TRADABLE FACTOR RISK PREMIA AND ORACLE TESTS OF ASSET PRICING MODELS

ALBERTO QUAINI, FABIO TROJANI AND MING YUAN*

First version: June 2022. This version: March 8, 2024

Abstract

Tradable factor risk premia are defined by the negative factor covariance with the Stochastic Discount Factor projection on returns. They are robust to misspecification or weak identification in asset pricing models, and they are zero for any factor weakly correlated with returns. We propose a simple estimator of tradable factor risk premia that enjoys the Oracle Property, i.e., it performs as well as if the weak or useless factors were known. This estimator not only consistently removes such factors, but also gives rise to reliable tests of asset pricing models. We study empirically a family of asset pricing models from the factor zoo and detect a robust subset of economically relevant and well-identified models, which are built out of factors with a nonzero tradable risk premium. Well-identified models feature a low factor space dimension and some degree of misspecification, which harms the interpretation of other notions of factor risk premia in the literature.

Keywords: Testing of asset pricing models, factor risk premia, useless and weak factors, factor selection, model misspecification, Oracle estimation and inference.

JEL classication: G12, C12, C13, C51, C52, C58.

*Alberto Quaini (email: quaini@ese.eur.nl) is with Erasmus University of Rotterdam, Erasmus school of Economics. Fabio Trojani (email: Fabio.Trojani@unige.ch) is with the University of Geneva, the AXA Chair in socio-economic risks at the University of Turin and a Senior Chair of the Swiss Finance Institute. Ming Yuan (email: Ming.Yuan@columbia.edu) is with Columbia University, Department of Statistics and Data Science Institute. The methods developed in this paper are implemented in the Rcpp package intrinsicFRP, available on CRAN. For useful comments, we thank Harjoat Bhamra, Andrea Buraschi, Frank Kleiberghen, Marcin Kacperczyk, Chu Zhang, Svetlana Bryzgalova, Dacheng Xiu, Cesare Robotti, Yacine Ait-Sahalia, Caio Almeida, Christian Julliard, Chen Zhou, Robin Lumsdaine, José Scheinkman, Mikkel Plagborg-Moller, Sofonias Korsaye, Rasmus Lönn, Gustavo Freire, Matthieu Gomez, Noémie Pinardon-Touati, Anastasija Tetereva, Onno Kleen, Maria Grith, Paolo Zaffaroni and participants of CFE 2022, EcoSta 2023, IWiFE 2023, ESSEC-CYU-WARWICK Econometrics Workshop 2023, FinEML 2023 and seminars at Columbia University, Princeton University, Imperial College Business School and Collegio Carlo Alberto. Ming Yuan was supported in part by the Columbia-CityU/HK collaborative project that is supported by the InnoHK Initiative, the Government of the HKSAR, and the AIFT Lab.

1 Introduction

Over the last decades, a myriad of factors and low-dimensional factor models has been suggested to explain the cross-section of asset returns, giving rise to a so-called zoo of asset pricing factors; see, e.g., Harvey et al. (2016). Most of the empirical evidence for these models is based on various versions of the so-called two-step cross-sectional estimation and inference methodology, in which (i) factor risk premia are estimated by a cross-sectional projection of average returns on estimated asset return betas and (ii) asset return betas are estimated by a time series regression of returns on factors; see, e.g., Fama and MacBeth (1973) and Shanken (1992).[1] However, this methodology has been shown to produce unreliable results in empirically relevant situations where an asset pricing model may be misspecified or a candidate risk factor may be at best only weakly correlated with returns. See Kan and Zhang (1999b), Shanken and Zhou (2007), Kleibergen (2009), Kan et al. (2013), Gospodinov et al. (2014), Kleibergen and Zhan (2015) and Kleibergen and Zhan (2020), among many others.

Misspecification arises when asset expected returns are not completely spanned by the vectors of factor betas in an asset pricing model, and it has important implications for the properties of risk premia and test statistics produced by two-step cross-sectional regression approaches. First, misspecification gives rise to non zero pricing errors that impact the asymptotic distribution of estimated factor risk premia and derived statistics. Second, misspecification can lead to a definition of (pseudo) factor risk premia in two-step cross-sectional regression approaches that is to some extent arbitrary, as it depends on the weighting scheme chosen to penalize pricing errors in the second step of the methodology.

The presence in an asset pricing model of risk factors that may be at best only weakly correlated with returns gives rise to the so-called weak (or no) identification problem. This issue implies a cross-sectional projection of expected returns on asset return betas that does not uniquely identify a factor risk premium. Weak (or no) identification has important implications for the reliability of estimated risk premia and test statistics produced by two-step cross-sectional regression approaches, as it may spuriously favour statistically weak factors or even completely useless ones, i.e., factors that are uncorrelated with returns. These undesirable features may be further exacerbated by asset pricing models in which misspecification and weak

[1] This methodology is still the workhorse underlying a large body of empirical finance research. Inference with it is based on standard asymptotic assumptions, including a fix number of test assets and an increasing time dimension.

(or no) identification coexist. This creates important challenges also for a well-founded model comparison, as intuitively higher dimensional models may naturally tend to be less misspecified, but at the same time they may also more likely incur into weak or no identification issues, due to potential redundancies in some of the model's factor exposures; see, e.g., Kleibergen and Zhan (2022).

This paper addresses the testing of asset pricing models in the joint presence of misspecification and weak (or no) identification, based on a well-established, economically motivated, notion of a tradable factor risk premium. A tradable factor risk premium is defined by the negative factor covariance with the projection of any stochastic discount factor (SDF) on the space of asset returns. Importantly, tradable factor risk premia are well-defined whenever the standard no arbitrage condition of a finite maximum asset Sharpe ratio holds, and they are zero by construction for any factor that is only weakly correlated with returns. Therefore, they are naturally suited objects to develop inference procedures in potentially misspecified or weakly identified asset pricing models.

We show that a consistent finite sample selection of factors that are at best only weakly correlated with returns is easily achievable using tradable factor risk premia, together with a reliable and efficient inference on the risk premia of all other factors in an asset pricing model. To achieve this, we exploit the fact that the sample version of tradable factor risk premia follows an asymptotically Gaussian distribution under quite general assumptions, which cover models potentially including factors weakly correlated with returns. Crucially, this distribution may imply a distortion of estimated risk premia in weakly-identified models, but only for the risk premia of factors that are weakly correlated with returns.

We exploit these key properties to build a convenient Oracle estimator of tradable factor risk premia, which overcomes the potential distortions generated by weak factors in an asset pricing model. This estimator is given in closed-form as a penalized minimum distance correction of sample tradable risk premia and improves on the latter in two main directions. First, it consistently identifies the set of factors that are weakly correlated with returns. Second, it yields a uniform (asymptotically) efficient inference for the tradable risk premia of all other factors in an asset pricing model, as if these factors were known ex-ante. By design, our Oracle estimator directly produces a valid inference for the risk premia of factors that are not weakly correlated with returns. Alternatively, it can be used to perform a consistent preliminary screening of factors weakly correlated with returns, in order to reliably apply in a second step standard cross-sectional inference methodologies to the risk premia of all other factors in an asset pricing model.

Our Oracle estimation and inference methodology for tradable factor risk premia allows us to build an easily applicable, coherent diagnostics framework for studying the asset pricing properties of a broad class of models in presence of weak factors. For this purpose, we form randomized families of models rom the factor zoo, in order to single out a subset of well-identified models consisting exclusively of factors that are neither useless nor weak. From this set, we pin down the resulting distribution of factor selections and factor risk premia across well-identified models. In this way, we produce a global summary of key properties of asset pricing models, including (i) their identification and correct specification features, (ii) their factor dimension and composition, and (iii) the risk compensation of individual factors.

In our empirical analysis, we first find that our methodology is indeed effective in mitigating identification problems resulting from useless or weak factors. This is illustrated by Figure 1 in Section 6, showing that the identification properties of a random asset pricing model from the factor zoo are substantially improved after applying our Oracle factor selection.

Second, we detect a robust set of economically relevant, well-identified models from the factor zoo, which are all built out of a relatively small set of factors having a nonzero tradable risk premium. This is illustrated by Figures 2–4 in Section 6, which provide global summaries of (i) the typical factor space dimensions of well-identified models and (ii) the set of asset pricing factors appearing in these models. Figure 2 confirms that the most frequent factor space dimensions of well-identified models are rather low and concentrating around three-four factors. Figure 3–4 document instead the existence of a robust set of factors most frequently appearing in well-identified models and having a non zero tradable risk premium. They include market, size, intermediaries capital ratio, market with a hedged unpriced component, a mispricing factor of anomalies that firms' managements can affect directly, a long-term behavioral factor, a liquidity factor, a conservative minus aggressive factor and a composite equity issuance factor.

Third, we document that the low factor space dimensions of well-identified models are typically also associated with a degree of misspecification. This fact has important implications for the interpretation of other established definitions of a factor risk premium in the literature, which we show can be problematic. This is so, because these notions may induce excessively model-dependent risk premium estimates and inferences in presence of untradable factor risks. For instance, Figure 5 of Section 6 documents that interpreting the risk premia of factors intermediaries capital ratio and market with a hedged unpriced component is particularly challenging, because they produce weakly statistically significant and economically implausibly variable risk premium estimates, across well-identified models and choices of test assets.

The rest of the paper is structured as follows. After a review of the related literature, Section 3 defines tradable factor risk premia and clarifies their main properties. In Section 4, we introduce our Oracle estimator of tradable factor risk premia. We next derive its asymptotic properties, under varying assumption on the (potentially vanishing) correlation between factors and returns. Furthermore, we explain our factor selection and inference procedure based on tradable risk premia. Section 5 summarizes the main findings regarding the finite sample properties of our methodology, based on extensive Monte Carlo simulation evidence reported in the Online Appendix. Section 6 presents our empirical study of the factor zoo, while Section 7 concludes and provides directions for future research.

2 Literature review

Our work is motivated by the widely established two-step cross-sectional estimation and inference framework, introduced in Fama and MacBeth (1973) and studied formally by Shanken (1992), for testing low-dimensional linear asset pricing models. This framework has been the workhorse of an uncountable number of earlier empirical finance studies and is the methodology underlying several important recent studies.[2]

Our contribution originates from the key difficulties caused by a model's weak identification and/or misspecification for properly applying two-step cross-sectional estimation and inference methods.[3] Compared to existing proposals in the literature for addressing these problems, our approach is inherently different and built on the desirable properties of tradable factor risk premia for testing weakly identified and potentially misspecified asset pricing models.

We follow Balduzzi and Kallal (1997) and define a tradable factor risk premium as the negative factor covariance with the (tradable) projection of any SDF on the space of asset returns. They exploit this factor risk premium

[2]See Asness et al. (2019), Daniel et al. (2020), Daniel et al. (2020), Langlois (2020), Gonçalves (2021), Ehsani and Linnainmaa (2022) and Gormsen and Lazarus (2023), among many others.

[3]See again Kan and Zhang (1999a), Kan and Zhang (1999b), Kleibergen (2005), Kleibergen (2009), Gospodinov et al. (2014),Gospodinov et al. (2017), Kleibergen and Zhan (2020) and Kleibergen and Zhan (2022).

definition to show that any SDF pricing a candidate factor satisfies a tighter variance bound than the Hansen and Jagannathan (1991) bound, in order to build sharper misspecification tests for asset pricing models. In contrast, we take tradable factor risk premia as naturally relevant objects of interest in weakly identified and potentially misspecified models, in order to obtain reliable Oracle inference and factor selection procedures for these settings.

As emphasized by Breeden et al. (1989), Huberman et al. (1987), Balduzzi and Robotti (2008), and Balduzzi and Robotti (2010), among many others, working with tradable factor risk premia in well-identified asset pricing settings can have a number of advantages.[4] However, it is important to recognize that conventional estimators of these risk premia are not immune from the effects of a model's weak identification. First, in these settings factor mimicking portfolio betas are not identified, which threatens inference based on two-pass cross-sectional regression methods also when factors are replaced by mimicking portfolio returns (Kleibergen and Zhan (2018)). Second, even though the expected returns of mimicking portfolios of weak factors are zero by definition, standard inference based on sample tradable risk premia may indicate a non zero premium spuriously.

We investigate these features in greater depth and show (i) that the asymptotic distribution of sample tradable risk premia is Gaussian also under a model's weak identification, (ii) that the distribution of individual factor sample risk premia is unaffected by the joint distribution of other factors with asset returns, (iii) that a spurious inference arises asymptotically only for the risk premia of weakly factors, and (iv) that the resulting estimator's mean square error is unchanged under a model's weak identification. These key properties allow us to adjust the inference distortions generated by weak factors with a relatively simple Oracle estimator, which is defined by a closed-form minimum distance correction of sample tradable risk premia. Crucially, besides the consistent selection of factors that are neither useless nor weak, this estimator also provides a uniform efficient inference for the tradable risk premia of all other factors in an asset pricing model.

The implications of weak identification and misspecification for two-step cross-sectional estimation and inference approaches have been studied in a series of important contributions. Kan et al. (2013) introduce misspecification-robust factor risk premia, defined as the minimizers of cross-

[4]They include, e.g., the possibility of applying two-step cross-sectional regression methods with nontradable factors replaced by factor mimicking portfolio returns, a straighforward interpretation and computation of factor risk premia as the expected returns of factor mimicking portfolios, enhanced factor correlations with test asset returns, or a potentially improved finite sample estimation and inference.

sectional pricing errors under a corresponding metric, within a well-identified but potentially misspecified asset pricing model. They then derive inference procedures for risk premia and goodness of fit statistics that are robust to misspecification. Gospodinov et al. (2014) further consider models including some useless factors. They show that estimated factor risk premia follow a nonstandard asymptotic distribution, which implies spuriously positive assessments of pricing performance and useless factor risk premium significance, whenever a model misspecification and/or identification failure is ignored. To overcome these problems, they propose a sequential testing procedure that first consistently eliminates the useless factors, in order to restore the standard inference for the risk premia of not useless factors and for the null of a correct model specification. Bryzgalova (2015) proposes instead a penalized two-step cross-sectional Oracle estimator of misspecification robust risk premia. This estimator consistently removes the useless/weak factors and provides at the same time an (asymptotically) efficient inference for the risk premia of all other factors, which is however not uniform.[5]

We add to this literature with a different approach, which fully turns to good use the fact that in weakly identified and potentially misspecified models tradable risk premia are always well-defined. This property directly implies asymptotic distributions of sample tradable risk premia that are less affected by the presence of weak factors, and hence more practicable, allowing us to build easily applicable Oracle estimators with good factor selection and efficiency properties, both asymptotically and in finite samples. Contrary to Gospodinov et al. (2014), in our approach the consistent selection of all weak/useless factors is performed at once, and the associated (conditional) inference on the tradable risk premia of other factors is automatically uniformly efficient.[6] Differently from Bryzgalova (2015), we consciously aim for Oracle estimation and inference of tradable risk premia, as we find that the nonstardard asymptotic properties of misspecification-robust risk premium estimators also harm Oracle estimators built on them, when some factors are weak.[7]

[5] As we clarify in Section 5, her factor risk premium estimator arises from a penalized weighted minimum distance correction of misspecification robust risk premia, in which a data-driven penalty consistently eliminates all factors that are weak/useless.

[6] Alternatively, after our consistent selection of weak/useless factors it is also possible to apply two-step cross-sectional estimation and inference approaches to produce valid (conditional) inference on the misspecification-robust risk premia of all other factors in a model.

[7] These nonstandard asymptotic properties make a uniformly efficient Oracle inference for misspecification-robust risk premia hardly feasible. Such lack of uniformity may partly explain why these Oracle estimators produce a poor finite sample factor selection and

The above literature focuses on procedures performing a consistent selection of weak/useless factors, and at the same time providing an efficient inference for the risk premia of other factors. In contrast, Kleibergen (2009), Kleibergen and Zhan (2018), and Kleibergen and Zhan (2020) build on various extensions of the Gibbons-Ross-Shanken statistic to form valid inference procedures for all factor risk premia in potentially weakly-identified, but correctly specified, models.[8] Kleibergen and Zhan (2022) make use of the objective function of a continuous updating two-step cross-sectional regression estimator, in order to construct size correct Lagrange multiplier tests of misspecification-robust factor risk premia under weak identification.[9]

We contribute to this literature with novel uniformly efficient Oracle inference procedures for tradable factor risk premia. Compatibly with the findings in Kleibergen and Zhan (2022), our empirical evidence shows that controlling for the complex interaction between misspecification and weak-identification is key for producing a reliable inference on factor risk premia. Our approach achieves such control by exploiting the desirable properties of tradable factor risk premia, which allow us to perform at the same time a consistent selection of the weak/useless factors and an efficient inference on the risk premia of the other factors. With this approach, we further isolate a robust set of well-identified asset pricing models from the factor zoo, built out exclusively of factors with a nonzero tradable risk premium. In parallel, we find that other established definitions of a factor risk premium in the literature may be hard to interpret in such settings, due to the intrinsic misspecification of well-identified models.

Our approach naturally complements also recent approaches testing asset pricing models in high-dimensional environments. Giglio and Xiu (2021) build on an exact linear asset pricing model defined by a small set of pervasive unobservable factors, which are consistently estimated with standard methods.[10] The risk premia of the unobservable factors are estimated with a cross-sectional regression of assets' average returns onto the estimated unobservable factors' exposures. Finally, the risk premium of any observable factor is estimated as the risk premium of the observable factor's mimick-

inference for a number of relevant factor models in our Monte Carlo simulations.

[8]By construction, these procedures deliver empty factor risk premium confidence sets when a model is misspecified.

[9]They further offer a diagnostics criterion measuring the relative strength of misspecification and identification, which indicates when a candidate risk premium is interpretable as a factor risk premium.

[10]Equivalently, this setting assumes identical minimum variance SDF projections on the spaces spanned by returns and pervasive factors, respectively, i.e., a correctly specified linear asset pricing model built only with the pervasive factors.

ing portfolio built from the estimated unobservable factors. In such high-dimensional environments, out methodology offers a consistent selection of observable factors that are neither useless nor weak with respect to the unobserved pervasive factors. It does so, by applying our Oracle estimator of tradable factor risk premia to test assets given by the estimated pervasive factors. Coherently with our theory, the tradable risk premium of any observable factor is then given by the negative factor covariance with the minimum variance SDF projection of the space of unobservable pervasive factors. This risk premium is well-defined and economically clearly interpretable, independent of whether the minimum variance SDF projections on the space spanned by asset returns and the space spanned by the pervasive factors coincide.

