Inferences about Portfolio and Stochastic Discount Factor Mean Variance Frontiers*

Francisco Peñaranda

UPF, Ramón Trias Fargas 25-27, 08005 Barcelona, Spain.

<francisco.penaranda@upf.edu>

Enrique Sentana

CEMFI, Casado del Alisal 5, E-28014 Madrid, Spain.

<sentana@cemfi.es>

October 2011

Abstract

We propose GMM procedures that consistently estimate the mean-variance frontiers for returns and SDFs and the weights of the portfolios that belong to them, and derive analytically and computationally simple joint confidence regions. We discuss efficiency gains obtained by exploiting asset pricing, tangency or spanning restrictions, and study the associated overidentification tests. We systematically exploit the duality of return and SDF frontiers so that our estimators, confidence regions and tests apply to both of them. We also analyse in detail the situation in which a researcher only has data on excess returns.

Keywords: Asset Pricing, GMM, Mean-Variance Frontiers, Tangency, Spanning.

JEL: G11, G12, C12, C13.

*This is a substantially revised and extended version of Penaranda and Sentana (2004). We would like to thank audiences at Indiana University for very useful comments and suggestions. Of course, the usual caveat applies. Financial support from the Spanish Ministry of Science and Innovation through grants ECO 2008-03066 (Peñaranda) and ECO 2008-00280 (Sentana) is gratefully acknowledged. Peñaranda also aknowledges the financial support of the Barcelona GSE and the Government of Catalonia.

1 Introduction

Mean-variance portfolio analysis is widely regarded as the cornerstone of modern investment theory. Despite its simplicity, and the fact that almost six decades have elapsed since Markowitz published his seminal work on the theory of portfolio allocation under uncertainty (Markowitz (1952)), it remains the most widely used asset allocation method. There are several reasons for its popularity. First, it provides a very intuitive assessment of the relative merits of alternative portfolios, as their risk and expected return characteristics can be compared in a two-dimensional graph. Second, return mean-variance frontiers (RMVF) are spanned by only two funds, a property that simplifies their calculation and interpretation. Finally, mean-variance analysis is fully compatible with expected utility maximisation if we assume Gaussian or elliptical distributions for asset returns (see e.g. Chamberlain (1983a), Owen and Rabinovitch (1983) and Berk (1997)), or if the mutual fund separation conditions hold (see Ross (1978)).

In turn, the stochastic discount factor mean-variance frontier (SMVF) introduced by Hansen and Jagannathan (1991) represented a major breakthrough in the way financial economists look at data on asset returns to discern which asset pricing theories are not empirically falsified. One of the main advantages of this frontier, though, is that one can compute it without a specification of the preferences of the representative agent because its focus are the mean-variance constraints that financial markets data imposes on asset pricing models (see e.g. Campbell, Lo and MacKinlay (1997) or Cochrane (2001) for advanced textbook treatments). Somewhat remarkably, it turns out that both frontiers are intimately related, as they effectively summarise the information contained in the first and second moments of asset payoffs [1]

However, those moments are unknown, and there is a substantial body of literature proposing different ways of reducing the sampling uncertainty involved in the estimation of mean variance frontiers, or at least bringing it to the forefront. In fact, the sampling uncertainty in estimated expected returns is so large that several authors have forcefully raised some doubts about the usual practice of computing mean-variance frontiers by simply replacing expected returns, variances and covariances by their sampling counterparts.

Simultaneously, there is also a large literature which focuses on statistical tests of the restrictions on mean-variance frontiers implied by various plausible assumptions, such as asset pricing models, or the tangency and spanning constraints that arise in mutual

[1] In line with most of the literature, in this paper we do not consider SDF frontiers that impose positivity of the SDF. See Hansen and Jagannathan (1991) for details.

fund performance evaluation (see De Roon and Nijman (2001) for a recent survey) or the assessment of portfolio diversification gains. Despite the fact that originally these two literatures were intertwined (see for example Jobson and Korkie (1980, 1982)), more recently the testing and estimation literatures have often been disconnected.

In this context, the contribution of this paper is threefold:

1. We propose GMM-based procedures that allow us to consistently estimate the frontiers and the weights of the portfolios that belong to them, as well as to derive joint confidence regions that provide analytically tractable and computationally simple alternatives to the Monte Carlo methods considered by Jorion (1992) and Michaud (1998) among others.

2. We explain how to achieve efficiency gains in estimating those frontiers by exploiting theoretically motivated restrictions, such as those derived from asset pricing models or other commonly used assumptions like tangency or spanning.

3. We exploit the integration of estimation and testing implicit in GMM, and study the associated overidentification tests, which can be formally understood as parametric tests of the null hypothesis that the additional restrictions are satisfied.

In addition, we follow Peñaranda and Sentana (2010a,b) in providing a unifiying approach that applies at three different levels:

a. We exploit the duality of the RMVF and SMVF so that our estimators, confidence regions and tests are not necessarily tied down to the specific properties of either frontier.

b. We compare our proposed tests to the extant tests, and show that they are all asymptotically equivalent under the null and compatible Pitman sequences of local alternatives, despite the fact that in some cases the number of parameters and moment conditions can be different.

c. We show that by using single-step GMM procedures such as the Continuously Updated (CU) version in Hansen, Heaton and Yaron (1996), we can make all the different overidentification tests numerically identical.

It is important to emphasise that all our results are obtained under fairly weak assumptions on the distribution of asset returns. In particular, in no way do they require that asset returns are independent or identically distributed as Gaussian or elliptical random vectors.

We complete our theoretical analysis by considering three special cases in which the RMVF and/or the SMVF take special forms. In particular, we analyse in detail the situation in which a researcher only has data on excess returns. We also discuss briefly in the appendix the less realistic contexts in which a constant payoff is included among the original vector of asset returns, or the expected returns of all the assets are identical.

The rest of the paper is organised as follows. In section 2 we derive the necessary theoretical background. Then, we consider frontiers for zero-cost portfolios in section 3, and devote section 4 to the case of gross returns. Finally, we summarise our conclusions in section 5. Proofs of propositions and auxiliary results are gathered in the appendix.

2 Theoretical background

In this section, we first describe the representing portfolios introduced by Chamberlain and Rothschild (1983), which we then use to characterise the RMVF and SMVF. We focus most of our discussion on the case of gross returns in the absence of a safe asset, although at the end we briefly consider the special situation in which all available assets are arbitrage (zero-cost) portfolios, which is often encountered in practice in working with excess returns.

2.1 Cost and Mean Representing Portfolios

Consider an economy with a finite number N of risky assets whose random payoffs $\mathbf{x} = (x_1, \ldots, x_N)$ are defined on an underlying probability space. Importantly, these assets can be either primitive, like stocks and bonds, or mutual funds managed according to some specific active portfolio strategy (see the discussion in chapter 8 of Cochrane (2001)). Let $E(\mathbf{x})$ and $E(\mathbf{x}\mathbf{x}')$ denote the first and second uncentred moments of those payoffs, which we assume are bounded. Thus, $x_i \in L^2$ $(i = 1, \ldots, N)$, which is the collection of all random variables defined on the underlying probability space with bounded second moments. We also assume that the covariance matrix of the N asset payoffs, $V(\mathbf{x})$, has full rank, which implies that none of the original assets is either riskless or redundant, and that it is not possible to generate a riskless portfolio from $\mathbf{x}$ other than the trivial one. Finally, we assume that the cost of the assets, $C(\mathbf{x})$, is not proportional to the vector of expected payoffs $E(\mathbf{x})$.[2]

Let $p = w_1x_1 + \ldots + w_Nx_N = \mathbf{w}'\mathbf{x}$ denote the payoffs to a portfolio of the N primitive assets with weights given by the vector $\mathbf{w} = (w_1, w_2, \ldots, w_N)'$. There are at least three characteristics of portfolios in which investors are interested: their cost, the expected value of their payoffs, and their variance, which will be given by $C(p) = \mathbf{w}'C(\mathbf{x})$, $E(p) = \mathbf{w}'E(\mathbf{x})$ and $V(p) = \mathbf{w}'V(\mathbf{x})\mathbf{w}$, respectively. Let $\mathcal{P}$ be the set of the payoffs from all possible portfolios of the N original assets, which is given by the linear span of $\mathbf{x}$, $\langle\mathbf{x}\rangle$. Within this set, two subsets deserve special attention: the set of all unit cost portfolios $\mathcal{P}(1) = \{p \in \mathcal{P} : C(p) = 1\}$, whose payoffs can be directly understood as returns per unit invested; and also the set of all zero cost, or arbitrage portfolios $\mathcal{P}(0) = \{p \in \mathcal{P} : C(p) = 0\}$. In this sense, note that any non-arbitrage portfolio can be transformed into a unit-cost portfolio by simply scaling its weights by its cost. For example, we can define $R_1 = x_1/C(x_1)$ as the gross return on the first asset provided that $C(x_1) \neq 0$. Similarly, if we partition $\mathbf{x}$ as $(x_1, \mathbf{x}_{-1})$, where $\mathbf{x}_{-1}$ is of dimension $n = N - 1$, it is clear that $\mathcal{P}(0)$ coincides with the linear span of $\mathbf{r}$, $\langle\mathbf{r}\rangle$, where $\mathbf{r} = \mathbf{x}_{-1} - R_1C(\mathbf{x}_{-1})$ can be understood as the vector of payoffs on the last n risky assets in excess of the first one. In empirical work, excess returns are often computed by subtracting from gross returns a purportedly riskless asset, but the representation of $\mathcal{P}(0)$ as the linear span of $\mathbf{r}$ remains valid irrespective of the ordering of the original assets as long as $C(x_1) \neq 0$. More generally, we can express any arbitrary $p = \mathbf{w}'\mathbf{x}$ as

$$p = w_1R_1 + \mathbf{w}_r'\mathbf{r}, \tag{1}$$

which means that we can effectively assume that $\mathbf{x}$ is given by $(R_1, \mathbf{r})$ without loss of generality. Thus, we can equate the weight on R_1, w_1, with the portfolio cost, $C(p)$, because $C(\mathbf{x})$ will coincide with the first column of the identity matrix of order N, $\mathbf{e}_1$.

Since $\mathcal{P}$ is a closed linear subspace of L^2, it is also a Hilbert space under the mean square inner product, $E(xy)$, and the associated mean square norm $\sqrt{E(x^2)}$, where $x, y \in L^2$. Such a topology allows us to define the least squares projection of any $y \in L^2$ onto $\mathcal{P}$ as:

$$E(y\mathbf{x}')E^{-1}(\mathbf{x}\mathbf{x}')\mathbf{x}, \tag{2}$$

which is the element of $\mathcal{P}$ that is closest to y in the mean square norm. In this context, we can formally understand $C(.)$ and $E(.)$ as linear functionals that map the elements of $\mathcal{P}$ onto the real line. The expected value functional is always continuous on L^2, while our

[2] Given their limited empirical relevance, we postpone to the appendix the discussion of the special cases in which either there is a constant payoff or expected payoffs and costs are proportional.

full rank assumption on $V(\mathbf{x})$ implies that $E(\mathbf{x}\mathbf{x}')$ has full rank too, which is tantamount to the law of one price, and consequently that the cost functional is also continuous on $\mathcal{P}$. The Riesz representation theorem then implies that there exist two unique elements of $\mathcal{P}$ that represent these functionals over $\mathcal{P}$ (see Chamberlain and Rothschild (1983)). In particular, the uncentred cost and mean representing portfolios, p^* and p°, respectively, will be such that:

$$C(p) = E(p^*p) \quad \text{and} \quad E(p) = E(p^\circ p) \quad \forall p \in \mathcal{P}.$$

It is then straightforward to show that

$$\begin{aligned} p^* &= C(\mathbf{x}')[E(\mathbf{x}\mathbf{x}')]^{-1}\mathbf{x}, \\ p^\circ &= E(\mathbf{x}')[E(\mathbf{x}\mathbf{x}')]^{-1}\mathbf{x}. \end{aligned} \tag{3}$$

If $\mathcal{P}$ included a constant unit payoff, then p° would coincide with it. But even though it does not, it follows from (2) and (3) that p^0 is the projection of 1 onto $\mathcal{P}$, which in financial markets parlance simply means that the mean representing portfolio is the portfolio that "mimics" the safe asset with the minimum "tracking error". To give a similar economic interpretation to p^*, it is convenient to recall that a stochastic discount factor, m say, is any scalar random variable defined on the same underlying probability space which prices assets in terms of their expected cross product with it. We can again use (2) to interpret p^* as the projection of any m onto $\mathcal{P}$, i.e. as the portfolio that best mimics stochastic discount factors. In addition, since $C(1) = E(1 \cdot m) = c$ say, the expected value of m defines the shadow price of a unit payoff.

Chamberlain and Rothschild (1983) show that an alternative valid topology on $\mathcal{P}$ can be defined with covariance as inner product and standard deviation as norm when there is not a constant payoff in $\mathcal{P}$. Consequently, we could also represent the two functionals by means of two alternative centred representing portfolios, þ^* and þ° in $\mathcal{P}$, such that

$$C(p) = Cov(\text{þ}^*, p) \quad \text{and} \quad E(p) = Cov(\text{þ}^\circ, p) \quad \forall p \in \mathcal{P}.$$

Not surprisingly,

$$\begin{aligned} \text{þ}^* &= C(\mathbf{x}')[V(\mathbf{x})]^{-1}\mathbf{x} = p^* + [1 - E(p^\circ)]^{-1}C(p^\circ)p^\circ, \\ \text{þ}^\circ &= E(\mathbf{x}')[V(\mathbf{x})]^{-1}\mathbf{x} = [1 - E(p^\circ)]^{-1}p^\circ. \end{aligned} \tag{4}$$

2.2 SDF and Portfolio Mean-Variance Frontiers for Gross Returns

The elements of the SMVF, or Hansen and Jagannathan (1991) frontier, are those admissible SDF's with the lowest variance for a given mean. Therefore, they formally solve the programme

$$\min_{m\in L^2} V(m) \quad s.t. \quad E(m) = c \in \mathbb{R}^+, \quad E(m\mathbf{x}) = C(\mathbf{x}). \tag{5}$$

If there were a constant payoff then the reciprocal of its gross return will pin down the single element of the SMVF frontier. But even though no safe asset exists, we can find the elements of the SMVF by solving the above programme for any notional safe return $c^{-1} \geq 0$. As shown by Hansen and Jagannathan (1991), the solution to (5) can be expressed as

$$m_1^{MV}(c) = p^* + \alpha(c)(1 - p^\circ) = \alpha(c) + þ^* - cþ^\circ, \tag{6}$$

where

$$\alpha(c) = \frac{c - E(p^*)}{1 - E(p^\circ)} = c\,[1 + E(þ^\circ)] - E(þ^*), \tag{7}$$

which shows that they are all shifted portfolios spanned by p^* and $1 - p^\circ$ (or $þ^\circ$ and $(1{-}þ^\circ)$) alone.

It is then easy to show that

$$Var[m_1^{MV}(c)] = E(p^{*2}) + \frac{[c - E(p^*p^\circ)]^2}{1 - E(p^{\circ 2})} - c^2 = V(þ^*) - 2cov(þ^\circ, þ^*)c + V(þ^\circ)c^2. \tag{8}$$

Hereinafter, we shall usually refer to the function $\{c, Var[m_1^{MV}(c)]\}$ as the SMVF for gross returns, which is a parabola in mean-variance space, although sometimes we will consider instead the function $\{c, \sqrt{Var[m_1^{MV}(c)]}\}$, which is a hyperbola in mean-standard deviation space.

