Duality in Mean-Variance Frontiers with Conditioning Information*

Francisco Peñaranda

Queens College CUNY, 65-30 Kissena Blvd., Flushing, NY 11367, US.

<francisco.penaranda@qc.cuny.edu>

Enrique Sentana

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

<sentana@cemfi.es>

Revised: March 2016

Abstract

Portfolio and stochastic discount factor mean-variance frontiers are usually regarded as dual objects. However, the Hansen and Richard (1987) and Gallant, Hansen and Tauchen (1990) unconditional frontiers are not dual unless some strong conditions hold. We characterise the objects that are always dual to those frontiers, which are not generally proper SDFs or returns. We avoid the common practice of parametrically specifying conditional moments of returns, estimating instead the frontiers with easily implementable sieve methods, which have a managed portfolio interpretation. We empirically assess the validity of SDFs with constant risk prices and the relevance of predictability for portfolio choice.

Keywords: Asset Pricing, Conditional Moment Restrictions, Dynamic Portfolio Strategies, Representing Portfolios, Sieve Minimum Distance, Stochastic Discount Factors.

JEL: C58, G11, G12

*We would like to thank Caio Almeida, Torben Andersen, Tony Berrada, Xiaohong Chen, Wayne Ferson, Lars Hansen, Jens Jackwerth, Ravi Jagannathan, Weiping Li and Cesare Robotti, as well as seminar audiences at Kellogg, UPF, EFA (Ljubljana), ESEM (Budapest), the Nonlinear and Financial Econometrics Conference in honour of A. Ronald Gallant (Toulouse), the Midwest Finance Association Meeting (Orlando) and the Recent Developments in Financial Econometrics and Empirical Finance Conference (Essex) for useful comments and discussions. The suggestions of two anonymous referees have also improved the paper. Julio Crego and Julio Gálvez provided able research assistance for the Monte Carlo simulations. Of course, the usual caveat applies. Financial support from the Spanish Government through grants ECO 2008-03066 and 2011-25607 (Peñaranda) and ECO 2008-00280, 2011-26342 and 2014-59262 (Sentana) is gratefully acknowledged. The first author is also very grateful to SanFI for its funding.

1 Introduction

Mean-variance analysis continues to be widely used in economics and finance, with applications that cover such key issues as portfolio choice, asset pricing tests and performance evaluation. In fact, finance students nowadays learn that there is not just one, but two types of mean-variance frontiers: one for portfolios due to Markowitz (1952), and another one for stochastic discount factors (SDFs) due to Hansen and Jagannathan (1991). They learn that the first frontier characterises the risk-return trade-offs that investor face, and the second one the mean-variance constraints that financial markets data imposes on asset pricing models.[1]

Students also learn that asset returns are predictable, if not in mean at least in variance, and that investors can exploit this fact to their advantage by using conditional distributions in designing their portfolio strategies. For instance, an investor can not only choose a portfolio strategy with constant (value) weights, but also define a trading strategy as a function of yield spreads. As a result, more advanced students learn that there are different versions of the return and SDF mean variance frontiers, depending on the information used in their construction. Frontiers for such active strategies were introduced by Hansen and Richard (1987) in the case of portfolios, and Gallant, Hansen and Tauchen (1990) for SDFs, and were subsequently revisited by Ferson and Siegel (2001, 2003, 2009), Bekaert and Liu (2004) and Abhyankar, Basu and Stremme (2007).[2]

The first contribution of this paper is to characterise the precise relationship between mean-variance frontiers across both these dimensions, namely type (i.e. portfolio vs SDF) and information. This is an important issue because portfolio and stochastic discount factor frontiers are usually regarded as dual objects (in the sense that every element in one frontier is believed to be perfectly correlated with one element in the other frontier) to the extent that sometimes researchers use one type of frontier to answer questions that arise more naturally in the other type. For example, De Santis (1995) and Bekaert and Urias (1996) assess the gains for a mean-variance investor from internationally diversifying her portfolio by testing if the restrictions that domestic market data imposes on asset pricing models are strengthened by the inclusion of data on foreign assets. Similarly, Bekaert and Liu (2004) mention in their concluding remarks that

[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.

[2] See Bansal, Dahlquist, and Harvey (2004) and Brandt and Santa-Clara (2006) for adaptations of the unconditional mean-variance frontier for returns to account for intertemporal hedging in multiperiod portfolio problems.

one interesting application of their SDF framework could be the computation of optimal portfolios. On the other hand, Cochrane (2001, sec. 21.1) uses unconditional Sharpe ratios of traded assets to infer the volatility of the SDF required to explain the equity premium puzzle with a consumption-based asset pricing model. Likewise, Ferson and Siegel (2003) use mean-variance efficient returns to construct SDF bounds.

Nevertheless, the widely cited duality result in Hansen and Jagannathan (1991) applies to their specific set-up: unconditional moments of fixed-weight portfolios based on a given vector of asset payoffs, and it does not usually hold more generally. In particular, their result does not automatically apply to the unconditional return frontier (RF) introduced by Hansen and Richard (1987), which for each level of expected return provides the highest lower bound on the variance of any portfolio whose weights may depend on conditioning information but whose price is always one, and the unconditional SDF frontier (SF) introduced by Gallant, Hansen and Tauchen (1990), which yields the highest lower bound on the variance of SDFs that correctly price any portfolio whose weights may also depend on conditioning information.

We show that these frontiers, which have substantial empirical interest because they rely on unconditional moments, are not dual unless the position of certain frontier portfolios in conditional mean-standard deviation space does not depend on the values of the variables in the information set. Given that these strong conditions are unlikely to hold in practice, empirical researchers willing to take into account conditioning information should be careful, and focus on the type of frontier that is really relevant for the particular question they want to address.

Further, we characterise the SDF-like random variables that are always dual to the elements of the unconditional RF, as well as the return-like random variables that are always dual to the elements of the unconditional SF. Specifically, the dual to the unconditional RF is the extended SF, which provides the highest lower bound on the variance of those univariate random variables that price on average any portfolio whose weights may depend on conditioning information but whose cost is constant. Since these are not necessarily valid SDFs because they may not price correctly portfolios whose cost is a function of the information available at the time of trading, the dual to the unconditional RF does not deliver sufficiently tight constraints on asset pricing models. In turn, the dual of the unconditional SF is the extended RF, which for each level of expected return yields the highest lower bound on the variance of any portfolio with weights that may depend on conditioning information but whose price is only one on average. Therefore, the dual to the unconditional SF does not deliver interesting risk-return trade-offs because its elements are either infeasible or leave money on the table.

Figure 1 summarises our theoretical analysis, which also includes the passive frontiers in Hansen and Jagannathan (1991) and the conditional frontiers in Hansen and Richard (1987) and Gallant, Hansen and Tauchen (1990).

<FIGURE 1>

Our second contribution is a computationally simple yet efficient frontier estimation procedure, which avoids likely misspecifications associated to parametric assumptions on the first two moments of the joint distribution of asset returns given the agents' information set. Specifically, we rely on sieve methods, and in particular on the sieve minimum distance (SMD) semiparametric procedures proposed by Ai and Chen (2003) to estimate the different frontiers in Figure 1 and make inferences about them. Chen and Ludvigson (2009), Nagel and Singleton (2011), and Chen, Favilukis and Ludvigson (2013) are other recent examples of the use of sieve methods in empirical finance.

Importantly, we explicitly relate the mean-variance frontiers that such a procedure generates to the popular empirical strategy of approximating the effect of conditioning information by constructing passive RFs and SFs from managed portfolios, i.e. portfolios whose scale is a function of some variables in the econometrician's information set, as suggested by Hansen and Jagannathan (1991).[3] In doing so, we show that the use of managed portfolios is not necessarily an ad-hoc procedure because they can form the basis of an estimation method with proper statistical foundations. In practice, we choose sieves that have three important implications: the SMD objective function is numerically equivalent to a GMM criterion, the implied non-linear managed portfolios span the linear ones that are common in empirical work, and the weighting matrix is positive definite by construction. In this regard, we work with a continuously updated criterion, which has several valuable invariance properties, and a potentially improved finite sample behaviour.

Our third contribution is the combination of our theoretical results and econometric methods to empirically explore two questions of substantive interest related to the ways in which conditioning information sharpens return and SDF frontiers.

First, we formally test whether the SDF of the popular linear factor pricing model that

[3] These payoffs are also known as "multiplicative" or "scaled" returns (see section 8.1 in Cochrane (2001)).

assumes constant risk prices on the three Fama-French factors belongs to the unconditional SF. Although the evidence against the null is relatively weak if we focus on the linear payoff space of those pricing factors scaled by the price earnings ratio and the default spread, we clearly reject when we use sieves. In simple economic terms, the rejection that we find implies that the risk prices of the Fama-French model, which we can identify with the coefficients of its candidate SDF, should not be time-invariant. We obtain the same conclusion when we study their new five-factor model in Fama and French (2015).

Second, we investigate whether investors would effectively expand their mean-variance opportunity set by considering active portfolio strategies instead of restricting their choices to asset combinations with constant weights. The null hypothesis that the unconditional RF coincides with the standard Markowitz frontier is clearly rejected for both sieve and linear managed portfolios. Therefore, there is added value in exploiting conditioning information for investors who choose portfolios on the unconditional RF.

The rest of the paper is organised as follows. Section 2 studies unconditional frontiers with conditioning information, obtains precise duality conditions and introduces their dual counterparts. Next, we explain how to construct the different frontiers by means of sieve managed portfolios in section 3. Then, we introduce our inference procedures in section 4, and apply them to the empirical questions previously described. Finally, we summarise our conclusions in section 5. Appendix A deals with some important special cases, while proofs and auxiliary results are relegated to a supplemental appendix.

2 Duality Relationships for Unconditional Mean-Variance Frontiers

2.1 Information Structure and Active Portfolio Strategies

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. These payoffs may correspond to stocks, bonds, derivatives, mutual funds, hedge funds, etc. To incorporate conditional information, we closely follow Hansen and Richard (1987), where further details can be found.

Specifically, we assume that there are three important dates in this economy: the decision, trading, and payoff dates. Investors design ex ante portfolio strategies at the decision date which may depend on the information that they will observe at the trading date. Finally, they receive payoffs at the final date. Let G denote the investors' information at the trading date. We will typically think of G as containing one or more signals that are informative about future asset payoffs. We denote the set of all random variables that are measurable with respect to G by I.

In this context, we denote the first two conditional moments of the primitive payoffs and their conditional costs by

$$E(\mathbf{x}|G), \qquad E(\mathbf{x}\mathbf{x}'|G), \qquad C(\mathbf{x}|G), \tag{1}$$

respectively, all of which belong to I. To avoid a trivial uninformative set up, we assume that not all these random variables are degenerate. We also assume that the diagonal elements of $E(\mathbf{x}\mathbf{x}'|G)$ are uniformly bounded with probability one (a.s.), so that a fortiori all the elements of $\mathbf{x}$ belong to L^2, which is the collection of all random variables defined on the underlying probability space with bounded unconditional second moments. Regarding the covariance matrix of $\mathbf{x}$, $Var(\mathbf{x}|G)$, we initially assume its smallest eigenvalue is uniformly bounded away from 0 a.s., which implies that none of the primitive assets is either conditionally riskless or redundant, and moreover, that it is not possible to generate a conditionally riskless portfolio from $\mathbf{x}$ other than the trivial one.

Although we deliberately allow asset prices $C(\mathbf{x}|G)$ to depend on the values of the signals, there are two important examples of payoffs whose costs are non-random: gross returns, which are payoffs with unit prices, and excess returns or arbitrage portfolios, which are payoffs of zero cost. For simplicity, though, we exclude the possibility that all primitive assets are arbitrage portfolios by assuming that the vector $C(\mathbf{x}|G)$ has at least one entry different from 0 a.s. We also assume that not all expected payoffs are conditionally proportional to their prices with a common factor of proportionality. In this way, we implicitly rule out those situations in which all conditionally expected returns are the same.[4]

We denote the unconditional counterparts to (1) as

$$E(\mathbf{x}) = E[E(\mathbf{x}|G)], \qquad E(\mathbf{x}\mathbf{x}') = E[E(\mathbf{x}\mathbf{x}'|G)], \qquad C(\mathbf{x}) = E[C(\mathbf{x}|G)],$$

which are now real numbers instead of random variables. Following Hansen and Richard (1987), we will sometimes use the term pseudo-prices to refer to average costs.

[4]The special cases of a riskless asset, zero-cost portfolios, and equal expected returns can also be analysed in our set up, but for pedagogical reasons we postpone their discussion to Appendix A. In the case of a riskless asset, in particular, we show that the geometric interpretation of duality in terms of Sharpe ratios in Hansen and Jagannathan (1991) only applies to their specific set-up (i.e. unconditional moments of passive strategies), so that one must again be careful in extending their result to other contexts.

As we said before, investors can condition their portfolios weights on the information they know they will have at the time of trading, which is given by G. For instance, investors may prefer different portfolios depending on whether yield spreads at the trading date are high or low. Consequently, they can construct portfolio strategies with payoffs $p = \mathbf{x}'\mathbf{w}$, where the portfolio weights $\mathbf{w} \in I$. In what follows, we will refer to the conditional span of $\mathbf{x}$ as the payoff space P. In this context, $\mathbf{w} \in I$ indicates an *active* portfolio strategy, while a vector of constant weights $\mathbf{w} \in \mathbb{R}^N$ indicates a *passive* portfolio.