3 Factor risk premia

Let $\boldsymbol{R}_t = [R_{1t}, \ldots, R_{Nt}]'$ and $\boldsymbol{F}_t := [F_{1t}, \ldots, F_{Kt}]'$, with $K \leq N$, be a vector of excess returns and a vector of candidate asset pricing factors observed at times $t = 1, \ldots, T$. The joint vector $\boldsymbol{Y}_t := [\boldsymbol{R}_t', \boldsymbol{F}_t']'$ has moments partitioned as:

$$\mathbb{E}\left[\boldsymbol{Y}_t\right] = \begin{bmatrix} \boldsymbol{\mu}_R \\ \boldsymbol{\mu}_F \end{bmatrix} \quad ; \quad \mathbb{C}\mathrm{ov}\left[\boldsymbol{Y}_t, \boldsymbol{Y}_t\right] = \begin{bmatrix} \boldsymbol{V}_R & \boldsymbol{V}_{RF} \\ \boldsymbol{V}_{FR} & \boldsymbol{V}_F \end{bmatrix} .$$

Given a set of test asset returns and candidate asset pricing factors, the main question studied by empirical asset pricers is whether some factors may reflect sources of systematic risk that are priced in the cross-section of asset returns. In order to answer this question empirically, one crucially needs both an appropriate factor risk premium definition and reliable inference techniques for testing hypotheses about factor risk premia.

Existing notions of factor risk premia all rely on the basic asset pricing equation implied by the choice of a corresponding Stochastic Discount Factor (SDF) projection M_t:

$$\boldsymbol{\lambda}(M) := -\mathbb{C}ov[\boldsymbol{F}_t, M_t] \ . \tag{1}$$

In all cases, SDF projection M_t is defined as some suitable approximation of a corresponding unknown SDF in an arbitrage-free market. Such SDF projections can be quite different with respect to several important features, including, e.g., the type of risks they span and the optimality criteria used to build them. More importantly for our purposes, most of these projections directly depend on the properties of an asset pricing model and are not well-defined in not or weakly identified models.

3.1 Tradable factor risk premia

In order to define tradable factor risk premia, the only assumption needed is the nonredundancy of asset returns.

Assumption 1. Covariance matrix $\boldsymbol{V}_R$ is positive definite.

Assumption 1 is standard and ensures a finite upper bound on the largest Sharpe ratio of any excess return portfolio. This bound is attained by the maximum Sharpe ratio portfolio with weights:

$$\boldsymbol{w} = \boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ . \tag{2}$$

Portfolio (2) identifies the tradable minimum variance projection of any stochastic discount factor (SDF) onto the space spanned by excess returns and constant payoffs:[11]

$$M_t^* := 1 - \boldsymbol{\mu}_R'\boldsymbol{V}_R^{-1}(\boldsymbol{R}_t - \boldsymbol{\mu}_R) \ . \tag{3}$$

By construction, this projection satisfies the basic risk premium identity:

$$\mathbb{E}[\boldsymbol{R}_t] = -\,\mathbb{C}\text{ov}[\boldsymbol{R}_t, M_t^*] \ .$$

Consistently with these pricing properties, a tradable factor risk premium is the risk premium assigned to a candidate asset pricing factor by SDF projection (3).

Definition 1. Given Assumption 1, the tradable risk premium of factor vector $\boldsymbol{F}_t$ is defined by:

$$\boldsymbol{\lambda}^* := -\,\mathbb{C}\text{ov}[\boldsymbol{F}_t, M_t^*] = \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ .$$

The tradable factor risk premium in Definition 1 is a special case of general risk premium definition (1) with $M_t := M_t^*$ and it satisfies a number of desirable properties listed below.

(i) It is well-defined whenever Assumption 1 holds. In particular, it remains well-defined when covariance matrix $\boldsymbol{V}_{FR}$ does not have a full rank, i.e., $\text{rank}(\boldsymbol{V}_{FR}) < K$. This feature emerges in presence of useless or weak factors or, more generally, when covariances of different factors with returns are collinear.

[11] The normalization $\mathbb{E}[M_t^*] = 1$ is without loss of generality since we work with excess returns.

(ii) It gives rise to individual factor risk premia depending exclusively on the joint factor distribution with test asset returns, but not on the joint distribution of test asset returns with other factors in an asset pricing model.

(iii) It offers a straightforward economic interpretation as it coincides with the expectation of the least squares projection of factors on the span of excess returns:

$$\boldsymbol{\lambda}^* = \mathbb{E}[\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{R}_t] \ ,$$

i.e., tradable factor risk premia are interpretable as the (tradable) risk premia of factor replicating portfolios.

(iv) It uniquely assigns a zero risk premium to any factor having a zero or a vanishing population covariance with all excess returns, i.e., any useless or weak factor.

(v) It is independent of the misspecification properties of a factor asset pricing model, because it is directly defined via the SDF projection on the space of returns.

(vi) When factors are not redundant ($\boldsymbol{V}_F$ is positive definite), it equals the negative factor covariance with an associated linear SDF projection on the factor space. Precisely, we obtain $\boldsymbol{\lambda}^* = -\mathbb{C}\mathrm{ov}[\boldsymbol{F}_t, M_t(\boldsymbol{\gamma}^*)]$ for an SDF projection $M_t(\boldsymbol{\gamma}^*) := 1 - \boldsymbol{\gamma}^{*\prime}(\boldsymbol{F}_t - \boldsymbol{\mu}_F)$ such that:

$$\boldsymbol{\gamma}^* = \boldsymbol{V}_F^{-1}\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ . \tag{4}$$

(vii) It is invariant to linear invertible transformations of excess returns, i.e., to a simple repackaging of the test assets.[12]

3.2 Alternative notions of factor risk premia

Tradable factor risk premia differ from other commonly used notions of factor risk premia in the literature, because of the different SDF projection underlying them. Indeed, as explained above tradable risk premia are defined by a linear SDF projection on the return space, which correctly prices all test assets. In contrast, other prominent notions of factor risk premium

[12] A detailed discussion of this important invariance property is given in Kandel and Stambaugh (1995).

are defined by a SDF projection on the factor space, which can imply some degree of asset mispricing in misspecified models.

To clarify these differences within a single coherent framework, we borrow from the definition of a misspecification-robust factor risk premium in Kan et al. (2013):

$$\lambda(\boldsymbol{\gamma}_{\boldsymbol{W}}) := -\,\mathbb{C}\mathrm{ov}[\boldsymbol{F}_t, M_t(\boldsymbol{\gamma}_{\boldsymbol{W}})] \ , \tag{5}$$

where $M_t(\boldsymbol{\gamma}_{\boldsymbol{W}}) := 1 - \boldsymbol{\gamma}_{\boldsymbol{W}}'(\boldsymbol{F} - \boldsymbol{\mu}_F)$ is an associated SDF projection on the factor space. This projection is defined by factor loadings that minimize pricing errors under the metric induced by some $N \times N$ symmetric and positive-definite matrix $\boldsymbol{W}$:

$$\boldsymbol{\gamma}_{\boldsymbol{W}} \in \underset{\boldsymbol{\gamma} \in \mathbb{R}^K}{\operatorname{argmin}} \ \mathbb{E}[M_t(\boldsymbol{\gamma})\boldsymbol{R}_t]'\boldsymbol{W}\mathbb{E}[M_t(\boldsymbol{\gamma})\boldsymbol{R}_t] \ . \tag{6}$$

It follows that factor risk premium (5) is another special case of general factor risk premium definition (1). However, it is crucial to emphasize that factor risk premium (5) is identified if and only if factor loadings (6) are identified. Clearly, this happens if and only if covariance matrix $\boldsymbol{V}_{FR}$ has full rank, in which case the factor loadings are uniquely given by:

$$\boldsymbol{\gamma}_{\boldsymbol{W}} = (\boldsymbol{V}_{FR}\boldsymbol{W}\boldsymbol{V}_{RF})^{-1}\boldsymbol{V}_{FR}\boldsymbol{W}\boldsymbol{\mu}_R \ . \tag{7}$$

Therefore, in identified asset pricing models one obtains following standard two-step weighted least squares expression for factor risk premium (5):

$$\boldsymbol{\lambda}(\gamma_{\boldsymbol{W}}) = \boldsymbol{V}_F\boldsymbol{\gamma}_{\boldsymbol{W}} = (\boldsymbol{\beta}'\boldsymbol{W}\boldsymbol{\beta})^{-1}\boldsymbol{\beta}'\boldsymbol{W}\boldsymbol{\mu}_R \ , \tag{8}$$

where $\boldsymbol{\beta} = \boldsymbol{V}_{RF}\boldsymbol{V}_F^{-1}$ is the $N \times K$ matrix of factor betas.[13] The next example illustrates in a simplified setting the basic identification issues underlying two-step weighted least squares factor risk premia and the relation to tradable factor risk premia.

Example 1. Consider a two-factor asset pricing model with factors F_1, F_2 having proportional covariances with test asset returns: $\mathbb{C}\mathrm{ov}[F_{2t}, \boldsymbol{R}_t] = m\,\mathbb{C}\mathrm{ov}[F_{1t}, \boldsymbol{R}_t]$ for some $m > 0$. In this case, $N \times 2$ matrix $\boldsymbol{V}_{RF}$ has column rank equal to 1. As a consequence, 2×2 matrices $\boldsymbol{V}_{FR}\boldsymbol{W}\boldsymbol{V}_{RF}$ and $\boldsymbol{\beta}'\boldsymbol{W}\boldsymbol{\beta}$ are singular and both factor loading (7) and two-step factor risk

[13] Here, the factor risk premium induced by Fama and MacBeth (1973) two-pass procedures emerges for an identity weighting matrix $\boldsymbol{W} = \boldsymbol{I}$.

premium (8) are not identified. In contrast, the tradable factor risk premium $\boldsymbol{\lambda}^* = \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R$ is still well-defined and naturally incorporates the covariance redundancies implied by the two factors:

$$\lambda_2^* = \boldsymbol{V}_{F_2R}\boldsymbol{V}_R^{-1}\boldsymbol{\mu} = m\boldsymbol{V}_{F_1R}\boldsymbol{V}_R^{-1}\boldsymbol{\mu} = m\lambda_1^* \ .$$

Since factor mimicking portfolio returns $\tilde{\boldsymbol{F}}_t := \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{R}_t$ are still well defined when factor covariances are redundant, one may be tempted to solve the identification issue using the two-step factor risk premia of the factor mimicking portfolios. Unfortunately, this approach does not solve either the identification issue arising for two-step factor risk premia. This is so, because the covariance matrix $\boldsymbol{V}_{\tilde{F}} = \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF}$ of the factor mimicking portfolio returns is singular and hence the factor mimicking portfolio betas are not identified. As a consequence, in unidentified models tradable factor risk premia cannot be recovered from the two-step factor risk premia of factor mimicking portfolios.

As Example 1 shows, the identification issue of two-step factor risk premia arises when covariance matrix $\boldsymbol{V}_{FR}$ is not full rank, because then SDF projection coefficient (6) is not uniquely defined. In contrast, SDF projection coefficient (4) underlying tradable factor risk premia is always well-defined, independent of the rank properties of $\boldsymbol{V}_{FR}$. The next proposition adds insight into the relation and the differences between SDF projection coefficients (4) and (6).

Proposition 3.1. *Let Assumption 1 be satisfied and $\boldsymbol{V}_F$ be positive definite. Then, SDF projection coefficient (4) is equivalently given by:*

$$\boldsymbol{\gamma}^* = \underset{\boldsymbol{\gamma}\in\mathbb{R}^K}{\operatorname{argmin}} \ \mathbb{E}[M_t(\boldsymbol{\gamma})\boldsymbol{R}_t]'\boldsymbol{V}_R^{-1}\mathbb{E}[M_t(\boldsymbol{\gamma})\boldsymbol{R}_t] + \boldsymbol{\gamma}'\boldsymbol{A}\boldsymbol{\gamma} \ , \tag{9}$$

with the symmetric positive semi-definite matrix:

$$\boldsymbol{A} := \boldsymbol{V}_F - \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF} \ .$$

Proposition 3.1 states that SDF projection coefficient $\boldsymbol{\gamma}^*$ can be obtained by solving a minimum pricing error problem with weighting matrix $\boldsymbol{W} = \boldsymbol{V}_R^{-1}$, which is complemented by a quadratic penalty $\boldsymbol{\gamma}'\boldsymbol{A}\boldsymbol{\gamma}$. This penalty is larger precisely when some factors are less accurately replicated by their factor mimicking portfolios and thus naturally penalizes situations where the variance covariance matrix $\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF}$ of the factor mimicking portfolios may be singular due to a model non identification.

Note that the penalty in equation (9) vanishes precisely when factors are spanned by test asset returns (i.e., they are factor mimicking portfolio returns) and the asset pricing model is identified, in which case $\boldsymbol{\gamma}^* = \boldsymbol{\gamma}_{\boldsymbol{V}_R^{-1}}$ and $\boldsymbol{\lambda}^* = \boldsymbol{\lambda}(\boldsymbol{\gamma}_{\boldsymbol{V}_R^{-1}})$. If instead the model is identified but not all factors are spanned, then equation (8) gives:[14]

$$\boldsymbol{\lambda}(\boldsymbol{\gamma}_{\boldsymbol{V}_R^{-1}}) = \boldsymbol{V}_F(\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF})^{-1}\boldsymbol{\lambda}^* \ .$$

This identity further implies that in an identified model the tradable risk premium of any spanned factor $F_k = \boldsymbol{V}_{F_k R}\boldsymbol{V}_R^{-1}\boldsymbol{R}$ is identical to its two step factor risk premium: $\lambda_k^* = \lambda_k(\boldsymbol{\gamma}_{\boldsymbol{V}_R^{-1}})$.[15] Therefore, misspecification-robust risk premia differ from tradable risk premia by the fact that they assign to not spanned factors a rescaled risk premium, which is adjusted to the larger variance covariance matrix of not spanned factors relative to the variance-covariance matrix of factor mimicking portfolio returns.

Having clarified the desirable properties of tradable factor risk premia when working with potentially unidentified models, relative to other common notions in the literature, we study in the next section the asymptotic properties of a convenient class of tradable factor risk premium estimators, under varying assumptions on the data generating process for returns and potentially useless or weak asset pricing factors. Building on these estimators, we then set up our Oracle tests of asset pricing models.

4 Inference for tradable factor risk premia

Definition 1 naturally motivates a direct estimator of tradable factor risk premia, which is given by the sample version of equation (1):[16]

$$\hat{\boldsymbol{\lambda}} := \hat{\boldsymbol{V}}_{FR}\hat{\boldsymbol{V}}_R^{-1}\hat{\boldsymbol{\mu}}_R \ , \tag{10}$$

[14] In models that are additionally correctly specified, this identity extends to factor risk premia implied by Fama-MacBeth two-step regression approaches, given their equivalence to Kan et al. (2013) misspecification-robust factor risk premia in such settings.

[15] Denoting by $\boldsymbol{e}_k$ the $k-$unit vector, this property follows by noting that:

$$\lambda_k(\boldsymbol{\gamma}_{\boldsymbol{V}_R^{-1}}) = \boldsymbol{e}_k'\boldsymbol{V}_F(\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF})^{-1}\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu} = \boldsymbol{e}_k'\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu} = \lambda_k^* \ ,$$

since by construction $\boldsymbol{e}_k'\boldsymbol{V}_F = \boldsymbol{e}_k'\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF}$.

[16] All the methods developed in this section are implemented in the Rcpp package intrinsicFRP, available on CRAN.

with the sample mean and sample covariance matrix estimators given by:

$$\hat{\boldsymbol{\mu}} := [\hat{\boldsymbol{\mu}}_R', \hat{\boldsymbol{\mu}}_F']' := \frac{1}{T}\sum_{t=1}^{T} \boldsymbol{Y}_t \ ,$$

and

$$\hat{\boldsymbol{V}} := \begin{bmatrix} \hat{\boldsymbol{V}}_R & \hat{\boldsymbol{V}}_{RF} \\ \hat{\boldsymbol{V}}_{FR} & \hat{\boldsymbol{V}}_F \end{bmatrix} := \frac{1}{T}\sum_{t=1}^{T} (\boldsymbol{Y}_t - \hat{\boldsymbol{\mu}})(\boldsymbol{Y}_t - \hat{\boldsymbol{\mu}})' \ .$$

In the sequel, we show that estimator (10) has a number of useful properties, which make it a convenient initial estimator for building tests of potentially misspecified and not necessarily identified asset pricing models. The first obvious property follows directly from definition (10):

$$\hat{\boldsymbol{\lambda}} = \frac{1}{T}\sum_{t=1}^{T} \hat{\boldsymbol{V}}_{FR}\hat{\boldsymbol{V}}_R^{-1} R_t \ ,$$

i.e., sample tradabable risk premia are the average returns of the estimated factor mimicking portfolios. This feature directly implies that the sample tradable risk premium of any factor spanned by the test asset returns is automatically the factor's average return. Conversely, if a factor is spanned by a set of asset returns and it appears as weak or useless with respect to a given set of test assets, then it can deliver a nonzero risk premium (expected excess return) only if it has exposure to sources of systematic risks uncorrelated with test asset returns.

4.1 Properties of sample tradable factor risk premia

Further important and less direct properties are introduced and discussed below. To this end, we prove in Section 4.1.1 that estimator (10) always has a standard asymptotically Gaussian distribution, even in the presence of useless factors uncorrelated with returns, or, more generally, weak factors having a correlation with returns that vanishes sufficiently fast with the sample size.[17] Furthermore, we show in Section 4.1.2 that in presence of weak factors having return correlations vanishing at a slow rate,[18] the uniform asymptotic distribution of estimator (10) is Gaussian, but dependent on an asymptotic bias term. Crucially, this bias only affects the estimated

[17] I.e., at rate at least $1/T^{\alpha}$ for some $\alpha > 1/2$.

[18] I.e., at rate $1/\sqrt{T}$.

risk premia of the weak factors with slowly vanishing return correlations. Therefore, the uniform asymptotically Gaussian distribution of estimator (10) for the estimated risk premia all other factors is entirely unaffected by the presence of weak factors.

By exploiting the asymptotic properties of sample tradable risk premia, we next build in Section 4.2 an Oracle estimator of tradable factor risk premia. This estimator is given in closed-form by a penalized minimum distance correction of sample tradable risk premia, which corrects the potential biases generated by weak factors with slowly vanishing return correlations. We show that our Oracle estimator improves on sample tradable risk premia in two directions. First, by ensuring a consistent finite-sample selection of factors that are not weak or useless. Second, by giving rise to a uniform efficient asymptotic distribution for the estimated risk premia of all other factors, which is independent of the potential presence of useless or weak factors in an asset pricing model. Therefore, our Oracle estimator naturally produces a valid asymptotic inference framework for testing potentially misspecified and not or only weakly identified asset pricing models.