As we mentioned in the introduction, in this paper we are interested in estimating these curves, and in making inferences about them. Importantly, note that the coefficients of the parabola (8) depend exclusively on the uncentred second moments of the uncentred representing portfolios, or on the variances and covariances of the centred representing portofios. In fact, the frontier is linear in those parameters, which will simplify inferences. In addition, we are also interested in making inferences about the weights of the different assets in $m_1^{MV}(c)$, which are also simple affine functions of their weights on the representing portfolios in view of (6).

In turn, the elements of the RMVF, or Markowitz (1952) frontier, are those unit-cost portfolios that have the lowest variance for a given mean. Therefore, they formally solve the programme

$$\min_{p\in\mathcal{P}(1)} V(p) \quad s.t. \quad E(p) = \nu \in \mathbb{R}. \tag{9}$$

As shown by Chamberlain and Rothschild (1983), the solution to (9) can be written as

$$R^{MV}(\nu) = \left[\frac{E(p^{\circ 2}) - E(p^*p^\circ)\nu}{E(p^{*2})E(p^{\circ 2}) - E^2(p^*p^\circ)}\right] p^* + \left[\frac{E(p^{*2})\nu - E(p^*p^\circ)}{E(p^{*2})E(p^{\circ 2}) - E^2(p^*p^\circ)}\right] p^\circ$$
$$= \left[\frac{V(\text{þ}^\circ) - cov(\text{þ}^*, \text{þ}^*)\nu}{V(\text{þ}^*)V(\text{þ}^\circ) - cov^2(\text{þ}^*, \text{þ}^\circ)}\right] \text{þ}^* + \left[\frac{V(\text{þ}^*)\nu - cov(\text{þ}^*, \text{þ}^\circ)}{V(\text{þ}^*)V(\text{þ}^\circ) - cov^2(\text{þ}^*, \text{þ}^\circ)}\right] \text{þ}^\circ. \qquad (10)$$

which is also spanned by p^* and p°, or þ^* and þ°.

Hence, we will have that

$$Var[R^{MV}(\nu)] = \frac{1}{E(p^{*2})E(p^{\circ 2}) - E^2(p^*p^\circ)}\left[E(p^{*2})\nu^2 - 2E(p^*p^\circ)\nu + E(p^{\circ 2})\right] - \nu^2$$
$$= \frac{1}{V(\text{þ}^*)V(\text{þ}^\circ) - cov^2(\text{þ}^*, \text{þ}^\circ)}\left[V(\text{þ}^*)\nu^2 - 2cov(\text{þ}^*, \text{þ}^\circ)\nu + V(\text{þ}^\circ)\right]. \qquad (11)$$

We shall refer to the function $\{\nu, Var[R^{MV}(\nu)]\}$ as the RMVF, which is another parabola in mean-variance space, although sometimes will consider the related function $\{\nu, \sqrt{Var[R^{MV}(\nu)]}\}$, which is a hyperbola in mean-standard devation space.

Importantly, note once again that the coefficients of the parabola (11) depend exclusively on the uncentred second moments of the uncentred representing portfolios, or on the variances and covariances of the centred representing portfofios. Similarly, the weights of the different assets in $R^{MV}(\nu)$ are also simple linear functions of their weights on the representing portfolios in view of (10).

Finally, it is worth mentioning that those linear combinations of the centred representing portfolios $\tau^*\text{þ}^* + \tau^\circ\text{þ}^\circ$ such that $\tau^* = 1$ reflect the risky component of some $m_1^{MV}(c)$, whereas those others in which $\tau^*V(\text{þ}^*) + \tau^\circ cov(\text{þ}^*,\text{þ}^\circ) = 1$ correspond to some $R^{MV}(\nu)$. This fact corroborates the well-known duality between the RMVF and SMVF highlighted by Hansen and Jagannathan (1991). Specifically, any $m_1^{MV}(c)$ such that $V(\text{þ}^*) - cov(\text{þ}^*,\text{þ}^\circ)c \neq 0$ can be translated into a $R^{MV}(\nu)$, and any $R^{MV}(\nu)$ such that $V(^\circ) - cov(^*,^*)\nu \neq 0$ can be translated into a $m_1^{MV}(c)$ (see Appendix C of Peñaranda and Sentana (2011) for further details on those exceptions).

2.3 SDF and Portfolio Mean-Variance Frontiers for excess returns

Let us now study the special situation in which all primitive assets are arbitrage portfolios, so that the relevant payoff space is $\mathcal{P}(0)$, a closed linear subspace of $\mathcal{P}$ that inherits its Hilbert space structure. Although the centred and uncentred cost representing portfolios in $\mathcal{P}(0)$ are both 0, the respective mean representing portfolios will be

$$\begin{aligned} a^\circ &= E(\mathbf{r}')[E(\mathbf{rr}')]^{-1}\mathbf{r}, \\ \eth^\circ &= E(\mathbf{r}')[V(\mathbf{r})]^{-1}\mathbf{r} = [1-E(a^\circ)]^{-1}a^\circ. \end{aligned}$$

Interestingly, a° coincides with the residual from the projection of p° onto p^*, which is not surprising given that $\mathcal{P}(0)$ is the orthogonal complement of $\langle p^* \rangle$ on $\mathcal{P}$. Note also that the weights of the centred mean representing portfolio coincide with the most frequent textbook presentation of the weights of the tangency portfolio.

In this context, we can show that the elements SMVF based on arbitrage portfolios only will be given by

$$m_0^{MV}(c) = \frac{c}{1-E(a^\circ)}(1-a^\circ) = c\{1-[\eth^\circ - E(\eth^\circ)]\}, \tag{12}$$

so that they are spanned by a single "fund". Not surprisingly, their variance will be

$$Var[m_0^{MV}(c)] = \sigma_0^2(c) = c^2\frac{E(a^\circ)}{1-E(a^\circ)} = c^2 E(\eth^\circ), \tag{13}$$

which is a perfect square in c that depends on a single parameter given by the second moment of the uncentred mean representing portfolio, or the variance of the centred one. In addition, the weights of the different assets in $m_0^{MV}(c)$ are proportional to the weights of the mean representing portfolios in view of (12).

We shall refer to the function $[c, \sigma_0^2(c)]$ as the SMVF for arbitrage portfolios, which is a parabola in mean-variance space, although sometimes we will consider instead the function $[c, \sigma_0(c)]$, which is a half line starting from the origin in mean-standard deviation space.

On the other hand, although the RMVF cannot be defined over $\mathcal{P}(0)$, we can construct the arbitrage (i.e. zero-cost) mean variance frontier (AMVF), whose elements will be of the form

$$r^{MV}(\mu) = \mu\frac{1}{E(a^\circ)}a^\circ = \mu\frac{1}{E(\eth^\circ)}\eth^\circ. \tag{14}$$

Not surprisingly,

$$V[r^{MV}(\mu)] = \frac{1-E(a^\circ)}{E(a^\circ)}\mu^2 = \frac{1}{E(\eth^\circ)}\mu^2, \tag{15}$$

which confirms the duality between the SDF and portfolio frontiers because the maximum (squared) Sharpe ratio in $\mathcal{P}(0)$, which is given by $\mu^2/V[r^{MV}(\mu)] = E(\eth^\circ)$, is equal to $\sigma_0^2(c)/c^2$ (see Hansen and Jagannathan (1991)). As is well known, the function $\{\mu, V[r^{MV}(\mu)]\}$ will be a parabola tangent to the origin in mean-variance space, while the related function $\{\mu, \sqrt{V[r^{MV}(\mu)]}\}$ will be a reflected straight line in mean-standard deviation space. Finally, the weights of the different assets in $r^{MV}(\mu)$ are proportional to their weights on the mean representing portfolios in view of (14).

3 Estimation of mean-variance frontiers with excess returns

For pedagogical reasons, we first study the estimation of mean-variance frontiers for SDF's and portfolios when the only payoffs available to the researcher are excess returns, leaving the discussion of more general payoffs for section 4. We initially consider the situation in which there are no a priori restrictions on the frontiers, and then discuss how to increase efficiency by incorporating either spanning or asset pricing restrictions. To avoid duplicities, most of our theoretical discussions will focus on SMVFs, but the graphs will show AMVFs.

3.1 Unrestricted estimation

Expression (13) implies that the SMVF is linear in the single unknown parameter

$$\theta = E(\eth^\circ) = [1 - E(a^\circ)]^{-1} E(a^\circ),$$

which we interpreted before as the maximum (square) Sharpe ratio attainable, whereas expression (15) implies that the AMVF frontier is linear in its reciprocal, $\eta = \theta^{-1}$. Therefore, given a vector of n excess returns $\mathbf{r}$, we can estimate both frontiers from the following exactly identified system of $n+1$ moment conditions:

$$E\begin{bmatrix} \mathbf{r}\mathbf{r}'\boldsymbol{\phi}^\circ - \mathbf{r} \\ \mathbf{r}'\boldsymbol{\phi}^\circ - \mu^\circ \end{bmatrix} = \mathbf{0}, \tag{16}$$

where $\mu^\circ = \theta/(1+\theta) = 1/(1+\eta)$ identifies $E(a^\circ) = E(a^{\circ 2})$, and $\boldsymbol{\phi}^\circ$ the portfolio weights of this uncentred mean representing portfolio. Alternatively, we could work with the analogous $n+1$ moment conditions for the centred representing portfolios

$$E\begin{bmatrix} \mathbf{r}(\mathbf{r}'\boldsymbol{\varphi}^\circ - \theta) - \mathbf{r} \\ \mathbf{r}'\boldsymbol{\varphi}^\circ - \theta \end{bmatrix} = \mathbf{0}, \tag{17}$$

where θ identifies $E(\eth^\circ) = Var(\eth^\circ)$ and $\boldsymbol{\varphi}^\circ$ the corresponding portfolio weights.[3] Systems (16) and (17) are equivalent in the sense that they provide the same numerical

[3] Alternatively, we could consider the $2n$ moment conditions

$$E\begin{bmatrix} \mathbf{r}(\mathbf{r}-\boldsymbol{\mu})'\boldsymbol{\varphi}^\circ - \mathbf{r} \\ \mathbf{r}-\boldsymbol{\mu} \end{bmatrix} = \mathbf{0},$$

estimate of θ and the same standard error through the Delta method.

Under standard regularity conditions (see e.g. Newey and McFadden (1994)), the resulting GMM estimator of θ will converge in probability to its true value, and the same applies to the weights of the mean representing portfolios. Therefore, the GMM estimators of $Var[m_0^{MV}(c)]$ and $V[r^{MV}(\mu)]$ will also converge in probability to their population counterparts for fixed c and μ. Further, we can easily show that GMM estimators of the entire SMVF and AMVF will converge uniformly to their population analogues over any finite range. Specifically, in the case of the SMVF frontier we will have that

$$\sup_{c\in[\underline{c},\bar{c}]}\left|c^2\hat{\theta}-c^2\theta\right|=\left(\sup_{c\in[\underline{c},\bar{c}]}c^2\right)\left|\hat{\theta}-\theta\right|=o_p\left(1\right).$$

Despite the uniform consistency, though, the SMVF and AMVF frontiers are subject to substantial sample variability (see Jobson and Korkie (1980) for some early results and references). To emphasise the importance of sampling uncertainty in this context, we have conducted the following simulation experiment. We have assumed that investors have access to six arbitrage portfolios, whose excess returns roughly replicate the distribution of the 6 Fama and French portfolios formed on size and book-to-market (see Appendix A for further details on the experimental design we have used). Then we simulate forty years of monthly data many times, and compute the mean-variance frontiers. Figure 1a presents part of the ensemble of AMVFs thus obtained, while Figure 1b includes the weights of the first asset on the estimated optimal portfolio.

<Figures 1a and 1b>

As can be seen from these pictures, if one did not take into account sampling uncertainty, one would form very different optimal portfolios depending on the sample, and would reach rather different conclusions about the available risk-return trade-offs. In this respect, there is a clear tendency to reach overly optimistic conclusion about the mean-variance trade-offs that investors really face. This result is confirmed by Figure 2, which presents a kernel density estimate of the sample maximum Sharpe ratio.

<Figure 2>

---

where $\boldsymbol{\mu}$ is the vector of risk premia. This system identifies $\boldsymbol{\varphi}^{\circ}$ and $\boldsymbol{\mu}$ directly, but θ indirectly as $\boldsymbol{\mu}'\boldsymbol{\varphi}^{\circ}$. Although single-step GMM procedures will yield numerically identical results, we prefer to use (17) because it involves fewer parameters for $n>1$.

Given the importance of sampling uncertainty, and its potentially misleading nature, it is convenient to provide joint confidence intervals for the different quantities involved. Simultaneous confidence intervals are important because investors or researchers will not be necessarily concerned with just one point on the frontiers.

Let us start with $\sigma^2(c)$ for K values of c, $(c_1, c_2, ..., c_K)$ say. To obtain the relevant joint confidence region, we need the asymptotic distribution of the estimators of the variance of the SMVF elements corresponding to those values (see Mittnik and Zadrozny (1993) for a related approach in the context of impulse response functions). Given expression (13), it is clear that

$$\sqrt{T}\begin{bmatrix} \hat{\sigma}^2(c_1) - \sigma^2(c_1) \\ \vdots \\ \hat{\sigma}^2(c_K) - \sigma^2(c_K) \end{bmatrix} \xrightarrow{d} N\left[0, \upsilon \begin{pmatrix} c_1^2 \\ \vdots \\ c_K^2 \end{pmatrix} \begin{pmatrix} c_1^2 & \dots & c_K^2 \end{pmatrix}\right],$$

where υ denotes the asymptotic variance of $\sqrt{T}(\hat{\theta} - \theta)$, which is the only parameter estimator involved. The singularity of this asymptotic distribution implies that regardless of how big K is, the joint confidence region for $[\sigma^2(c_1), \sigma^2(c_2), \dots \sigma^2(c_K)]$ will be the unidimensional line segment

$$\left[\frac{\sum_{k=1}^{K} c_j^2[\hat{\sigma}^2(c_k) - \sigma^2(c_k)]}{\sum_{k=1}^{K} c_j^4}\right]^2 \le \frac{\upsilon}{T} Q(1-\alpha; 1), \tag{18}$$

where $Q(1-\alpha;1)$ is the $1-\alpha$ quantile of a chi-square random variable with 1 degree of freedom. Figure 3 presents this interval for two different values of c, $(= .9, .95)$. As expected, it is centred around the estimated values of $\sigma^2(.9)$ and $\sigma^2(.95)$ (in green), and in this case it contains the corresponding true values (in black).