Finally, it will be convenient for our purposes to express an arbitrary active portfolio $p = \mathbf{x}'\mathbf{w}$ as

$$p = Rw + \mathbf{r}'\mathbf{w}_{-1}, \qquad \mathbf{r} = \mathbf{x}_{-1} - RC(\mathbf{x}_{-1}|G), \tag{2}$$

where the subscript -1 means that we have deleted the first element of the corresponding vector, R is the gross return on the first asset (which we can assume has a non-zero price without loss of generality), and the vector $\mathbf{r}$ contains the remaining asset payoffs transformed into excess returns. Thus, we can establish a direct connection between the weight on R and the portfolio cost because $C(p|G) = w$.

2.2 Representing Portfolios and Stochastic Discount Factors

Hansen and Richard (1987) introduce a conditional analogue to a standard Hilbert space based on the mean square inner product, $E(xy|G)$, and the associated mean square norm $\sqrt{E(x^2|G)}$, where x, y belong to the conditional analogue to L^2. Such a topology allows them to define the conditional least squares projection of any y onto P as

$$E(y\mathbf{x}'|G)E^{-1}(\mathbf{x}\mathbf{x}'|G)\mathbf{x}, \tag{3}$$

which is the element of P that is closest to y in the conditional mean square norm.

In this context, we can formally understand $C(\cdot|G)$ and $E(\cdot|G)$ as conditionally continuous linear functionals that map the elements of P onto I. The expected value functional is always conditionally continuous on the conditional analogue to L^2 by a conditional version of the Markov inequality. Similarly, our full rank assumption on $Var(\mathbf{x}|G)$ implies that $E(\mathbf{x}\mathbf{x}'|G)$ has full rank too, and consequently, that the cost functional is also conditionally continuous on P, which is tantamount to the law of one price. A conditional version of the Riesz representation theorem then implies that there exist two unique elements of P that represent these conditional functionals over P.[5] As we shall see, the corresponding representing portfolios will be the basis of the SDF mean-variance frontiers discussed in the next section.

The conditional mean and cost representing portfolios, p° and p^*, respectively, will be such that:

$$E(p|G) = E(p^\circ p|G), \qquad C(p|G) = E(p^* p|G), \quad \forall p \in P. \tag{4}$$

It is then straightforward to show that

$$p^\circ = \mathbf{x}' E^{-1}(\mathbf{x}\mathbf{x}'|G) E(\mathbf{x}|G), \qquad p^* = \mathbf{x}' E^{-1}(\mathbf{x}\mathbf{x}'|G) C(\mathbf{x}|G). \tag{5}$$

If P included the conditionally (and unconditionally) safe payoff $x_0 = 1$, then p° would coincide with it. But even though it does not, it follows from (3) that we can interpret p° as the portfolio that "mimics" the safe asset with the minimum "tracking error". We can also use (3) to interpret p^* as the conditional projection of any valid SDF onto P. As is well known, an SDF is any scalar random variable m in the conditional analogue to L^2 that prices all conceivable payoffs in terms of their expected cross product with it. More formally,

$$E(mp|G) = C(p|G), \quad \forall p \in P. \tag{6}$$

Equivalently, admissible SDFs are fully characterised by the condition

$$E(m\mathbf{x}|G) = C(\mathbf{x}|G).$$

In addition, since $C(x_0|G) = E(1 \cdot m|G)$, the conditionally expected value of m defines the shadow price of the unit payoff.

Expression (5) may suggest that one would need the conditional moments of returns to obtain the representing portfolios above, and thereby, the mean-variance frontiers. However, Hansen and Richard (1987) show that representing portfolios and SDFs can be defined in terms of unconditional moments too. Specifically, the law of iterated expectations implies that p° and p^* also represent unconditional means and average costs on the active payoff space P, so that:

$$E(p) = E(p^\circ p), \qquad C(p) = E(p^* p), \quad \forall p \in P. \tag{7}$$

Similarly, we could also define SDFs as those m that give the right pseudo-price for any conceivable p, i.e.

$$E(mp) = C(p), \quad \forall p \in P.$$

[5] Chamberlain and Rothschild (1983) introduced mean and cost representing portfolios to study unconditional mean-variance analysis in infinite dimensional payoff spaces in which information plays no role. Hansen and Richard (1987) extended their results to conditioning information.

Therefore, there is no loss of information in moving from pricing to pseudo-pricing, but only as long as we focus on the whole of P, and not simply on a subset. We will exploit this result in section 3 to compute the representing portfolios without a parametric model of conditional moments.

From (p^*, p°), we can construct a pair of constant cost portfolios that will play a crucial role in the definition of return frontiers in the next section. The first one is the *gross* return associated to the cost representing portfolio

$$R^* = p^*/C(p^*|G) = p^*/E(p^{*2}|G), \tag{8}$$

which has the minimum second moment among all possible unit cost portfolios. If we express the vector $\mathbf{x}$ in terms of a gross return R and $N-1$ excess returns $\mathbf{r}$, as in (2), then we can write (8) as

$$R^* = R - r^*, \tag{9}$$

where

$$r^* = \mathbf{r}'E^{-1}(\mathbf{rr}'|G)E(\mathbf{r}R|G)$$

is the arbitrage portfolio that represents the expected cross-product with R in the space of zero-cost portfolios.

The second portfolio is the mean representing portfolio in the space of zero-cost portfolios,

$$r^\circ = \mathbf{r}'E^{-1}(\mathbf{rr}'|G)E(\mathbf{r}|G).$$

which is the *excess* return that "mimics" the safe asset with the minimum "tracking error". This portfolio achieves the maximum conditional Sharpe ratio among all arbitrage portfolios, namely $\sqrt{E(r^\circ|G)/[1-E(r^\circ|G)]}$.

Given that we can construct r° as the following portfolio of (p^*, p°)

$$r^\circ = p^\circ - p^*C(p^\circ|G)/C(p^*|G) = p^\circ - R^*E(p^*|G), \tag{10}$$

we can combine (8) and (10) to trivially recover (p^*, p°) from (R^*, r°) as

$$p^* = R^*/E(R^{*2}|G),$$
$$p^\circ = r^\circ + p^*E(R^*|G).$$

Importantly, though, the linear transformation relating both portfolio pairs generally depends on the available information, and therefore its practical implementation requires active strategies.

As we will see in section 2.4, the duality between return and SDF frontiers crucially depends on the active or passive nature of this relationship.

2.3 Unconditional Return and SDF frontiers

The first column of Figure 1 refers to those active portfolio strategies that are optimal with respect to conditional first and second moments, which we characterise in detail in supplemental appendix C. They are called Conditional Return and SDF Frontiers (RF and SF respectively), and they are dual objects in the sense that every element in one frontier has perfect conditional correlation with one element in the other frontier, with two exceptions whose geometric interpretation we also provide in the same appendix.

In this section we focus instead on those active portfolio strategies that are optimal with respect to unconditional moments, which correspond to columns 2 and 3 in Figure 1. At first sight, it might seem odd to study unconditional moments when we think of active strategies whose weights depend on conditioning information. However, in many practical situations the observer of the agents' decisions only has access to an information set that is much coarser than the agents' information set. The performance evaluation of a portfolio manager by means of the first two unconditional moments of her returns is a typical example of the use of unconditional return frontiers by an outside evaluator who may not have access to the proprietary strategies followed by the manager.

Hansen and Richard (1987) define the Unconditional RF as the *highest* lower bound on the variance for each level of expected return that can be achieved by portfolios with weights that may depend on conditioning information, but whose price is always one. Thus, the unconditional RF will be given by the set of active portfolio strategies that solve the problem

$$\min_{p\in P} E(p^2) \qquad s.t. \qquad E(p) = \nu \in \mathbb{R}, \qquad C(p|G) = 1. \tag{11}$$

Hansen and Richard (1987) show that the gross returns that solve (11) correspond to the following passive portfolio of (R^*, r°):

$$p_U(\nu) = R^* + \omega_U(\nu) r^\circ, \tag{12}$$

where the constant $\omega_U(\nu)$ guarantees that the constraint $E[p_U(\nu)] = \nu$ is satisfied and the unit weight on R^* guarantees that $p_U(\nu)$ has unit cost. As expected, the unconditional RF is a hyperbola in unconditional mean-standard deviation space for returns.

In turn, Gallant, Hansen and Tauchen (1990) define the Unconditional SF as the *highest* lower bound on the unconditional variance of the SDFs defined in (6), which correctly price all the active portfolios in P. Thus, the unconditional SF will be given by the set of scalar random variables that solve the optimisation problem

$$\min_{m \in L^2} E(m^2) \qquad s.t. \qquad E(m) = c \in \mathbb{R}, \qquad E(m\mathbf{x}|G) = C(\mathbf{x}|G). \tag{13}$$

Gallant, Hansen and Tauchen (1990) go on to show that the solution to (13) can be written as a constant plus a passive portfolio of (p^*, p°):

$$m_U(c) = p^* + \varpi_U(c)[1 - p^\circ], \tag{14}$$

where the constant $\varpi_U(c)$ guarantees that the constraint $E[m_U(c)] = c$ is satisfied while the unit weight on p^* guarantees the correct pricing of payoffs. Not surprisingly, the unconditional SF also takes the shape of a hyperbola in unconditional mean-standard deviation space for SDFs.

Given the close analogy between these unconditional frontiers and the conditional frontiers in the first column of Figure 1, one is tempted to conclude that the unconditional RF and SF are also dual objects, in the sense that every element in one frontier has perfect unconditional correlation with another element in the other frontier. However, this is not true in general, as the next simple example illustrates. Natural candidates for duality would be p^*, which belongs to the unconditional SF, and its return R^*, which belongs to the unconditional RF. However, they do not have perfect unconditional correlation unless $E(R^{*2}|G)$ is constant. The next section characterises the conditions required to obtain pairs of points in those frontiers with perfect correlation.

2.4 Duality Conditions for Unconditional Frontiers

The relationship between the unconditional RF and SF is easier to understand if we respectively re-write (12) and (14) as

$$p_U(\nu) = [E(R^{*2}|G) - \omega_U(\nu)E(R^*|G)]p^* + \omega_U(\nu)p^\circ \tag{15}$$

and

$$m_U(c) = [E(p^{*2}|G) - \varpi_U(c)E(p^*|G)]R^* + \varpi_U(c)(1 - r^\circ). \tag{16}$$

On this basis, we can characterise the duality between the unconditional RF and SF:

Proposition 1 *Let ν and c denote some specific means for the unconditional RF and SF, respectively, whose elements are characterised in (12) and (14). Then:*

1. *An element of the unconditional RF has perfect unconditional correlation with some element of the unconditional RF if and only if*

$$E(R^{*2}|G) - \omega_U(\nu)E(R^*|G)$$

is a nonzero constant, while an element of the unconditional SF has perfect unconditional correlation with some element of the unconditional RF if and only if

$$E(p^{*2}|G) - \varpi_U(c)E(p^*|G)$$

is a nonzero constant.

2. *Two elements of the unconditional RF have perfect unconditional correlation with two elements of the unconditional SF if and only if*

$$E(R^{*2}|G) = E(R^{*2}) \text{ and } E(R^*|G) = E(R^*),$$

or equivalently

$$E(p^{*2}|G) = E(p^{*2}) \text{ and } E(p^*|G) = E(p^*),$$

in which case the entire frontiers will be dual.

The first part of Proposition 1 shows that duality at a specific point requires an affine relationship with constant coefficients between the first two conditional moments of p^* or equivalently, the first two conditional moments of R^*. In turn, the second part of Proposition 1 states that the unconditional RF and SF are fully dual if and only if the first two conditional moments of p^* and R^* are constant, in which case the location of p^* and R^* in conditional mean-variance space will be constant too. Strictly speaking, though, there will be two duality exceptions: the element on the unconditional RF for which $E(R^{*2}) - \omega_U(v)E(R^*) = 0$, and the element on the unconditional SF for which $E(p^{*2}) - \varpi_U(c)E(p^*) = 0$. Nevertheless, in both cases we can establish a link between an element of one frontier and the asymptotes of the other. See supplemental appendix C for a detailed explanation of this link in the context of conditional frontiers.

In contrast, the location of p° and r° in conditional mean-variance space is unaffected by the conditions in this proposition. In other words, the duality restrictions constrain the time-variation of the minimum second moment portfolio in conditional mean-variance space but not the time-variation of the maximum Sharpe ratio (and hence the slopes) of the conditional RF.

Still, given that the strong conditions in Proposition 1.2 are unlikely to hold in practice, empirical researchers who wish to take into account conditioning information should be careful, and focus on the type of frontier that is really relevant for the particular question they want to address, either investors' risk-return trade-offs or constraints on asset pricing models.

2.5 Extended Return and SDF frontiers

As we have seen in the previous section, the elements of the unconditional RF and SF are not generally dual to each other, so it is interesting to characterise the SDF-like random variables that are always dual to the elements of the unconditional RF, as well as the return-like random variables that are always dual to the elements of the unconditional SF. Not very surprisingly, the dual variables also solve mean-variance problems based on unconditional moments, but with milder restrictions.

Let us start with the dual to the unconditional SF. Expression (16) implies that the weight of its elements on R^* is not generally constant, which motivates the definition of a mean-variance problem with a weaker cost constraint.

Let us define extended returns as portfolios with unitary average cost,[6] so that

$$C(p) = E(g) = 1,$$

where

$$g = C(p|G) = C(\mathbf{x}|G)'\mathbf{w}.$$

Similarly, we can also define extended arbitrage portfolios as those that satisfy $C(p) = E(g) = 0$.