4.1.1 Asymptotic distribution with useless factors

Let the vectors of excess return and factor innovations at time t be denoted by:

$$\begin{aligned} \bar{\boldsymbol{R}}_t &:= \boldsymbol{R}_t - \boldsymbol{\mu}_R \ , \\ \bar{\boldsymbol{F}}_t &:= \boldsymbol{F}_t - \boldsymbol{\mu}_F \ . \end{aligned}$$

The next proposition establishes the asymptotic properties of tradable factor risk premium estimator (10), under the general assumption of factor and excess return stationarity and ergodicity.[19]

Proposition 4.1. *Let Assumption 1 hold and $\mathbf{Y}_t$ be jointly stationary and ergodic with finite fourth moments. Let further matrix:*

$$\boldsymbol{\Sigma} := \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{s=1}^{T} \mathbb{E}[\boldsymbol{h}_t \boldsymbol{h}_s'] \ , \tag{11}$$

where

$$\boldsymbol{h}_t := \bar{\boldsymbol{F}}_t \bar{\boldsymbol{R}}_t' \boldsymbol{V}_R^{-1} \boldsymbol{\mu}_R - \boldsymbol{V}_{FR} \boldsymbol{V}_R^{-1} \bar{\boldsymbol{R}}_t \bar{\boldsymbol{R}}_t' \boldsymbol{V}_R^{-1} \boldsymbol{\mu}_R + \boldsymbol{V}_{FR} \boldsymbol{V}_R^{-1} \bar{\boldsymbol{R}}_t \ , \tag{12}$$

[19] This assumption, together with the assumption of finite fourth moments of random vector $\boldsymbol{Y}_t$, is needed to invoke a suitable central limit theorem for jointly stationary and ergodic random vector $((\boldsymbol{Y}_t - \mathbb{E}[\boldsymbol{Y}_t])', \mathrm{vec}(\boldsymbol{Y}_t \boldsymbol{Y}_t' - \boldsymbol{V})')'$.

be positive definite. It then follows:

$$\sqrt{T}(\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^*) \to_d \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}) \ . \tag{13}$$

In Proposition 4.1, the asymptotic covariance matrix (11) depends on three sources of estimation uncertainty, which are present in the sum on the right hand side of equation (12). They capture the uncertainty generated by the estimation of population moments $\boldsymbol{V}_{FR}$, $\boldsymbol{V}_R$ and $\boldsymbol{\mu}_R$ in the definition of tradable factor risk premia.

Several useful properties of asymptotic distribution (13) are worth noticing. First, the marginal asymptotic distribution of sample tradable risk premia only depends on the joint factor distribution with asset returns, but not on the joint distribution with other factors. This property does not hold in general, e.g., for factor risk premium estimators implied by two-step cross-sectional regression approaches. Second, in Proposition 4.1 the $K \times N$ matrix $\boldsymbol{V}_{FR}$ is not required to have a full rank. Therefore, asymptotic distribution (13) is valid also in presence of useless factors and one can directly test with statistics having this asymptotic distribution the null hypothesis of a zero tradable risk premium. This is a sharp difference with, e.g., the approach in Gospodinov et al. (2014), in which useless factors have first to be consistently eliminated with a sequential testing procedure, before testing risk premium hypotheses involving factors that are not useless.

4.1.2 Uniform asymptotic distribution with weak factors

The stationarity assumption in Proposition 4.1 is not suited to cover weak factor settings, in which some factors may exhibit a non degenerate covariance structure with returns in finite samples but may give rise to a degenerate covariance structure as the sample size grows. Such settings have been shown in Kleibergen (2009), among many others, to be empirically relevant for many macro-based asset pricing factors, as well as for potentially data-mined factors from the factor zoo (Harvey et al. (2016)).

Studying the properties of the asymptotic distribution of sample tradable risk premia in presence of weak factors is straightforward with our approach. This is so, because weak factor features are easily embedded as a special case of a sequence of asymptotically vanishing local alternatives on population tradable risk premia. These local alternatives are parametrized using sequences of factor covariance matrices with returns, which may differ from their population limit $\boldsymbol{V}_{FR}$.[20]

[20] Equivalently, nearly-singular weak factor settings are usually defined in the literature

Assumption 2. The sample size-dependent joint covariance matrix of factors and returns is given by:

$$\boldsymbol{V}^{(T)} = \begin{bmatrix} \boldsymbol{V}_R & \boldsymbol{V}_{RF}^{(T)} \\ \boldsymbol{V}_{FR}^{(T)} & \boldsymbol{V}_F \end{bmatrix} ,$$

where submatrix $\boldsymbol{V}_{FR}^{(T)}$ is of the form:

$$\boldsymbol{V}_{FR}^{(T)} := \boldsymbol{V}_{FR} + \boldsymbol{\Gamma}/T^{\alpha} + \boldsymbol{\Delta}/\sqrt{T} ,$$

for some $\alpha > 1/2$ and corresponding $K \times N$ matrices $\boldsymbol{\Gamma}$, $\boldsymbol{\Delta}$.

In Assumption 2, matrices $\boldsymbol{\Gamma}$ and $\boldsymbol{\Delta}$ parametrize the way how covariance matrix $\boldsymbol{V}_{FR}^{(T)}$ converges to its (possibly reduced-rank) limit $\boldsymbol{V}_{FR}$. We denote by

$$\mathcal{S} := \{k \in \{1, \ldots, K\} \ : \ \boldsymbol{V}_{F_k R} \neq \boldsymbol{0}\} , \tag{14}$$

the set indexing factors that are neither weak nor useless. Analogously, for any vector $\boldsymbol{v} \in \mathbb{R}^K$ (matrix $\boldsymbol{M} \in \mathbb{R}^{K \times N}$), let $\boldsymbol{v}_{\mathcal{S}}$ ($\boldsymbol{M}_{\mathcal{S}}$) be the subvector (submatrix) consisting only of the rows of $\boldsymbol{v}$ ($\boldsymbol{M}$) with index in $\mathcal{S}$.

With this notation, weak or useless factors emerge in Assumption 2 simply when set $\mathcal{S}^c$ is not empty, i.e., some row of matrix $\boldsymbol{V}_{FR}$ is a zero vector. Furthermore, given a factor with index $k \in \mathcal{S}^c$, such factor is weak with rate $1/2$ (not larger than α) if and only if $\boldsymbol{\Delta}_k \neq \boldsymbol{0}$ ($\boldsymbol{\Gamma}_k \neq \boldsymbol{0}$). Therefore, nonzero matrices $\boldsymbol{\Delta}_{\mathcal{S}^c}$ and $\boldsymbol{\Gamma}_{\mathcal{S}^c}$ indeed model weak factors having slowly and fastly vanishing covariances with returns, respectively.

A related, but conceptually different interpretation applies in Assumption 2 for cases where $\boldsymbol{\Delta}_{\mathcal{S}} \neq 0$. Indeed, since $\mathcal{S}$ indexes factors that are neither useless nor weak, the tradable risk premium of these factors admits an economically founded interpretation as the risk premium of a not degenerate factor mimicking porfolio:

$$\boldsymbol{\lambda}_{\mathcal{S}}^* = \boldsymbol{V}_{F_{\mathcal{S}} R} \boldsymbol{V}_R^{-1} \boldsymbol{\mu}_R . \tag{15}$$

Accordingly, submatrix $\boldsymbol{\Delta}_{\mathcal{S}}$ in Assumption 2 incorporates a sequence of local alternative with respect to a null hypothesis of the form (15):

$$\boldsymbol{\lambda}_{\mathcal{S}}^{*(T)} := \boldsymbol{V}_{F_{\mathcal{S}} R}^{(T)} \boldsymbol{V}_R^{-1} \boldsymbol{\mu}_R = \boldsymbol{\lambda}_{\mathcal{S}}^* + \frac{1}{\sqrt{T}} \boldsymbol{\Delta}_{\mathcal{S}} \boldsymbol{V}_R^{-1} \boldsymbol{\mu}_R . \tag{16}$$

---

using sequences $\boldsymbol{\beta}^{(T)} := \boldsymbol{V}_{RF}^{(T)} \boldsymbol{V}_F^{-1}$ of matrices of factor exposures converging to their population limit at a corresponding rate.

Along such a sequence of local alternatives, it is highly desirable that our Oracle inference methodology for tradable factor risk premia uniformly exhibits power. Therefore, in the next proposition we establish the uniform asymptotic distribution of tradable factor risk premium estimator (10) over sequences of local alternatives depicted in Assumption 2, under the general assumption of factor and excess return near-epoch dependence.[21]

Proposition 4.2. *Let Assumptions 1 and 2 hold, $\boldsymbol{Y}_t$ be near-epoch dependent with finite fourth moments, and regularity conditions detailed in Davidson (1992) hold. Let further matrix $\boldsymbol{\Sigma}$ defined in Proposition 4.1 be positive definite. It then follows:*

$$\sqrt{T}(\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^*) \to_d \mathcal{N}\left(\boldsymbol{b}(\boldsymbol{\Delta}), \boldsymbol{\Sigma} - \boldsymbol{b}(\boldsymbol{\Delta})\boldsymbol{b}'(\boldsymbol{\Delta})\right) \ , \tag{17}$$

where

$$\boldsymbol{b}(\boldsymbol{\Delta}) := \boldsymbol{\Delta}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ .$$

Proposition 4.2 first shows that weak factors with return covariances vanishing at fast rate $1/T^{\alpha}$ with $\alpha > 1/2$ leave completely unaffected the asymptotic distribution of sample tradable factor risk premia.[22] This property is remarkable, because such weak factor features do instead crucially affect the asymptotic distribution of risk premium estimators based on two-step cross-sectional regression approaches.

A distinct situation emerges with respect to weak factors with slowly vanishing return covariances ($\boldsymbol{\Delta}_{\mathcal{S}^c} \neq 0$). Indeed, a comparison of asymptotic distributions (13) and (17) shows that the latter distribution displays both a noncentrality parameter $\boldsymbol{b}(\boldsymbol{\Delta})$ such that:

$$\begin{bmatrix} \boldsymbol{b}_{\mathcal{S}}(\boldsymbol{\Delta}) \\ \boldsymbol{b}_{\mathcal{S}^c}(\boldsymbol{\Delta}) \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Delta}_{\mathcal{S}}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \\ \boldsymbol{\Delta}_{\mathcal{S}^c}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \end{bmatrix} \ , \tag{18}$$

and a smaller covariance matrix because of component $-\boldsymbol{b}(\boldsymbol{\Delta})\boldsymbol{b}'(\boldsymbol{\Delta})$. Therefore, weak factors with slowly vanishing return covariances generate a bias $\boldsymbol{b}_{\mathcal{S}^c}(\boldsymbol{\Delta}) = \boldsymbol{\Delta}_{\mathcal{S}^c}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R$ in the asymptotic distribution of their sample tradable risk premia, which implies an asymptotic inference tending to spuriously

[21] This assumption, together with the assumption of finite fourth moments of random vector $\boldsymbol{Y}_t$ and regularity conditions detailed in Davidson (1992), is needed to invoke a suitable central limit theorem for jointly near-epoch dependent random vector $((\boldsymbol{Y}_t - \mathbb{E}[\boldsymbol{Y}_t])', \mathrm{vec}(\boldsymbol{Y}_t\boldsymbol{Y}_t' - \boldsymbol{V}^{(T)})')'$.

[22] This result directly follows from the fact that asymptotic distribution (17) is independent of matrix $\boldsymbol{\Gamma}$.

conclude for a non zero risk premium of these factors. Conversely, the non-centrality parameter $\boldsymbol{b}_{\mathcal{S}}(\boldsymbol{\Delta}) = \boldsymbol{\Delta}_{\mathcal{S}} \boldsymbol{V}_R^{-1} \boldsymbol{\mu}_R$ in asymptotic distribution (17) for the sample tradable risk premia of factors that are neither weak nor useless is independent of wether $\boldsymbol{\Delta}_{\mathcal{S}^c} = 0$ or $\boldsymbol{\Delta}_{\mathcal{S}^c} \neq 0$. Therefore, the resulting asymptotic power of inferences for the tradable risk premia of these factors is independent of whether some weak factors appear in an asset pricing model. Finally, we note that the asymptotic mean square error of this estimator for all factor risk premia is also independent of the presence of weak factors.

These features have important implications for the construction of Oracle tests of tradable risk premia in presence of weak factors. First, sample estimator (10) paired with a consistent selection of weak/useless factors can automatically achieve an efficient Gaussian distribution, uniformly over local alternatives in Assumption 2. Second, such factor selection procedure should aim at removing factors with a zero or vanishing covariance with test asset returns.

4.2 Oracle tradable factor risk premium estimation

This section introduces an Oracle tradable risk premium estimator that removes the useless and weak factors with probability growing with the sample size, and at the same achieves a uniform efficient asymptotic Gaussian distribution over local alternatives depicted in Assumption 2, independently of potential covariance distortions generated by weak factors. Borrowing the terminology from Fan and Li (2001), we define an Oracle tradable risk premium estimator as an estimator satisfying two key properties. First, it consistently selects in finite samples factors that are not weak or useless. Second, it implies an efficient asymptotic distribution for the estimated risk premia of the selected factors.

Thanks to the convenient asymptotic properties of sample tradable risk premia, Oracle tradable risk premium estimators can be built with a simple approach. Denote by

$$\widehat{\boldsymbol{\rho}} := [\widehat{\boldsymbol{\rho}}_1, \ldots, \widehat{\boldsymbol{\rho}}_K] := \widehat{\mathbb{C}\mathrm{or}}[\boldsymbol{R}_t, \boldsymbol{F}_t'] \ ,$$

the $N \times K$ matrix of sample correlations between returns and factors. We propose an Oracle tradable risk premium estimator built by means of a convenient minimum distance correction of sample tradable risk premia, in which estimated risk premia of factors having small sample correlations with all asset returns are shrank using a suitable data-driven penalty.

Definition 2. Given a penalty parameter $\tau_T > 0$, consider the penalized estimator defined by:

$$\check{\boldsymbol{\lambda}}^* := (\check{\lambda}_1^*, \ldots, \check{\lambda}_K^*)' := \underset{\boldsymbol{\lambda} \in \mathbb{R}^K}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}\|_2^2 + \tau_T \sum_{k=1}^{K} \frac{|\lambda_k|}{\|\widehat{\boldsymbol{\rho}}_k\|_2^2} \right\} . \tag{19}$$

Estimator (19) is defined by a penalized minimum Euclidean distance correction of sample tradable risk premia and belongs to the class of proximal estimators studied in Quaini and Trojani (2022). The optimization problem in equation (19) is solvable in closed-form and gives rise to the soft-thresholding formula:

$$\check{\lambda}_k^* = \operatorname{sign}(\widehat{\lambda}_k^*) \max \left\{ |\hat{\lambda}_k^*| - \frac{\tau_T}{\|\hat{\boldsymbol{\rho}}_k\|_2^2}, 0 \right\} \; ; \; k = 1, \ldots K \; . \tag{20}$$

Therefore, estimator $\check{\boldsymbol{\lambda}}^*$ implies a zero estimated risk premium for all factors associated with a sample tradable risk premium $\hat{\lambda}_k^*$ that is smaller in absolute value than scaled penalty parameter $\tau_T / \|\widehat{\boldsymbol{\rho}}_k\|_2^2$.

We next show that an appropriate choice of tuning parameter τ_T in equation (19) implies the Oracle property. Let the estimated set of priced factors implied by estimator $\check{\boldsymbol{\lambda}}^*$ be denoted by

$$\check{\mathcal{S}} := \{k \in \{1, \ldots, K\} \; : \; \check{\lambda}_k^* \neq 0\} \; . \tag{21}$$

Using this notation, we characterize in the next proposition the asymptotic distribution of estimator (19) and its Oracle property.

Proposition 4.3. *Let the assumptions of Proposition 4.2 hold. Let further the sequence of tuning parameter τ_T be such that $\tau_T \sqrt{T} \to 0$ and $\tau_T T \to \infty$ as $T \to \infty$. It then follows, as $T \to \infty$:*

1. *Consistent variable selection:*

$$\mathbb{P}(\check{\mathcal{S}} = \mathcal{S}) \to 1 \; ,$$

with $\mathcal{S}$ as defined in (14).

2. *Asymptotic distribution:*

$$\sqrt{T} \left(\check{\boldsymbol{\lambda}}_{\mathcal{S}}^* - \boldsymbol{\lambda}_{\mathcal{S}}^* \right) \to_d \mathcal{N}(\boldsymbol{b}_{\mathcal{S}}(\boldsymbol{\Delta}), \boldsymbol{\Sigma}_{\mathcal{S}\mathcal{S}} - \boldsymbol{b}_{\mathcal{S}}(\boldsymbol{\Delta}) \boldsymbol{b}_{\mathcal{S}}'(\boldsymbol{\Delta})) \; ,$$

with $\boldsymbol{b}_{\mathcal{S}}(\boldsymbol{\Delta})$ defined in (18) and

$$\boldsymbol{\Sigma}_{\mathcal{S}\mathcal{S}} := \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{s=1}^{T} \mathbb{E}[\boldsymbol{h}_{\mathcal{S}t} \boldsymbol{h}_{\mathcal{S}s}'] \; ,$$

where

$$\boldsymbol{h}_{\mathcal{S}t} := \boldsymbol{V}_{F_{\mathcal{S}}R}\boldsymbol{V}_R^{-1}\bar{\boldsymbol{R}}_t - \boldsymbol{V}_{F_{\mathcal{S}}R}\boldsymbol{V}_R^{-1}\bar{\boldsymbol{R}}_t\bar{\boldsymbol{R}}_t'\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R + \bar{\boldsymbol{F}}_{\mathcal{S}t}\bar{\boldsymbol{R}}_t'\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ .$$

A direct consequence of Proposition 4.3 is that proximal tradable risk premium estimator (19) is an Oracle estimator. Indeed, Property 1 states that the set of factors that are neither useless nor weak is correctly selected with a probability converging to one as the sample size grows. Property 2 further implies that the asymptotic distribution of the estimated risk premia of factors that are neither useless nor weak is identical to the asymptotic distribution of a sample tradable risk premium estimator knowing ex ante which factors are useless or weak, uniformly over the local alternatives in Assumption 2.[23] In this way, our Oracle estimator directly produces a valid inference for the tradable risk premia of factors that are not weakly correlated with returns. Alternatively, it can be used to perform a consistent preliminary screening of factors weakly correlated with returns, in order to reliably apply in a second step standard cross-sectional inference methodologies to the risk premia of all other factors in an asset pricing model.

Remark 1. The fact that the asymptotic distribution of our Oracle estimator in Proposition 4.3 is uniform over local alternatives in Assumption 2 is an important property and not in contradiction with the well-known impossibility of hard- and soft-thresholding estimators to produce uniform convergence at rate $1/\sqrt{T}$ and consistent model selection at the same time; See, e.g., Pötscher and Leeb (2009). In our approach, uniform convergence in distribution emerges over sequences of local alternatives of the form (16) for the tradable risk premia of factors that are neither weak not useless, whenever $\boldsymbol{\Delta}_{\mathcal{S}} \neq \mathbf{0}$. In contrast, along local alternatives depicting weak factor features ($\boldsymbol{\Delta}_{\mathcal{S}^c} \neq \mathbf{0}$), our Oracle estimator correctly performs a consistent factor selection, by uniformly estimating a zero tradable factor risk premium for all weak or useless factors, with a probability converging to one as the sample size grows.

[23] Such asymptotic distribution is reported in Proposition 4.1. The Oracle property of our proposed proximal tradable risk premium estimator is a natural consequence of the desirable properties of sample tradable risk premia. To ensure a consistent selection of factors that are not useless or weak, we exploit the fact that the sample correlations with returns of factors that are useless or weak are at most of order $1/\sqrt{T}$ in probability. Once a consistent factor selection is ensured, the efficiency of the asymptotic distribution for the estimated risk premia of the selected factors follows directly, because the marginal asymptotic distributions of sample tradable risk premia are independent of the properties of the other factors in an asset pricing model.