<Figure 3>

Given that we cannot represent this confidence interval in more than three dimensions, we find it convenient to map it onto the space $[c, \sigma^2(c)]$ by plotting in that space all the points that belong to the joint confidence interval (18) for a chosen value of α. Nevertheless, it is important to remark that such a plot fails to capture the strong dependence that exists between the different values of c. It turns out that it is rather easy to characterise the envelope of this set, as the only operation required is the projection of the joint interval (18) on to the $\sigma^2(c_1), \sigma^2(c_2), \ldots$ and $\sigma^2(c_K)$ axes. Straightforward algebra shows that the projection over $\sigma^2(c_1)$ will be given by

$$[\hat{\sigma}^2(c_1) - \sigma^2(c_1)]^2 \leq \frac{\upsilon c_1^4}{T} Q(1-\alpha;1)$$

irrespective of the values of $(c_2, ..., c_K)$. Therefore, the limits of the mapping of the joint interval (18) onto $[c, \sigma^2(c)]$ space will be

$$c^2 \left[\hat{\theta} \pm \sqrt{\frac{\upsilon}{T} Q(1-\alpha;1)} \right]. \tag{19}$$

Not surprisingly, the width of these limits increases with c, υ and α, and decreases with the sample size.

Interestingly, the region generated by (19) turns out to the right coverage in $[c, \sigma^2(c)]$ space too because

$$\lim_{T\to\infty} P \left\{ \begin{array}{c} c^2 \left[\hat{\theta} + \sqrt{\frac{\upsilon}{T} Q(1-\alpha;1)} \right] \leq c^2 E(\eth^\circ) \leq c^2 \left[\hat{\theta} + \sqrt{\frac{\upsilon}{T} Q(1-\alpha;1)} \right] \\ \forall c \in [\underline{c}, \bar{c}] \end{array} \right\} = 1 - \alpha.$$

This result trivially follows from the fact that the upper and lower bounds of (19) correspond to the maximum and minimum values of $\sigma^2(c) = c^2\theta$ that can be achieved within the $100(1-\alpha)\%$ asymptotic confidence interval for θ [4]

$$T \frac{(\hat{\theta} - \theta)^2}{\upsilon} \leq Q(1-\alpha;1).$$

In fact, (19) also coincides with the pointwise confidence bands for $\sigma^2(c)$ due to the presence of a single estimated parameter.

We can repeat the same analysis in $[c, \sigma(c)]$ space by simply taking the square root of the limits (19).[5]

[4] Not surprisingly, if we replace $\hat{\sigma}^2(c_k)$ and $\sigma^2(c_k)$ by $c^2\hat{\theta}$ and $c^2\theta$, respectively, then (18) collapses to the confidence interval for $\hat{\theta}$. For that reason, asymptotically equivalent regions for $\sigma^2(c)$ could be computed from confidence intervals for θ obtained by inverting either a distance metric test or a Lagrange multiplier test of the null hypothesis $H_0 : \theta = \bar{\theta}$. For example, in the former case we would rely on the values of $\bar{\theta}$ such that the GMM J statistic is below $Q(1-\alpha;1)$. In fact, the results in Newey and West (1987) imply that if we estimated (17) by two step GMM, then the confidence regions for θ obtained from the trinity of classical tests will be numerically identical for a given estimator of the long-run covariance matrix, and the same will be true for the resulting confidence regions. In general, though, the Distance metric-, Lagrange multiplier- and Wald-based confidence intervals will be different in finite samples.

[5] This method will break down in finite samples if the lower bound of (19) is negative. In contrast, if we use the Delta method, which implies that

$$\sqrt{T}\left(\sqrt{c^2\hat{\theta}} - \sqrt{c^2\theta}\right) \xrightarrow{d} N\left[0, \left(\frac{1}{2\sqrt{c^2\theta}}\right)^2 c^4 \upsilon\right],$$

Similarly, given that the variance of the elements of the AMVF frontier are linear in η, which is the reciprocal of the maximum (square) Sharpe ratio attainable, we can repeat the same exercise in terms of this parameter to study this other frontier instead. In this sense, Figure 4a presents a graph of the relevant limits (in red) using the same design for excess returns considered in the previous figures, together with the estimated (in green) and true (in black) AMVFs.

<Figures 4a and 4b>

Given that the weight of asset i on the SMVF is given by $-c\varphi_i^\circ = -c(1+\theta)\phi_i^\circ$, where φ_i° and ϕ_i are the i^{th} elements of $\boldsymbol{\varphi}^\circ$ and $\boldsymbol{\phi}^\circ$, respectively, while the corresponding weight for the AMVF will be $\mu\eta\varphi_i^\circ = \mu\eta(1+\eta)^{-1}\phi_i^\circ$, it is also straightforward to derive joint confidence regions for those weights for $(c_1, c_2, ..., c_K)$ or $(\mu_1, \mu_2, ..., \mu_K)$. Moreover, we can also map those joint confidence intervals into the relevant space, as shown in Figure 4b for the optimal weight of the first asset in the AMVF.

Finally, it would be easy to consider joint confidence intervals for the weights of two or more assets.[6]

3.2 Efficient estimation imposing spanning restrictions: Mean-variance efficiency tests

Let us partition $\mathbf{r}$ into two sets of portfolios $\mathbf{r}_1$ and $\mathbf{r}_2$ of dimensions n_1 and n_2, respectively, with $n = n_1 + n_2$, so that $\mathbf{r}' = (\mathbf{r}_1', \mathbf{r}_2')$. Sometimes theoretical or empirical considerations may suggest that the addition of $\mathbf{r}_2$ should not improve the investment opportunity set of investors. Equivalently, we may believe that the inclusion of $\mathbf{r}_2$ would not tighten the bounds on admissible SDFs. In both cases cases, we say that $\mathbf{r}_1$ spans the mean-variance frontiers generated from $\mathbf{r}_1$ and $\mathbf{r}_2$.

---

the confidence bands for standard deviations will be

$$\sqrt{c^2\hat{\theta}} \pm \frac{1}{2\sqrt{c^2\hat{\theta}}}\sqrt{\frac{c^4 \upsilon}{T} Q(1-\alpha;1)}.$$

We can interpret these bands as a first-order approximation to

$$\sqrt{c^2\hat{\theta} \pm \sqrt{\frac{c^4 \upsilon}{T} Q(1-\alpha;1)}},$$

which converge to the same limit as $T \to \infty$.

[6] Confidence intervals for relative weights are even simpler, as they do not depend on the value of the indices c or μ (see Jobson and Korkie (1980) and Britten-Jones (1999)).

Under the null hypothesis that this is the indeed case, there will be only one pair of mean-variance frontiers. Under the alternative, there will be two: the frontiers generated from $\mathbf{r}_1$ alone, and the ones generated from $\mathbf{r}$, which will only touch at the origin. The procedures discussed in the previous subsection allow us to estimate those unrestricted frontiers. The purpose of this section is to explain how we can efficiently estimate the common frontiers under the null, which despite first appearances, will generally be different from the unrestricted frontiers estimated on the basis of $\mathbf{r}_1$ alone.

Efficient estimation of curves is a somewhat unusual concept in econometrics. Given two alternative estimators of a given curve, we could say that one is more efficient than the other if loosely speaking some efficiency gains accrue in estimating any arbitrary vector of points on the curve. Alternatively, we could say that one curve estimator is more efficient than another curve estimator if their efficiency ranking is preserved for any linear functional (see Arellano, Hansen and Sentana (2011) for further discussion of these concepts). Since in our case the entire SMVF and AMVF depend exclusively on a single parameter estimator, both concepts trivially imply that more efficient "curve estimators" of the frontiers will be obtained by using more efficient estimators of θ.

In the GMM context described at the beginning of the previous section, the imposition of the null hypothesis of spanning on the weights of the uncentred moment conditions (16) gives rise to the overidentified system

$$E\left[\begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \\ 1 \end{pmatrix} \mathbf{r}_1' \boldsymbol{\phi}_1^\circ - \begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \\ \mu^\circ \end{pmatrix}\right] = \mathbf{0}. \tag{20}$$

The optimal GMM estimator of $\theta = \mu^\circ/(1-\mu^\circ)$ obtained from (20) will generally be more efficient than the corresponding estimator obtain from the unrestricted system (16) as long as the equality restriction $\boldsymbol{\phi}_2^\circ = \mathbf{0}$ holds (see Property 10.6 in Gourieroux and Monfort (1995)). Moreover, the results in Breusch et al (1999) imply that this estimator will also be generally more efficient than the one obtained from the just identified $n_1 + 1$ moment conditions

$$E\left[\begin{pmatrix} \mathbf{r}_1 \\ 1 \end{pmatrix} \mathbf{r}_1' \boldsymbol{\phi}_1^\circ - \begin{pmatrix} \mathbf{r}_1 \\ \mu^\circ \end{pmatrix}\right] = \mathbf{0}. \tag{21}$$

An exception arises in the following situation:

Proposition 1 *If $\mathbf{r}_t$ is an i.i.d. elliptical random vector with bounded fourth moments, and the null hypothesis of spanning is true, then:*
a) The asymptotic variance of the optimal GMM estimator of μ° obtained from (20), which

imposes the spanning constraint $\boldsymbol{\phi}_2^\circ = 0$, will coincide with the asymptotic variances of both the estimator obtained from (21), which will be given by

$$T^{-1}\left(\sum_{t=1}^{T}\mathbf{r}_{1t}\right)'\left(\sum_{t=1}^{T}\mathbf{r}_{1t}\mathbf{r}_{1t}'\right)^{-1}\left(\sum_{t=1}^{T}\mathbf{r}_{1t}\right),$$

and the estimator obtained from (16), which will be given by

$$T^{-1}\left(\sum_{t=1}^{T}\mathbf{r}_{t}\right)'\left(\sum_{t=1}^{T}\mathbf{r}_{t}\mathbf{r}_{t}'\right)^{-1}\left(\sum_{t=1}^{T}\mathbf{r}_{t}\right).$$

b) The asymptotic variance of the GMM estimator of $\boldsymbol{\phi}_1^\circ$ obtained from (16) will be larger (in the usual positive definite sense) than the asymptotic variance of the optimal GMM estimator based on (20), which in turn coincides with the asymptotic variance of

$$\left(\sum_{t=1}^{T}\mathbf{r}_{1t}\mathbf{r}_{1t}'\right)^{-1}\left(\sum_{t=1}^{T}\mathbf{r}_{1t}\right),$$

which is the GMM estimator obtained from (21).

This results extends Lemma 1 in Peñaranda and Sentana (2010b), who prove the asymptotic equivalence mentioned in part b. Trivially, part b) extends to the "estimators" of $\boldsymbol{\phi}_2^\circ$.

In order to gauge the efficiency gains in estimating the frontiers, we have repeated the simulation exercise described in the previous section, which satisfies by construction that the frontiers are a function of $\mathbf{r}_1$ only (see Appendix A for further details). The results for a specific simulation are reported in Figure 5. As expected, the unrestricted empirical frontier obtained from $\mathbf{r}$ (in blue) is always outside the frontier generated from $\mathbf{r}_1$ alone (in green) even though both frontiers are identical in the population. The differences between the frontier generated from $\mathbf{r}_1$ alone and from $\mathbf{r}$ imposing the spanning restriction $\boldsymbol{\phi}_2^\circ = \mathbf{0}$ (in red) are relatively small. Given the tendency of unrestricted frontiers to overestimate the true risk return trade-off, it is perhaps not surprising that the restricted frontiers are closer to the truth (in black).

<Figures 5a and 5b>

Similarly, Figures 6a and 6b illustrate the asymptotic efficiency gains obtained by imposing the spanning restriction $\boldsymbol{\phi}_2^\circ = \mathbf{0}$, which turn out to be relatively minor for the frontier itself, but rather large for the portfolio weights.

<Figures 6a and 6b>

Unfortunately, these efficiency gains come at a cost: if the spanning restrictions are wrong, then the frontier based on (20) will be inconsistently estimated. There is a huge literature on testing those restrictions, which usually comes under the heading of mean-variance efficiency test (see Sentana (2009) for a recent survey). The advantage of our GMM set-up is that we can readily use the overidentification test of the moment conditions (20) to test for spanning, since it coincides with the distance metric test of the null hypothesis $H_0 : \boldsymbol{\phi}_2^\circ = \mathbf{0}$.[7] As is well known, this test will have a limiting chi-square distribution with n_2 degrees of freedom under the null.

An analogous test could be based on the moment conditions that define the centred mean representing portfolio. In this respect, note that since $[E(\mathbf{r}\mathbf{r}')]^{-1}E(\mathbf{r}) = \{1 + E(\mathbf{r}')[V(\mathbf{r})]^{-1}E(\mathbf{r})\}^{-1}[V(\mathbf{r})]^{-1}E(\mathbf{r})$ by virtue of the Woodbury formula, $\boldsymbol{\phi}_2^\circ$ and $\boldsymbol{\varphi}_2^\circ$ will be proportional to each other, so that the null hypotheses are equivalent.

The most popular mean-variance efficiency tests by far, though, are the regression-based tests considered by Gibbons, Ross and Shanken (1989), and robustified against non-normality by MacKinlay and Richardson (1991). Specifically, their test would correspond to the overidentification test of the $n_2(n_1 + 1)$ moment conditions

$$E\left[\begin{pmatrix} 1 \\ \mathbf{r}_1 \end{pmatrix} \otimes (\mathbf{r}_2 - \mathbf{B}\mathbf{r}_1)\right] = \mathbf{0}. \tag{22}$$

Peñaranda and Sentana (2010b) show that all three approaches (namely, uncentred and centred representing portfolios and regression) are numerically equivalent when implemented by single-step methods such as CU-GMM. This fact also implies that the three approaches will be asymptotically equivalent when implemented by two-stage or iterated GMM, even though they will not be numerically equivalent in that case. For that reason, we shall not discuss (22) further, especially taking into account that those moment conditions cannot be used directly in the estimation of mean-variance frontiers.

[7] Alternatively, we could conduct a Wald test of the same null hypothesis from (17), or indeed a Lagrange multiplier test, all of which are asymptotically equivalent under the null and sequences of local alternatives. In this sense, it is interesting to note that the approach used by Britten-Jones (1999) to test the mean-variance efficiency of a given portfolio by looking at its weights can be easily cast in our GMM framework too, because the regression of a vector of ones onto the vector of excess returns gives the orthogonality conditions (20) that define the mean representing portfolio (see also Jobson and Korkie (1983) and Sentana (2009)).