By analogy with the unconditional RF discussed in the previous section, we define the Extended RF as the *highest* lower bound on the variance for each level of expected return that can be achieved by portfolios in P whose *pseudo price* is one. More formally, the extended RF is the set of portfolio strategies that solve the problem

$$\min_{p \in P} E(p^2) \qquad s.t. \qquad E(p) = \nu \in \mathbb{R}, \qquad C(p) = 1, \tag{17}$$

which is an unconditional mean-variance problem similar to (11), but in the space of extended returns. Then, we can show that:

Proposition 2 *The solution to program (17) is given by*

$$p_E(\nu) = g(\nu)R^* + \omega_E(\nu)r^\circ = \left[\frac{1-\omega_E(\nu)E(p^*)}{E(p^{*2})}\right]p^* + \omega_E(\nu)p^\circ, \tag{18}$$

$$g(\nu) = \frac{E(p^{*2}|G)}{E(p^{*2})} + \omega_E(\nu)\left[E(p^*|G) - \frac{E(p^{*2}|G)}{E(p^{*2})}E(p^*)\right], \tag{19}$$

where the constant $\omega_E(\nu)$ guarantees that the constraint $E[p_E(\nu)] = \nu$ is satisfied.

[6] Hansen and Richard (1987) refer to them as pseudo-returns.

The main difference between expressions (12) and (18) is that in the latter the cost $g(\nu)$ is not systematically one, although it is one on average.

The extended RF, which is also a hyperbola in mean-standard deviation space, will necessarily be to the left of the unconditional RF because (17) has the same objective function as (11) but with less demanding cost constraints. This can be clearly seen in the left panel of Figure 2, which uses the data and methodology of the empirical application in section 4.

<FIGURE 2>

Nevertheless, the relative position of the extended RF does not really reflect an improvement in investment opportunities relative to the unconditional RF because extended RF portfolios are either infeasible or leave money on the table.

The following result characterises the element-by-element duality between the extended RF and unconditional SF:

Proposition 3 *Let ν and c denote some specific means for the extended RF and unconditional SF, respectively, whose elements are characterised in (18) and (14). Then:*

1. *Any element of the unconditional SF such that $E(p^{*2}) - \varpi_U(c)E(p^*) \neq 0$ has perfect unconditional correlation with some element of the extended RF.*
2. *Any element of the extended RF such that $1 - \omega_E(\nu)E(p^{*2}) \neq 0$ has perfect unconditional correlation with some element of the unconditional SF.*

Intuitively, the duality between the unconditional SF and the extended RF derives from the fact that the second formula in (18) expresses $p_E(\nu)$ as having fixed-weights on (p^*, p°), exactly like the elements of the unconditional SF in (14). Proposition 3 shows that save for the two stated exceptions, every element on the unconditional SF is perfectly correlated with another element on the extended RF and vice versa. Nevertheless, it is important to emphasise once again that while the unconditional SF delivers the optimal constraints on asset pricing models, the extended RF is useless from the vantage point of an investor.

Let us now turn to the dual frontier to the unconditional RF. Expression (15) implies that the weight of the unconditional SF on p^* is not generally constant, which motivates the definition of a mean-variance problem with a weaker pricing constraint.

Let us focus on constant conditional cost portfolios by defining the restricted payoff space $P_c \subset P$ as

$$P_c = \{p \in P : C(p|G) = C(p)\},$$

which includes both gross and excess returns. Expression (2) allows us to clarify the constraint that a constant cost imposes on active strategies. Specifically, while the active payoff space P does not impose any constraint on the dependence of w and $\mathbf{w}_{-1}$ on the information in G, the constant-cost payoff space P_c imposes the restriction that w belongs to $\mathbb{R}$.

In this context, we define *extended* SDFs as those random variables $m \in L^2$ that price correctly on average any payoff that belongs to the constant-cost payoff space:

$$E(mp) = C(p), \quad \forall p \in P_c.$$

Given that (6) implies that proper SDFs satisfy an analogous condition for the richer set of payoffs in P, extended SDFs will not price correctly portfolios whose cost is not constant. The following lemma provides an equivalent characterisation for extended SDFs:

Lemma 1 *Extended SDFs are fully characterised by the condition*

$$\begin{aligned} E(m\mathbf{x}|G) &= hC(\mathbf{x}|G), \qquad h \in I \\ E(h) &= 1. \end{aligned}$$

Therefore, $E(m\mathbf{r}|G) = \mathbf{0}$ but $E(mR|G) = h$, so that the only assets that extended SDFs price correctly are zero cost portfolios. For all other assets, the ratios of extended SDFs' prices to actual prices will be equal across portfolios because h is a scalar random variable associated to m but not to $\mathbf{x}$.

By analogy with the unconditional RF discussed in section 2.3, we can now define the Extended SF as the *highest* lower bound on the variance of those univariate random variables that price correctly on average any portfolio of $\mathbf{x}$ whose weights may depend on conditioning information, *but whose cost is constant.* Using Lemma 1, we can formally characterise the extended SF as the set of scalar random variables m that solve the problem

$$\min_{m \in L^2} E(m^2) \qquad s.t. \qquad E(m) = c \in \mathbb{R}, \qquad E(m\mathbf{x}|G) = hC(\mathbf{x}|G). \tag{20}$$

Then, we can show:

Proposition 4 *The solution to program (20) is given by*

$$m_E(c) = h(c)p^* + \varpi_E(c)[1 - p^\circ] = \left[\frac{1 - \varpi_E(c)E(R^*)}{E(R^{*2})}\right] R^* + \varpi_E(c)[1 - r^\circ], \tag{21}$$

$$h(c) = \frac{E(R^{*2}|G)}{E(R^{*2})} + \varpi_E(c)\left[E(R^*|G) - \frac{E(R^{*2}|G)}{E(R^{*2})}E(R^*)\right], \tag{22}$$

where the constant $\varpi_E(c)$ guarantees that the constraint $E[m_E(c)] = c$ is satisfied.

The main difference between expressions (14) and (21) is that in the latter the mispricing factor $h(c)$ is not systematically one, only on average.

The extended SF, which is also a hyperbola in mean-standard deviation space, will necessarily be below the unconditional SF because (20) has the same objective function as (13) but with less demanding pricing constraints. This can be clearly seen in the right panel of Figure 2.

The following result characterises the element-by-element duality between the unconditional RF and extended SF:

Proposition 5 *Let ν and c denote some specific means for the unconditional RF and extended SF, respectively, whose elements are characterised in (12) and (21). Then:*

1. *Any element of the extended SF such that $1 - \varpi_E(c)E(R^*) \neq 0$ has perfect unconditional correlation with some element of the unconditional RF.*
2. *Any element of the unconditional RF such that $E(R^{*2}) - \omega_U(v)E(R^*) \neq 0$ has perfect unconditional correlation with some element of the extended SF.*

Intuitively, the duality between the unconditional SF and the extended RF derives from the fact that the second expression in (21) shows that $m_E(c)$ has fixed-weights on (R^*, r°), exactly like the elements of the unconditional RF in (12). Once again, Proposition 5 shows that save for the usual two exceptions, every element on the unconditional RF is perfectly correlated with another element on the extended SF and vice versa. Nevertheless, while the unconditional RF characterises the optimal unconditional risk-return trade-offs that an investor faces, the extended SF represents a loss of asset pricing information relative to the unconditional SF because it only provides suboptimal (but valid) constraints on asset pricing models.

We can also explicitly link the existence of a single dual point between the unconditional RF and SF in Proposition 1.1 to tangency between unconditional and extended frontiers. Specifically, the unconditional RF and extended RF are tangent if and only if there is a unique $\bar{\nu}$ such that $g(\bar{\nu}) = 1$, while the unconditional SF and extended SF are tangent if and only if there is a unique $\bar{c}$ such that $h(\bar{c}) = 1$. One tangency implies the other,[7] and both tangencies are dual points, unless the tangency point from which we start coincides with one of the duality exceptions described in Propositions 3 and 5.

In turn, the full duality of the unconditional RF and SF in Proposition 1.2 can be linked to the unconditional and extended frontiers being equal. Specifically, given that equality between

[7]Note that if G is given by a signal that can only take two values, then there will be at least one tangency. To see why, let us focus on $g(\nu)$, which can only take two values in this case. We can choose $\bar{\nu}$ such that $g(\bar{\nu}) = 1$ for one of the values of the signal. Given that $E[g(\nu)] = 1$ by construction, it follows that $g(\bar{\nu}) = 1$ for the other signal value. A similar argument applies to $h(c)$.

those frontiers is equivalent to

$$p^\circ \in P_c \text{ and } p^* \in P_c, \tag{23}$$

the necessary and sufficient condition for duality between the unconditional RF and SF is that the two representing portfolios that unconditionally span the unconditional SF must have constant cost.

Finally, Proposition B1 in supplemental appendix B characterises the points of minimum distance between the unconditional and extended RFs on the one hand, and the unconditional and extended SFs on the other, which are not generally dual to each other. Figure 2 displays the points of minimum distance for the depicted frontiers.

3 Mean-Variance Frontiers with Sieve Managed Portfolios

The unconditional mean-variance frontiers discussed in the previous section seem to require the correct specification of the first two conditional moments of asset returns because they are constructed from the representing portfolios (p^*, p°) and (R^*, r°), whose information-dependent weights are defined in (5), (8) and (10). As a result, it seems rather natural to parametrically specify those conditional moments, as Bekaert and Liu (2004) and Ferson and Siegel (2003) did (see supplemental appendix D for further details). However, parametric models are often restrictive and their results sensitive to deviations from the chosen specifications.

Estimating those first and second moments by means of semi- or non-parametric procedures should provide a more flexible and robust approach, but they effectively introduce a huge number of parameters, which can lead to numerical problems. For example, the semi-non-parametric method in Gallant, Hansen, and Tauchen (1990) requires the estimation of the conditional distribution of the vector of returns, a task which is particularly complicated with multiple assets. In addition, even if one could satisfactorily deal with the numerical problems, the resulting estimators may require very large samples to provide reliable inferences.

Given these difficulties, it is perhaps not surprising that many empirical studies rely on constant weight strategies of linear managed portfolios, an approach which is often regarded as an ad-hoc way of approximating columns 2 and 3 of Figure 1 (see chapter 8 in Cochrane (2001)). As we shall show below, though, suitably selected managed portfolios can provide the basis for an efficient estimation method with proper statistical foundations.

In addition, given that (p^*, p°) and (R^*, r°) are not only the building blocks of the unconditional SF and RF, respectively, but also of the conditional SF and RF in the first column of Figure 1, such managed portfolios will also be useful for researchers interested in conditional mean variance frontiers for SDFs and returns, as we explain in supplemental appendix C.

For the sake of clarity, though, we start with the last column of Figure 1, which is associated to frontiers that do not use information at all, and postpone to section 3.2 the discussion of those situations in which investors *implicitly* use the information available at the trading date in constructing their portfolio weights, which corresponds to the fourth and fifth columns of the same figure.

3.1 Passive Mean-Variance Frontiers for $\mathbf{x}$

Given a vector of asset payoffs $\mathbf{x}$, its unconditional linear span $\langle \mathbf{x} \rangle$, which is the space of *passive* portfolios with payoffs $p = \mathbf{x}'\mathbf{w}$, $\mathbf{w} \in \mathbb{R}^N$, will be a subset of the payoff space P. In this context, Hansen and Jagannathan (1991) define a frontier that we will label as the Passive SF for $\mathbf{x}$, which puts the *highest* variance bound on those univariate random variables that *pseudo price* any portfolio $p \in \langle \mathbf{x} \rangle$. More formally, such a passive SF is given by the set of scalar random variables that solve a variant of (13) in which the pricing conditions hold on average instead of conditionally. These random variables, though, are generally *passive* SDFs, and not necessarily valid SDFs, since they may not price correctly portfolios whose weights depend on information.

Hansen and Jagannathan (1991) also define a dual frontier to the passive SF discussed in the previous paragraph, which we will label as the Passive RF for $\mathbf{x}$, such that any element of the passive SF has perfect unconditional correlation with some element of the passive RF, with the usual two exceptions. Formally, such a passive RF will be made up of portfolio strategies that solve a problem analogous to (17) but defined over $\langle \mathbf{x} \rangle$ instead of P. However, the elements of this passive RF will generally be extended returns instead of returns since the cost constraint $C(p) = 1$ is stated as an average, while the prices of the asset payoffs under consideration, $C(\mathbf{x}|G)$, may depend on the information available at the time of trading.

This passive RF will differ from the usual Markowitz frontier for returns, unless $\mathbf{x}$ is effectively an $N \times 1$ vector of *constant cost* payoffs. In what follows, we assume that this is indeed the case to simplify the exposition. For the same reason, but without loss of generality, we will express the vector $\mathbf{x}$ in terms of a gross return R and an $(N-1) \times 1$ vector of excess returns $\mathbf{r}$, as in (2).

The passive RF for $(R, \mathbf{r}')$ will be a constrained version not only of the extended RF but also of the unconditional RF because any extended return in $\langle R, \mathbf{r}' \rangle$ will also be a return. Therefore, we will come across the extended RF, the unconditional RF and the Markowitz frontier as we go from left to right on $[\sqrt{Var(p)}, E(p)]$ space, as in the left panel of Figure 3, which is also constructed with the data used in the empirical application. Thus, we can understand the passive RF obtained from a constant cost $\mathbf{x}$ as providing a lower bound on the actual risk-return trade-offs that investors face, which are described by the unconditional RF, not the extended RF.

<FIGURE 3>

In turn, the passive SF for $(R, \mathbf{r}')$ will be a constrained version of the extended SF because $\langle \mathbf{x} \rangle$ is also a subset of P_c in this special case. Therefore, as we move upwards on $[\sqrt{Var(p)}, E(p)]$ space we will come across this passive SF, the extended SF and finally the unconditional SF, as the right panel of Figure 3 illustrates.