Remark 2. An obvious implication of Proposition 4.3 is that the asymptotic distribution of Oracle estimator $\check{\boldsymbol{\lambda}}^*$ under null hypothesis (15) is such that:

$$\sqrt{T}\,(\boldsymbol{\lambda}^*_{\check{\mathcal{S}}} - \boldsymbol{\lambda}^*_{\mathcal{S}}) \to_d \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathcal{S}\mathcal{S}}) \ .$$

Furthermore, while according to Proposition 4.3 the shrinkage of sample tradable risk premia implied by soft-thresholding formula (20) is asymptotically negligible for factors that are neither useless nor weak, it may give rise to a finite-sample bias. Such bias can be removed using a hard-thresholding approach, in which one directly estimates sample risk premia, but only for factors that have been selected in finite samples by our Oracle estimator, giving rise to a "relaxed" proximal estimator $\bar{\boldsymbol{\lambda}}^*$ such that:

$$\bar{\boldsymbol{\lambda}}^*_{\check{\mathcal{S}}} = \hat{\boldsymbol{\lambda}}^*_{\check{\mathcal{S}}} \text{ and } \bar{\boldsymbol{\lambda}}^*_{\check{\mathcal{S}}^c} = \mathbf{0} \ . \tag{22}$$

This hard-thresholding approach is founded theoretically by the fact that under the assumptions of Proposition 4.3 the relaxed proximal estimator is asymptotically equivalent to proximal estimator (19), uniformly over the set of local alternatives in Assumption 2. Due to its better finite-sample bias properties, we use it as our benchmark estimator in our Monte Carlo simulations and empirical analysis.

5 Monte Carlo simulation evidence

We study the finite-sample estimation and inference properties of sample tradable risk premium estimator (10), Oracle estimator (19), and its relaxed version introduced in Remark 2. To this end, we consider several Monte Carlo simulation settings (i) including correctly specified and misspecified factor models of increasing dimensions and (ii) featuring a varying degree of no or weak identification. In addition, we investigate the finite-sample factor selection properties induced by Oracle estimator (19).

5.1 Benchmark methodologies

We benchmark our methodology to other important estimators in the literature. As a first natural comparison, we consider the misspecification-robust estimation and inference approach in Kan et al. (2013) (KRS), based on factor risk premium definition (8) for a weighting matrix $\boldsymbol{W} = \boldsymbol{V}_R^{-1}$. This gives rise to the factor risk premium estimator:

$$\hat{\boldsymbol{\lambda}} := \hat{\boldsymbol{\lambda}}_{\hat{\boldsymbol{V}}_R^{-1}} := (\hat{\boldsymbol{\beta}}'\hat{\boldsymbol{V}}_R^{-1}\hat{\boldsymbol{\beta}})\hat{\boldsymbol{\beta}}\hat{\boldsymbol{V}}_R^{-1}\hat{\boldsymbol{\mu}}_R \ .$$

This estimator follows as a direct sample version of factor risk premium (8). Second, we document how our approach compares with respect to the Penalized two-step Fama-MacBeth (Pen-FM) estimation and inference approach introduced in Bryzgalova (2015). This approach relies on a penalized factor risk premium estimator defined by:

$$\check{\boldsymbol{\lambda}} := \check{\boldsymbol{\lambda}}_{\hat{\boldsymbol{V}}_R^{-1}} := \arg\min_{\boldsymbol{\lambda}\in\mathbb{R}^K} \left\{ (\hat{\boldsymbol{\lambda}} - \boldsymbol{\lambda})'\hat{\boldsymbol{\beta}}'\hat{\boldsymbol{V}}_R^{-1}\hat{\boldsymbol{\beta}}(\hat{\boldsymbol{\lambda}} - \boldsymbol{\lambda}) + \tau_T \sum_{k=1}^{K} \frac{|\lambda_k|}{||\hat{\boldsymbol{\beta}}_k||_1^d} \right\}, \quad (23)$$

where $\tau_T > 0$, $d > 0$ are tuning parameters and $||\hat{\boldsymbol{\beta}}_k||_1$ denotes the l_1-norm of the $k-$column in the $N \times K$ matrix $\hat{\boldsymbol{\beta}}$ of sample betas. By definition, this estimator is given by a penalized weighted minimum distance correction of misspecification-robust factor risk premium estimator $\hat{\boldsymbol{\lambda}}$. Since in unidentified models weighting matrix $\hat{\boldsymbol{W}} := \hat{\boldsymbol{\beta}}'\hat{\boldsymbol{V}}_R^{-1}\hat{\boldsymbol{\beta}}$ does not converge to a positive definite limit matrix, the penalty term in equation (23) is chosen to ensure that estimator $\check{\boldsymbol{\lambda}}$ gives rise asymptotically to a zero risk premium for all useless and weak factors. As shown in Proposition 2 of Bryzgalova (2015), estimator (23) consistently selects in finite samples factors that are neither weak nor useless. Moreover, it induces for all other factor risk premia an asymptotic distribution as efficient as the one of an Oracle misspecification-robust factor risk premium estimator, which knows ex-ante the factors that are neither useless not weak. Therefore, estimator (23) offers a second obvious benchmark for the finite-sample inference and factor selection properties of our approach.[24]

To select the tuning parameter τ_T of our proximal tradable risk premium estimator in a purely data-driven way, we make use of typical Cross Validation (CV) schemes and a Generalized Cross Validation (GCV; Wahba, 1990) scheme based on following criterion:

$$\text{GCV} = \frac{\left\| \hat{\boldsymbol{\mu}}_R - \hat{\boldsymbol{\beta}}^*_{\check{\mathcal{S}}} \check{\boldsymbol{\lambda}}^*_{\check{\mathcal{S}}} \right\|_2^2}{(1 - |\check{\mathcal{S}}|/T)^2} \ ,$$

where $\check{\mathcal{S}}$ is the active set estimated by $\check{\boldsymbol{\lambda}}^*$, $\hat{\boldsymbol{\beta}}^*_{\check{\mathcal{S}}} := \hat{\boldsymbol{V}}_{RF_{\check{\mathcal{S}}}}(\hat{\boldsymbol{V}}_{F_{\check{\mathcal{S}}}R}\hat{\boldsymbol{V}}_R^{-1}\hat{\boldsymbol{V}}_{RF_{\check{\mathcal{S}}}})^{-1}$,

[24] In contrast to our Oracle estimator, the Oracle property of estimator (23) does not hold uniformly over local alternative of the relevant factor risk premium hypotheses. In our implementation of this estimator, we use $d = 4$, consistently with the choice in the simulation section of Bryzgalova (2015), and we define the adaptive penalty weights based on sample betas. Results for different choices of tuning parameter d and penalty weights – e.g., based on correlations or partial correlations of factors and returns – are similar to the ones obtained below, and are available upon request.

and $|\check{\mathcal{S}}|$ is the cardinality of set $\check{\mathcal{S}}$, which is a proxy for a model's number of degrees of freedom.[25]

5.2 Summary of the key findings

For brevity of exposition, we report in this section a summary of the main findings emerging from our extensive Monte Carlo study. A detailed description of the Monte Carlo simulation settings considered and the corresponding results is given in Section 8.3 of the Online Appendix.

In the comparison with the KRS estimator, we find in general that the former estimator consistently exhibits a higher Mean Squared Errors (MSE) with respect to the sample, Oracle, and Relaxed estimators of tradable factor risk premia. This is expected and due to the well-known steep deterioration of KRS's estimation properties in the presence of weak and useless factors. Our Oracle estimator proves successful in employing shrinkage and factor selection to balance bias and variance more effectively, which makes it a viable factor risk premium estimator. The Relaxed estimator further refines the Oracle estimator by reducing its bias, especially in higher dimensional models. The comparison with the Pen-FM estimator further highlights clearly lower MSEs provided by our Oracle and Relaxed estimators of tradable factor risk premia in higher dimensional models, where it becomes apparent that accurately estimating misspecification-robust factor risk premia in presence of weak or useless factors is hardly feasible with the typical sample sizes available in practice.

When comparing the factor selection properties of our Oracle estimator and the Pen-FM estimator, we find that the former exhibits remarkable variable selection capabilities in essentially all models, by nearly exactly selecting factors that are neither weak nor useless. In contrast, the Pen-FM estimator delivers sufficient factor selection properties in relatively low dimensional models with no more than 3 factors (Model 1-3 of the Appendix), but it breaks down in higher dimensional models where it only rarely selects the correct set of factors that are neither weak nor useless.

When considering finite sample inference properties, we find that our Oracle estimator and our Relaxed estimator are quite successful in miti-

[25] When comparing our Oracle estimator to the Pen-FM estimator of Bryzgalova (2015), we additionally consider 5-fold Cross Validation (CV). Following, e.g., James et al. (2021), we select the tuning parameters of the penalized estimators in our analysis according to the "one-standard-error" rule: we select the most parsimonious model among the models with a CV/GCV score that is not more than one standard error above the CV/GCV score of the best model.

gating the well-known distortions caused by weak or useless factors for the inference produced by KRS estimators. In particular, the Relaxed estimator is quite successful also in reducing the finite sample bias associated with the shrinkage properties of the Oracle estimator. Both the Oracle and Relaxed estimators also exhibit higher power in detecting factors with non zero risk premia, due to the smaller finite-sample distortion and higher finite-sample efficiency of these estimators in estimating tradable factor risk premia when some factors are weak or useless. Importantly and consistently with our asymptotic results in Propositions 4.2–4.3, the finite sample inference for tradable risk premia of factors that are neither weak nor useless is unaffected by the presence of weak or useless factors in a model. This key robustness property is an additional desirable feature of our approach, which is not preserved for KRS misspecification-robust notions of factor risk premia and their estimators, as shown in Kan and Zhang (1999b), Kleibergen (2009) and Gospodinov et al. (2014), among others. As a consequence, also the Oracle inference provided by estimator (23) for misspecification-robust factor risk premia is intrinsically quite sensitive to an accurate finite sample selection of factors that are neither weak nor useless. However, as our simulations document, such an accurate selection is hardly feasible with the typical sample sizes available in empirical work for a number of theoretically and empirically relevant factor models.

6 Empirical analysis

We exploit our factor selection and inference methodology based on the Oracle proximal estimator of Definition 2, in order to explore and analyze a wide range of empirical asset pricing models generated from the factor zoo. In this way, we produce a theoretically founded, feasible diagnostics framework, which helps to characterize the main properties of asset pricing models from the factor zoo, in the potential presence of useless/weak factors. With this framework, we ideally aim to first isolate a relevant subset of low-dimensional, well-identified asset pricing models. Second, we aim to understand whether a robust subset of factors exists, which consistently appears in such subset of well-identified asset pricing models. Finally, if such a subset of factors exists, we can try to study which of these factors are priced.

6.1 Empirical setting

Our analysis is based on a universe of potential factors given by the 51 tradable and nontradable factors compiled in Bryzgalova et al. (2023).[26] As test assets, we consider the 25 portfolios sorted on size and book-to-market and the 17 industry portfolios. Section 8.5 of the Online Appendix reports results for test assets including (i) the 25 size and book-to-market and the 25 operating profitability and investment portfolio sorts, and (ii) the 25 size and book-to-market and a varying number of principal components extracted from 17 industry and 310 double-sorted portfolios. Our datasets consist of monthly observations of test asset excess returns and risk factors collected from October 1973 to December 2016.[27]

Using this dataset, we build randomized families of factor models having factor dimension growing from 1 to 10. To these families of models, we apply our factor selection methodology based on tradable factor risk premia. Intuitively, models of higher dimension are more likely to be weakly or not identified. In this respect, a main goal of our analysis is to understand whether our methodology is effective in mitigating these identification problems, exclusively by consistently pruning the useless and the weak factors, in a way that may ideally give rise to a selected subset of economically relevant and well-identified lower dimensional models.

In order to verify the identification properties of the various candidate models in our empirical study, we make use of the rank test proposed in Chen and Fang (2019).[28] We perform Chen and Fang (2019) test both before and after our consistent selection procedure for factors that are neither weak nor useless is applied. Precisely, let $\boldsymbol{\beta} = \boldsymbol{V}_{RF}\boldsymbol{V}_F^{-1}$ be the $N \times K$ matrix of factor betas of a set of candidate factors, before our selection procedure has been applied. With Chen and Fang (2019) test we first test the null hypothesis:

$$\mathcal{H}_0 : \ \mathrm{rank}(\boldsymbol{\beta}) \leq K - 1 \ ,$$

[26]Since tradable risk premia satisfy the identity $\boldsymbol{\lambda}^* = \boldsymbol{V}_F^{1/2}\,\mathbb{C}\mathrm{or}[\boldsymbol{F},\boldsymbol{R}]\boldsymbol{V}_R^{-1/2}\boldsymbol{\mu}_R$, we scale all factors so that they have a unit sample variance, in order not to spuriously select factors with a large variance.

[27]All data on test assets is sourced from the Kenneth French data library. As for risk factors, we are grateful to the authors of Bryzgalova et al. (2023) for making their data publicly accessible. Detailed descriptions of this factor universe can be found in their Table B.1 and Table IA.XIII from their Online Appendix.

[28]This test first employs the iterative Kleibergen and Paap (2006) rank test procedure to compute a first-step rank estimate. In a second step, it implements a bootstrap procedure that allows to directly test the null hypothesis of a reduced column rank in the matrix of factor betas of a candidate asset pricing model.

against the alternative hypothesis $\text{rank}(\boldsymbol{\beta}) = K$, using the corresponding matrix $\hat{\boldsymbol{\beta}} = \hat{\boldsymbol{V}}_{RF}\hat{\boldsymbol{V}}_F^{-1}$ of estimated factor betas. Let further $\boldsymbol{\beta}_{\mathcal{S}} = \boldsymbol{V}_{RF_{\mathcal{S}}}\boldsymbol{V}_{F_{\mathcal{S}}}^{-1}$ be the matrix of factor betas of factors that are neither useless nor weak. We then test the null hypothesis

$$\mathcal{H}_{0\mathcal{S}}: \ \text{rank}(\boldsymbol{\beta}_{\mathcal{S}}) \leq |\mathcal{S}| - 1 \ ,$$

against the alternative hypothesis $\text{rank}(\boldsymbol{\beta}_{\mathcal{S}}) = |\mathcal{S}|$, using the estimated cardinality $|\hat{\mathcal{S}}|$ of active set $\mathcal{S}$ and the corresponding matrix $\hat{\boldsymbol{\beta}}_{\hat{\mathcal{S}}} = \hat{\boldsymbol{V}}_{RF_{\hat{\mathcal{S}}}}\hat{\boldsymbol{V}}_{F_{\hat{\mathcal{S}}}}^{-1}$ of estimated factor betas after the factor selection has been performed.

By comparing the rank test results before and after performing the factor selection, we can finally analyze whether our methodology is effective in mitigating potential identification problems, exclusively by pruning the useless and the weak factors. Moreover, by inspecting the distribution of selected factors across randomized models, we can study whether a robust subset of economically relevant, well-identified low dimensional asset pricing models appears. Finally, we can study the pricing properties of asset pricing factors appearing in such well-identified low dimensional models.

6.2 Identification evidence pre- and post-factor screening

We first quantify the degree to which our consistent factor selection methodology can help improve the identification of a candidate asset pricing model. Intuitively, if a model's weak or no identification is exclusively due to the existence of some useless or weak factors, then our methodology is able to consistently select an associated set of well-identified submodels, each including exclusively factors that are neither useless nor weak. In such a situation, the rejection frequency of the null of no identification by Chen and Fang (2019) test should be substantially larger after applying our factor selection procedure over randomized factor models from the factor zoo. Moreover, after performing the factor selection, a potential identification problem should depend to a less extent on the initial dimension of an asset pricing model.

Figure 1 reports rejection frequencies of the null of no identification by Chen and Fang (2019) test, before and after applying our factor selection methodology, for different randomization procedures over factor models of dimension between 5 and 10 from the factor zoo.[29] Each null rejection is based on asymptotic critical values for a nominal size $\alpha = 0.05$. Note that this analysis is not aimed to directly identify a subset of well-identified

[29] Results for lower dimensions are reported in the supplemental appendix.

models from an initial larger set of models in a multiple hypothesis testing problem, but rather to quantify how often a test of identification applied to a randomly selected model would indicate an identification problem. Therefore, we do not introduce a correction for multiple testing at this stage, even though such a correction could be in principle developed using existing methodologies in the literature.

The upper panels of Figure 1 show that the pre-screening identification frequency decreases rapidly in the number of candidate factors of an asset pricing model. Moreover, this frequency is lower for the testing settings based on the lower dimensional set of test assets. For instance, identification frequencies decrease from about 0.95 (0.99) for single factor models, to about 0.05 for eight, nine and ten (ten) factor models, when test assets comprise the 25 size/book-to-maket (25 size/book-to-maket and the 17 industry) portfolios. This evidence is partly expected, as the probability of observing a reduced column rank in the matrix of factor betas of randomly selected models may be naturally higher for larger dimensional models and for testing frameworks based on less test assets. However, it crucially indicates that the identification problem may appear quite frequently already across randomly selected low dimensional models including no more than, e.g., five factors.

A clearly distinct picture emerges for the post-screening identification frequencies, which are uniformly higher than the pre-screening identification frequency, in many cases by a large extent. For instance, the post-screening identification frequency decreases from about 0.95 (0.99) for single factor models to about 0.65 (0.8) for ten factor models, when test assets comprise the 25 size/book-to-maket (25 size/book-to-maket and the 17 industry) portfolios. This evidence suggests that our factor screening methodology is indeed effective in selecting factors that are neither useless nor weak. Moreover, it indicates that after the consistent selection of these factors the weak or no identification problem emerges much less frequently across randomly selected models.

The lower panels of Figure 1 further consider model randomizations that always include the market factor in the initial model. Also in this case, the pre-screening identification frequencies feature a dramatic degradation as the asset pricing model size increases. In contrast, all post-screening identification frequencies are always very large and above 0.9 (0.95) even for ten factor models, when test assets comprise the 25 size/book-to-maket (25 size/book-to-maket and the 17 industry) portfolios. This additional evidence further confirms that our screening methodology is quite successful in selecting factors that are neither useless nor weak and thereby largely improving the identification properties of empirical asset pricing models.