3.3 Efficient estimation imposing a linear factor pricing model: Asset pricing tests

A closely related way to reduce the sampling uncertainty in the construction of mean-variance frontiers is to use an asset pricing model, which will impose some discipline on the estimators of expected returns (see Black and Litterman (1992), Pástor (2000), Pástor and Stambaugh (2000) or Tu and Zhou (2004) for related approaches). Although nothing prevents us from considering non-linear models, the standard approach in empirical finance is to model m as an affine transformation of some $k \leq n$ observable risk factors $\mathbf{f}$, even though this ignores that m must be positive with probability 1 to avoid arbitrage opportunities (see Hansen and Jagannathan (1991)). In this context, we can express the pricing equation as

$$E\left[(\lambda_0 - \boldsymbol{\lambda}'\mathbf{f})\,\mathbf{r}\right] = \mathbf{0} \tag{23}$$

for some real numbers $(\lambda_0, \boldsymbol{\lambda}')'$. As argued by Cochrane (2001) among others, we can in fact understand the spanning restrictions discussed in the previous section as imposing a linear factor pricing model in which the pricing factors $\mathbf{f}$ coincide with some excess returns $\mathbf{r}_1$, as in the CAPM or the Fama and French (1993) model. In general, though, $\mathbf{f}$ does not have to be a subset of $\mathbf{r}$.[8]

Although $\mathbf{r}$ only contains assets with 0 cost, which leaves the *scale* and *sign* of m undetermined, we would like our candidate SDF to price other assets with positive prices. Therefore, we require a *scale normalisation* to rule out the trivial solution $(\lambda_0, \boldsymbol{\lambda}')' = (0, \mathbf{0}')'$ (see Cochrane (2001, pp. 256-258)). For example, we could choose the popular asymmetric normalisations $\lambda_0 = 1$ or $E(m) = \lambda_0 - \boldsymbol{\lambda}'E(\mathbf{f}) = 1$.[9] For simplicity, we will follow the former normalisation in our exposition, although as shown by Kan and Robotti (2008) and Peñaranda and Sentana (2010b), the normalisation is generally inconsequential for single-step GMM methods. In this context, assuming that $\mathbf{f}$ and $\mathbf{r}$ do not share any common elements, we can add the pricing conditions (23) to the exactly identified moment

[8] In fact, it is possible to prove that the pricing constraint (23) can also be interpreted as spanning constraints, but this time with respect to the mimicking portfolios of $\mathbf{f}$, which are given by

$$\mathbf{r}'\boldsymbol{\phi}^{\circ} = \mathbf{r}'E^{-1}\left(\mathbf{r}\mathbf{r}'\right)E\left(\mathbf{r}\right) = \left[\mathbf{r}'E^{-1}\left(\mathbf{r}\mathbf{r}'\right)E\left(\mathbf{r}\mathbf{f}'\right)\right]\boldsymbol{\lambda}.$$

[9] Alternatively, we could choose the symmetric normalisation $\lambda_0^2 + \boldsymbol{\lambda}'\boldsymbol{\lambda} = 1$, together with a sign restriction on one of the nonzero coefficients. See Peñaranda and Sentana (2010b) for further details.

conditions (16), thereby obtaining.

$$E\begin{bmatrix} \mathbf{r}\,(1-\mathbf{f}'\boldsymbol{\lambda}) \\ \mathbf{r}\mathbf{r}'\boldsymbol{\phi}^{\circ}-\mathbf{r} \\ \mathbf{r}'\boldsymbol{\phi}^{\circ}-\mu^{\circ} \end{bmatrix}=\mathbf{0}, \tag{24}$$

where the unknown parameters are $(\boldsymbol{\lambda}', \boldsymbol{\phi}^{\circ\prime}, \mu^{\circ})$. In this way, we obtain more efficient estimators of the mean-variance frontiers that exploit the pricing equations (23).

As expected, the overidentifiying restriction test of the moment conditions (24) yields a valid asset pricing test, whose asymptotic distribution will be a chi-square with $n-k$ degrees of freedom. Peñaranda and Sentana (2010b) show that the uncentred and centred representing portfolios approaches are numerically equivalent to the corresponding regression approach when implemented by single-step methods such as CU-GMM. This fact also implies that the three approaches will be asymptotically equivalent when implemented by two-stage or iterated GMM, even though they will not be numerically equivalent in that case.

4 Estimation of mean-variance frontiers with returns

4.1 Unrestricted estimation

Let us now consider the alternative situation in which an empirical researcher has at her disposal at least one asset whose cost is different from 0. Expression (8) implies that the variance of the elements of the SMVF for gross returns will be given by

$$\begin{aligned}\sigma^2(c) &= Var[m^{MV}(c)] = \begin{pmatrix} 1 & -2c & c^2 \end{pmatrix}\begin{pmatrix} V(\text{þ}^*) \\ cov(\text{þ}^{\circ},\text{þ}^*) \\ V(\text{þ}^{\circ}) \end{pmatrix} \\ &= \begin{pmatrix} 1 & -2c & c^2 \end{pmatrix}\begin{pmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{pmatrix} = \mathbf{k}'(c)\,\boldsymbol{\theta},\end{aligned} \tag{25}$$

where $\boldsymbol{\theta}$ represents the three unknown parameters that we need to estimate. It is interesting to note that these parameters can be directly identified with the three letters commonly used in textbook treatments of mean-variance frontiers, as

$$
\begin{aligned}
\theta_1 &= V(\text{þ}^*) = C(\text{þ}^*) = \textsc{c}, \\
\theta_2 &= cov(\text{þ}^\circ, \text{þ}^*) = C(\text{þ}^\circ) = E(\text{þ}^*) = \textsc{a}, \\
\theta_3 &= V(\text{þ}^\circ) = E(\text{þ}^\circ) = \textsc{b}.
\end{aligned}
$$

Similarly, expression (11) shows that the variance of the elements of the RMVF are linear in a new set of parameters that correspond to the elements of the inverse of the second moment matrix of the uncentred representing portfolios, or the inverse of the variance matrix of the centred representing portfolios. As a result, we can express the RMVF as

$$
Var[R^{MV}(\nu)] = \begin{pmatrix} 1 & -2\nu & \nu^2 \end{pmatrix} \begin{pmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \end{pmatrix} = \mathbf{k}'(\nu)\,\boldsymbol{\eta}, \tag{26}
$$

where $\boldsymbol{\eta}$ represents the three unknown parameters that we need to estimate. Once again, these parameters are also related to the letters used in textbook treatments of mean-variance frontiers because

$$
\eta_1 = \textsc{b}/\textsc{d}; \quad \eta_2 = \textsc{a}/\textsc{d}; \quad \eta_3 = \textsc{c}/\textsc{d},
$$

where $\textsc{d}=\textsc{bc}\text{-}\textsc{a}^2$.[10]

Given that we saw in section 2 that we can assume without loss of generality that the vector of available asset payoffs is $\mathbf{x} = (R, \mathbf{r}')'$, with cost $\mathbf{e}_1$ (see (1)), we can estimate all the required parameters from the exactly identified system of moment conditions

$$
E\begin{bmatrix} \mathbf{x}\mathbf{x}'\boldsymbol{\phi}^* - \mathbf{e}_1 \\ \mathbf{x}\mathbf{x}'\boldsymbol{\phi}^\circ - \mathbf{x} \\ \mathbf{x}'\boldsymbol{\phi}^\circ - \nu^\circ \end{bmatrix} = \mathbf{0}, \tag{27}
$$

where the parameters to estimate are $(\boldsymbol{\phi}^{*\prime}, \boldsymbol{\phi}^{\circ\prime}, \nu^\circ)'$. The main advantage of decomposing $\mathbf{x}$ into a single return and a vector of n excess returns is that the first entry of $\boldsymbol{\phi}^*$ identifies $C(p^*) = E(p^{*2})$, while the first entry of $\boldsymbol{\phi}^\circ$ identifies $C(p^\circ) = E(p^* p^\circ)$. Finally, the additional parameter ν° identifies $E(p^\circ) = E(p^{\circ 2})$. On this basis, we can obtain $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$

[10] Therefore, it would be straightforward to compute from these parameters other objects such as the mean and variance of the minimum variance portfolio ($\textsc{a}/\textsc{c}$ and $1/\textsc{c}$, respectively), or the curvature of the frontier in mean-variance space ($\textsc{c}/\textsc{d}$), which is such that the reciprocal of its square root yields the slope of the asymptotes in mean-standard deviation space.

from the one-to-one mappings

$$\boldsymbol{\theta} = \left[1 - E\left(p^{\circ 2}\right)\right]^{-1} \begin{pmatrix} E(p^{*2})\left[1 - E\left(p^{\circ 2}\right)\right] + E^2(p^* p^\circ) \\ E(p^* p^\circ) \\ E\left(p^{\circ 2}\right) \end{pmatrix},$$

$$\boldsymbol{\eta} = \left[E(p^{*2})E\left(p^{\circ 2}\right) - E^2(p^* p^\circ)\right]^{-1} \begin{pmatrix} E(p^{*2})\left[1 - E\left(p^{\circ 2}\right)\right] + E^2(p^* p^\circ) \\ E(p^* p^\circ) \\ E\left(p^{\circ 2}\right) \end{pmatrix}.$$

Alternatively, we could start from the exactly identified set of moment conditions that defines the centred representing portfolios:

$$E\begin{bmatrix} \mathbf{x}(\mathbf{x}'\boldsymbol{\varphi}^* - \theta_3) - \mathbf{e}_1 \\ \mathbf{x}(\mathbf{x}'\boldsymbol{\varphi}^\circ - \theta_1) - \mathbf{x} \\ \mathbf{x}'\boldsymbol{\varphi}^\circ - \theta_1 \end{bmatrix} = \mathbf{0}, \tag{28}$$

and then obtain $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\eta}}$ from the estimators of $(\boldsymbol{\varphi}^*, \boldsymbol{\varphi}^\circ, \theta_1)$. In particular, $\theta_3 = V(\text{þ}^*) = C(\text{þ}^*)$ is identified by the first entry of $\boldsymbol{\varphi}^*$, which is associated to the only nonzero cost payoff, and $\theta_2 = cov(\text{þ}^\circ, \text{þ}^*) = C(\text{þ}^\circ)$ is identified by the first entry of $\boldsymbol{\varphi}^\circ$ for the same reason. Finally, we estimate $\theta_1 = V(\text{þ}^\circ) = E(\text{þ}^\circ)$ directly.

Under standard regularity conditions (see e.g. Newey and McFadden (1994)), the resulting GMM estimator of $\boldsymbol{\theta}$ will converge in probability to its true value, and the same applies to the weights of the mean representing portfolios $\boldsymbol{\phi}^*$ and $\boldsymbol{\phi}^\circ$. Therefore, the GMM estimators of $Var[m_1^{MV}(c)]$ and $V[R^{MV}(\nu)]$ will also converge in probability to their population counterparts for fixed c and ν. Further, we can use the Cauchy-Schwarz inequality to show that the GMM estimators of the entire SMVF and RMVF will converge uniformly to their population analogues over any finite range. Specifically,

$$\sup_{c \in [\underline{c}, \bar{c}]} \left| \mathbf{k}'(c)\,\hat{\boldsymbol{\theta}} - \mathbf{k}'(c)\,\boldsymbol{\theta} \right| = \sup_{c \in [\underline{c}, \bar{c}]} \left| \mathbf{k}'(c)\,(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \right|$$
$$\leq \left\{ \sup_{c \in [\underline{c}, \bar{c}]} \left[\mathbf{k}'(c)\,\mathbf{k}(c)\right]^{1/2} \right\} \left[(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})'(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})\right]^{1/2} = o_p(1).$$

As in the case of arbitrage portfolios, though, the SMVF and RMVF are subject to substantial sampling variability despite their uniform convergence. To emphasise yet again the importance of sampling uncertainty in this context, we have conducted a simulation experiment similar to the one reported in section 3, but this time assuming that investors also have access to US Tbills, whose gross returns are measured in real terms. The results are reported in Figure 7. Once again, if one did not take into account sampling uncertainty, one would form very different optimal portfolios depending on the sample, and would reach rather different conclusions about the mean-variance trade-offs.

<Figures 7a, 7b and 7c>

For that reason, we should again construct joint confidence regions. Let us start with K values of $\sigma^2(c)$. The main difference between gross and excess returns is that there are 3 unknown parameters as opposed to only 1. As a result, the asymptotic covariance matrix of the corresponding point estimators will be singular whenever we jointly consider more than 3 points. In particular, given that the SMVF is linear in $\boldsymbol{\theta}$ in view of (25), the asymptotic joint distribution will be

$$\sqrt{T}\begin{bmatrix} \hat{\sigma}^2(c_1) - \sigma^2(c_1) \\ \vdots \\ \hat{\sigma}^2(c_K)\,\sigma^2(c_K) \end{bmatrix} \xrightarrow{d} N\left\{\mathbf{0}, \begin{bmatrix} \mathbf{k}'(c_1) \\ \vdots \\ \mathbf{k}'(c_K) \end{bmatrix} \boldsymbol{\Upsilon} \begin{bmatrix} \mathbf{k}(c_1) & \dots & \mathbf{k}(c_K) \end{bmatrix}\right\},$$

where

$$\sqrt{T}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} N(\mathbf{0}, \boldsymbol{\Upsilon}),$$

with $\boldsymbol{\Upsilon}$ obtained by applying the Delta method to the GMM estimators of $(\boldsymbol{\phi}^{*\prime}, \boldsymbol{\phi}^{\circ\prime}, \nu^{\circ})'$ from (27). As a result, the joint confidence region in $[\sigma^2(c_1), \dots \sigma^2(c_K)]$ space for any $K \geq 3$ will be the three dimensional ellipsoid:

$$\begin{bmatrix} \hat{\sigma}^2(c_1)\text{-}\sigma^2(c_1) & \dots & \hat{\sigma}^2(c_K)\text{-}\sigma^2(c_K) \end{bmatrix} \mathbf{K}'(\mathbf{c})\,\boldsymbol{\Upsilon}^{-1}\mathbf{K}(\mathbf{c}) \begin{bmatrix} \hat{\sigma}^2(c_1)\text{-}\sigma^2(c_1) \\ \vdots \\ \hat{\sigma}^2(c_K)\text{-}\sigma^2(c_K) \end{bmatrix} \leq \frac{Q(1\text{-}\alpha;3)}{T}, \tag{29}$$

where

$$\mathbf{K}(\mathbf{c}) = \left\{\begin{bmatrix} \mathbf{k}(c_1) & \dots & \mathbf{k}(c_K) \end{bmatrix} \begin{bmatrix} \mathbf{k}'(c_1) \\ \vdots \\ \mathbf{k}'(c_K) \end{bmatrix}\right\}^{-1} \begin{bmatrix} \mathbf{k}(c_1) & \dots & \mathbf{k}(c_K) \end{bmatrix}$$

and $\mathbf{c} = (c_1, \dots, c_K)$.

Figure 8 presents this ellipsoid for three different values of c, $(= .9, .95, 1)$. As expected, it is centred around the estimated values of $\sigma^2(.9)$, $\sigma^2(.95)$ and $\sigma^2(1)$ (marked by a '+'), and in this case it contains the corresponding true values (market by a '*').

<Figure 8>

Given that we cannot represent this confidence region in more than three dimensions, once again we find it convenient to map it onto the space $[c, \sigma^2(c)]$ by plotting in that space all the points that belong to the joint confidence interval (29) for a chosen value of α. Nevertheless, it is important to remark that such a plot fails to capture the strong dependence that exists between the different values of c. It turns out that it is rather easy to characterise the envelope of this set, as the only operation required is the projection of the joint interval (29) on to the $\sigma^2(c_1), \sigma^2(c_2), \ldots$ and $\sigma^2(c_K)$ axes.

Lemma 1 *The projection of the ellipsoid (29) onto the $\sigma^2(c_1)$ axis is*

$$[\hat{\sigma}^2(c_1) - \sigma^2(c_1)]^2 \leq \frac{\mathbf{k}'(c_1)\,\boldsymbol{\Upsilon}\mathbf{k}(c_1)}{T} Q(1-\alpha; 3) \tag{30}$$

irrespective of the chosen values of $(c_2, ..., c_K)$.