3.2 From Passive to Unconditional Frontiers

3.2.1 From Passive to Unconditional SF

The payoff space $\langle \mathbf{x} \rangle$ when $\mathbf{x} = (R, \mathbf{r}')'$ may be too narrow relative to P, which is the relevant space of strategies available to investors. For that reason, Hansen and Jagannathan (1991) also relied on an alternative empirical approach based on the *linear* managed portfolios

$$\begin{pmatrix} 1 \\ \mathbf{z} \end{pmatrix} \otimes \mathbf{x}, \tag{24}$$

which scale the vector $\mathbf{x}$ by some variables $\mathbf{z}$ in I. Their approach corresponds to the penultimate column of Figure 1. These "multiplicative" or "scaled" returns are no longer proper returns since their true cost varies with the values of the signals. As far as the unconditional and extended frontiers discussed in the previous sections are concerned, though, the use of $\mathbf{x}$ or (24) leads to the same answer because they do not enlarge the payoff spaces P_c and P.

In contrast, the unconditional span of (24) nests $\langle \mathbf{x} \rangle$, but is nested by P. As a result, in $[\sqrt{Var(p)}, E(p)]$ space, a tighter passive RF constructed in this way will lie between the extended RF and the Markowitz frontier, and could cross the unconditional RF. Hence, such a passive RF is not a relevant object for an investor because it is not constructed from constant cost payoffs, neither can it be used to place a bound on the unconditional RF. Similarly, a sharpened passive SF constructed from (24) will be between the unconditional SF and the passive SF for $\mathbf{x}$ in $[E(m), \sqrt{Var(m)}]$ space, and might even cross the extended SF. Figure 3 displays these tighter passive frontiers in our data, which can be identified by the subscript 1.

Often, though, linear managed portfolios only provide a rough approximation to P. Nevertheless, if we were able to take into account all possible managed portfolios, then the corresponding unconditional span would be equal to P. A sharper passive RF constructed with all managed portfolios instead of the initial $\mathbf{x}$ would coincide with the extended RF in (17), while its dual passive SF would coincide to the unconditional SF in (13). In other words, the fifth column in Figure 1 would coincide with the third column. In practice, we could consider an increasing sequence of managed portfolios such that the fifth column in Figure 1 would converge to the third column. Obviously, the corresponding sequence of dual passive frontiers would cross the unconditional RF/extended SF at some point.

Our empirical strategy inspired by Ai and Chen (2003) uses sieve methods as a proper non-parametric procedure to achieve this goal. For a given original space, these methods rely on a sequence of less complex approximating spaces (see Chen (2007) for a survey of sieve methods in econometrics). Let $\mathbf{b}^{k_T}(\mathbf{z})$ denote a vector of known sieve basis functions (power series, splines, Fourier series, etc.) with the property that its linear combinations can approximate any square integrable real-value function of $\mathbf{z}$ as the smoothing parameter k_T increases. In this context, we can identify P with the (closure of the union over k_T of the) unconditional spans of

$$\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{x}, \tag{25}$$

so that these sieve managed portfolios identify (p^*, p°) through the unconditional moments (7).[8] Specifically, p° can be identified as $(\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{x})'\boldsymbol{\varphi}^{k_T}$ from the just identified moment conditions

$$E\{(\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{x})[(\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{x})'\boldsymbol{\varphi}^{k_T} - 1]\} = \mathbf{0}, \tag{26}$$

which are equivalent to the conditional moments that identify p° in (4) as k_T grows. A similar argument applies to p^*. Thus, it is indeed possible to reproduce the conditional and unconditional RF and SF to any desired degree of accuracy by means of suitably selected managed portfolios without modelling the first two conditional moments of asset returns. In this regard, the main numerical advantage of our proposed procedure is that since the approximating spaces are characterised by a finite number of parameters, sieve methods effectively reduce the estimation problem to a parametric one. Nevertheless, the quality of the approximation that can be

[8] In the usual situation in which $\mathbf{b}^{k_T}(\mathbf{z})$ spans $(1, \mathbf{z}')$, the PSF frontier constructed from (25) cannot be below the one obtained from (24). This does not need be the case when we construct the unconditional SF from a model of conditional moments. For example, the SDF bounds obtained from linear managed portfolios by Ferson and Siegel (2003) are higher than their estimates of the unconditional SF and extended SF, which they attribute to sampling error or misspecification.

realistically achieved in practice will depend on the sample size.

3.2.2 From Passive to Unconditional RF

Imagine now that we restrict the managed portfolios that we use to have constant cost, which corresponds to the fourth column in Figure 1. As we mention in section 2.5, portfolios in P_c must have a constant weight on R, while the weights on $\mathbf{r}$ may belong to I. This motivates the approximation of P_c by means of the unconditional span of the vector $[R, (1, \mathbf{z}') \otimes \mathbf{r}']$.[9]

Once again, we can exploit sieve methods to construct a payoff space based on R and the arbitrage portfolios

$$\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{r}, \tag{27}$$

so that we can identify P_c with the (closure of the union over k_T of the) unconditional spans $\langle R, \mathbf{b}^{k_T}(\mathbf{z})' \otimes \mathbf{r}' \rangle$. Therefore, we could also express (R^*, r°) in terms of unconditional moments. In particular, we could identify r° from moment conditions analogous to (26) with the vector $\mathbf{r}$ replacing $\mathbf{x}$.

As expected, a tightened passive RF obtained in that way will converge to the unconditional RF in (11) as k_T grows, while its dual passive SF will converge to the extended SF in (20). In other words, the fourth column in Figure 1 will coincide with the second column in the limit.

In summary, the use of unrestricted managed portfolios and fixed-weight frontiers yields a relevant object when applied to SDF frontiers, but not when applied to portfolio frontiers, which should be based on managed portfolios of constant cost instead. In effect, this conclusion reflects the lack of duality between the unconditional frontiers discussed in section 2.

3.3 Relationship between Passive and Unconditional Frontiers

We conclude this section with two results on the relationship between passive and unconditional frontiers for a given vector of payoffs, which will investigate in section 4.3. The following proposition makes use of the representing portfolios to formalise the conditions for the fixed-weight and active frontiers to share a single element.

Proposition 6 *Given a vector of payoffs* $\mathbf{x} = (R, \mathbf{r}')'$*:*

1. *The unconditional RF shares an element with the Markowitz frontier if and only if there is an* $\omega \in \mathbb{R}$ *such that* $r^* - \omega r^\circ$ *has fixed-weights on the payoffs* $\mathbf{r}$*.*

[9] These linear managed portfolios are used by Bansal, Dahlquist, and Harvey (2004); see Brandt and Santa-Clara (2006) for a related approach.

2. *The unconditional SF shares an element with the Hansen-Jagannathan frontier based on returns if and only if in addition Proposition 1.1 holds, in which case there will be a* $\varpi \in \mathbb{R}$ *such that* $p^* - \varpi p^\circ$ *has fixed-weights on the payoffs* $\mathbf{x}$.

In other words, the unconditional frontiers for active portfolios will share an element with their fixed-weight counterparts if and only if there are linear combinations of the portfolios that unconditionally span the unconditional RF and SF with constant weights on the original payoffs. In addition, tangency between the unconditional SF and the Hansen-Jagannathan bounds implies tangency between the unconditional RF and the Markowitz frontier, but the converse is not generally true. In this regard, note that the duality condition in Proposition 1.1 is precisely the extra condition that combined with the first part of Proposition 6 implies tangency in the unconditional SF, and hence the existence of an SDF with constant coefficients that prices active strategies. This is due to the fact that the unconditional SF has a random weight on $R^* = R - r^*$ (see (16)). Thus, the passivity of the weight on R automatically imposes the duality condition 1.1

Figure 4 illustrates Proposition 6.2 with an example in which the unconditional SF shares an element with the Hansen-Jagannathan frontier based on returns. As we have just explained, this result also implies that the duality condition between unconditional frontiers must hold at that particular point, which in turn means that the unconditional RF shares an element with the Markowitz frontier. For ease of comparison, we keep the passive RF and SF in Figure 3.

<FIGURE 4>

In economic terms, tangency in the unconditional SF means that there is an SDF with constant risk prices that can price any active strategy. In contrast, tangency in the unconditional RF is less relevant because it only means that there is one (and only one) optimal portfolio with constant weights on the original payoffs.

We can extend the previous proposition to the case where fixed-weight and active frontiers coincide.

Proposition 7 *Given a vector of payoffs* $\mathbf{x} = (R, \mathbf{r}')'$*:*

1. *The unconditional RF and the Markowitz frontier are equal if and only if* r^* *and* r° *have fixed-weights on the payoffs* $\mathbf{r}$.
2. *The unconditional SF and the Hansen-Jagannathan frontier based on returns are equal if and only if in addition Proposition 1.2 holds, in which case* p^* *and* p° *will have fixed-weights on the payoffs* $\mathbf{x}$.

Thus, the unconditional frontiers for active portfolios will coincide with their fixed-weight counterparts if and only if the portfolios that span the unconditional RF and SF have constant weights. In the first case, the conditional RF is generated by active strategies only through a time-varying choice of the weight on r°. Again, if the unconditional SF coincides with the Hansen-Jagannathan frontier based on returns, then the unconditional RF will also coincide with the Markowitz frontier, the extra condition being simply the duality condition in Proposition 1.2.

Figure 5 illustrates Proposition 7.1 with an example in which the unconditional RF and the Markowitz frontier are equal. Once again, we simplify the comparison by keeping the same passive RF and SF as in Figure 3.

<FIGURE 5>

Finally, it is important to emphasise that the conditions in Propositions 6 and 7 do not imply that conditioning information plays no role, or indeed that there is not predictability in the first and second moments of returns. Even if the more restrictive conditions in Proposition 7.2 hold, the maximum conditional Sharpe ratio can change freely over time, as explained after Proposition 1. Those conditions simply mean that the investor can attain an optimal risk-return trade-off by means of a simple passive strategy. In the next section, we will further illustrate the different concepts with US stock data.

4 The empirical relevance of conditioning information

Despite hundreds of papers over three decades, the evidence on the predictability of the levels of financial returns remains controversial (see Spiegel (2008) and the references therein). In fact, there is not even agreement about the predictability channel among those who believe in it (see Bansal, Kiku and Yaron (2012) and Beeler and Campbell (2012)). In contrast, there is much stronger evidence on time variation in volatilities at daily frequencies, but at the same time the extent to which those effects are relevant at lower frequencies, such as monthly or quarterly, is less clear.

In this section, we would like to answer two questions which are related but not identical to the predictability of the first and second moments of asset returns.

Linear factor pricing models with constant weights on some traded factors are often used for pricing a broad cross-section of US stock portfolios (see e.g. Cochrane (2001, sec. 20.2)).

Our first question is whether they can also correctly value portfolios whose weights depend on the information set. In this regard, it is worth recalling that the SDF implied by those models will belong by construction to the passive SF generated from the pricing factors. Therefore, for those models to correctly price active strategies, it must be the case that this passive SF and the unconditional SF are tangent at a single point, as shown in Proposition 6.2.

The second question that we would like to answer is whether the use of conditioning information adds value in portfolio choice. We do so by testing the null hypothesis that the unconditional RF coincides with the standard Markowitz frontier, as stated in Proposition 7.1.

4.1 Econometric methodology

Propositions 6 and 7 provide conditional moment restrictions that we can exploit for conducting inferences which explicitly take into account sampling uncertainty. Specifically, the second part of Proposition 6 on tangency between the passive SF for $\mathbf{x}$ and the unconditional SF is equivalent to the existence of an unconditional SF element (14) that can be expressed as $\varpi + \mathbf{x}'\boldsymbol{\varphi}$ for a scalar ϖ and some vector $\boldsymbol{\varphi}$. Therefore, we can test this passive tangency at some unspecified point by means of the conditional moments

$$E[\mathbf{x}(\varpi + \mathbf{x}'\boldsymbol{\varphi}) - \mathbf{e}_1|\mathbf{z}] = \mathbf{0}, \tag{28}$$

where we have written the information set in terms of the vector of predictors $\mathbf{z}$ and exploited the fact that the true cost of $\mathbf{x}$ is $\mathbf{e}_1 = (1, \mathbf{0}')'$ when $\mathbf{x} = (R, \mathbf{r}')'$.

Similarly, the first part of Proposition 7 implies the existence of vectors $\boldsymbol{\phi}^*$ and $\boldsymbol{\phi}^\circ$ such that

$$E\left[\begin{array}{c} \mathbf{r}(\mathbf{r}'\boldsymbol{\phi}^* - R) \\ \mathbf{r}(\mathbf{r}'\boldsymbol{\phi}^\circ - 1) \end{array} \middle| \mathbf{z}\right] = \mathbf{0}, \tag{29}$$

where $\mathbf{r}'\boldsymbol{\phi}^*$ and $\mathbf{r}'\boldsymbol{\phi}^\circ$ yield the representing portfolios r^* and r° that span the unconditional RF (see (12) and (9)), which have constant weights under the null.

We deal with these conditional moment conditions by means of the optimally weighted sieve minimum distance (SMD) semiparametric procedures introduced by Ai and Chen (2003), which are both easy to implement and intuitive. Let $\mathbf{h}(\mathbf{y}; \boldsymbol{\theta})$ denote a vector of influence functions such that we can express the conditional moment restrictions as

$$E[\mathbf{h}(\mathbf{y}; \boldsymbol{\theta})|\mathbf{z}] = \mathbf{0},$$

where $\mathbf{y}$ represents the data and $\boldsymbol{\theta}$ the unknown parameters, which are finite in both (28) and (29). Importantly, the linearity of those moment conditions guarantees the identifiability of $\boldsymbol{\theta}$ in view of our assumptions on the non-redundant nature of the asset payoffs in $\mathbf{x}$.