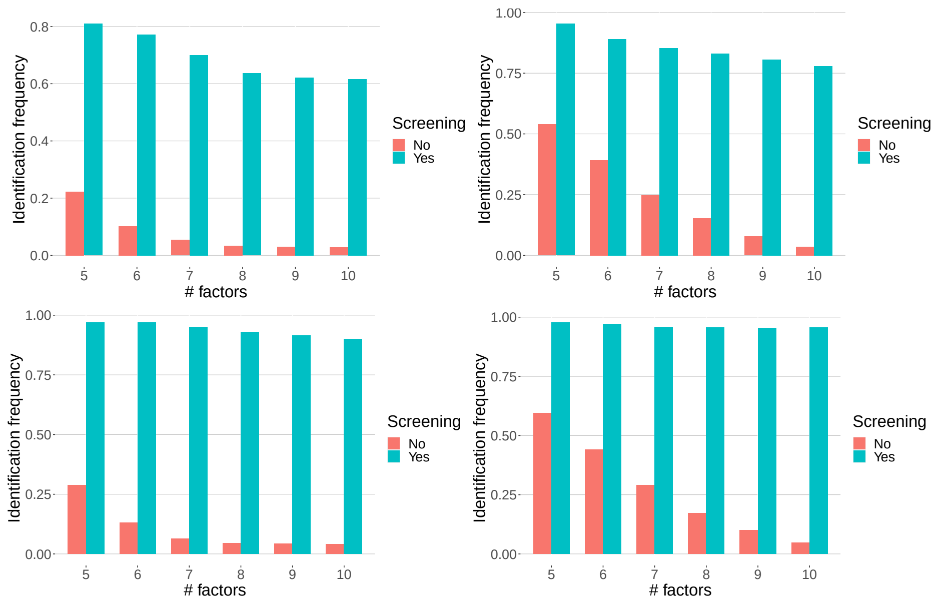

Figure 1: **Model identification frequencies**: Frequency of null hypothesis (6.1) rejections under a significance level $\alpha = 0.05$ by the Chen and Fang (2019) test, across randomized factor models including 5–10 initial factors. The red (blue) bars indicate model identification frequencies before (after) having applied our Oracle factor selection. The upper (lower) panels report the model identification frequencies for selections with no (the market as) ubiquitous factor across models. The left (right) panels report the results for test assets comprising the 25 size/book-to-market (the 25 size/book-to-market and the 17 industry) portfolios. For models having 1 (2) randomized factors, we consider all 52 (1326) possible model combinations. For models including more than 2 randomized factors, we consider 10'000 random factor combinations.

6.3 Pre- and post-screening factor space dimension

While the post-screening dimension of an asset pricing model is by construction not larger than the pre-screening dimension, the different pre- and post-screening identification properties documented above have to be related to strictly lower model dimensions after screening. This feature may imply that the typical dimension of an identified asset pricing models from the factor zoo is relatively low. To understand this issue properly, we study in

Figure 2 the distribution of post-screening factor space dimensions emerging across randomized asset pricing models.

The upper panels of Figure 2 show that, as expected, the post-screening factor space dimension increases with the initial dimension of an asset pricing model. More importantly, regardless of the pre-screening dimension, lower post-selection factor space dimensions occur with higher frequencies than higher dimensions. For instance, for settings with five initial factors, we find that models with at most one, at most two, and more than four selected factor(s) occur around 30% (60%), 50% (73%) and 30% (16%) of the cases, respectively, when the test assets comprise the 25 size/book-to-market (the 25 size/book-to-market and the 17 industry) portfolios. The post-screening factor dimension further shrinks uniformly for model randomizations that always include the market factor. For instance, for settings with five initial factors, we find that models with at most one, at most two, and more than four selected factor(s) occur 33% (30%), 77% (73%) and 6% (8%) of the cases, respectively, when the test assets comprise the 25 size/book-to-market portfolios (the 25 size/book-to-market and the 17 industry portfolios).[30]

Overall, we conclude that the most likely factor space dimensions of identified asset pricing models from the factor zoo are relatively low, and even more so for models including the market factor. As a consequence, standard tests of asset pricing models based on two-step cross-sectional regression methodologies are likely to be harmed by a weak identification problem already when testing low-dimensional models with, e.g., between five and seven factors.

[30]Analogous evidence is obtained in Section 8.5 of the Online Appendix for test assets including the 25 size/book-to-market portfolios and 25 portfolios sorted on operating profitability and investment.

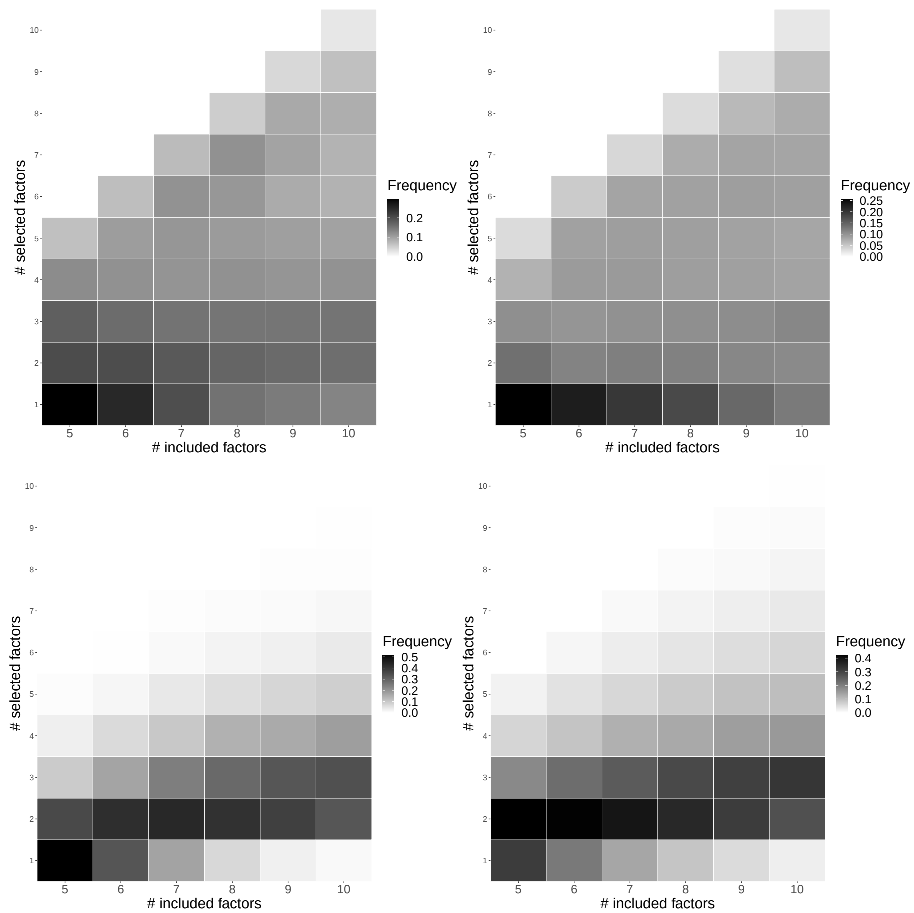

Figure 2: **Post-selection model size**: Frequency of post-selection model dimensions, i.e., number of factors selected by our Oracle factor selection, across randomized factor models including 5–10 initial factors. The red (blue) bars indicate model identification frequencies before (after) having applied our Oracle proximal factor selection. The upper (lower) panels report the model identification frequencies for selections with no (the market as) ubiquitous factor across models. The left (right) panels report the results for test assets comprising the 25 size/book-to-market (the 25 size/book-to-market and the 17 industry) portfolios.

6.4 Oracle factor selection and factor types

The previous evidence of a low post-screening factor space dimension suggests that the majority of identified asset pricing models from the factor zoo is relatively low-dimensional. Therefore, we next ask whether we can detect asset pricing factors that most consistently appear in the subset of identified low-dimensional models.

Figure 3 illustrates the marginal selection frequencies of individual risk factors determined by our screening procedure. These frequencies indicate how often each risk factor is selected from randomized initial models that include the particular factor, and they are plotted against the model's dimension. The upper panels of Figure 3 show that factors such as the market, the SMB, the ICR in He et al. (2017), and the market factor with a hedged unpriced component in Daniel et al. (2020) (MKT_star) are those with the highest and very consistent marginal selection frequencies across randomized models and test asset choices. For instance, when using the 25 size/book-to-market and 17 industry portfolios, the marginal selection frequencies of these factors when randomizing 10-factor models are all 100%, except those for factor SMB which has a selection frequency of 93.4%. Analogously, when using the 25 size/book-to-market portfolios, these selection frequencies are 99.4% and 92.9% for SMB and MKT_star, respectively, and 100% for the other factors. Overall, we conclude that this set of factors is unlikely to generate an identification problem when included in an asset pricing model from the factor zoo.

A second tier of factors display high, but less consistent, marginal selection frequency across randomized models and choices of test assets. These factors may be crowd out in some identified models by other factors more strongly co-moving with test asset returns. When using the 25 size/book-to-market and 17 industry portfolios as test assets, there are just a few of these factors, namely the long-term reversal factor in Jegadeesh and Titman (2001) (LTRev), the liquidity factor in Pástor and Stambaugh (2003) (LIQ_NT), nondurable consumption (NONDUR) and the systematic skeweness factor in Langlois (2020) (SKEW). When using just the 25 size/book-to-market portfolios as test assets, additional such factors are the SMB factor with a hedged unpriced component in Daniel et al. (2020) (SMB_star), the long-term behavioral factor in Daniel et al. (2020) (BEH_FIN), the profitability factor measured by gross profits-to-assets in Novy-Marx (2013) (GR_PROF), the mispricing factor in Stambaugh and Yuan (2017) (MGMT), and the distress risk factor in Campbell et al. (2008) (DISSTR).

An even sharper evidence arises in the lower panels of Figure 3, for model randomizations that always include the market factor. Here, factors including SMB, ICR, MKT_star and MGMT are those having by far the highest and most consistent marginal selection frequencies, with, e.g., selection frequencies above 97% when randomizing ten-factor asset pricing models. Second tier factors with high, but less consistent, marginal selection frequencies include BEH_FIN, the composite equity issuance in Daniel and Titman (2006) (COMP_ISSUE) and the conservative minus aggressive in Fama and French (2015) (CMA). For randomizations of ten-factor models, these factors have marginal selection frequencies of 89.1% (85.4%), 69.3% (89.2%) and 58% (73.4%), respectively, when using the 25 size/book-to-market (25 size/book-to-market and 17 industry) portfolios as test assets. A similar, and to some extent stronger, evidence is obtained in Section 9 of the Online Appendix for test assets including the 25 size/book-to-market portfolios and 25 portfolios sorted on operating profitability and investment. Here, all these factors have marginal selection frequencies above 95% when randomizing ten-factor asset pricing models.

Overall, we conclude that models including robust asset pricing factors MKT, SMB, ICR, MKT_star, and MGMT, together with some of the factors BEH_FIN, COMP_ISSUE and CMA, are natural well-identified benchmarks for explaining the cross-section of test assets' expected returns.

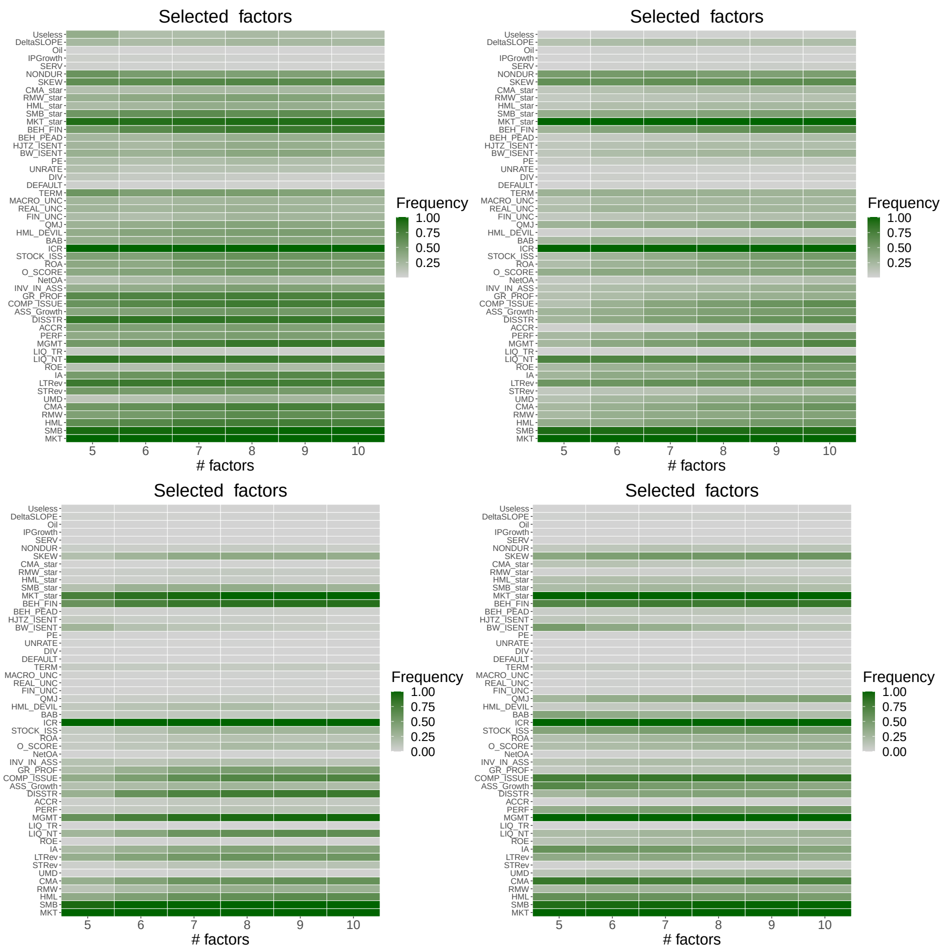

Figure 3: **Factor selection frequencies**: Selection frequencies of individual factors determined through our Oracle factor screening procedure, observed across randomized factor models containing the specific factor along with 4–9 additional factors. The upper (lower) panels report the frequencies of the post-selection model dimensions for selections with no (the market as) ubiquitous factor across models. The left (right) panels report the results for test assets comprising the 25 size/book-to-market (the 25 size/book-to-market and the 17 industry) portfolios.

6.5 Tradable factor risk premia

Given the above evidence regarding the set of well-identified candidate models and corresponding factors for explaining the cross-section of asset returns, we next expore the associated tradable factor risk premia. Exploiting the properties of our Oracle estimator in Definition 2, we can test whether a factor gives rise to a tradeable risk premium component. Indeed, recall that conditionally on a factor being identified as neither useless nor weak by our Oracle factor selection, we can still make use of the marginal asymptotic distributions from Proposition 4.3 to test the null hypothesis of a zero individual tradeable risk premium. Since these asymptotic distributions are independent of the dimension of the factor space in an asset pricing model, this inference is also invariant across randomized models.

Figure 4 reports the resulting inference, with estimated tradable factor risk premia and their confidence intervals ranked according to a factor's selection frequency. In the set of robust asset pricing factors MKT, SMB, ICR, MKT_star, and MGMT with highest selection frequencies, we find that they all exhibit a statistically significant tradable risk premium at the 10% significance level, for both sets of test assets. The tradeable risk premium of factor CMA is as well significant for both sets of test assets, while the risk premium of factor COMP_ISSUE is significant for the 25 size/book-to-market and 17 industry portfolios. In contrast, factor BEH_FIN, is (marginally) insignificant at the 10% significance level for both sets of test assets.

An analogous, and to some extent stronger, evidence is obtained in Section 9 of the Online Appendix for test assets including the 25 size/book-to-market portfolios and 25 portfolios sorted on operating profitability and investment. Here, all tradeable risk premia of the above factors are significant at the 10% level and each of these factors is selected in more than 97% of randomized ten-factor models.

Overall, this evidence points to a robust subset of economically relevant and well-identified low dimensional models from the factor zoo, which can be built out of selected factors giving rise to a nonzero tradable risk premium.

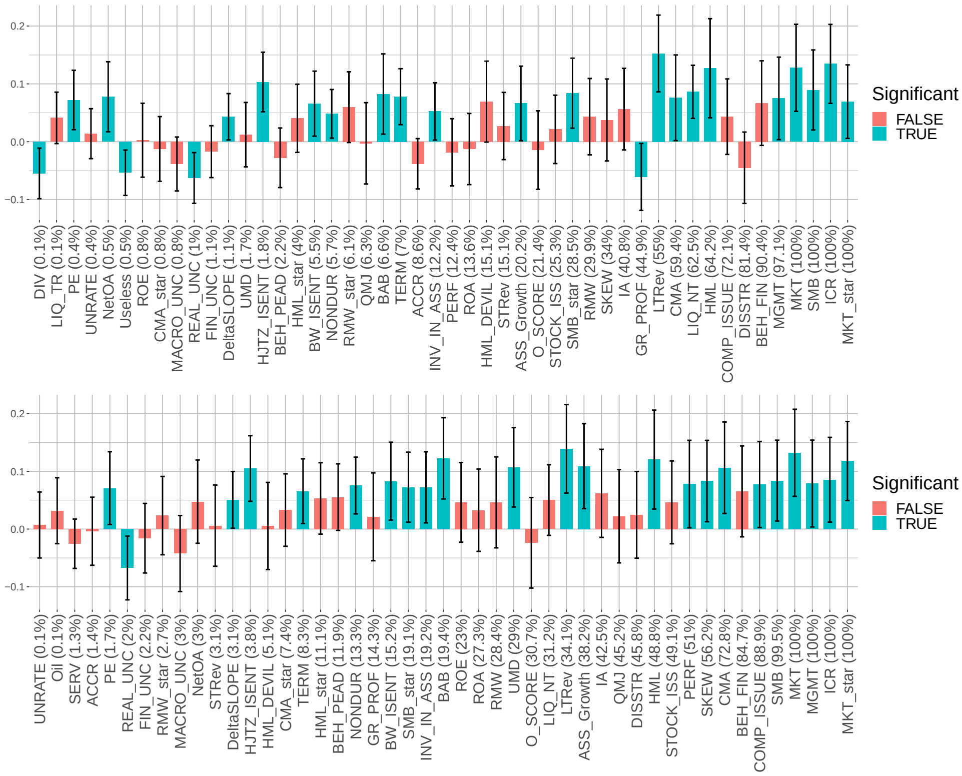

Figure 4: **Tradable factor risk premia**: Point estimates and confidence intervals at confidence level 90% of sample tradable factor risk premia. Factors are ordered by their selection frequency (conditional on the factor being present in the initial model) in randomized 10-factor models always including the market, reported in parenthesis, and only factors with positive selection frequency are retained. The upper (lower) panel reports the results obtained with test assets comprising the 25 size/book-to-market (25 size/book-to-market and the 17 industry portfolios).

6.6 Misspecification-robust factor risk premia

Given the above evidence regarding the set of low-dimensional well-identified candidate models, we next explore the properties of misspecification-robust factor risk premia estimated by two-step cross-sectional regression methods. In doing so, we build on the fact that our Oracle factor selection ensures a good identification of asset pricing models after screening.

Recall that a factor has different misspecification-robust and tradable risk premia only when it is partly unspanned by test asset returns. Furthermore, the latter risk premium depends in general on the factor composition of a model. Therefore, misspecification-robust risk premia in general imply different factor risk premium estimators across asset pricing models. Figures 5–6 summarize the composite post screening effects of unspanned factor risks for such factor risk premium estimates, across well-identified and potentially misspecified models from the factor zoo.

We find that while factors MKT, SMB, MGMT, CMA and COMP_ISSUE imply a relatively robust statistical significance of estimated risk premia, factors like ICR, MKT_star, and to some extent BEH_FIN, give rise to a uniformly weak significance. In addition, factors such as ICR, MKT_star or CMA give rise to a large variability of estimated risk premia across models, which is hardly interpretable economically. This evidence emerges despite the above findings that all these factors consistently appear in the subset of well-identified asset pricing models with a statistically significant tradable risk premium.