Therefore, the limits of the mapping of the joint interval (29) onto $[c, \sigma^2(c)]$ space will be

$$\mathbf{k}'(c)\,\hat{\boldsymbol{\theta}} \pm \sqrt{\frac{\mathbf{k}'(c)\,\boldsymbol{\Upsilon}\mathbf{k}(c)}{T} Q(1-\alpha; 3)}.$$

Not surprisingly, the width of these limits increases with $\boldsymbol{\Upsilon}$ and α, and decreases with the sample size. Given the singular nature of (29), in effect the resulting plot simply depicts the set of SMVF that correspond to the $100(1-\alpha)\%$ confidence interval for $\hat{\boldsymbol{\theta}}$,[11] which is given by[12]

$$T(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})'\boldsymbol{\Upsilon}^{-1}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \leq Q(1-\alpha; 3).$$

Importantly, these limits no longer coincide with the point by point confidence intervals. The reason is that for a fixed c, the asymptotic marginal distribution of the estimator of $\sigma^2(c)$ will be

$$\sqrt{T}\mathbf{k}'(c)\,(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} N[\mathbf{0}, \mathbf{k}'(c)\,\boldsymbol{\Upsilon}\mathbf{k}(c)].$$

[11] Specifically, it is easy to show that the upper limit of (30) is the solution to the programme

$$\max_{\boldsymbol{\theta}} \mathbf{k}'(c)\,\boldsymbol{\theta} \qquad s.t. \qquad T(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})'\boldsymbol{\Upsilon}^{-1}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \leq Q(1-\alpha; 3),$$

while the lower limit coincides with the corresponding minimisation.

[12] Not surprisingly, if we replace $\hat{\sigma}^2(c_k)$ and $\sigma^2(c_k)$ by $\mathbf{k}'(c_k)\,\hat{\boldsymbol{\theta}}$ and $\mathbf{k}'(c_k)\,\boldsymbol{\theta}$, respectively, then (29) collapses to the confidence interval for $\hat{\boldsymbol{\theta}}$. In this regard, note that the determinant of

$$\begin{pmatrix} \mathbf{k}(c_1) & \mathbf{k}(c_2) & \mathbf{k}(c_3) \end{pmatrix}$$

is $2(c_1 - c_2)(c_1 - c_3)(c_2 - c_3)$. Therefore, we can invert this matrix to get the implied $\boldsymbol{\theta}$ using three different variance values.

Hence, the pointwise confidence interval will look like (30), but with $Q(1-\alpha;1)$ instead of $Q(1-\alpha;3)$, so they will be narrower (see Bansal, Dahlquist and Harvey (2004) and Abhyankar, Basu and Stremme (2007)).

Similarly, given that the variance of the elements of the RMVF frontier are linear in $\boldsymbol{\eta}$, we can repeat the same exercise in terms of these parameters to study this other frontier instead. In this sense, Figures 9a and 9b present a graph of the limits (in red) and the pointwise confidence bands (in blue) using the same design for gross returns considered in the previous figures.

<Figures 9a, 9b and 9c>

Given that the weights of the traded part of the SMVF is given by $\boldsymbol{\varphi}^* - c\boldsymbol{\varphi}^\circ$ in view of (6), it is also straightforward to derive joint confidence regions for those weights for $(c_1, c_2, ..., c_K)$. The same applies to the weights of the RMVF for $(\nu_1, \nu_2, ..., \nu_K)$, which can obtained from the textbook formula

$$(\eta_1\boldsymbol{\varphi}^* - \eta_2\boldsymbol{\varphi}^\circ) + (\eta_3\boldsymbol{\varphi}^\circ - \eta_2\boldsymbol{\varphi}^*)\,\nu$$

in view of (10). Therefore, regardless of the frontier the weights for a particular asset, i say, depend on two parameters only: $(\varphi_i^*, \varphi_i^\circ)$ in the SDF case, and $(\eta_1\varphi_i^* - \eta_2\varphi_i^\circ, \eta_3\varphi_i^\circ - \eta_2\varphi_i^*)$ in the portfolio case. Therefore, the joint confidence regions for the weights of the i^{th} asset will be two-dimensional ellipses for any $K \geq 2$. As before, we can also map those joint confidence regions into the relevant space, as shown in Figure 9c for the weight of the second asset in the RMVF, which corresponds to the first excess return.

Finally, it would be straightforward to consider joint confidence intervals for the weights of two or more assets.

4.2 Efficient estimation imposing a common point: Tangency tests

Let us again partition the available assets into two sets of payoffs $\mathbf{x}_1$ and $\mathbf{x}_2$ of dimensions N_1 and N_2, respectively, with $N = N_1 + N_2$, so that $\mathbf{x}' = (\mathbf{x}_1', \mathbf{x}_2')$. We want to compare the SMVF and RMVF frontiers generated by $\mathbf{x}_1$ alone with the ones generated by the whole of $\mathbf{x}$, where in line with the previous section, we assume that $\mathbf{x}_1$ contains at least one asset of non-zero cost. In general, when we consider both $\mathbf{x}_1$ and $\mathbf{x}_2$, the RMVF frontier will shift to the left because the available risk-return trade-offs improve, while the SMVF frontier will rise because there is more information in the data about the underlying SDF. However, this is not always the case. In particular, we say that $\mathbf{x}_1$ spans the SMVF and/or RMVF generated from $\mathbf{x}$ when the original and extended frontiers coincide. We shall return to this case in section 4.4. Unlike in the case of excess returns, though, a third, and last, possibility arises, namely, that the original and extended frontiers touch at a single point. Although it is common in the literature to refer to this situation as "intersection", we prefer to use the word "tangency" because the frontiers are never secant to each other, as the word "intersection" may suggest. As expected, the duality between the SMVF and RMVF means that tangency in one of them implies tangency in the other.[13]

Given that the elements of the SMVF can be written as (a constant plus) a portfolio of the cost and mean representing portfolios (see expression (6)), there will be tangency on the SMVF if and only if there is a $\dot{c}$ such that

$$\boldsymbol{\phi}^* - \left[\frac{\dot{c} - E(p^*)}{1 - E(p^\circ)}\right] \boldsymbol{\phi}^\circ = \boldsymbol{\varphi}^* - \dot{c}\boldsymbol{\varphi}^\circ \tag{31}$$

has zero weights on $\mathbf{x}_2$. Similarly, expression (10) implies there will be tangency on the RMVF if and only if there is a $\dot{\nu}$ such that

$$\begin{aligned}&\left[E(p^{\circ 2}) - E(p^* p^\circ)\dot{\nu}\right] \boldsymbol{\phi}^* + \left[E(p^{*2})\dot{\nu} - E(p^* p^\circ)\right] \boldsymbol{\phi}^\circ \\ = \;& \left[V(\mathfrak{p}^\circ) - cov(\mathfrak{p}^*, \mathfrak{p}^\circ)\dot{\nu}\right] \boldsymbol{\varphi}^* + \left[V(\mathfrak{p}^*)\dot{\nu} - cov(\mathfrak{p}^*, \mathfrak{p}^\circ)\right] \boldsymbol{\varphi}^\circ \end{aligned} \tag{32}$$

has zero weights on $\mathbf{x}_2$. Therefore, if we want to estimate the SMVF and RMVF frontiers exploiting tangency, then we simply have to impose these restrictions in the just identified moment conditions (27), or in the system (28) that identifies the centred representing portfolios. A particularly convenient way of doing so is by writting

$$\boldsymbol{\phi}^* = \begin{pmatrix} \boldsymbol{\phi}_1^* \\ \boldsymbol{\phi}_2^* \end{pmatrix} = \begin{pmatrix} \lambda_0 \boldsymbol{\phi}_1^\circ - \boldsymbol{\lambda}_1 \\ \lambda_0 \boldsymbol{\phi}_2^\circ \end{pmatrix},$$

where in view of (7) we can interpret

$$\lambda_0 = \left[\frac{\dot{c} - E(p^*)}{1 - E(p^\circ)}\right]$$

and

$$\boldsymbol{\lambda}_1 = \lambda_0 \boldsymbol{\phi}_1^\circ - \boldsymbol{\phi}_1^*,$$

[13]Save in the two duality exceptions discussed in appendix C of Peñaranda and Sentana (2011), in which tangency in one frontier implies common asymptotes for the other one.

the restriction being

$$\boldsymbol{\lambda}_2 = \lambda_0 \boldsymbol{\phi}_2^{\circ} - \boldsymbol{\phi}_2^{*} = \mathbf{0}.$$

As a result, we can work with the overidentified system:

$$E \begin{bmatrix} \mathbf{x}\left(\lambda_0 - \mathbf{x}_1' \boldsymbol{\lambda}_1\right) - \mathbf{e}_1 \\ \mathbf{x}\mathbf{x}'\boldsymbol{\phi}^{\circ} - \mathbf{x} \\ \mathbf{x}'\boldsymbol{\phi}^{\circ} - \nu^{\circ} \end{bmatrix} = \mathbf{0}, \tag{33}$$

which is linear in the unknown parameters $(\lambda_0, \boldsymbol{\lambda}_1', \boldsymbol{\phi}^{\circ\prime}, \nu^{\circ})'$.[14]

We can also estimate the SMVF and RMVF frontiers corresponding to $\mathbf{x}_1$ subject to the tangency restriction by using the overidentified system of moment conditions

$$E \begin{bmatrix} \mathbf{x}\left(\lambda_0 - \mathbf{x}_1' \boldsymbol{\lambda}_1\right) - \mathbf{e}_1 \\ \mathbf{x}_1\mathbf{x}_1'\boldsymbol{\phi}_1^{\circ} - \mathbf{x}_1 \\ \mathbf{x}_1'\boldsymbol{\phi}_1^{\circ} - \nu_1^{\circ} \end{bmatrix} = \mathbf{0}. \tag{34}$$

Since we are using optimal GMM in all cases, the imposition of the equality restrictions (31) or (32) should generally lead to efficiency gains relative to the unrestricted estimation of both frontiers. Figures 10a and 10b present the restricted and unrestricted frontiers imposing the tangency constraint for a given simulation, as well as the corresponding confidence intervals. Finally, Figure 10c looks at portfolio weights.

<Figures 10a, 10b and 10c>

As in previous cases, we can test the tangency restrictions by means of the overidentified restriction test of (33), or its centred counterpart, both of which will have $N_2 - 1$ degrees of freedom under the null. As expected, these tests coincide with the corresponding distance metric tests of the null hypothesis $H_0 : \boldsymbol{\lambda}_2 = \mathbf{0}$.

Once again, though, the most popular tangency test in empirical finance is the regression test considered by Gibbons (1982), Kandel (1984) and Shanken (1985, 1986). Although these authors discussed likelihood ratio and F-tests under the assumption that the conditional distribution of $\mathbf{x}_2$ given $\mathbf{x}_1$ is multivariate normal with an affine mean and

[14] If $N_1 = 1$ then there will be a singularity if $\dim(\mathbf{x}_1) = 1$. To solve it, we can impose

$$\nu^{\circ} = \phi_1^{\circ 2} / \phi_1^{*}.$$

and skip the condition for ν° (see Penaranda and Sentana (2010a) for a more formal treatment of optimal GMM with this type of singularities).

a constant covariance matrix, it is straightforward to obtain robust GMM versions (see section 5.3.2 of Campbell, Lo and MacKinlay (1997)). Specifically, if we exploit the fact that the first element of $\mathbf{x}_1$ is the only payoff with a non-zero cost, the distance metric version of the regression-based tangency tests will be the overidentification restriction test of the $N_2(N_1+1)$ normal equations

$$E\left\{\left(\begin{array}{c}1\\ \mathbf{x}_1\end{array}\right)\otimes[\mathbf{x}_2-\mathbf{b}_1(R_1-\dot{c}^{-1})-\mathbf{B}\mathbf{r}_1]\right\}=\mathbf{0}, \tag{35}$$

where $(\mathbf{b}_1, \mathbf{B})$ denotes the matrix of regression coefficients of $\mathbf{x}_2$ onto $\mathbf{x}_1$, and $\dot{c}^{-1}$ is the unknown expected value of the zero-beta return corresponding to the tangency portfolio. It is tedious but otherwise straightforward to show that this regression test is asymptotically equivalent under the null and compatible sequences of local alternatives to the overidentification restriction test of (33). In addition, it is also possible to prove that they will be numerically identical for single-step GMM estimators such as Continuously Updated GMM.

The main advantage of the test based on (33) relative to the tangency test based on (35) is that the moment conditions under the null are linear in the parameters, which simplifies the computations considerably. In addition, given that (35) is not particularly useful for the purposes of estimating mean-variance frontiers, we shall not discuss it any further.[15]

Finally, in some important cases we could be interested in studying the tangency of two frontiers at a pre-specified point. For instance, Chen and Knez (1996) propose to test the tangency of the two SMVF frontiers at p^* in order to evaluate the performance of a single investment fund with gross returns R_2 with respect to the "benchmarks" included in $\mathbf{R}_1$. Specifically, they test $\boldsymbol{\phi}_2^*=0$ on the basis of

$$E(\mathbf{x}\mathbf{x}_1'\boldsymbol{\phi}_1^*-\mathbf{e}_1)=\mathbf{0},$$

which are the moment conditions that implicitly define the uncentred cost representing portfolio (see also De Roon and Nijman (2001), who provide a survey of performance evaluation, and its relationship to tangency tests).

More generally, if we want to test for tangency between SMVFs at a given value of $\dot{c}$, we can either impose the linear restriction

$$\boldsymbol{\varphi}_2^*=\dot{c}\boldsymbol{\varphi}_2^\circ$$

[15] See Beaulieu, Dufour and Khalaf (2007) for alternative tangency tests with exact finite sample distributions.

on the moment conditions (28), or else impose the non-linear restriction

$$\phi_1^\circ = \dot{c} + \lambda_0(1 - \nu^\circ)$$

on the moment conditions (33), where ϕ_1° denotes the first entry of $\boldsymbol{\phi}^\circ$.

Testing for tangency between two nested RMVFs at a given value of $\dot{\nu}$ is somewhat more involved. One possibility would be to impose the non-linear restriction

$$(\eta_1 - \eta_2\dot{\nu})\,\boldsymbol{\varphi}_2^* + (\eta_3\dot{\nu} - \eta_2)\,\boldsymbol{\varphi}_2^\circ = \mathbf{0}$$

on the moment conditions (28). Another possibility would be as follows. Let τ^* denote the reciprocal of the first element of $-\boldsymbol{\lambda}_1$. If we multiply the first block of the moment conditions in (33) by this parameter, we end up with

$$E\left[\mathbf{x}\mathbf{x}_1'\begin{pmatrix} 1 \\ \boldsymbol{\omega}_1 \end{pmatrix} - (\tau^*\mathbf{e}_1 + \tau^\circ\mathbf{x})\right] = \mathbf{0}, \tag{36}$$

where we have called $\tau^\circ = \lambda_0\tau^*$ and $(1, \boldsymbol{\omega}_1')' = -\boldsymbol{\lambda}_1\tau^*$. If we then impose the constraint

$$\tau^\circ = \left(\frac{\dot{\nu} - \tau^*\phi_1^\circ}{\nu^\circ}\right),$$

the moment conditions (36) will identify $\boldsymbol{\omega}_1$ as the weights on $\mathbf{r}_1$ of a tangency portfolio in the RMVF with mean $\dot{\nu}$ because $\nu^\circ = E(p^\circ)$ and $\phi_1^\circ = E\,(p^*)$. Therefore, we can test for tangency at $\dot{\nu}$ by using the overidentification test of the system

$$E\left\{\begin{array}{c} \mathbf{x}\,(R_1 + \mathbf{r}_1'\boldsymbol{\omega}_1) - \tau^*\mathbf{e}_1 - [(\dot{\nu} - \tau^*\phi_1^\circ)/\nu^\circ]\mathbf{x} \\ \mathbf{x}\mathbf{x}'\boldsymbol{\phi}^\circ - \mathbf{x} \\ \mathbf{x}'\boldsymbol{\phi}^\circ - \nu^\circ \end{array}\right\} = \mathbf{0}. \tag{37}$$

In fact, if we set $\dot{\nu} = 0$ in these moment conditions, the duality of the RMVF and SMVF would imply that we would obtain the test for tangency at the minimum of the SMVF considered by DeSantis (1995). Similarly, we could set $\dot{c} = 0$ to test for tangency at the minimum variance portfolio in the RMVF.