Let $\mathbf{b}^{k_T}(\mathbf{z})$ denote the vector of sieve basis functions introduced in section 3, with the property that a linear combination can approximate the conditional mean of $\mathbf{h}(\mathbf{y};\boldsymbol{\theta})$ as the smoothing parameter k_T increases. Following Ai and Chen (2003), we first project $\mathbf{h}(\mathbf{y};\boldsymbol{\theta})$ onto the linear span of $\mathbf{b}^{k_T}(\mathbf{z})$ as follows

$$\mathbf{g}(\mathbf{z};\boldsymbol{\theta}) = E[\mathbf{h}(\mathbf{y};\boldsymbol{\theta})\mathbf{b}^{k_T}(\mathbf{z})']\boldsymbol{\Lambda}^{-1}\mathbf{b}^{k_T}(\mathbf{z}), \qquad \boldsymbol{\Lambda} = E[\mathbf{b}^{k_T}(\mathbf{z})\mathbf{b}^{k_T}(\mathbf{z})'].$$

Then, we estimate the unknown parameters by minimising the sample analogue to the criterion function

$$\mathcal{J} = E[\mathbf{g}(\mathbf{z};\boldsymbol{\theta})'\boldsymbol{\Omega}^{-1}(\mathbf{z})\mathbf{g}(\mathbf{z};\boldsymbol{\theta})], \tag{30}$$

where the optimal weighting matrix

$$\boldsymbol{\Omega}(\mathbf{z}) = Var[\mathbf{h}(\mathbf{y};\boldsymbol{\theta})|\mathbf{z}]$$

guarantees that the resulting estimator of $\boldsymbol{\theta}$ will attain the semiparametric efficiency bound. In this context, consistency is ensured if asymptotically the sieves are dense in the relevant space.

Importantly, we can express the SMD criterion (30) as

$$\mathcal{J}=E[\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{h}(\mathbf{y};\boldsymbol{\theta})]'E[\boldsymbol{\Lambda}^{-1}\mathbf{b}^{k_T}(\mathbf{z})\mathbf{b}^{k_T}(\mathbf{z})'\boldsymbol{\Lambda}^{-1} \otimes \boldsymbol{\Omega}^{-1}(\mathbf{z})]E[\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{h}(\mathbf{y};\boldsymbol{\theta})], \tag{31}$$

which suggests that $\mathcal{J}$ may also be interpreted as a GMM criterion for managed portfolios for any given k_T. In particular, $\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{h}(\mathbf{y};\boldsymbol{\theta})$ corresponds to sieve managed portfolios with payoffs (25) in the context of the SF-related moments (28), while the relevant sieve managed portfolios become (27) in the case of the unconditional RF-related moments (29).

As for the weighting matrix in (31), we can show that under some conditions on $\mathbf{b}^{k_T}(\mathbf{z})$ and $\boldsymbol{\Omega}(\mathbf{z})$ that our choice of sieves will satisfy, we will have that:

$$E[\boldsymbol{\Lambda}^{-1}\mathbf{b}^{k_T}(\mathbf{z})\mathbf{b}^{k_T}(\mathbf{z})'\boldsymbol{\Lambda}^{-1} \otimes \boldsymbol{\Omega}^{-1}(\mathbf{z})] = \left\{E[\mathbf{b}^{k_T}(\mathbf{z})\mathbf{b}^{k_T}(\mathbf{z})' \otimes \boldsymbol{\Omega}(\mathbf{z})]\right\}^{-1},$$

in which case we can interpret T times the sample counterpart to $\mathcal{J}$ as a GMM overidentifying restrictions statistic. Nevertheless, by using SMD we mitigate the numerical problems associated to the inversion of large, poorly conditioned matrices that plague standard GMM procedures. Intuitively, the reason is that the size of the vector $\mathbf{g}(\mathbf{z};\boldsymbol{\theta})$ and the matrix $\boldsymbol{\Omega}(\mathbf{z})$ that appear in

(30) depend on the number of assets in $\mathbf{x}$, but not on the dimension of $\mathbf{z}$ or the sieve that we use, which only affect the size of $\mathbf{\Lambda}$.

An additional numerical advantage of the Ai and Chen (2003) estimation procedure is that it only involves a finite number of parameters for any given sample size. At the same time, by explicitly recognising that the number of parameters would increase slowly with the sample size, sieves methods are more flexible and robust that classical parametric methods which assume fixed, finite-dimensional parameter spaces regardless of the sample size.

In practice, we first map each of our predictors onto the interval $[0, 1]$ by means of their empirical probability integral transform (PIT). Then we construct $\mathbf{b}^{k_T}(\mathbf{z})$ as the Kronecker product of the linearly independent elements of the B-splines of order 1 with knot vector $(0, 0.5, 1^+)$ and order 2 with knot vector $(-0.5, 0, 0.5, 1, 1.5)$ for each predictor. Importantly, the implied sieve managed portfolios span the standard linear managed portfolios obtained from the PITs of each element of $\mathbf{z}$. Higher order splines would eventually approximate any active strategy whose weights are smooth functions of the predictors, but at the cost of introducing a much larger number of parameters. We also use the same B-splines to generate the payoff spaces and the cost and mean representing portfolios underlying the unconditional and extended frontiers in Figure 2 and 3, as explained at the end of section 3.2.

Similarly, we estimate $\mathbf{\Omega}(\mathbf{z})$ using the Kronecker product of the same B-splines of order 1 for each predictor to guarantee that $\hat{\mathbf{\Omega}}(\mathbf{z})$ will be positive semidefinite for all values of $\mathbf{z}$ by construction. Nevertheless, in our empirical applications we will also report a GMM counterpart to (31) that uses the inverse of $Var[\mathbf{b}^{k_T}(\mathbf{z}) \otimes \mathbf{h}(\mathbf{y}; \boldsymbol{\theta})]$ as weighting matrix to check that our results do not depend on this particular choice of $\mathbf{\Omega}(\mathbf{z})$. In what follows, we shall refer to this alternative procedure as the standard GMM approach.

As usual, we have two possibilities to deal with the fact that we do not know the true $\boldsymbol{\theta}$. Either we use some initial consistent estimator of $\boldsymbol{\theta}$ and iterate to obtain k-step SMD estimators, or we explicitly take into account in the criterion function the dependence of the weighting matrix on the parameter values, along the lines of the single-step continuously updated (CU) GMM estimator of Hansen, Heaton and Yaron (1996). Although this estimator is often more difficult to compute than two-step and iterated estimators, particularly in linear models, an important advantage is that (30) becomes numerically invariant to normalisation, bijective reparametrisations and parameter-dependent linear transformations of the conditional moment conditions. In contrast, these properties do not necessarily hold for two-step or iterated SMD.

Hansen, Heaton and Yaron (1996) showed the relevance of these invariance properties in testing asset pricing models using GMM. Newey and Smith (2004) confirmed the advantages of CU- over two-step GMM by going beyond the usual first-order asymptotic equivalence results. More recently, Peñaranda and Sentana (2015) conducted a detailed simulation experiment which shows that GMM asymptotic theory provides a reliable guide for the CU version of the J test when the moment conditions hold. For those reasons, all the empirical results reported below have been computed with a CU version of SMD.

Finally, we follow Donald, Imbens, and Newey (2003) in calculating the p-value of the overidentifying restrictions test statistic associated to (31) from a chi-square distribution with degrees of freedom equal to the difference between the number of moment conditions and the number of parameters. These authors formally show that this statistic, standardised by subtracting its mean and dividing by its standard deviation, converges to a normal distribution when the number of observations and basis functions converge to infinity at suitable rates. Nevertheless, they have a preference for a chi-squared asymptotic approximation because, among other motivations, it is correct for fixed k_T.[10] In any event, we explicitly study the finite sample reliability of their asymptotic approximation in the Monte Carlo experiments reported in supplemental appendix E.5.

4.2 Data

We initially focus our analysis on the three Fama and French factors for US stocks, which we have obtained from Ken French's Data Library (see his web page, as well as Fama and French (1993) for further details). We use monthly data from January 1952 to December 2012 (732 observations), so that our sample begins soon after the Treasury - Federal Reserve Accord whereby the Fed stopped its wartime pegging of interest rates. As in previous sections, we will express the payoff vector $\mathbf{x}$ in terms of a gross return R, which we identify with the US market portfolio, MK, and a vector of excess returns $\mathbf{r}$ associated to the portfolios that capture size and value effects, so that:

$$(R, \mathbf{r}')' = \begin{pmatrix} R_{MK} & r_{SMB} & r_{HML} \end{pmatrix}', \tag{32}$$

[10] Strictly speaking, their theoretical results are developed for *i.i.d.* data, while our influence functions $\mathbf{h}(\mathbf{y};\boldsymbol{\theta})$ are martingale difference sequences under the null hypotheses that we study. Still, under suitable additional regularity conditions they should apply to our case too.

where SMB means long/short in small/large capitalisation stocks, and HML long/short in high/low book-to market ones.

We work with two prominent predictors: the US (cyclically adjusted) price earnings ratio, and the default spread.[11] The former is taken from Robert Shiller's web page, and the latter is constructed from FRED data (from yields on AAA and BAA-rated bonds). These predictors are among the ones that Goyal and Welch (2008) and Campbell and Thompson (2008) considered in their analysis of mean predictability. In what follows, the vector $\mathbf{z}$ will denote these two predictors once they have been mapped onto the interval $[0, 1]$ by means of their empirical PIT.

The first thing we do is to check that our sieve procedure is indeed able to predict the Fama-French returns. Table 1 reports predictability tests in the first moment of MK, SMB and HML, which we assess by means of the conditional moment restrictions

$$E(\mathbf{x} - \boldsymbol{\mu}|\mathbf{z}) = \mathbf{0}, \tag{33}$$

where $\boldsymbol{\mu}$ is a vector of real parameters.

<TABLE 1>

As can be seen in Panel A, we cannot reject the constancy of the first conditional moment of R_{MK} or r_{SMB} with affine functions of $\mathbf{z}$, but we clearly reject with B-splines. In contrast, we cannot find predictability in the first moment of r_{HML}. The standard CU-GMM tests in Panel B yield similar conclusions.

Still, it is worth emphasising again that the results in Proposition 6 and 7 are compatible with predictability in levels of $\mathbf{x}$. For that reason, our empirical methodology puts the emphasis directly on portfolio weight predictability instead of first or second moment predictability.

4.3 Empirical results

4.3.1 Validity of an SDF with constant risk prices

As we mentioned at the beginning of section 4, linear factor pricing models with constant weights are often used for pricing a broad cross-section of US stock portfolios. The purpose of our exercise is to test whether they can also correctly value portfolios whose weights depend on the information set. Given that the SDF implied by those models belongs to the passive SF generated from the pricing factors, it must be the case that this passive SF and the unconditional SF are tangent at a single point for those models to correctly price active strategies.

[11] With two predictors, our choice of sieves implies that $\dim[b^{k_T}(\mathbf{z})] = 16$.

Table 2 shows the results of the tests of the null hypothesis that the unconditional SF shares one element with the Hansen-Jagannathan frontier based on returns, which is formally stated in Proposition 6.2. As explained at the beginning of section 4.1, this null gives rise to the conditional moments (28), which can be regarded as conditional counterparts to the unconditional moments in the tangency tests of Peñaranda and Sentana (2011).

<TABLE 2>

As can be seen, we reject that an SDF with constant weights on the Fama-French factors is able to price actively managed portfolios of those factors. Therefore, we do not expect a fortiori that such a passive SDF would be able to price a richer cross-section of active strategies constructed from size, book-to-market, and momentum sorted portfolios. In simple economic terms, the rejection that we find implies that the risk prices of the Fama-French model, which we can identify with the coefficients of its candidate SDF, cannot be time-invariant.

However, if we restrict the weights of the managed portfolios to be affine in $\mathbf{z}$, then we do not reject at the 1% level with SMD in Panel A (or at the 5% with standard GMM in Panel B). We can confirm these differential results by formally testing the relevance of extending the managed portfolios from $(1, \mathbf{z}')'$ to $\mathbf{b}^{k_T}(\mathbf{z})$. The distance metric test of this null hypothesis has a p-value of 0.3% under the maintained hypothesis of passive tangency.

To further illustrate the lack of duality between the unconditional RF and SF, we also test the hypothesis that the unconditional RF shares an element with the Markowitz frontier. Given that Proposition 6.1 implies that passive tangency on the unconditional RF is equivalent to the existence of a scalar ω and a vector $\boldsymbol{\phi}$ such that $R - r^* + \omega r^\circ = R - \mathbf{r}'\boldsymbol{\phi}$, we rely on the conditional moments

$$E[\mathbf{r}(\omega + \mathbf{r}'\boldsymbol{\phi} - R)|\mathbf{z}] = \mathbf{0}. \tag{34}$$

The results in Table 2 show that the evidence against this null is weaker, especially for linear managed portfolios; see Ferson and Siegel (2009) for related evidence.

4.3.2 The new Fama-French five-factor model

Fama and French (2015) have recently developed an improved version of their three factor model that adds two factors: RMW, long/short in robust/weak profitability stocks, and CMA, long/short in conservative/aggressive stocks, which they identify as low/high investment firms. Their rationale for those factors is that they capture the profitability and investment patterns in stock returns beyond the usual size and value effects. Given that there is growing interest in this model in the empirical literature, we subject it to the same testing procedure used in the previous subsection.