Overall, we conclude that the misspecification of identifiable models from the factor zoo generates a challenge for interpreting notions of risk premia allocating a non zero price to unspanned factor risks. In contrast, tradable factor risk premia are economically-founded in general and can be studied empirically based on our Oracle inference framework.

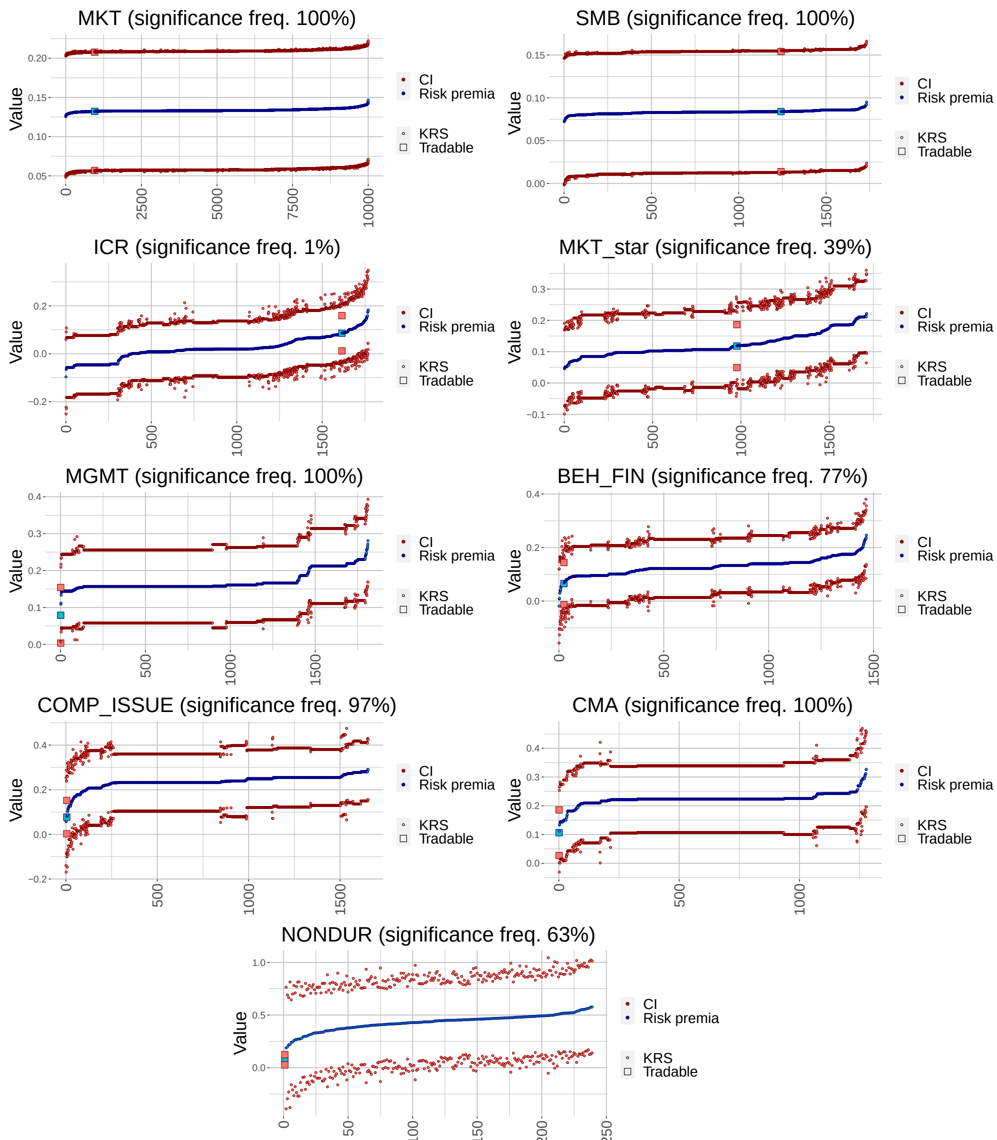

Figure 5: **Factor risk premia and** 90% **confidence intervals**: Misspecification-robust (represented by dots) and tradable (represented by squares) factor risk premia and corresponding 90% confidence intervals of a selection of factors over various randomized 10-factor models always including the market. The factor significance frequencies (conditional on the factor being present in the initial model) are reported in parenthesis, next to the factor's label. The results are obtained with test assets comprising the 25 size/book-to-market and the 17 industry portfolios.

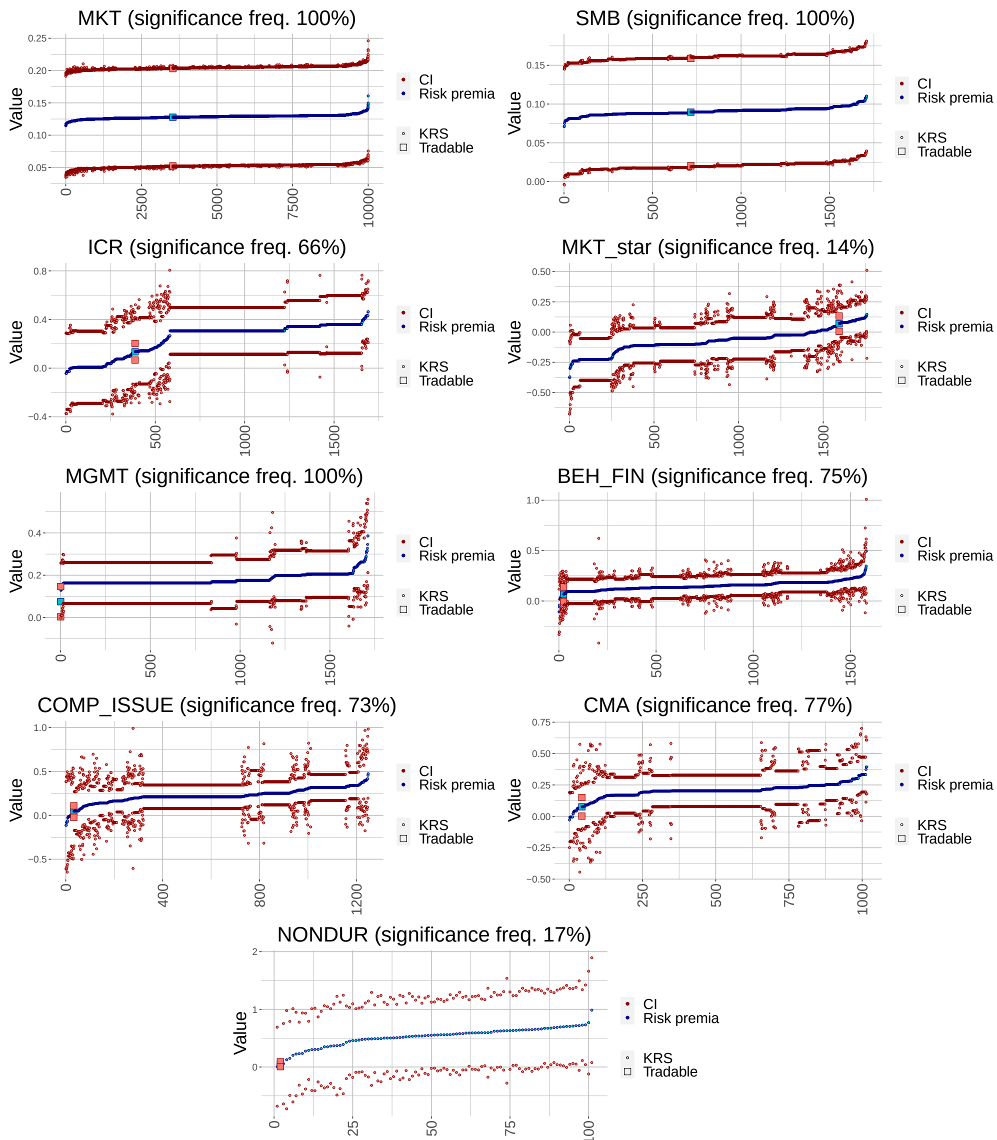

Figure 6: **Factor risk premia and** 90% **confidence intervals**: Misspecification-robust (represented by dots) and tradable (represented by squares) factor risk premia and corresponding 90% confidence intervals of a selection of factors over various randomized 10-factor models always including the market. The factor significance frequencies (conditional on the factor being present in the initial model) are reported in parenthesis, next to the factor's label. The results are obtained with test assets comprising the 25 size/book-to-market portfolios.

7 Conclusions

This paper addresses the challenge of testing misspecified and weakly (or not) identified asset pricing models. Existing research has emphasized the limitations of two-step testing procedures for factor risk premia defined by the negative factor covariance with a Stochastic Discount Factor (SDF) projection on the factor space. In such settings, misspecification introduces non-zero pricing errors that affect the asymptotic distribution of estimated factor risk premia, while weak (or no) identification hinders the unique identification of the candidate SDF projection, and, consequently, the associated factor risk premia. In contrast, our inference approach relies on the well-established concept of a tradable factor risk premium, which is defined by the negative factor covariance with the SDF projection on the asset return space. We show that tradable factor risk premia give rise to a robust and well-defined framework for Oracle inference in potentially misspecified or weakly identified asset pricing models.

Our methodology consistently identifies the set of factors that are weakly correlated with returns and simultaneously enables a reliable and efficient inference on the risk premia of all other factors in an asset pricing model. We achieve this by leveraging the properties of a simple sample version of tradeable factor risk premia, which exhibits an asymptotically Gaussian distributions even in models including factors weakly correlated with returns. Since this distribution may imply some asymptotic biases only for the estimated risk premia of weakly correlated factors, we develop an Oracle estimator of tradeable factor risk premia that overcomes these distortions.

Our Oracle estimator is derived through a closed-form penalized minimum distance correction of sample tradable risk premia. It improves upon existing estimators by producing a consistent identification of the set of weak factors and simultaneously generating a uniform efficient asymptotic distribution for the estimated risk premia of other factors. By design, our Oracle estimator directly provides a valid inference for the tradable risk premia of factors that are not weakly correlated with returns. Alternatively, it can be used for a consistent preliminary screening of potentially weakly correlated factors, which enhances the identification of a model and hence facilitates the reliable application of efficient two-step cross-sectional tests of asset pricing models in a subsequent step of the analysis.

We make use of our tractable Oracle estimation and inference methodology for tradable factor risk premia, in order to build a coherent and easily applicable framework for studying the asset pricing properties of a broad class of factor models from the factor zoo. For this purpose, we first form randomized families of candidate models of largest factor space dimension, in order to single out a set of well-identified submodels consisting exclusively of factors that are neither useless nor weak. Based on this set, we pin down the properties of the resulting distribution of factors and factor risk premia across well-identified models, for various benchmark choices of test assets in the literature.

Our empirical analysis reveals that the proposed factor selection procedure greatly enhances model identification. This allows us to detect a robust subset of economically relevant and well-identified models that are all built out of factors with a nonzero tradable risk premium. Such factors include market, size, intermediaries capital ratio, market with a hedged unpriced component, a mispricing factor, a long-term behavioral factor, a liquidity factor, a conservative minus aggressive factor and a composite equity issuance factor. At the same time, we find that the relatively low factor space dimension of well-identified models is associated with some degree of misspecification. This feature harms the interpretation of other established notions of a factor risk premium in the literature, which may allocate a non zero price to factor risks that are not tradable.

References

Asness, C. S., A. Frazzini, and L. H. Pedersen (2019). Quality minus junk. *Review of Accounting Studies 24*(1), 34–112.

Attouch, H. (1984). *Variational convergence for functions and operators*, Volume 1. Pitman Advanced Publishing Program.

Balduzzi, P. and H. Kallal (1997). Risk premia and variance bounds. *The Journal of Finance 52*(5), 1913–1949.

Balduzzi, P. and C. Robotti (2008). Mimicking portfolios, economic risk premia, and tests of multi-beta models. *Journal Business and Economic Statistics 26*(3), 354–368.

Balduzzi, P. and C. Robotti (2010). Asset pricing models and economic risk premia: A decomposition. *Journal of Empirical Finance 17*(1), 54–80.

Breeden, D. T., M. R. Gibbons, and R. H. Litzenberger (1989). Empirical tests of the consumption-oriented capm. *The Journal of Finance 44*(2), 231–262.

Bryzgalova, S. (2015). Spurious factors in linear asset pricing models. *LSE manuscript 1*(3).

Bryzgalova, S., J. Huang, and C. Julliard (2023). Bayesian solutions for the factor zoo: We just ran two quadrillion models. *The Journal of Finance 78*(1), 487–557.

Campbell, J. Y., J. Hilscher, and J. Szilagyi (2008). In search of distress risk. *The Journal of finance 63*(6), 2899–2939.

Chen, Q. and Z. Fang (2019). Improved inference on the rank of a matrix. *Quantitative Economics 10*(4), 1787–1824.

Daniel, K., D. Hirshleifer, and L. Sun (2020). Short-and long-horizon behavioral factors. *The review of financial studies 33*(4), 1673–1736.

Daniel, K., L. Mota, S. Rottke, and T. Santos (2020). The cross-section of risk and returns. *The Review of Financial Studies 33*(5), 1927–1979.

Daniel, K. and S. Titman (2006). Market reactions to tangible and intangible information. *The Journal of Finance 61*(4), 1605–1643.

Davidson, J. (1992). A central limit theorem for globally nonstationary near-epoch dependent functions of mixing processes. *Econometric theory 8*(3), 313–329.

Ehsani, S. and J. T. Linnainmaa (2022). Factor momentum and the momentum factor. *The Journal of Finance 77*(3), 1877–1919.

Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of financial economics 116*(1), 1–22.

Fama, E. F. and J. D. MacBeth (1973). Risk, return, and equilibrium: Empirical tests. *Journal of political economy 81*(3), 607–636.

Fan, J. and R. Li (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association 96*(456), 1348–1360.

Frazzini, A. and L. H. Pedersen (2014). Betting against beta. *Journal of financial economics 111*(1), 1–25.

Geyer, C. J. (1994). On the asymptotics of constrained m-estimation. *The Annals of statistics*, 1993–2010.

Giglio, S. and D. Xiu (2021). Asset pricing with omitted factors. *Journal of Political Economy 129*(7), 1947–1990.

Gonçalves, A. S. (2021). The short duration premium. *Journal of Financial economics 141*(3), 919–945.

Gormsen, N. J. and E. Lazarus (2023). Duration-driven returns. *The Journal of Finance 78*(3), 1393–1447.

Gospodinov, N., R. Kan, and C. Robotti (2014). Misspecification-robust inference in linear asset-pricing models with irrelevant risk factors. *The Review of Financial Studies 27*(7), 2139–2170.

Gospodinov, N., R. Kan, and C. Robotti (2017). Spurious inference in reduced-rank asset-pricing models. *Econometrica 85*(5), 1613–1628.

Hansen, L. P. and R. Jagannathan (1991). Assessing specification errors in stochastic discount factors models. *Journal of finance 52*, 557–590.

Harvey, C. R., Y. Liu, and H. Zhu (2016). … and the cross-section of expected returns. *The Review of Financial Studies 29*(1), 5–68.

He, Z., B. Kelly, and A. Manela (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics 126*(1), 1–35.

Hou, K., C. Xue, and L. Zhang (2015). Digesting anomalies: An investment approach. *The Review of Financial Studies 28*(3), 650–705.

Huberman, G., S. Kandel, and R. F. Stambaugh (1987). Mimicking portfolios and exact arbitrage pricing. *The Journal of Finance 42*(1), 1–9.

James, G., D. Witten, T. Hastie, R. Tibshirani, G. James, D. Witten, T. Hastie, and R. Tibshirani (2021). Statistical learning. *An introduction to statistical learning: with applications in R*, 15–57.

Jegadeesh, N. and S. Titman (2001). Profitability of momentum strategies: An evaluation of alternative explanations. *The Journal of finance 56*(2), 699–720.

Kan, R., C. Robotti, and J. Shanken (2013). Pricing model performance and the two-pass cross-sectional regression methodology. *The Journal of Finance 68*(6), 2617–2649.

Kan, R. and C. Zhang (1999a). Gmm tests of stochastic discount factor models with useless factors. *Journal of Financial Economics 54*(1), 103–127.

Kan, R. and C. Zhang (1999b). Two-pass tests of asset pricing models with useless factors. *the Journal of Finance 54*(1), 203–235.

Kandel, S. and R. F. Stambaugh (1995). Portfolio inefficiency and the cross-section of expected returns. *The Journal of Finance 50*(1), 157–184.

Kleibergen, F. (2005). Testing parameters in gmm without assuming that they are identified. *Econometrica 73*(4), 1103–1123.

Kleibergen, F. (2009). Tests of risk premia in linear factor models. *Journal of econometrics 149*(2), 149–173.

Kleibergen, F. and R. Paap (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of econometrics 133*(1), 97–126.

Kleibergen, F. and Z. Zhan (2015). Unexplained factors and their effects on second pass r-squared's. *Journal of Econometrics 189*(1), 101–116.

Kleibergen, F. and Z. Zhan (2018). Identification-robust inference on risk premia of mimicking portfolios of non-traded factors. *Journal of Financial Econometrics 16*(2), 155–190.

Kleibergen, F. and Z. Zhan (2020). Robust inference for consumption-based asset pricing. *The Journal of Finance 75*(1), 507–550.

Kleibergen, F. and Z. Zhan (2022). Misspecification and weak identification in asset pricing. Technical report, Working Paper.

Langlois, H. (2020). Measuring skewness premia. *Journal of Financial Economics 135*(2), 399–424.

Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of financial economics 108*(1), 1–28.

Pástor, L. and R. F. Stambaugh (2003). Liquidity risk and expected stock returns. *Journal of Political economy 111*(3), 642–685.

Pötscher, B. M. and H. Leeb (2009). On the distribution of penalized maximum likelihood estimators: The lasso, scad, and thresholding. *Journal of Multivariate Analysis 100*, 2065–2082.

Quaini, A. and F. Trojani (2022). Proximal estimation and inference. *arXiv preprint arXiv:2205.13469*.

Shanken, J. (1992). On the estimation of beta-pricing models. *The review of financial studies 5*(1), 1–33.

Shanken, J. and G. Zhou (2007). Estimating and testing beta pricing models: Alternative methods and their performance in simulations. *Journal of Financial Economics 84*(1), 40–86.

Stambaugh, R. F. and Y. Yuan (2017). Mispricing factors. *The review of financial studies 30*(4), 1270–1315.

Wahba, G. (1990). *Spline models for observational data*. SIAM.

Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association 101*(476), 1418–1429.

8 ONLINE APPENDIX

This online appendix is structured as follows. Section 8.1 contains the proofs of the mathematical statements in the main text. Section 8.2 introduces and discusses the exact beta representation of expected returns based on tradable factor risk premia. The simulation analysis in the main text is detailed in Section 8.3, while Section 8.4 assesses the performance of our methodologies based on tradable factor risk premia in the simulation settings of Gospodinov et al. (2014). Finally, Section 8.5 reports further empirical analysis and figures complementing the empirics in the main text.