Obviously, all these tests will have N_2 degrees of freedom when the tangency point is known, since there is one parameter less to estimate.

4.3 Efficient estimation imposing a linear factor pricing model: Asset pricing tests

We can again consider using an asset pricing model in order to reduce the sampling variability of estimated mean-variance frontiers. If we model the true SDF m as an affine transformation of some $k \leq N$ observable risk factors $\mathbf{f}$, then we can express the pricing equation as

$$E[\mathbf{x}(\lambda_0 - \boldsymbol{\lambda}'\mathbf{f})] = \mathbf{e}_1$$

for some real numbers $(\lambda_0, \boldsymbol{\lambda}')'$. In fact, we can understand the tangency restrictions discussed in the previous section as imposing a linear factor pricing model in which the pricing factors $\mathbf{f}$ coincide with some traded payoffs $\mathbf{x}_1$. However, in general $\mathbf{f}$ does not have to be a subset of $\mathbf{x}$. Unlike in the case of arbitrage portfolios, though, the scaling of m is no longer an issue in the presence of at least one asset whose cost is not zero, which means that λ_0 and $\boldsymbol{\lambda}$ can be separately identified. Assuming that $\mathbf{f}$ and $\mathbf{x}$ do not share any common elements, we get

$$E\begin{bmatrix} \mathbf{x}(\lambda_0 - \boldsymbol{\lambda}'\mathbf{f}) - \mathbf{e}_1 \\ \mathbf{x}\mathbf{x}'\boldsymbol{\phi}^* - \mathbf{e}_1 \\ \mathbf{x}\mathbf{x}'\boldsymbol{\phi}^\circ - \mathbf{x} \\ \mathbf{x}'\boldsymbol{\phi}^\circ - \nu^\circ \end{bmatrix} = \mathbf{0}, \tag{38}$$

where the unknown parameters are $(\lambda_0, \boldsymbol{\lambda}', \boldsymbol{\phi}^{*\prime}, \boldsymbol{\phi}^{\circ\prime}, \nu^\circ)'$. In this way, we should generally obtain more efficient estimators of the mean-variance frontiers that exploit the pricing equations.[16]

As before, we can use the overidentifying restriction test of system (38) to test the asset pricing restrictions.

4.4 Efficient estimation imposing spanning restrictions: Spanning tests

As we mentioned before, we say that $\mathbf{x}_1$ spans the SMVF and/or RMVF generated from $\mathbf{x}$ when the original and extended frontiers coincide. Given expression (10), this will happen if and only if neither the cost nor the mean representing portfolios depend on the vector of payoffs $\mathbf{x}_2$. Under the null hypothesis, there will be one pair of MV frontiers. Under the alternative, there will be two: those generated from $\mathbf{x}_1$ alone, and those generated from $\mathbf{x}$. We can estimate the pairs of unrestricted frontiers by using the

[16] Once again, the asset pricing constraint can still be interpreted as tangency, but with respect to the factor mimicking portfolios. That is, there is a linear combination of the representing portfolios that depends on the factor mimicking portfolios only

$$\mathbf{x}'(\boldsymbol{\phi}^* - \lambda_0\boldsymbol{\phi}^\circ) = \mathbf{x}'E^{-1}(\mathbf{x}\mathbf{x}')[\mathbf{e}_1 - \lambda_0 E(\mathbf{x})] = \left[\mathbf{x}'E^{-1}(\mathbf{x}\mathbf{x}')E(\mathbf{x}\mathbf{f}')\right]\boldsymbol{\lambda}.$$

procedures discussed at the beginning of section 4. On the other hand, we can estimate the common mean-variance frontiers by imposing the restrictions $H_0 : \boldsymbol{\phi}_2^* = \boldsymbol{\phi}_1^\circ = \mathbf{0}$ on the set of moment conditions (27). The resulting system of overidentified moment conditions will be

$$E \begin{bmatrix} \mathbf{x}\mathbf{x}_1'\boldsymbol{\phi}_1^* - \mathbf{e}_1 \\ \mathbf{x}\mathbf{x}_1'\boldsymbol{\phi}_1^\circ - \mathbf{x} \\ \mathbf{x}_1'\boldsymbol{\phi}_1^\circ - \nu^\circ \end{bmatrix} = \mathbf{0}. \tag{39}$$

The optimal GMM estimator obtained from (39) will generally be more efficient than the corresponding estimator obtain from the unrestricted system (27) as long as the spanning restrictions hold. Moreover, the results in Breusch et al (1999) imply that this estimator will also be generally more efficient than the one obtained from the just identified $2N_1 + 1$ moment conditions

$$E \begin{bmatrix} \mathbf{x}_1\mathbf{x}_1'\boldsymbol{\phi}_1^* - \mathbf{e}_1 \\ \mathbf{x}_1\mathbf{x}_1'\boldsymbol{\phi}_1^\circ - \mathbf{x}_1 \\ \mathbf{x}_1'\boldsymbol{\phi}_1^\circ - \nu^\circ \end{bmatrix} = \mathbf{0}. \tag{40}$$

As in the case of excess returns, though, an exception to this rule arises in the case of *i.i.d.* elliptical returns.

Proposition 2 *If $\mathbf{x}_t$ is an i.i.d. elliptical random vector with bounded fourth moments, and the null hypothesis of spanning is true, then:*
a) The asymptotic variance of the optimal GMM estimator of ν° obtained from (39), which imposes the spanning constraints $\boldsymbol{\phi}_2^* = \boldsymbol{\phi}_1^\circ = \mathbf{0}$, will coincide with the asymptotic variances of both the estimator obtained from (40), and the estimator obtained from (27).
b) The asymptotic variance of the GMM estimators of $\boldsymbol{\phi}_1^*$ and $\boldsymbol{\phi}_1^\circ$ obtained from (27) will be larger (in the usual positive definite sense) than the asymptotic variance of the optimal GMM estimators based on (39), which in turn coincides with the asymptotic variance of the GMM estimator obtained from (40).

This results extends Lemma 1 in Peñaranda and Sentana (2010b), who prove the asymptotic equivalence mentioned in part b. Trivially, part b) extends to the "estimators" of $\boldsymbol{\phi}_2^*$ and $\boldsymbol{\phi}_2^\circ$.

As usual, the advantage of our GMM set-up is that we can readily use the overidentification test of the moment conditions (39) to test for spanning, since it coincides with the distance metric test of the null hypothesis $H_0 : \boldsymbol{\phi}_2^* = \boldsymbol{\phi}_1^\circ = \mathbf{0}$. Such a spanning test was proposed by Peñaranda and Sentana (2010a), who also considered a centred representing portfolio counterpart.

However, the most popular mean-variance spanning tests are the regression-based tests considered by Huberman and Kandel (1987). Given that the first element of $\mathbf{x}_1$ is the only asset with a non-zero cost, their test would correspond to the overidentification test of the $N_2(N_1+1)$ moment conditions

$$E\left[\begin{pmatrix} 1 \\ \mathbf{x}_1 \end{pmatrix} \otimes (\mathbf{x}_2 - \mathbf{B}\mathbf{r}_1)\right] = \mathbf{0}, \tag{41}$$

Peñaranda and Sentana (2010a) show that the overidentification test based on the moment conditions (39), their centred representing portfolios counterparts, and the regression version obtained from (41) can be made numerically identical by using single step methods such as CUE. More generally, these authors also show that all these tests are asymptotically equivalent under the null and compatible sequences of alternatives when implemented by two-stage or interated GMM, even though they will not be numerically equivalent in that case.

5 Summary and directions for future research

The contribution of this paper is threefold:

1. We propose GMM-based procedures that allow us to consistently estimate mean-variance frontiers for returns and stochastic discount factors and the weights of the portfolios that belong to them, as well as to derive joint confidence regions that provide analytically tractable and computationally simple alternatives to the Monte Carlo methods considered by Jorion (1992) and Michaud (1998) among others.

2. We explain how to achieve efficiency gains in estimating those frontiers by exploiting theoretically motivated restrictions, such as those derived from asset pricing models or other commonly used assumptions like tangency or spanning.

3. We exploit the integration of estimation and testing implicit in GMM, and study the associated overidentification tests, which can be formally understood as parametric tests of the null hypothesis that the additional restrictions are satisfied.

In addition, we follow Peñaranda and Sentana (2010a,b) in providing a unifiying approach that applies at three different levels:

a. We exploit the duality of the RMVF and SMVF so that our estimators, confidence regions and tests are not necessarily tied down to the specific properties of either frontier.

b. We compare our proposed tests to the extant tests, and show that they are all asymptotically equivalent under the null and compatible Pitman sequences of local alternatives, despite the fact that in some cases the number of parameters and moment conditions can be different.

c. We show that by using single-step GMM procedures such as the Continuously Updated (CU) version in Hansen, Heaton and Yaron (1996), we can make all the different overidentification tests numerically identical.

However, we have not explicitly considered the implications of mutual fund separation for the estimation of mean-variance frontiers. In fact, the only additional restriction in a RMVF context is that the residual of the theoretical regression of $\mathbf{x}_2$ on $\mathbf{x}_1$ must not only be orthogonal to $\mathbf{x}_1$, but also mean independent (see e.g. Chamberlain (1983) or Ferson, Foerster, and Keim (1993)).

We have not discussed either conditional versions of the RMVF or SMVF (see Hansen and Richard (1987) and Gallant, Hansen and Tauchen (1990), respectively). Given that Hansen and Richard (1987) derive conditional analogues to the centred and uncentred representing portfolios, our unifying approach provides a rather natural starting point to look at this problem. However, since the weights of the conditional mean and cost RP portfolios will generally be functions of the relevant information set, we should again consider conditional moment restrictions, as opposed to the unconditonal ones that we have discussed. Given the practical relevance of all these issues, they constitute obvious avenues for further research.

References

Abhyankar, A., D. Basu, and A. Stremme (2007): "Portfolio efficiency and discount factor bounds with conditioning information: an empirical study", *Journal of Banking & Finance* 31, 419-437.

Arellano, M., L.P. Hansen and E. Sentana (2011): "Underidentification?", forthcoming in the *Journal of Econometrics*.

Bansal, R., M. Dahlquist, and C.R. Harvey (2004): "Dynamic strategies and portfolio choice", NBER WP 10820.

Beaulieu, M.C., J.M. Dufour, and L. Khalaf (2007a): Testing mean-variance efficiency in CAPM with possibly non-gaussian errors: an exact simulation-based approach", *Journal of Business and Economic Statistics* 25, 398-410.

Black, F. and R. Litterman (1992): "Global portfolio optimization", *Financial Analysts Journal* 48, September-October.

Breeden, D.T. (1979): "An intertemporal asset pricing model with stochastic consumption and investment opportunities", *Journal of Financial Economics* 7, 265-296.

Breusch, T., H. Qian, P. Schmidt and D. Wyhowski (1999): "Redundancy of moment conditions", *Journal of Econometrics* 91, 89-111.

Britten-Jones, M. (1999): "The sampling error in estimates of mean-variance efficient portfolio weights", *Journal of Finance* 54, 655-671.

Campbell, J.Y., A.W. Lo, and A.C. MacKinlay (1997): *The Econometrics of Financial Markets*, Princeton University Press.

Chamberlain, G. (1983a): "A characterization of the distributions that imply mean-variance utility functions", *Journal of Economic Theory* 29, 185-201.

Chamberlain, G. (1983b): "Funds, factors, and diversification in arbitrage pricing models", *Econometrica* 51, 1305-1323.

Chamberlain, G. and M. Rothschild (1983): "Arbitrage, factor structure, and mean-variance analysis on large asset markets", *Econometrica* 51, 1281-1304.

Chen, N-F., R. Roll, and S.A. Ross (1986): "Economic forces and the stock market, *Journal of Business* 59, 383-403.

Chen, Z. and P.J. Knez, (1996): "Portfolio performance measurement: Theory and applications", *The Review of Financial Studies* 9, 511-555.

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

De Santis, G. (1995): "Volatility bounds for stochastic discount factors: tests and implications from international financial markets", mimeo, University of Southern California.

De Roon, F.A. and T.E. Nijman (2001). "Testing for mean-variance spanning: a survey", *Journal of Empirical Finance* 8-2, 111-156.

Fama, E.F., and K.R. French (1993): "Common risk factors in the returns on stock and bonds", *Journal of Financial Economics* 33, 3-56.

Fang, K.-T., S. Kotz and K.-W. Ng (1990): *Symmetric Multivariate and Related Distributions*, Chapman and Hall.

Ferson, W E., S.R. Foerster and D.B. Keim (1993): "General tests of latent variable models and mean-variance spanning" *Journal of Finance* 48, 131-156.

Gallant, A.R., L.P. 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, 141-179.

Gibbons, M. (1982): "Multivariate tests of financial models: a new approach", *Journal of Financial Economics* 10, 3-27.

Gibbons, M.R., S.A. Ross and J. Shanken (1989): "A test of the efficiency of a given portfolio", *Econometrica* 57, 1121-1152.

Gouriéroux, C. and Monfort, A. (1995): *Statistics and Econometric models*, vols. I and II, Cambridge.

Hansen, L.P. (1982): "Large sample properties of generalized method of moments estimators", *Econometrica* 50, 1029-1054.

Hansen, L.P., J. Heaton and A. Yaron (1996): "Finite sample properties of some alternative GMM estimators", *Journal of Business and Economic Statistics* 14, 262-280.

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

Hansen, L.P., and S.F. Richard (1987): "The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models", *Econometrica* 55, 587-613.

Huberman, G. and S. Kandel (1987). "Mean-variance spanning", *Journal of Finance* 42, 873-888.

Jobson, J.D., and B. Korkie (1980): "Estimation for Markowitz efficient portfolios", *Journal of the American Statistical Association* 75, 544-554.

Jobson, J.D. and B. Korkie (1982): "Potential performance and tests of portfolio efficiency", *Journal of Financial Economics* 10, 433–466.

Jobson, J.D. and B. Korkie (1983): "Statistical inference in two-parameter portfolio theory with multiple regression software", *Journal of Financial and Quantitative Analysis* 18, 189-197.

Jorion, P. (1992): “Portfolio optimization in practice, *Financial Analysts Journal* 48, January/February.

Kan, R. and C. Robotti (2008): “Specification tests of asset pricing models using excess returns”, *Journal of Empirical Finance* 15, 816–38.

Kandel, S. (1984): “The likelihood ratio test statistic of mean-variance efficiency without a riskless asset”, *Journal of Financial Economics* 13, 575-592.

Lo, A.W. (2002): “The statistics of Sharpe ratios”, *Financial Analyst Journal* July/August, 36-52.