The new factors are available through Ken French's Data Library from July 1963 onwards. Accordingly, our empirical results cover the period that goes from that month until December 2014 (618 observations).

The results in Table 3 reject the null hypothesis that there is an SDF with constant weights on the five Fama-French factors which can correctly price active strategies on those factors. Once again, this suggests that the risk prices of this model should probably be time-varying too. In this regard, our result complements the conclusions of the Gibbons, Ross and Shanken (1989) test reported by Fama and French (2015), which indicates that their five-factor model is unable provide a valid SDF for the cross-section of stock returns that they consider.

<TABLE 3>

4.3.3 Relevance of the Markowitz frontier

Let us turn to our second empirical question, namely, whether the use of conditioning information adds value in portfolio choice. We answer that question by testing the null hypothesis that the unconditional RF coincides with the standard Markowitz frontier, which is constructed with passive portfolios. This corresponds to the first part of Proposition 7, which we test by means of the conditional moments (29), as discussed in section 4.1.

Table 4 shows that the null hypothesis is clearly rejected for both sieve and linear managed portfolios (the standard GMM tests in Panel B are qualitatively similar). Therefore, there is added value in exploiting conditioning information for investors that choose portfolios on the unconditional RF. Our results confirm those in Bansal, Dahlquist and Harvey (2004), who find that managed portfolios can significantly improve the mean-variance trade-off achievable with fixed-weight portfolios only. These findings are perhaps not entirely surprising in view of the fact that passive investors and fund managers who actively engage in market timing typically have different information sets and different abilities or resources to process their common information.

<TABLE 4>

For completeness, we also test the null hypothesis that the unconditional SF is spanned by the passive SF, stated in Proposition 7.2. This is equivalent to the existence of vectors $\boldsymbol{\varphi}^*, \boldsymbol{\varphi}^\circ$

such that

$$E\left[\begin{array}{c}\mathbf{x}\mathbf{x}'\boldsymbol{\varphi}^{*}-\mathbf{e}_1\\ \mathbf{x}(\mathbf{x}'\boldsymbol{\varphi}^{\circ}-1)\end{array}\middle|\mathbf{z}\right]=\mathbf{0},$$

which resemble the unconditional spanning tests in Peñaranda and Sentana (2012). Once again, we clearly reject the null, which is not surprising because passive spanning of the unconditional SF requires passive spanning of the unconditional SF; see Proposition 7.2.

5 Conclusions

Portfolio and stochastic discount factor frontiers are usually regarded as dual objects to the extent that sometimes researchers use one type of frontier to answer questions that arise more naturally in the other type. Nevertheless, the widely cited duality result in Hansen and Jagannathan (1991) does not usually hold when one explicitly exploits the potential predictability of the first two moments of asset returns in designing portfolio strategies, as Hansen and Richard (1987) did in the case of portfolios and Gallant, Hansen and Tauchen (1990) for SDFs.

In this regard, our first theoretical contributions is to derive the precise restrictions under which the unconditional SF and RF that they proposed are dual, which will happen when the position in conditional mean-standard deviation space of the cost representing portfolio and its gross return is independent of the variables in the information set. In general, though, empirical researchers taking into account conditioning information should focus on the type of frontier that is really relevant for the particular question they want to address.

Another theoretical contribution is to explicitly characterise the random variables for which the dual frontiers to the unconditional RF and SF provide the sharpest possible mean-variance bounds. Specifically, while the unconditional RF delivers interesting one period risk-return trade-offs, the extended SF does not deliver equally interesting constraints on asset pricing models because it may not price correctly portfolios whose cost is a function of the conditioning information. Therefore, the unconditionally efficient bound advocated by Ferson and Siegel (2003), which at best coincides with the extended SF frontier if the researcher uses the correct conditional model, provides a suboptimal SDF bound.

Likewise, while the unconditional SF delivers interesting constraints on asset pricing models, the extended RF does not deliver interesting risk-return trade-offs because its elements are either infeasible or leave money on the table. This means that the Bekaert and Liu (2004) SDF framework cannot be used to compute optimal portfolios without a substantial rethink.

Our methodological contribution is a computationally simple yet efficient procedure for the estimation of all the different frontiers in Figure 1. Specifically, we rely on the sieve minimum distance semiparametric procedures proposed by Ai and Chen (2003) to estimate the frontiers and make inferences about them. Importantly, we explicitly relate the mean-variance frontiers such a procedure generates to the alternative empirical strategy of approximating the effect of conditioning information by constructing passive RFs and SFs from managed portfolios.

Our empirical exercises combine our theoretical results and econometric methods to explore two questions related to the ways in which conditioning information sharpens return and SDF frontiers. First, we test whether the SDF of the popular Fama and French (1993) linear factor pricing model belongs to the unconditional SF. Although we clearly reject the null, the evidence is weaker if we focus on the linear payoff space of those pricing factors scaled by the price earnings ratio and the default spread. In simple economic terms, the rejection that we find implies that the risk prices of the Fama-French model cannot be time-invariant. Our conclusions are identical when we study the new five-factor model in Fama and French (2015).

Second, we use the same data set to investigate whether investors would effectively expand their mean-variance opportunity set by considering active portfolio strategies. The null hypothesis that the unconditional RF coincides with the standard Markowitz frontier is clearly rejected for both sieve and linear managed portfolios. While our Monte Carlo simulation exercises indicate some size distortions, they by no means overturn the empirical conclusion that there is added value in exploiting conditioning information.

Finally, although mean-variance analysis is still a common tool in portfolio choice, asset pricing tests and performance evaluation, one relevant extension of our work would be the introduction of higher order moments in our analysis (as Chabi-Yo (2008) and Almeida and Garcia (2016)). From a methodological perspective, other relevant extensions would be (i) an automatic data-driven choice of the order of the sieve in the context of mean-variance frontiers with conditioning information (see Donald, Imbens, and Newey (2009) for such results with *i.i.d.* data); (ii) an exploration of resampling methods that do not require a parametric model for the conditional distribution of asset returns; and (iii) an extension of our estimation and testing framework to situations in which the influence functions contain non-parametric components under the null along the lines of Chen and Pouzo (2015). We are currently exploring several of these interesting extensions.

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.

Ai, C., and X. Chen (2003): "Efficient estimation of models with conditional moment restrictions containing unknown functions", *Econometrica* 71, 1795-1843.

Almeida, C. and R. Garcia (2016): "Economic implications of nonlinear pricing kernels", forthcoming in *Management Science*.

Balduzzi, P. and C. Robotti (2008): "Mimicking portfolios, economic risk premia, and tests of multi-beta models", *Journal of Business & Economic Statistics* 26, 354-368.

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

Bansal, R., Kiku, D. and Yaron, A. (2012): "An empirical evaluation of the long-run risks model for asset prices", *Critical Finance Review* 1, 183–221.

Beeler, J. and Campbell, J.Y. (2012): "The long-run risks model and aggregate asset prices: an empirical assessment", *Critical Finance Review* 1, 141-182.

Bekaert, G., and R.J. Hodrick (1992): "Characterizing predictable components in excess returns on equity and foreign exchange markets", *Journal of Finance* 47, 467-509.

Bekaert, G.,and J. Liu (2004): "Conditioning information and variance bounds on pricing kernels", *Review of Financial Studies* 17, 339-378.

Bekaert, G., and M.S. Urias (1996): "Diversification, integration and emerging market closed-end funds", *Journal of Finance* 51, 835-869.

Brandt, M.W., and P. Santa-Clara (2006): "Dynamic portfolio selection by augmenting the asset space", *Journal of Finance* 61, 2187-2217.

Campbell, J.Y. and S.B. Thompson (2008): "Predicting excess stock returns out of sample: can anything beat the historical average?", *Review of Financial Studies* 21, 1509-1531.

Chabi-Yo, F. (2008): "Conditioning information and variance bounds on pricing kernels with higher-order moments: theory and evidence", *Review of Financial Studies* 21, 181-231.

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

Chen, X. (2007): "Large sample sieve estimation of semi-nonparametric models", in J. Heckman and E. Leamer (eds.), *Handbook of Econometrics* Vol. VI, Elsevier.

Chen, X., and S.C. Ludvigson (2009): "Land of addicts? An empirical investigation of habit-based asset pricing models", *Journal of Applied Econometrics* 24, 1057-1093.

Chen, X., J. Favilukis, and S.C. Ludvigson (2013): "An estimation of economic models with recursive preferences", *Quantitative Economics* 4, 39–83.

Chen, X., and D. Pouzo (2015): "Sieve quasi likelihood ratio inference on semi/nonparametric conditional moment models", *Econometrica* 83, 1013–1079.

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.

Donald, S.G., G.W. Imbens, and W.K. Newey (2003): "Empirical likelihood estimation and consistent tests with conditional moment restrictions", *Journal of Econometrics* 117, 55–93.

Donald, S.G., G.W. Imbens, and W.K. Newey (2009): "Choosing instrumental variables in conditional moment restriction models", *Journal of Econometrics* 152, 28–36.

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.

Fama, E.F., and K.R. French (2015): "A five-factor asset pricing model", *Journal of Financial Economics* 116, 1-22.

Ferson, W.E., and A.F. Siegel (2001): "The efficient use of conditioning information in portfolios", *Journal of Finance* 56, 967-982.

Ferson, W.E., and A.F. Siegel (2003): "Stochastic discount factor bounds with conditioning information", *Review of Financial Studies* 16, 567-595.

Ferson, W.E., and A.F. Siegel (2009): "Testing portfolio efficiency with conditioning information", *Review of Financial Studies* 22, 2735-2758.

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

Goyal, A., and I. Welch (2008): "A comprehensive look at the empirical performance of equity premium prediction", *Review of Financial Studies* 21, 1455-1508.

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.

Jagannathan, R. (1996): "Relation between the slopes of the conditional and unconditional mean-standard deviation frontiers of asset returns", in S. Saito, K. Sawaki, and K. Kubota (eds.) *Modern Portfolio Theory and its Applications*, Center for Academic Societies, Osaka.

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

Nagel, S., and K.J. Singleton (2011): "Estimation and evaluation of conditional asset pricing models", *Journal of Finance* 66, 873–909.

Newey, W.K. and R.J. Smith (2004): "Higher order properties of GMM and generalized empirical likelihood estimators", *Econometrica* 72, 219-255.

Peñaranda, F. (2014): "Understanding portfolio efficiency with conditioning information", forthcoming in *Journal of Financial and Quantitative Analysis*.

Peñaranda, F. and E. Sentana (2011): "Inferences about portfolio and stochastic discount factor mean-variance frontiers", mimeo, CEMFI.

Peñaranda, F., and E. Sentana (2012): "Spanning tests in return and stochastic discount factor mean-variance frontiers: a unifying approach", *Journal of Econometrics* 170, 303-324.

Peñaranda, F., and E. Sentana (2015): "A unifying approach to the empirical evaluation of asset pricing models", *Review of Economics and Statistics* 97, 412-435.

Spiegel, M. (2008): "Forecasting the equity premium: where we stand today", *Review of Financial Studies* 24, 1453-1454.

A Appendix: Special Cases

There are three special cases in which mean-variance frontiers adopt a simple form. One such case occurs when there is a safe asset. The second case arises when all expected payoffs are conditionally proportional to their prices, with a common scalar factor of proportionality. The final one occurs when all the primitive assets are arbitrage portfolios.

A.1 The Riskless Asset Case

Imagine that investors have access to a set of assets $\mathbf{y}$ that includes not only the original risky asset payoffs in $\mathbf{x}$, but also the safe payoff $x_0 = 1$, so that $\mathbf{y}' = (x_0, \mathbf{x}')$. In this context, the conditional span of $\mathbf{y}$, which we denote by Q, will be an enlarged payoff space such that $Q \supset P$. On this basis, we can define the conditionally safe return and the extended return associated to the riskless asset as

$$R_0 = \frac{1}{C(x_0|G)} \in I, \qquad S_0 = \frac{1}{C(x_0)} \in \mathbb{R} \tag{A1}$$

respectively. The safe asset is unconditionally riskless when $C(x_0|G) = C(x_0)$, so that $R_0 = S_0$.

The conditional mean and cost active representing portfolios in the payoff space Q will be

$$q^\circ = 1, \qquad q^* = p^* + \left[\frac{C(x_0|G) - E(p^*|G)}{E(1 - p^\circ|G)}\right](1 - p^\circ), \tag{A2}$$

respectively. Note that q° is trivially the conditional projection of x_0 onto Q, and hence the corresponding residual will be 0. On the other hand, q^* is the conditional projection of any valid SDF onto Q, which obviously coincides with $m_C[C(x_0|G)]$ (see (C4)).

In the rest of this section we describe in detail the different mean-variance frontiers that one can construct, with a special emphasis on their shape, the relationship between frontiers with and without a safe asset, and a geometrical interpretation of duality by means of Sharpe ratios.

The representation of the conditional RF in (C2) is still valid after the introduction of a safe payoff if we simply replace p° and p^* in (5) with q° and q^* in (A2), respectively. As expected, the elements of the conditional RF lie along two straight lines in $[\sqrt{Var(p|G)}, E(p|G)]$ space for each possible signal value in G. Moreover, those two lines intersect on the vertical axis at R_0.

In addition, we can choose the conditional mean ν such that the weight of the conditional RF on the conditionally safe payoff x_0 will be identically 0 for every possible signal realisation, which implies that it will be equal to the conditional RF without a safe asset $p_C(\nu)$ at that point. This shared element is usually referred to as the *tangency portfolio*.