8.1 Proofs of the mathematical results in the main text

Proof of Proposition 3.1. The objective function of optimization problem (9) is given by:

$$\begin{aligned} f(\boldsymbol{\gamma}) &:= \mathbb{E}[M_t(\boldsymbol{\gamma})\boldsymbol{R}_t]'\boldsymbol{V}_R^{-1}\mathbb{E}[M_t(\boldsymbol{\gamma})\boldsymbol{R}_t] + \boldsymbol{\gamma}'\boldsymbol{A}\boldsymbol{\gamma} \\ &= \boldsymbol{\mu}_R'\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R + \boldsymbol{\gamma}'\boldsymbol{V}_F\boldsymbol{\gamma} - 2\boldsymbol{\gamma}'\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ . \end{aligned}$$

Then, under the assumption that $\boldsymbol{V}_F$ is positive definite, the necessary and sufficient first order condition imply:

$$\boldsymbol{\gamma}^* = \boldsymbol{V}_F^{-1}\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ .$$

□

Proof of Proposition 4.1. Given the assumptions on stochastic process $\boldsymbol{Y}_t = [\boldsymbol{F}_t', \boldsymbol{R}_t']'$, terms $\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1}$, $\hat{\boldsymbol{V}}_{FR} - \boldsymbol{V}_{FR}$ and $\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R$ are all $O_{\mathrm{Pr}}(1/\sqrt{T})$. Therefore, we can write:[31]

$$\begin{aligned} \hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^* =& \hat{\boldsymbol{V}}_{FR}\hat{\boldsymbol{V}}_R^{-1}\hat{\boldsymbol{\mu}}_R - \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \\ =& \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}(\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R) + \boldsymbol{V}_{FR}(\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1})\boldsymbol{\mu}_R \\ &+ (\hat{\boldsymbol{V}}_{FR} - \boldsymbol{V}_{FR})\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R + O_{\mathrm{Pr}}(1/T) \ . \end{aligned}$$

Moreover:[32]

$$\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1} = \boldsymbol{V}_R^{-1}(\boldsymbol{V}_R - \hat{\boldsymbol{V}}_R)\boldsymbol{V}_R^{-1} + O_{\mathrm{Pr}}(1/T) \ .$$

[31] Here, the term $O_{\mathrm{Pr}}(1/T)$ collects $(\hat{\boldsymbol{V}}_{FR} - \boldsymbol{V}_{FR})(\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1})(\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R)$, $(\hat{\boldsymbol{V}}_{FR} - \boldsymbol{V}_{FR})(\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1})\boldsymbol{\mu}_R$, $(\hat{\boldsymbol{V}}_{FR} - \boldsymbol{V}_{FR})\boldsymbol{V}_R^{-1}(\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R)$ and $\boldsymbol{V}_{FR}(\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1})(\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R)$.

[32] Here, the term $O_{\mathrm{Pr}}(1/T)$ consists in $(\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1})(\boldsymbol{V}_R - \hat{\boldsymbol{V}}_R)\boldsymbol{V}_R^{-1}$.

This further gives, up to terms of order $O_{\mathrm{Pr}}(1/T)$:

$$\begin{aligned}\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^* &= \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}(\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R) - \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}(\hat{\boldsymbol{V}}_R - \boldsymbol{V}_R)\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \\ &\quad +(\hat{\boldsymbol{V}}_{FR} - \boldsymbol{V}_{FR})\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \\ &= \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}(\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R) - \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\hat{\boldsymbol{V}}_R\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R + \hat{\boldsymbol{V}}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ .\end{aligned}$$

Finally, note that:

$$\hat{\boldsymbol{V}}_R = \frac{1}{T}\sum_{t=1}^{T}\bar{\boldsymbol{R}}_t\bar{\boldsymbol{R}}_t' - (\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R)(\hat{\boldsymbol{\mu}}_R - \boldsymbol{\mu}_R)' = \frac{1}{T}\sum_{t=1}^{T}\bar{\boldsymbol{R}}_t\bar{\boldsymbol{R}}_t' + O_{\mathrm{Pr}}(1/T) \ .$$

Similarly:

$$\hat{\boldsymbol{V}}_{FR} = \frac{1}{T}\sum_{t=1}^{T}\bar{\boldsymbol{F}}_t\bar{\boldsymbol{R}}_t' + O_{\mathrm{Pr}}(1/T) \ .$$

Overall, we thus obtain:

$$\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^* = \frac{1}{T}\sum_{t=1}^{T}\boldsymbol{h}_t + O_{\mathrm{Pr}}(1/T) \ ,$$

where

$$\boldsymbol{h}_t := \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\bar{\boldsymbol{R}}_t - \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\bar{\boldsymbol{R}}_t\bar{\boldsymbol{R}}_t'\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R + \bar{\boldsymbol{F}}_t\bar{\boldsymbol{R}}_t'\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ ,$$

defines a stationary and ergodic process with zero expectation and finite second moments. Using the Central Limit Theorem for stationary and ergodic processes, we conclude:

$$\sqrt{T}(\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^*) \to_d \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}) \ ,$$

where

$$\boldsymbol{\Sigma} = \lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{\infty}\sum_{m=1}^{\infty}\mathbb{E}[\boldsymbol{h}_t\boldsymbol{h}_m'] \ .$$

This concludes the proof. □

Proof of Proposition 4.2. To prove the statement of the proposition, we follow the same steps as in the proof of Proposition 4.1, noting that under the given assumptions on stochastic process $\boldsymbol{Y}_t = [\boldsymbol{F}_t', \boldsymbol{R}_t{}']'$, quantities $\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}_R$, $\hat{\boldsymbol{V}}_R^{-1} - \boldsymbol{V}_R^{-1}$ and $\hat{\boldsymbol{V}}_{FR} - \boldsymbol{V}_{FR}$ are still all $O_{\mathrm{Pr}}(1/\sqrt{T})$. Therefore,

$$\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^* = \frac{1}{T}\sum_{t=1}^{T}\boldsymbol{h}_t + O_{\mathrm{Pr}}(1/T) \ ,$$

where

$$\boldsymbol{h}_t = \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\bar{\boldsymbol{R}}_t - \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\bar{\boldsymbol{R}}_t\bar{\boldsymbol{R}}_t'\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R + \bar{\boldsymbol{F}}_t\bar{\boldsymbol{R}}_t'\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ .$$

However the distribution of random vector $\boldsymbol{h}_t$ is different from the one obtained in Proposition 4.1. Indeed,

$$\boldsymbol{h}_t = \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\bar{\boldsymbol{R}}_t - \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}(\bar{\boldsymbol{R}}_t\bar{\boldsymbol{R}}_t' - \boldsymbol{V}_R)\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ + (\bar{\boldsymbol{F}}_t\bar{\boldsymbol{R}}_t' - \boldsymbol{V}_{FR})\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R \ ,$$

where, for any $\alpha > 1/2$:

$$\begin{aligned}\bar{\boldsymbol{F}}_t\bar{\boldsymbol{R}}_t' - \boldsymbol{V}_{FR} &= \bar{\boldsymbol{F}}_t\bar{\boldsymbol{R}}_t' - \boldsymbol{V}_{FR}^{(T)} + \boldsymbol{\Gamma}'/T^{\alpha} + \boldsymbol{\Delta}/\sqrt{T} \\ &= \bar{\boldsymbol{F}}_t\bar{\boldsymbol{R}}_t' - \boldsymbol{V}_{FR}^{(T)} + \boldsymbol{\Delta}/\sqrt{T} + o_{\mathrm{Pr}}(1/\sqrt{T}) \ .\end{aligned}$$

Therefore, using the Central Limit Theorem for near-epoch dependent processes, we obtain:

$$\frac{1}{\sqrt{T}}\sum_{t=1}^{T}\boldsymbol{h}_t \to_d \mathcal{N}\left(\boldsymbol{b}(\boldsymbol{\Delta}), \boldsymbol{\Sigma} - \boldsymbol{b}(\boldsymbol{\Delta})\boldsymbol{b}(\boldsymbol{\Delta})'\right) \ ,$$

where $\boldsymbol{b}(\boldsymbol{\Delta}) := \boldsymbol{\Delta}\boldsymbol{V}_R^{-1}\boldsymbol{\mu}_R$. This concludes the proof. □

Proof of Proposition 4.3. Under the assumptions of Proposition 4.2, it follows:

$$\sqrt{T}(\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^*) \to_d \boldsymbol{\xi} \sim \mathcal{N}\left(\boldsymbol{b}(\boldsymbol{\Delta}), \boldsymbol{\Sigma} - \boldsymbol{b}(\boldsymbol{\Delta})\boldsymbol{b}(\boldsymbol{\Delta})'\right) \ .$$

Moreover,

$$\sqrt{T}(\check{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^*) = \arg\min_{\boldsymbol{u}\in\mathbb{R}^K}\{g_T(\boldsymbol{u}) + q_T(\boldsymbol{u})\} \ ,$$

where

$$g_T(\boldsymbol{u}) := \frac{1}{2}\left|\left|\sqrt{T}\left(\hat{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^*\right) - \boldsymbol{u}\right|\right|_2^2 \ ,$$

and

$$q_T(\boldsymbol{u}) := T\tau_T\left(f_T\left(\boldsymbol{\lambda}^* + \frac{\boldsymbol{u}}{\sqrt{T}}\right) - f_T(\boldsymbol{\lambda}^*)\right) \ ,$$

with $f_T(\boldsymbol{\lambda}^*) := \sum_{k=1}^{K}|\lambda_k^*|/\left\|\hat{\boldsymbol{\rho}}_k\right\|_2^2$. Therefore,

$$g_T(\boldsymbol{u}) \to_d \frac{1}{2}||\boldsymbol{\xi} - \boldsymbol{u}||_2^2 =: g_0(\boldsymbol{u}) \ ,$$

uniformly on compact sets. Moreover, using the functional form of the Adaptive Lasso penalty f_T and the fact that, under the assumptions of Proposition 4.2, $\hat{\boldsymbol{\rho}}_{\mathcal{S}} \to_{\Pr} \boldsymbol{\rho}_{\mathcal{S}} \neq \mathbf{0}$, $\hat{\boldsymbol{\rho}}_{\mathcal{S}^c} \to_{\Pr} \mathbf{0}$ and that $\hat{\boldsymbol{\rho}}_{\mathcal{S}^c} = O_{\Pr}(1/\sqrt{T})$:

$$q_T(\boldsymbol{u}) \to_{\Pr} q_0(\boldsymbol{u}) := \begin{cases} +\infty & \boldsymbol{u}_{\mathcal{S}^c} \neq \mathbf{0} \\ 0 & else \end{cases} , \tag{24}$$

in epigraph. Using (Attouch, 1984, Thm. 2.15) and Geyer (1994), we thus obtain:

$$g_T(\boldsymbol{u}) + q_T(\boldsymbol{u}) \to_d g_0(\boldsymbol{u}) + q_0(\boldsymbol{u}) ,$$

in epigraph, and:[33]

$$\sqrt{T}(\check{\boldsymbol{\lambda}}^* - \boldsymbol{\lambda}^*) \to_d \underset{\{\boldsymbol{u}:\boldsymbol{u}_{\mathcal{S}^c}=\mathbf{0}\}}{\operatorname{argmin}} \{g_0(\boldsymbol{u})\} = \begin{pmatrix} \boldsymbol{\xi}_{\mathcal{S}} \\ \mathbf{0} \end{pmatrix} .$$

Finally, simple computations yield $\mathbb{E}[\boldsymbol{\xi}_{\mathcal{S}}] = \boldsymbol{b}_{\mathcal{S}}(\boldsymbol{\Delta})$, as shown in (18), and

$$\mathbb{V}\mathrm{ar}[\boldsymbol{\xi}_{\mathcal{S}}] = \boldsymbol{\Sigma}_{\mathcal{S}\mathcal{S}} - \boldsymbol{b}_{\mathcal{S}}(\boldsymbol{\Delta})\boldsymbol{b}'_{\mathcal{S}}(\boldsymbol{\Delta}) ,$$

where

$$\boldsymbol{\Sigma}_{\mathcal{S}\mathcal{S}} := \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \sum_{m=1}^{T} \mathbb{E}[\boldsymbol{h}_{t\mathcal{S}} \boldsymbol{h}'_{m\mathcal{S}}] .$$

This proves the second statement in Proposition 4.3. The first statement follows from the first statement using standard proofs of model selection consistency in Zou (2006) for the Adaptive Lasso penalty. This concludes the proof. □

8.2 Tradable beta representation of expected returns

Tradable factor risk premia naturally give rise to an exact beta representation of expected returns in correctly specified, but not necessarily identified, models.

Proposition 8.1. *Let Assumption 1 hold. Then, following statements are equivalent:*

[33] The equality in equation (24) holds after a reordering of the vector components.

1. *There exists a valid factor SDF of the form*

$$M_F(\boldsymbol{\gamma}) := 1 - \boldsymbol{\gamma}'(\boldsymbol{F}_t - \boldsymbol{\mu}_F)$$

for some $\boldsymbol{\gamma} \in \mathbb{R}^K$*.*

2. *Following tradable beta-representation of expected returns holds:*

$$\boldsymbol{\mu}_R = \boldsymbol{V}_{RF}\boldsymbol{\gamma} = \boldsymbol{\beta}^*\boldsymbol{\lambda}^* \ ,$$

for a $N \times K$ *matrix of mimicking portfolio return betas given by:*

$$\boldsymbol{\beta}^* := \boldsymbol{V}_{RF}(\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF})^+ \ , \tag{25}$$

where $\boldsymbol{A}^+$ *denotes the Moore–Penrose inverse of a matrix* $\boldsymbol{A}$*.*

Proposition 8.1 states that in correctly specified models there always exists an exact beta-representation of expected returns, based on tradable factor risk premia $\boldsymbol{\lambda}^*$ and $N \times K$ matrix $\boldsymbol{\beta}^*$ of tradable factor betas. Note that this representation is always valid, irrespective of the identification properties of an asset pricing model. Indeed, recall that $\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF}$ is the variance covariance matrix of the factor mimicking portfolio returns. Therefore, in identified models $\boldsymbol{\beta}^*$ is the uniquely given $N \times K$ matrix such that:

$$\boldsymbol{\beta}^*\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF} = \boldsymbol{V}_{RF} \ ,$$

i.e., the matrix of factor betas of mimicking portfolio returns. More generally, in unidentified models there always exists a multiplicity of $N \times K$ matrices $\tilde{\boldsymbol{\beta}}$ such that:

$$\tilde{\boldsymbol{\beta}}\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF} = \boldsymbol{V}_{RF} \ ,$$

i.e., without additional assumptions the factor betas of portfolio mimicking returns are not identified. In such settings, definition (25) identifies a unique matrix of betas of portfolio factor mimicking returns, which (i) satisfies a tradable beta representation of expected returns and (ii) implies the smallest beta exposures under the Euclidean metric.

Proof of Proposition 8.1. $M_F(\boldsymbol{\gamma})$ is a valid SDF if and only if $\boldsymbol{\mu}_R = \boldsymbol{V}_{RF}\boldsymbol{\gamma}$, which implies:

$$\boldsymbol{\lambda}^* = \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF}\boldsymbol{\gamma} \ .$$

Furthermore, using the definition of generalized inverse, it follows:

$$\boldsymbol{V}_{FR} = \boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF}\boldsymbol{\beta}^{*\prime} \ .$$

Together, this gives:

$$\boldsymbol{\mu}_R = \boldsymbol{V}_{RF}\boldsymbol{\gamma} = \boldsymbol{\beta}^*\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF}\boldsymbol{\gamma} = \boldsymbol{\beta}^*\boldsymbol{\lambda}^* \ .$$

Conversely, if $\boldsymbol{\mu}_R = \boldsymbol{\beta}^*\boldsymbol{\lambda}^*$, it follows:

$$\boldsymbol{\mu}_R = \boldsymbol{\beta}^*\boldsymbol{\lambda}^* = \boldsymbol{V}_{RF}(\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF})^+\boldsymbol{\lambda}^* =: \boldsymbol{V}_{RF}\boldsymbol{\gamma} \ ,$$

with $\boldsymbol{\gamma} := (\boldsymbol{V}_{FR}\boldsymbol{V}_R^{-1}\boldsymbol{V}_{RF})^+\boldsymbol{\lambda}^*$. This concludes the proof. $\square$

8.3 Monte Carlo simulation evidence

We simulate four linear asset pricing models with varying dimension and factor properties:[34]

- Model 1 is a 1-factor model with no useless or weak factor;
- Model 2 (Model 3) is a 3-factor model with two useless (weak) factors and one factor that is neither useless nor weak;
- Model 4 is a higher dimensional model with eleven useless factors, one weak factor and three factors that are neither useless nor weak.

The two weak factors in Model 3 have covariances with returns that are fastly and slowly vanishing at rate $1/T^{3/4}$ and $1/\sqrt{T}$, respectively, while the single weak factor in Model 4 has slowly vanishing covariances at rate $1/\sqrt{T}$.

For each of the above models, we generate a correctly specified and a misspecified version of the model. Returns on the test assets and factors are drawn from a multivariate normal distribution

$$\begin{bmatrix} \boldsymbol{R}_t \\ \boldsymbol{F}_t \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu}_R \\ \boldsymbol{\mu}_F \end{bmatrix}, \begin{bmatrix} \boldsymbol{V}_R & \boldsymbol{V}_{RF} \\ \boldsymbol{V}_{FR} & \boldsymbol{V}_F \end{bmatrix} \right) \ .$$

Moments in this distribution are calibrated to match those of following set of factors and test asset excess returns:

[34]In Online Appendix 8.4, we further report results for the Monte Carlo simulation settings considered in Gospodinov et al. (2014) of a 2-factor model, where the useless factor has mean zero and unit variance, while the factor that is neither useless nor weak has moments calibrated to the market factor.

- The five Fama-French factors, momentum, the q-factors in Hou et al. (2015), the intermediary factors of He et al. (2017), the betting-against-beta factors in Frazzini and Pedersen (2014), and the real per capita nondurable consumption growth;
- The 25 double-sorted portfolios constructed on size and book-to-market and 17 industry portfolios.

Data for the moment calibration are obtained from the Kenneth French data library and consist of monthly observations from January 1970 to October 2021 for the considered set of assets and factors, which gives rise to a panel including 624 observations across time. Accordingly, in our Monte Carlo simulation we produce $10'000$ random samples consisting each of 624 time series observations of factors and excess returns.

Regarding factors that are neither weak nor useless, in Model 1–3, the single such factor is calibrated to the market factor. In Model 4, they are calibrated to the three available factors implying the highest sum of squared sample correlations with test asset returns. These are the size factor (R_{ME}), the expected growth factor (R_{EG}) and the profitability factor (R_{ROE}).

Regarding factors that are weak or useless, their means and variances are calibrated to the SMB and HML factors in Model 2–3, and to the twelve factors implying the lowest sum of square sample correlations with test asset returns in Model 4. For these factors, means and marginal variances are calibrated from the data, while the covariances between factors are set to zero. The factor covariances with asset returns are set to zero when simulating useless factors and appropriately scaled by $1/T^{3/4}$ or by $1/\sqrt{T}$ when simulating weak factors with fastly and slowly vanishing covariances with returns.