MacKinlay, A.C. and M.P. Richardson (1991): “Using generalized method of moments to test mean-variance efficiency”, *Journal of Finance* 46, 511-527.

Markowitz, H. (1952): “Portfolio selection”, *Journal of Finance* 7, 77-99.

Merton, R.C. (1980): “On estimating the expected return on the market: an exploratory investigation”, *Journal of Financial Economics* 8, 323-361.

Michaud, R.O. (1998): *Efficient Asset Management*, Oxford University Press.

Mittnik, S. and P.A. Zadrozny (1993): “Asymptotic distributions of impulse responses, step responses, and variance decompositions of estimated linear dynamic models”, *Econometrica* 61, 857-870.

Newey, W. K. and D. L. McFadden (1994). “Large sample estimation and hypothesis testing”, in R. F. Engle and (eds.), *Handbook of Econometrics*, vol. IV, 2111-2245, Elsevier.

Newey, W.K. and K.D. West (1987): “Hypothesis testing with efficient method of moments estimation”, *International Economic Review* 28, 777-787.

Owen, J. and R. Rabinovitch (1983): On the class of elliptical distributions and their applications to the theory of portfolio choice, *Journal of Finance* 58, 745-752.

Pástor, L. (2000): Portfolio selection and asset pricing models, *Journal of Finance* 55, 179-223.

Pástor, L. and R.F. Stambaugh (2000): Comparing asset pricing models: an investment perspective, *Journal of Financial Economics* 56, 335-381.

Peñaranda, F. and E. Sentana (2004): “Mean-variance tangency tests”, mimeo, CEMFI.

Peñaranda, F., and E. Sentana (2010a): “Spanning tests in return and stochastic discount factor mean-variance frontiers: a unifying approach”, forthcoming in *Journal of Econometrics*.

Peñaranda, F. and Sentana, E. (2010b): “A unifying approach to the empirical evaluation of asset pricing models”, mimeo, CEMFI.

Peñaranda, F. and Sentana, E. (2011): "Duality in mean-variance frontiers with conditioning information", mimeo, CEMFI.

Ross, S.A. (1978): "Mutual fund separation in financial theory - The separating distributions", *Journal of Economic Theory* 17, 254-286.

Sentana, E. (2009): "The econometrics of mean-variance efficiency tests: a survey", *Econometrics Journal* 12, pp. C65-C101.

Shanken, J. (1985): "Multivariate tests of the zero-beta CAPM", *Journal of Financial Economics* 14, 327-348.

Shanken, J. (1986): "Testing portfolio efficiency when the zero-beta rate is unknown: A note", *Journal of Finance* 41, 269–276.

Tu, J., and G. Zhou (2004): "Data-generating process uncertainty: what difference does it make in portfolio decisions?, *Journal of Financial Economics* 72, 385-421.

Figures 1a and 1b: Ensemble of AMVFs and optimal weigths on the first asset

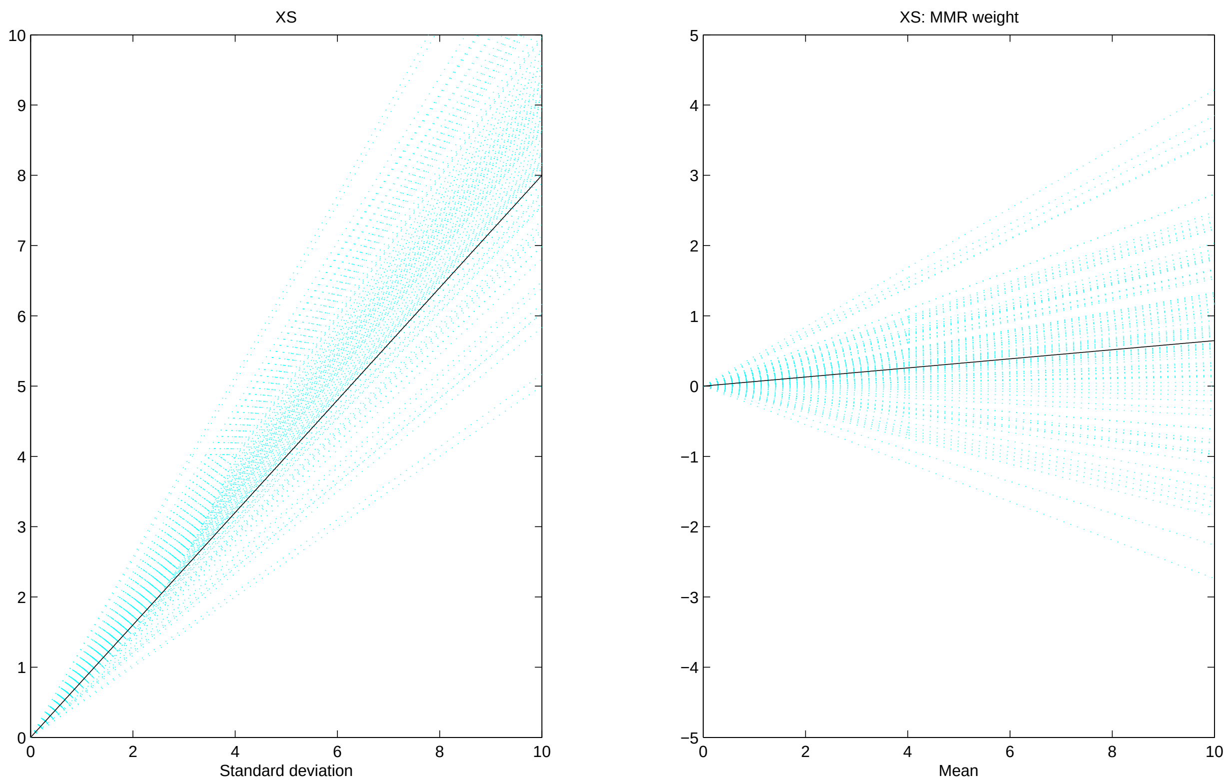

Figure 2: Sampling distribution of maximum Sharpe ratio

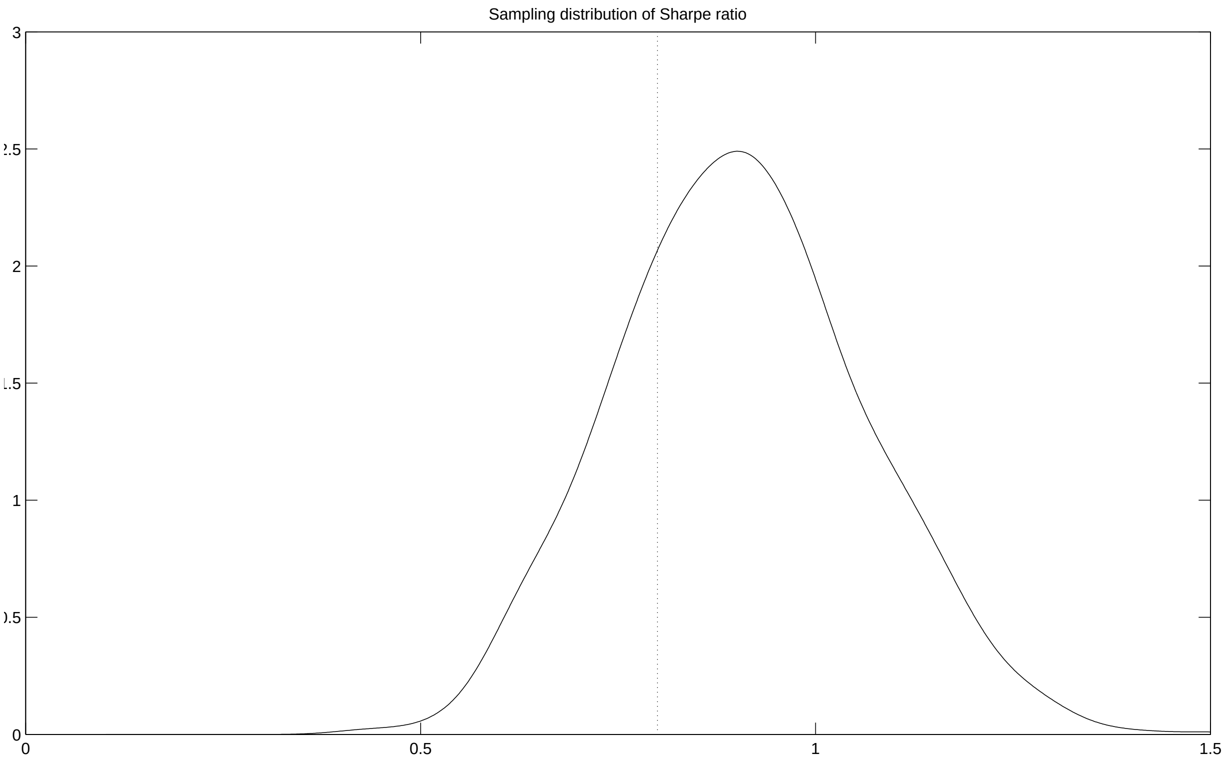

Figure 3: Joint confidence region for two points on the SMVF for excess returns

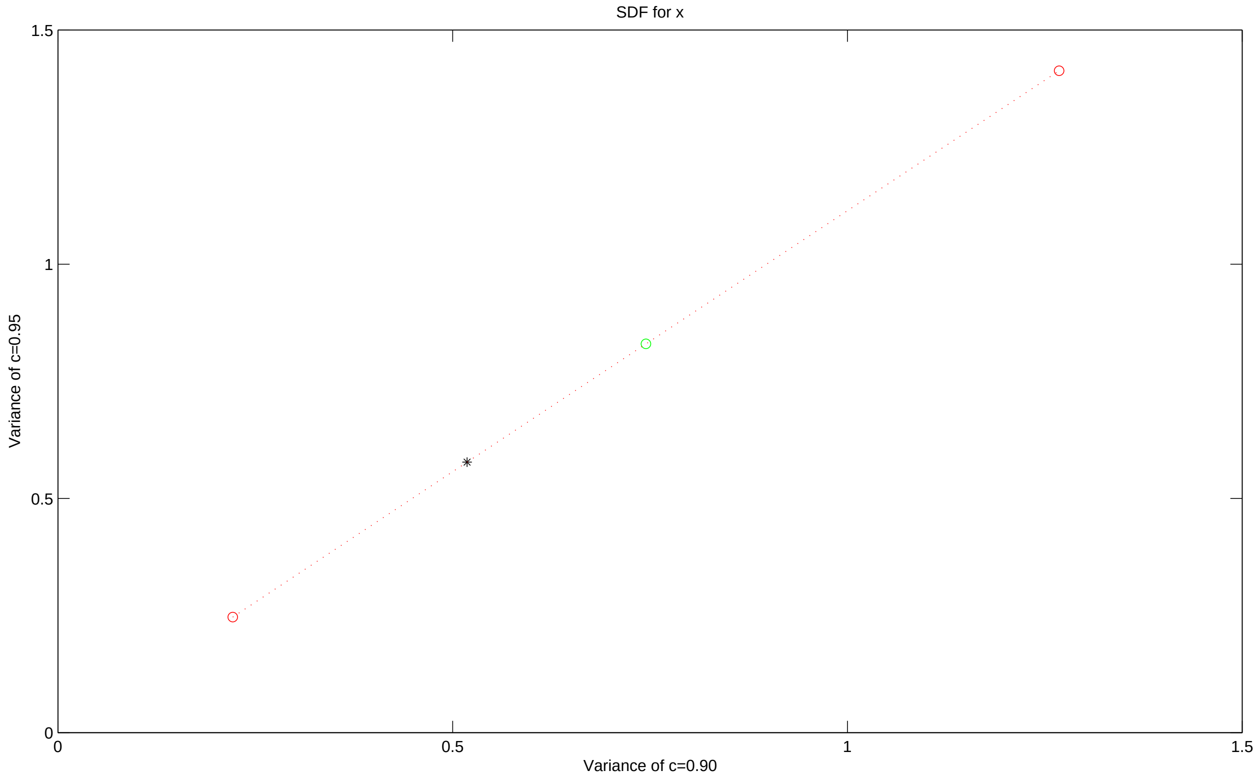

Figures 4a and 4b: Confidence regions for AMVF and optimal weights on the first asset

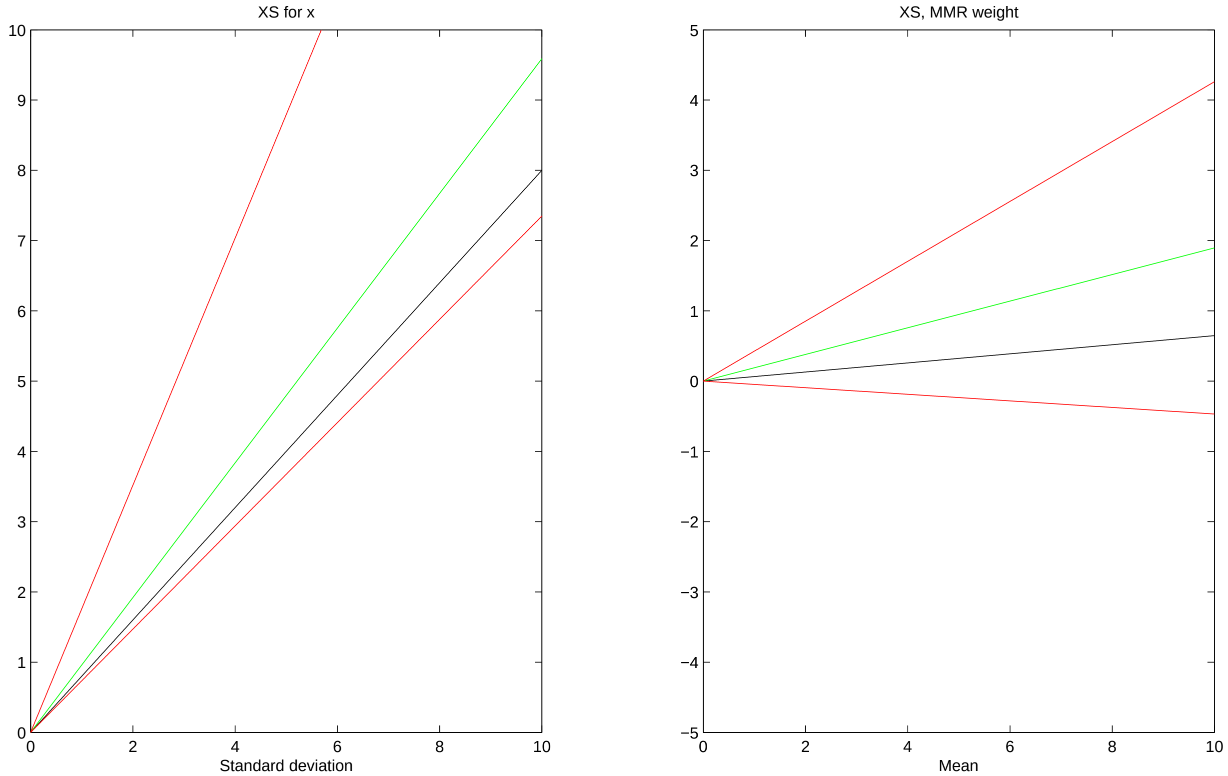

Figures 5a and 5b: Restricted and unrestricted AMVF and optimal weights on the first asset