Similarly, the elements of the conditional SF solve the same problem as in (C3) with the additional pricing restriction $E(mx_0|G) = C(x_0|G)$. Hence, the only conditional mean that we can choose is $c = C(x_0|G)$ and the conditional SF will be given by the singleton $q^* = m_C[C(x_0|G)]$, which belongs to the conditional SF without a safe payoff.

In this context, the duality between the conditional RF and SF is trivial since the latter is fully traded and its corresponding return will be

$$S^* = q^*/E(q^{*2}|G). \tag{A3}$$

Alternatively, we can illustrate the duality between the conditional RF and SF by adapting the geometrical argument on mean-standard deviation spaces that relates the passive RF and SF in Hansen and Jagannathan (1991); see e.g. Figure 5.4 in Cochrane (2001). Specifically, the optimal *conditional* Sharpe ratio on the conditional RF will be equal to the slope of the ray that joins the origin with the single point on the conditional SF, so that

$$\frac{|E(S^*|G) - R_0|}{Var^{1/2}(S^*|G)} = \frac{Var^{1/2}(q^*|G)}{E(q^*|G)}.$$

Let us turn to unconditional frontiers. Again, the representation of the unconditional RF in (12) is still valid after the introduction of a safe payoff if we simply replace p° and p^* with q° and q^* in (A2), respectively.

In this context, we find two facts that contradict the textbook analysis of mean-variance frontiers with a safe asset. First, if R_0 is random then it does not belong to the unconditional RF, as Hansen and Richard (1987) showed. As a result, there is not a unique optimal risk-return trade-off on the unconditional RF unless R_0 is also unconditionally riskless, in which case the unconditional RF will indeed consist of two straight lines in $[\sqrt{Var(p)}, E(p)]$ space that intersect on the vertical axis at $R_0 = S_0$. Second, we can show that there is no *tangency portfolio* irrespective of whether $R_0 = S_0$, because the risky component of the elements of the augmented unconditional RF will not be conditionally proportional to the returns that conform the original unconditional RF. As Peñaranda (2014) proves, this is due to the fact the conditional RF tangency portfolio does not belong to the unconditional RF in general. Therefore, the Sharpe ratios that Bekaert and Liu (2004) and Abhyankar, Basu and Stremme (2007) consider must be interpreted with some care, as they relate to *passive* strategies that combine an unconditionally riskless asset (traded or fictitious) with a portfolio on the unconditional RF of risky assets alone.

Therefore, those Sharpe ratios underestimate the maximum unconditional risk-return trade-off that can be achieved through *active* portfolio strategies.

The elements of the unconditional SF solve the same problem as in (13), but with the additional pricing restriction $E(mx_0|G) = C(x_0|G)$. Once again, the only c for which the mean and pricing constraints are compatible is $C(x_0)$. As a result, the unconditional SF will also be the singleton q^*. However, this portfolio does not generally coincide with any $m_U(c)$ (see (14)) because its weight on $1 - p^\circ$ will be random even in the case of an unconditionally riskless asset.

Regarding extended frontiers, the elements of the extended RF solve the same problem as in (17), except that p is allowed to belong to the enlarged payoff space Q. The extended RF without a safe asset does not generally share any point with the extended RF with a safe asset, which consists of two straight lines in $[\sqrt{Var(p)}, E(p)]$ space that cross on the vertical axis at S_0 regardless of whether the riskless asset is unconditionally safe or not.

As expected, the single element of the unconditional SF q^* defined in (A2) has a dual element on the extended RF, which is given by

$$S_e^* = q^*/E(q^{*2}).$$

In addition, there is a clear connection between slopes of the return and SDF frontiers because both elements are unconditionally proportional, which means that the constant pseudo-Sharpe ratio (based on the unconditional moments of S_e^* in excess of the "safe" extended return S_0) of the elements of the extended RF will be equal to the slope of a ray from the origin to the single element of the unconditional SF, so that

$$\frac{|E(S_e^*) - S_0|}{Var^{1/2}(S_e^*)} = \frac{Var^{1/2}(q^*)}{E(q^*)}.$$

However, the left hand side of the foregoing expression will not be a proper Sharpe ratio even if $R_0 = S_0$ because S_e^* is not a proper return.

To construct the extended SF, we must first define the subspace of constant-cost portfolios $Q_c \supset P_c$, and obtain the extended representing portfolios q_e° and q_e^* in that subspace. The elements of the extended SF solve the same problem as in (20), but with the additional "pricing" constraint $E(mx_0|G) = hC(x_0|G)$. Nevertheless, this pricing constraint is not generally enough to pin down a particular c, and hence the extended SF will contain infinite points. However, when there is an unconditionally riskless asset, extended SDFs must price a unit payoff correctly on average, in which case the extended SF will be given by the single point q_e^* such that $E(q_e^*) =$

$C(x_0)$. In either case, the extended SF with and without a safe asset will not generally share any points.

As for the duality between the extended SF and the unconditional RF, it is easy to see that the extended SF will always be fully traded, and moreover, that its return will be S^* (see (A3)), which belongs to the unconditional RF. Given that

$$q_e^* = S^*/E(S^{*2}),$$

it is not surprising that

$$\frac{|E(S^*) - S_0|}{Var^{1/2}(S^*)} = \frac{Var^{1/2}(q_e^*)}{E(q_e^*)},$$

which means that the pseudo-Sharpe ratio of S^* is equal to the slope of a ray from the origin to q_e^* in $[E(m), \sqrt{Var(m)}]$ space. This pseudo-Sharpe ratio was already defined by Jagannathan (1996), who related it to the Sharpe ratio of the conditional RF. Nevertheless, his analysis requires that the safe asset is unconditionally riskless, in which case $|E(S^*) - S_0| / Var^{1/2}(S^*)$ will be a proper Sharpe ratio. More recently, Ferson and Siegel (2009) develop a portfolio efficiency test whose interpretation in terms of the unconditional RF also requires implicitly that the safe asset is unconditionally riskless, as pointed out by Peñaranda (2014).

Still, the difference between the pseudo-Sharpe ratios of S^* and S_e^* implies that one must be careful in extending to unconditional frontiers of actively managed portfolios the geometrical relationship obtained by Hansen and Jagannathan (1991) in terms of pseudo-Sharpe ratios of passive portfolios. In particular, such a relationship does not hold between the elements of the unconditional RF and SF, which simply reflects the fact that these two frontiers are not dual.

If there is an unconditionally riskless asset, then the pseudo-Sharpe ratio of S^* is bounded above by the pseudo-Sharpe ratio of S_e^*, which means that a bound on the volatility of SDFs obtained from S^* might be too low, and a pseudo-Sharpe ratio obtained from q^* might be too high. As a result, the intertemporal marginal rate of substitution in consumption of a specific CCAPM may look volatile enough from the perspective of S^* even though it would be insufficiently volatile to match q^* (cf. Cochrane (2001, sect 21.1).

Finally, we can also show that the elements of the passive RF for $\mathbf{y}$ will lie in $[\sqrt{Var(p)}, E(p)]$ space along two straight lines that cross on the vertical axis at a point with mean S_0 defined in (A1), regardless of whether the riskless asset is unconditionally safe or not. Moreover, we can show that the passive SF for $\mathbf{y}$ is a singleton because $c = C(x_0)$ is the only choice compatible with the associated pricing constraints.

A.2 Prices Proportional to Expected Payoffs

This situation is typically linked to the equilibrium of an economy with a risk-neutral agent, but it also arises when $N = 1$, an example used by Ferson and Siegel (2003) and Bekaert and Liu (2004) to differentiate their papers. Intuitively, the approach used by Ferson and Siegel (2003) to obtain SDF bounds cannot exploit the existence of conditioning information when $N = 1$ because the elements of the URF in (11) are constrained to have constant (unit) cost.

Either way, $p^{\circ} = kp^{*}$, with $k \in I$, so the geometry of the return and SDF frontiers will be the mirror image of the safe asset case. In particular, while the main implication of the existence of a safe asset was that $1 - q^{\circ} = 0$, with the additional feature that $1 - q_e^{\circ} = 0$ if the safe asset asset was unconditionally riskless, the main implication now is that r° defined in (10) will be 0, with the additional feature that r_e° defined in (B2) will also be 0 if expected payoffs are unconditionally proportional to their prices, i.e. if $k \in \mathbb{R}$.

In this context, the conditional RF will be given by the single element R^{*}, which was defined in (8). On the other hand, the risky part of the elements of the conditional SF can be obtained by conditionally scaling R^{*}. As a result, for each signal value the conditional SF will be represented by two straight lines in $[E(m|G), \sqrt{Var(m|G)}]$ space that touch at the horizontal axis when $c = k^{-1}$. The duality between the straight lines that characterise $m_C(c)$ and the point $p_C(\nu)$ relies on the fact that the return corresponding to the traded part of any $m_C(c)$ is always R^{*}.

A similar type of duality applies for the pairs unconditional RF/extended SF and extended RF/unconditional SF. Specifically, the unconditional RF will be given by the same single point R^{*} for the reasons explained when we discussed the conditional RF in the presence of a riskless asset. Further, the extended SF will now be given by two straight lines in $[E(m), \sqrt{Var(m)}]$ space that touch the horizontal axis at $c = E(k^{-1})$ because the scaling of R^{*} is non-random.

In contrast, there are no relevant changes in the unconditional SF and the extended RF with respect to the general case. However, if $k \in \mathbb{R}$, then the unconditional SF will be given by two straight lines in $[E(m), \sqrt{Var(m)}]$ space that touch the horizontal axis at $c = k^{-1}$, and the extended RF will be the single point R_e^{*} defined in (B1) with $E(R_e^{*}) = E(R^{*}) = k$.

The situation is slightly different when we consider passive frontiers for $\mathbf{x}$, which again do not suffer any relevant changes unless $k \in \mathbb{R}$ or $N = 1$. It is only in these circumstances that we find the mirror image situation to the safe asset case, in that the passive RF will collapse to a single point, while the passive SF will be given by two straight lines in $[E(m), \sqrt{Var(m)}]$ space.

A.3 Zero-cost Portfolios

Let us finally study the situation in which all primitive assets are arbitrage portfolios, so that $C(\mathbf{x}|G) = \mathbf{0}$. This case is quite common in empirical work, as asset payoffs are routinely transformed into excess returns in the presence of a (conditionally) riskless asset. From the point of view of mean-variance frontiers, the main implication of dealing with arbitrage portfolios is that the active cost representing portfolios defined in (5) is zero. Therefore, there is one-fund spanning in every frontier and consequently, all of them can be represented by straight lines that start from the origin in the appropriate mean-standard deviation space.

More specifically, since the cost of any portfolio of $\mathbf{x}$ is 0 in this case, the portfolio frontiers problems can be defined as usual (see problems (C1), (11) and (17)) after dropping the cost constraints. In other words, each problem consists now in minimising the second moment of portfolios given a constraint on their first moment. As a result, the unconditional and extended RFs coincide in this context since their only difference is the cost constraint. The conditional RF is constructed by a conditional scaling of p° and the unconditional RF by an unconditional scaling, while the passive RF for $\mathbf{x}$ would scale the passive counterpart to p°.

Interestingly, if the N arbitrage portfolios under analysis correspond to the excess returns of N risky assets over an *unconditionally* riskless asset, the slope of the unconditional/extended RF discussed in the previous paragraph will coincide with the slope of the unconditional RF discussed in the safe asset section, which combines the original N risky returns and the unconditionally safe asset. Therefore, the maximum unconditional Sharpe ratios attainable in both situations will also be the same, and will exceed the unconditional Sharpe ratios in Bekaert and Liu (2003) and Abhyankar, Basu and Stremme (2007) mentioned in the same section.

On the other hand, the pricing constraints of the SDF frontiers (see problems (C3), (13) and (20)) imply that any valid SDF must be orthogonal to $\mathbf{x}$. Moreover, since $P = P_c$ in this context, the unconditional and extended SFs will also coincide. The conditional SF is constructed by a conditional scaling of $1 - p^\circ$, and the unconditional SF by an unconditional scaling, while the passive SF for $\mathbf{x}$ would scale 1 minus the passive counterpart to p°. This means that we need a normalisation of candidate SDFs in testing their validity with excess returns; see Cochrane (2001, pages 256-258), Balduzzi and Robotti (2008) or Peñaranda and Sentana (2015). Finally, we can also add managed portfolios of zero cost to this set-up, as Bekaert and Hodrick (1992) did to estimate the slope of the passive SF.

Table 1: Tests of First Moment Predictability of the Fama-French Factors

	Sieve	Linear
Panel A. SMD		
MK	0.5	5.8
SMB	0.8	7.0
HML	20.3	87.7
Joint	0.1	7.9
Panel B. Standard GMM		
MK	0.8	7.9
SMB	0.2	4.9
HML	57.7	86.7
Joint	0.0	6.0

Note: Overidentifying restrictions tests of the conditional moments (33). Panel A and B display the CU-SMD and the CU-GMM test, respectively. The first column assesses orthogonality with respect to B-splines constructed from the predictors, while the last column orthogonality with respect to linear terms. For each test, the p-value (%) is shown.

Table 2: Tests of Passive Tangency on the Unconditional RF and SF

	Sieve	Linear
Panel A. SMD		
RF	2.5	18.8
SF	0.0	1.7
Panel B. Standard GMM		
RF	10.8	15.0
SF	0.2	5.6

Note: Overidentifying restrictions tests of the conditional moments (34) and (28). Panel A and B display the CU-SMD and the CU-GMM test, respectively. The first column studies managed portfolios constructed from B-splines, while the last column studies linear portfolios. There are two lines, the upper one studies constant cost managed portfolios (associated to the unconditional RF), while the lower one studies unrestricted managed portfolios (associated to the unconditional SF). For each test, the p-value (%) is shown.

Table 3: Tests of Passive Tangency on the Unconditional RF and SF

with the five Fama-French factors

	Sieve	Linear
Panel A. SMD		
RF	0.0	0.0
SF	0.0	0.0
Panel B. Standard GMM		
RF	0.7	0.2
SF	0.5	0.3

Note: Overidentifying restrictions tests of the conditional moments (34) and (28) when the vector of excess returns includes the two additional factors. Panel A and B display the CU-SMD and the CU-GMM test, respectively. The first column studies managed portfolios constructed from B-splines, while the last column studies linear portfolios. There are two lines, the upper one studies constant cost managed portfolios (associated to the unconditional RF), while the lower one studies unrestricted managed portfolios (associated to the unconditional SF). For each test, the p-value (%) is shown.

Table 4: Tests of Passive Spanning of the Unconditional RF and SF

	Sieve	Linear
Panel A. SMD		
RF	0.0	0.6
SF	0.0	0.0
Panel B. Standard GMM		
RF	0.3	4.2
SF	0.0	0.0

Note: Overidentifying restrictions tests of the conditional moments (29) and (**??**). Panel A and B display the CU-SMD and the CU-GMM test, respectively. The first column studies managed portfolios constructed from B-splines, while the last column studies linear portfolios. There are two lines, the upper one studies constant cost managed portfolios (associated to the unconditional RF), while the lower one studies unrestricted managed portfolios (associated to the unconditional SF). For each test, the p-value (%) is shown.

Figure 1: Mean-Variance Frontiers across Type and Information

Decreasing use of conditioning information →

	Active Strategies (Explicit use of information)			Passive Strategies (Implicit use of information)		
	Conditional Moments	Unconditional Moments		Managed Portfolios		Returns
		Constant Cost	Random Cost	Constant Cost	Random Cost	
Porfolio Frontiers	Hansen-Richard (1987) **CRF**	Hansen-Richard (1987) **URF**	*Extended Return Frontier* **ERF**	**URF$\leftarrow$PRF**	**ERF$\leftarrow$PRF**	Markowitz (1952) **PRF**
Duality ↕				Hansen-Jagganathan (1991)		
SDF Frontiers	Gallant et al. (1990) **CSF**	*Extended SDF Frontier* **ESF**	Gallant et al. (1990) **USF**	**ESF$\leftarrow$PSF**	**USF$\leftarrow$PSF**	Hansen-Jagganathan (1991) **PSF**

Note: Columns are arranged by a decreasing use of conditioning information, while for each column the last two rows couple the appropriate dual frontiers. The Conditional Return Frontier (CRF) and Conditional SDF Frontier (CSF) in the 1st column, which are constructed from active portfolio strategies, are dual. The Unconditional Return Frontier (URF) in the 2nd column and the Unconditional SDF Frontier (USF) in the 3rd column are subsets of the CRF and CSF, but they are not dual. Their duals are the Extended SDF Frontier (ESF) and the Extended Return Frontier (ERF), respectively. There are several dual pairs of Passive Return Frontiers (PRFs) and Passive SDF Frontiers (PSFs) constructed as fixed-weight combinations of different managed portfolios. The 6th column is obtained without managed portfolios. With all the relevant managed portfolios, the PRF/PSF converge to the ERF/USF (5th column), but if only constant cost managed portfolios are used they converge to the URF/ESF (4th column)

Figure 2: Unconditional and Extended Frontiers for Returns and SDFs

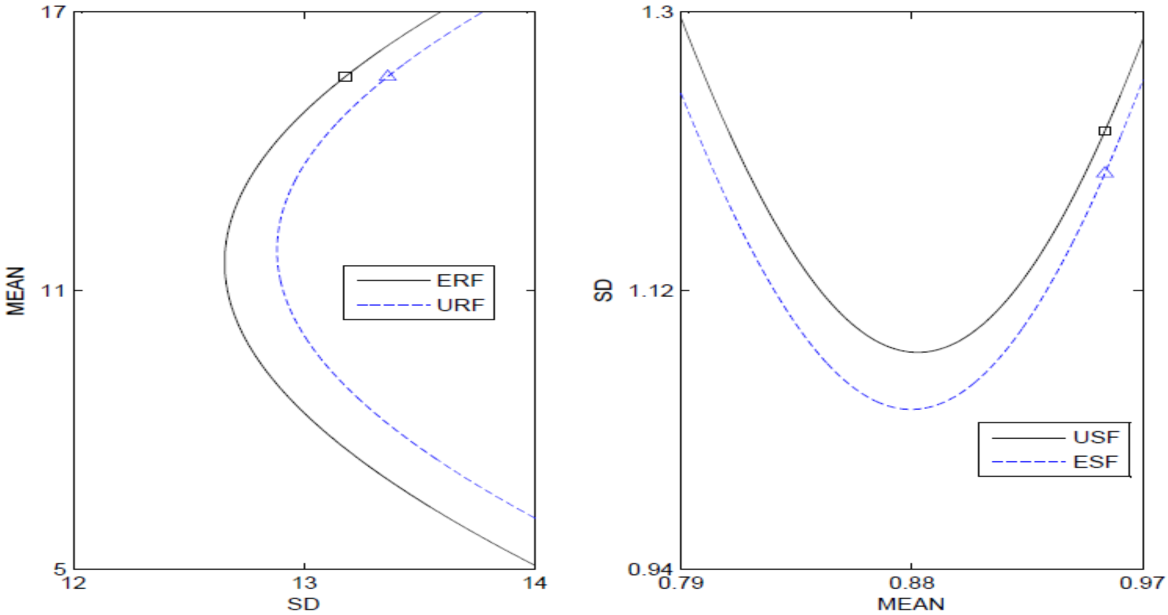

Note: The left mean-variance diagram displays the Unconditional and Extended Return Frontiers (URF and ERF, respectively). The right one displays the Unconditional and Extended SDF Frontiers (USF and ESF, respectively). The square and triangle represent the minimum distance points between extended and unconditional frontiers. Means and variances are annualised.

Figure 3: Unconditional, Extended and Passive Frontiers for Returns and SDFs

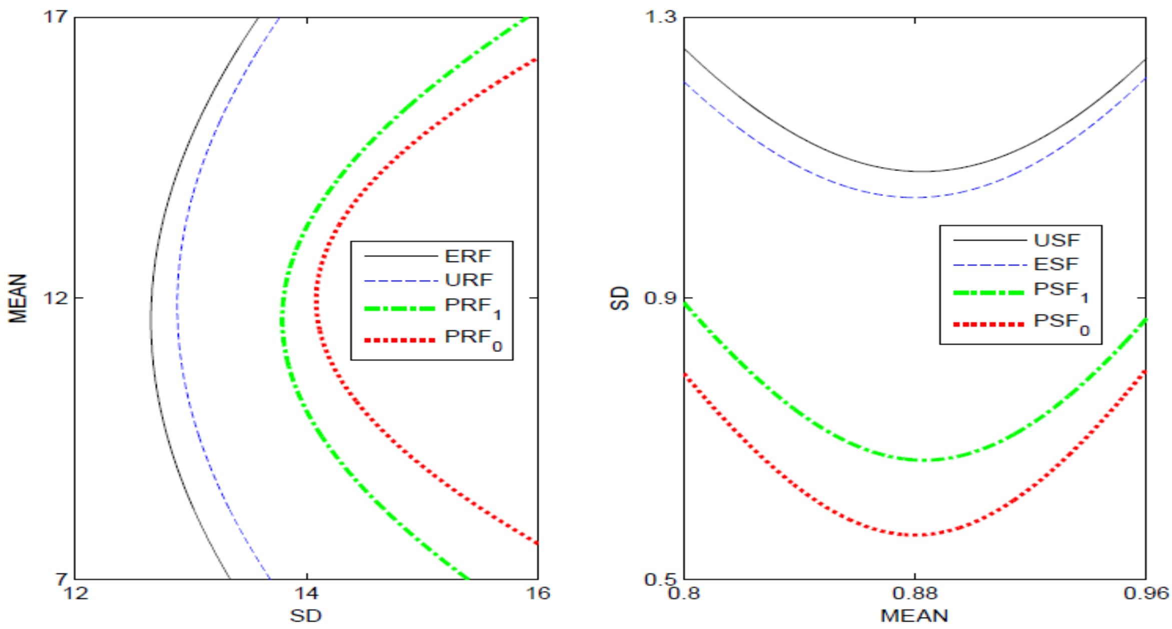

Note: The left mean-variance diagram displays the Unconditional and Extended Return Frontiers (URF and ERF, respectively). The right one displays the Unconditional and Extended SDF Frontiers (USF and ESF, respectively). PRF_0 and PSF_0 denote the passive frontiers for $\mathbf{x}$, while PRF_1 and PSF_1 denote the passive frontiers for $(1, \mathbf{z}')' \otimes \mathbf{x}$. Means and variances are annualised.

Figure 4: Passive Tangency on the Unconditional SF

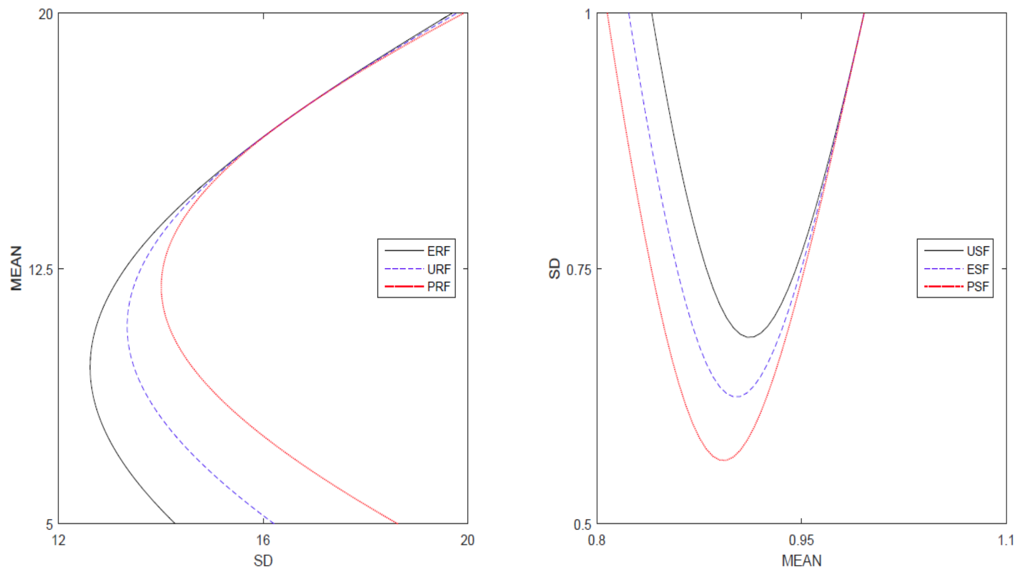

Note: The left mean-variance diagram displays the Unconditional and Extended Return Frontiers (URF and ERF, respectively). The right one displays the Unconditional and Extended SDF Frontiers (USF and ESF, respectively). The passive frontiers, PRF and PSF, are also displayed. Proposition 6.2 states the conditions that guarantee that the USF shares an element with the Hansen-Jagannathan frontier based on returns (PSF). Means and variances are annualised.

Figure 5: Passive Spanning of the Unconditional RF

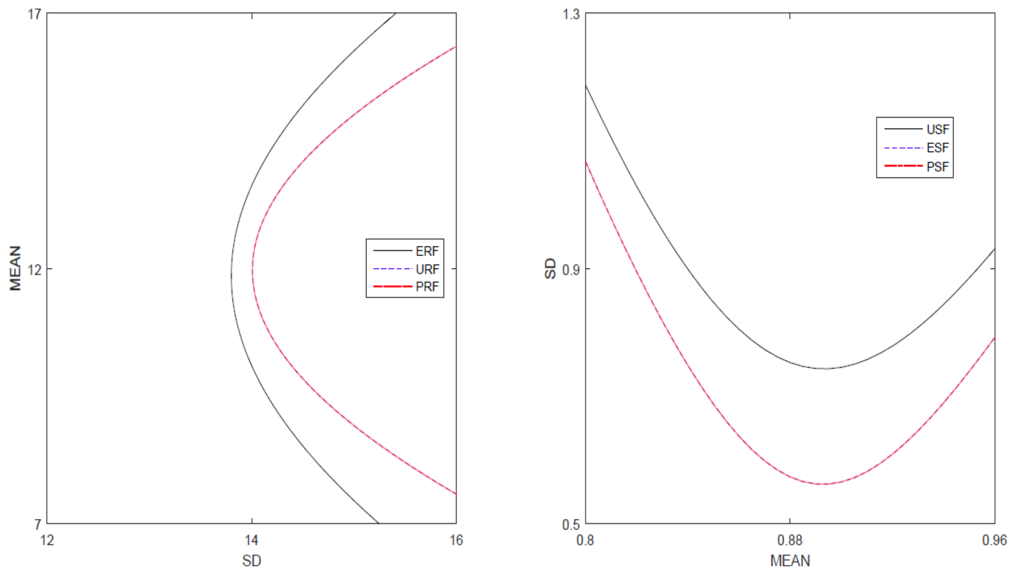

Note: The left mean-variance diagram displays the Unconditional and Extended Return Frontiers (URF and ERF, respectively). The right one displays the Unconditional and Extended SDF Frontiers (USF and ESF, respectively). The passive frontiers, PRF and PSF, are also displayed. Proposition 7.1 states the conditions that guarantee that the URF and the Markowitz frontier (PRF) are equal. Means and variances are annualised.