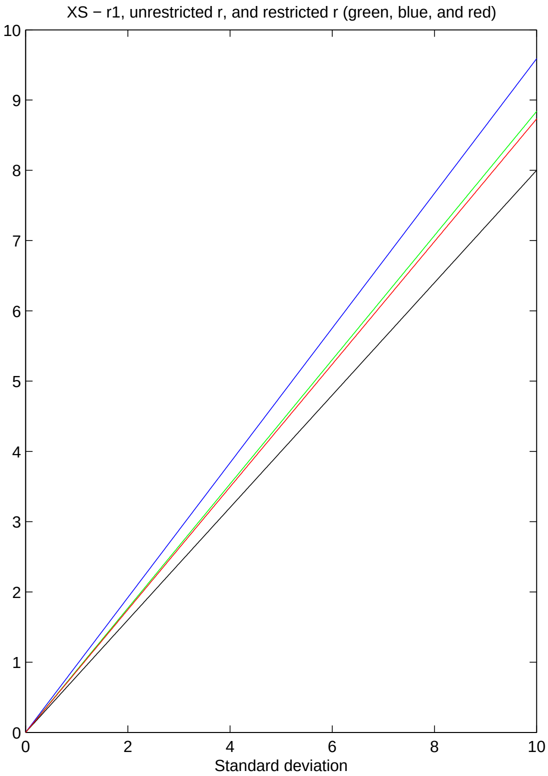

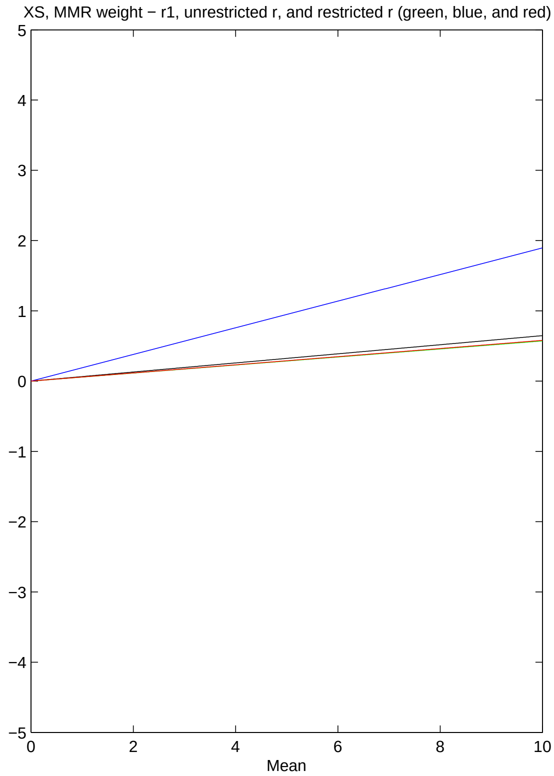

Figures 6a and 6b: Efficiency gains from imposing spanning restrictions on excess returns

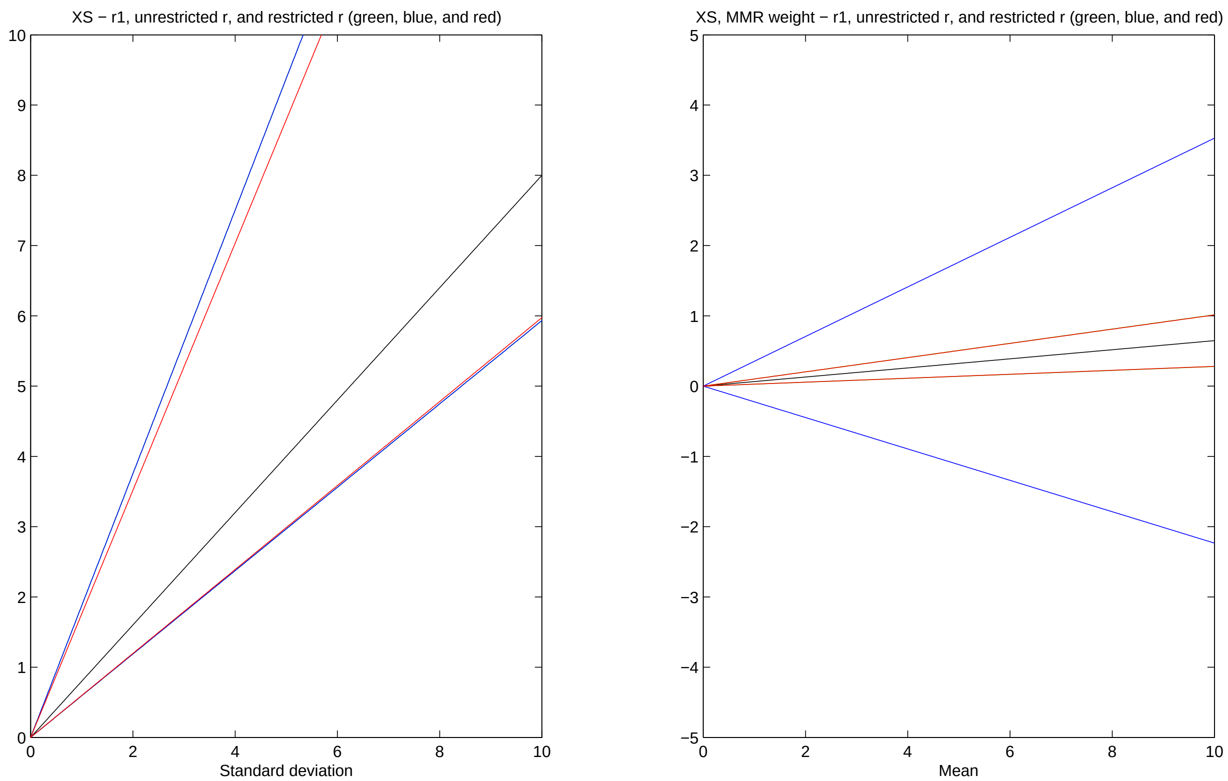

Figure 7: Ensemble of mean-variance frontiers for returns, and optimal weights on the second asset

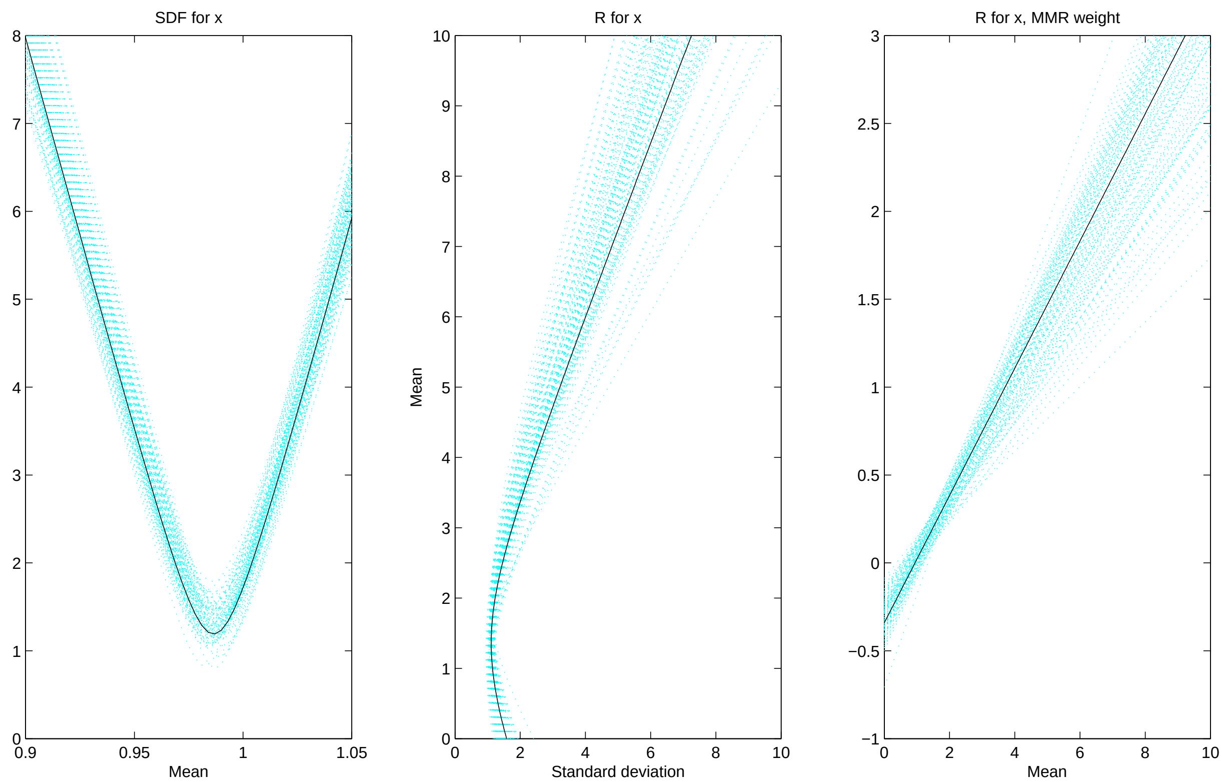

Figure 8: Joint confidence region for three points on the SMVF for gross returns

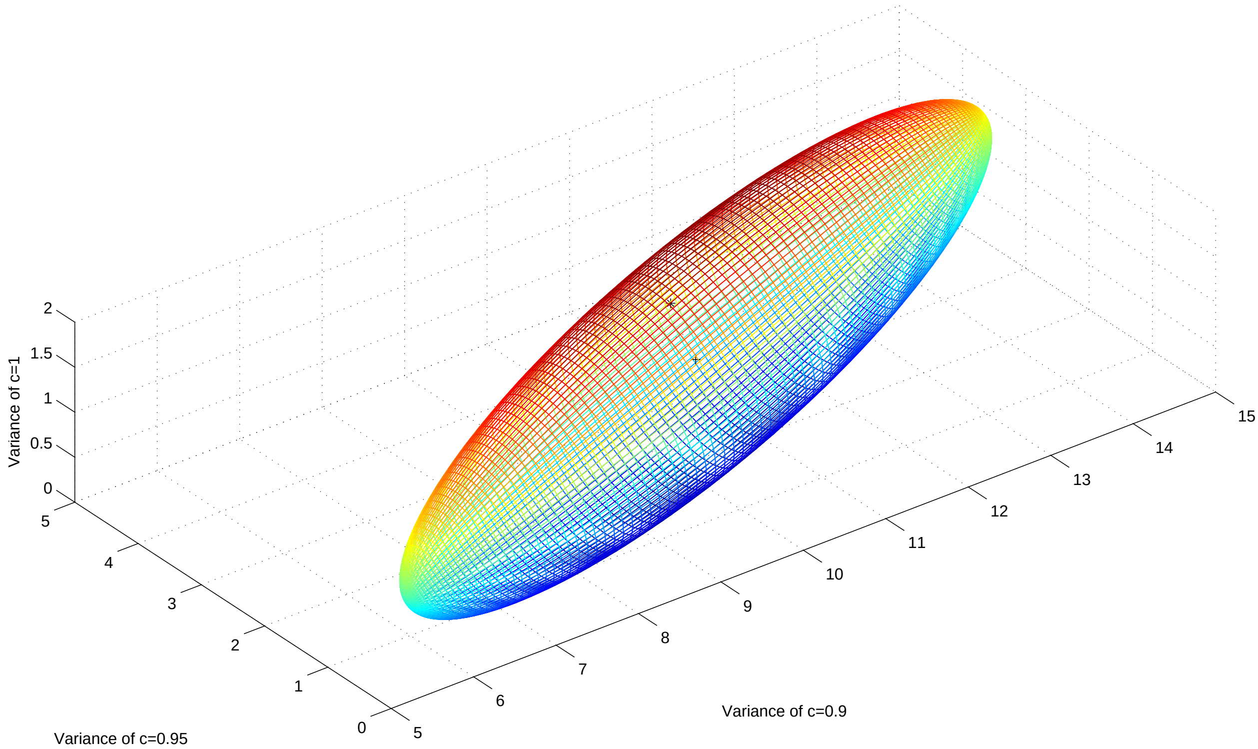

Figures 9a, 9b and 9c: Confidence limits for SMVF, RMVF and optimal weights on the second asset

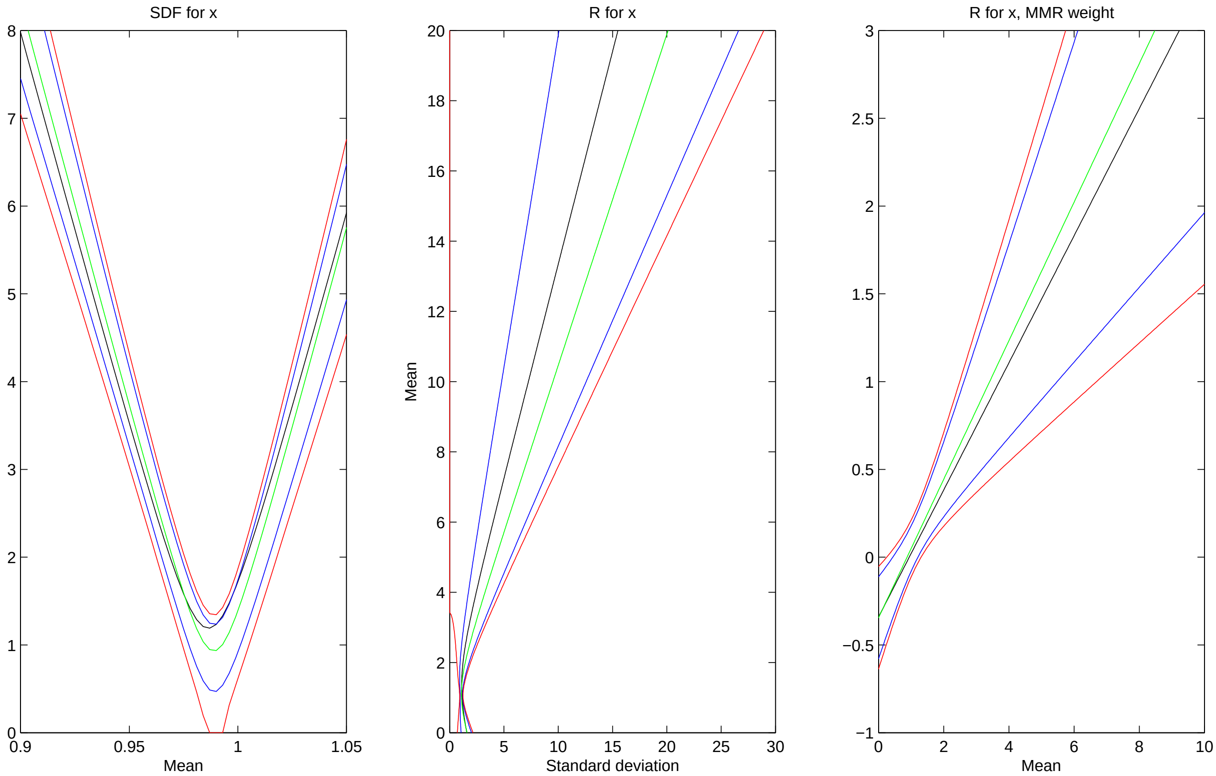

Figure 10a, 10b and 10c: Restricted and unrestricted SMVF, RMVF and optimal weights on the second asset