Our simulations cover both correctly specified and misspecified models. Following, e.g., Gospodinov et al. (2014), the misspecified version of each model is simulated by setting the expected asset return vector $\boldsymbol{\mu}_R$ equal to the sample asset returns means in the data. Instead, correctly specified models are simulated by setting

$$\boldsymbol{\mu}_R = \boldsymbol{\beta}_{(\mathcal{S})}\boldsymbol{\lambda}_{\mathcal{S}}^{KRS} = \boldsymbol{\beta}^*\boldsymbol{\lambda}^* \ .$$

Here, $\boldsymbol{\lambda}_{\mathcal{S}}^{KRS}$ ($\boldsymbol{\beta}_{(\mathcal{S})}$) is the vector (matrix collecting the columns) of misspecification-robust factor risk premia (factor betas) for the subset of factors that are not useless or weak.[35] All these parameters act as population parameters in our

[35] Recall from Proposition 8.1 that when the model is correctly specified, expected returns are equivalently given by $\boldsymbol{\beta}^*\boldsymbol{\lambda}^*$, where $\boldsymbol{\beta}^*$ ($\boldsymbol{\lambda}^*$) is the matrix (vector) of factor mimicking portfolios return betas (tradable factor risk premia).

simulations, and are appropriately calibrated to the data. In Models 2–4, which include useless or weak factors, we set $\boldsymbol{\lambda}^{KRS}_{\mathcal{S}^c} = \mathbf{0}$ for the factor risk premia of factors that are weak or useless, as Kan et al. (2013) risk premia are not identified in these models.

8.3.1 Finite sample bias-variance tradeoff and factor selection

We start by studying the finite sample bias-variance tradeoffs emerging for the KRS factor risk premium estimator and for our sample, Oracle and Relaxed estimators of tradable factor risk premia. For each simulated model, Figures 7–10 report the decomposition of the estimators' MSE into its squared bias and variance components.[36] Overall we observe that, with the exception of Model 1 where the estimators' MSE are almost identical, the KRS estimator's MSE is considerably higher than the one of sample, Oracle and relaxed estimators of tradable factor risk premia. This feature is a direct consequence of the fact that the KRS risk premia of useless and weak factors are not estimated consistently, which gives rise to a substantially larger estimator variance in the box plots of Figures 11–16. When further comparing our sample and Oracle estimators of tradable factor risk premia, the latter estimator gives rise to some degree of bias for the risk premia of factors that are not weak nor useless. This feature is a natural consequence of the shrinkage properties of this estimator. This shrinkage effect also implies that the tradable risk premia of useless and weak factors are estimated with a clearly lower variance, which results in a lower overall MSE with respect to the MSE of sample tradable factor risk premia. Finally, the relaxed estimator successfully removes the finite-sample bias implied by the Oracle estimator, while giving rise to a higher (lower) overall mean square error in misspecified (correctly specified) models.

Given the problematic finite-sample properties of misspecification-robust factor risk premium estimator (5.1) in weak or not identified models, we next ask how much its Pen-FM version (23) improves on them. Figure 18 reports the corresponding estimators' MSE decompositions for Models 2 and 4. Overall, we observe that Pen-FM estimator (23) nicely improves on estimator (5.1) in the low-dimensional Model 2, by providing comparable MSEs as our Oracle estimator. In contrast, for the higher dimensional Model 4, it still implies large MSEs, which exhibit a large bias component under both cross-validation schemes used. On the other hand, our Oracle estimator gives rise to clearly lower MSEs and virtually no bias when us-

[36] The bias of each estimator is computed with respect to the associated population factor risk premium definition.

ing 5-fold cross-validation. This evidence indicates that estimating precisely misspecification-robust factor risk premia in weakly or not identified asset pricing models may be a challenge also for Oracle-type estimators of these risk premia.

Figure 19 summarizes the finite sample selection properties of our Oracle estimator of tradable factor risk premia, based on a Generalized cross-validation scheme. To this end, we report the Monte Carlo simulation frequency with which the active set estimated by our Oracle estimator equals or contains, respectively, the unknown active subset of factors that are neither useless nor weak. Overall, we find that our methodology produces a quite reasonable factor selection. Indeed, we find that the finite sample factor selection is virtually exact in Model 1. In Model 2–3 it is correct in more than 98% (around 75-80%) of the cases in correctly specified (misspecified) models, and it virtually always contains the unknown active subset of factors that are neither useless nor weak. Even in the large dimensional Model 4, a correct factor selection is obtained in 95% (74%) of the cases for correctly specified (misspecified) models, while the unknown active subset of factors that are neither useless nor weak is contained in the estimated active set in 95% of the cases

Figure 20 sets the above evidence in relation to the finite sample factor selection properties of Oracle estimator (23) of misspecification-robust factor risk premia. Coherently with the above MSE results, we find that for both cross-validation schemes used this estimator produces a satisfactory factor selection in low-dimensional Model 2. In contrast, it gives rise often to incorrect factor selections in higher dimensional Model 4. On the other hand, we find that 5-fold cross-validation further improves the factor selection properties of our Oracle estimator under Generalized cross-validation, by providing a virtually perfect selection for all models considered.

In summary, we conclude that a reliable estimation and factor selection based on misspecification-robust factor risk premia may be hardly achievable in weak or not identified models, even when using Oracle-type estimators of these risk premia. Nonetheless, a trustworthy Oracle estimation and factor selection can still be produced based on tradable factor risk premia.

8.3.2 Finite-sample inference

We next document the finite-sample inference properties implied by KRS factor risk premium estimators and by the sample, Oracle and Relaxed estimators of tradable factor risk premia. For the estimators of tradable risk premia, we obtain a feasible finite sample inference based on the asymptotic standard error formulas implied by Propositions 4.1 and 4.3. For the KRS estimator, we rely on the misspecification-robust standard error formula implied by Proposition 2 of the internet appendix of Kan et al. (2013).

Figure 21 reports coverage probabilities for the event that a candidate risk premium parameter λ^* in a neighborhood of true population risk premium λ_0 is contained in a finite sample 95% confidence interval implied by the different methods. Therefore, when $\lambda^* = \lambda_0$, one minus this probability is the finite sample size of a test for testing null hypothesis $\lambda = \lambda_0$. Conversely, when $\lambda^* \neq \lambda_0$, one minus this probability gives the finite sample power of a test for testing null hypothesis $\lambda = \lambda^*$ against alternative hypothesis $\lambda = \lambda_0$. Since we obtain a similar and coherent evidence for all simulated models, we focus for brevity in the sequel on the evidence produced for the correctly specified versions of Model 4.

A first general pattern is that, relative to misspecification-robust risk premia, sample tradable risk premium estimators imply an improved ability to reject the hypothesis of a non-zero risk premium for weak and useless factors. This feature is reflected in the coverage probabilities in the lower panels of Figure 21. Indeed, while the size behavior of both tests under the null of a zero risk premium for the useless factor gives rise to a conservative empirical size lower than the nominal one, the null hypothesis of a non zero tradable risk premium is correctly rejected with a rapidly increasing probability using sample tradable risk premia, as the distance between the null of a nonzero risk premium and the alternative of a zero risk premium increases. In contrast, the same null hypothesis for KRS misspecification-robust risk premia is rarely rejected, even for quite large absolute risk premia under the null. The well-known conservative behavior of the inference implied by KRS misspecification-robust risk premia for testing the null of a zero risk premium for the useless factors, is the foundation of the sequential testing procedure proposed in Gospodinov et al. (2014) for consistently eliminating these factors from an asset pricing model. In our methodology, the consistent elimination of all useless and weak factors is performed jointly by our Oracle estimator of tradable factor risk premia. This feature is directly visible in the lower panels of Figure 21, by the fact that such estimator produces in virtually almost all simulations a zero estimated tradable risk premium.

The evidence in the upper panels of Figure 21 suggests that the inference regarding the tradable risk premia of factors that are not weak nor useless is unaffected by the presence of weak or useless factors in an asset pricing model. This feature holds exactly asymptotically for our estimators, as shown in Propositions 4.1 and 4.3. In contrast, Gospodinov et al. (2014) show that the same feature does not apply to estimators of KRS misspecification-robust risk premia.[37] Indeed, as displayed in the upper panels of Figure 21, the finite sample inference implied by sample tradeable risk premia is well-behaved and sharper than the one implied by misspecification-robust risk premia, in the sense that their coverage probabilities are decaying faster as we move away from the corresponding population risk premium. The coverage probabilities of our Oracle estimator present a similar sharp behavior, but they are shifted with respect to the population tradable risk premium due to the finite sample bias induced by the shrinkage effect. As shown in the figure, this finite sample bias can be reduced by means of the Relaxed estimator.

8.3.3 Monte Carlo simulation evidence: figures

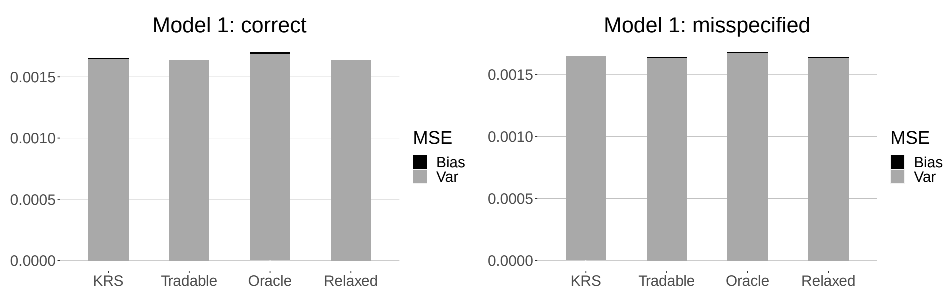

Figure 7: **Estimators' MSE decomposition, Model 1**: the figure presents the MSE bias-variance decomposition of KRS, tradable, Oracle and Relaxed factor risk premia estimators in simulation Model 1, for both correctly specified (left panel) and misspecified (right panel) version.

[37]This is the reason why they propose in the first place a sequential testing procedure for consistently eliminating the useless factors from an asset pricing model.

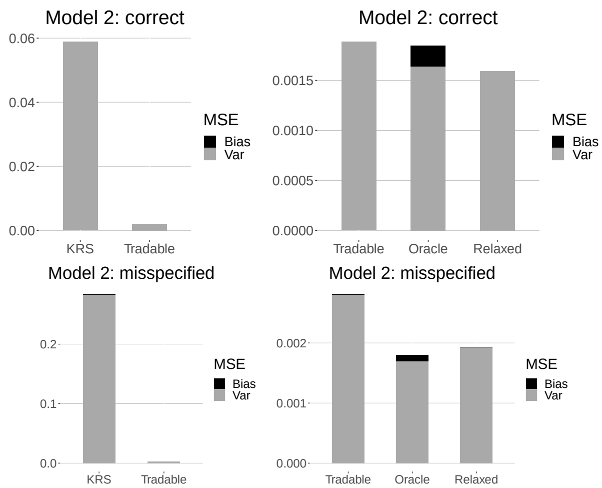

Figure 8: **Estimators' MSE decomposition, Model 2**: the figure presents the MSE bias-variance decomposition of KRS (left panels), tradable (in both left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators in simulation Model 2, for both correctly specified (upper panels) and misspecified (lower panel) version.

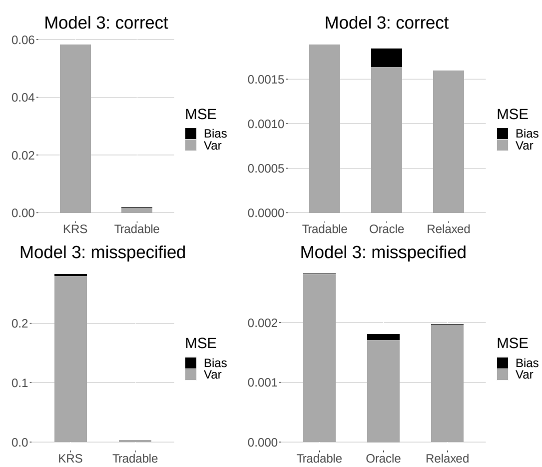

Figure 9: **Estimators' MSE decomposition, Model 3**: the figure presents the MSE bias-variance decomposition of KRS (left panels), tradable (in both left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators in simulation Model 3, for both correctly specified (upper panels) and misspecified (lower panels) version.

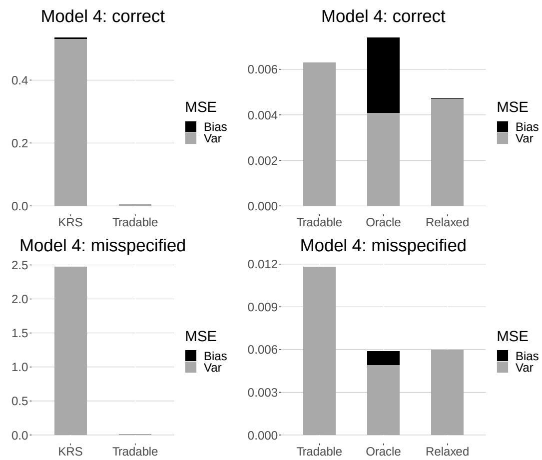

Figure 10: **Estimators' MSE decomposition, Model 4**: the figure presents the MSE bias-variance decomposition of KRS (left panels), tradable (in both left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators in simulation Model 4, for both correctly specified (upper panels) and misspecified (lower panels) version.

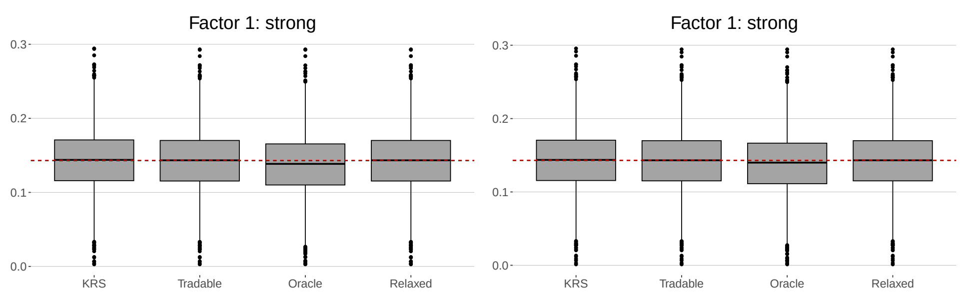

Figure 11: **Estimators' simulated distribution, Model 1**: the figure presents the simulated distribution of KRS, tradable, Oracle and Relaxed factor risk premia estimators of the strong factor of simulation Model 1, for both correctly specified (left panel) and misspecified (right panel) version. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

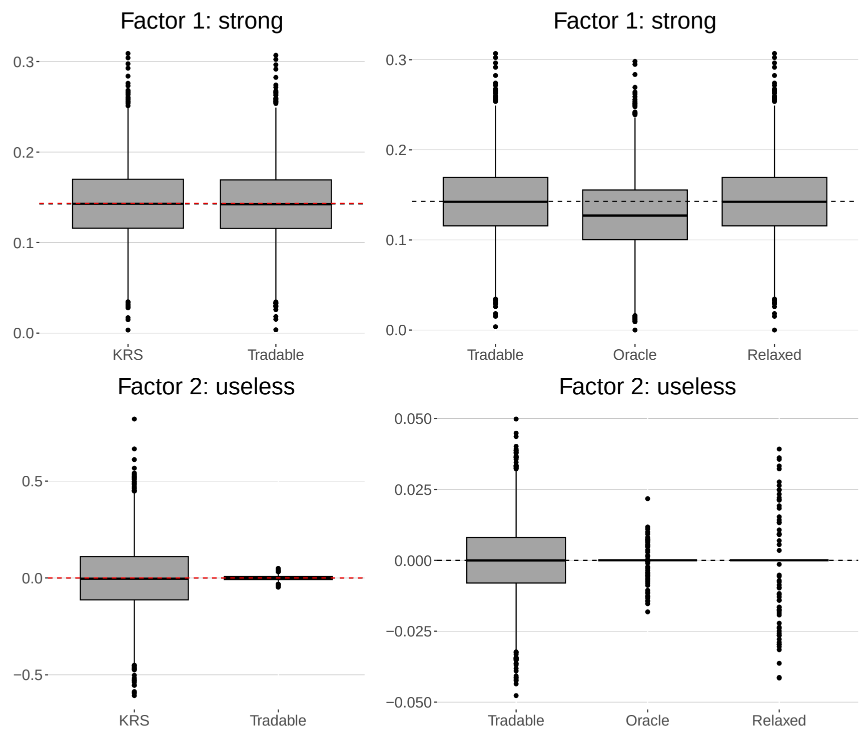

Figure 12: **Estimators' simulated distribution, Model 2 (correctly specified)**: the figure presents the marginal simulated distributions of KRS (left panels), tradable (left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators of the strong (upper panel) and one of the two useless (lower panel) factors of the correctly specified version of simulation Model 2. We omit the figure of the second useless factor as it is qualitatively identical to the one of the first useless factor. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

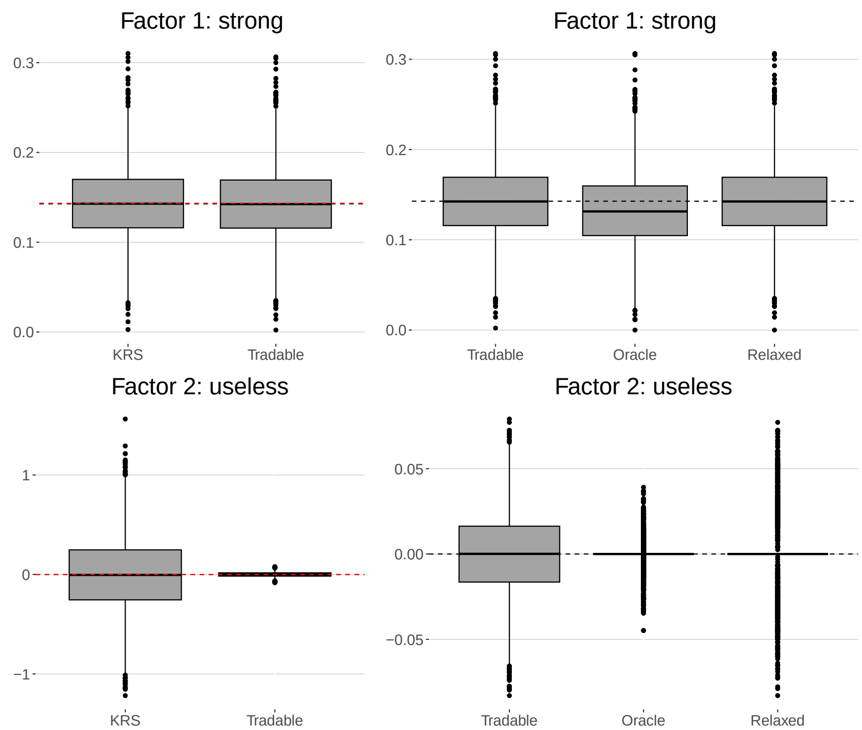

Figure 13: **Estimators' simulated distribution, Model 2 (misspecified)**: the figure presents the marginal simulated distributions of KRS (left panels), tradable (left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators of the strong (upper panel) and one of the two useless (lower panel) factors of the misspecified version of simulation Model 2. We omit the figure of the second useless factor as it is qualitatively identical to the one of the first useless factor. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

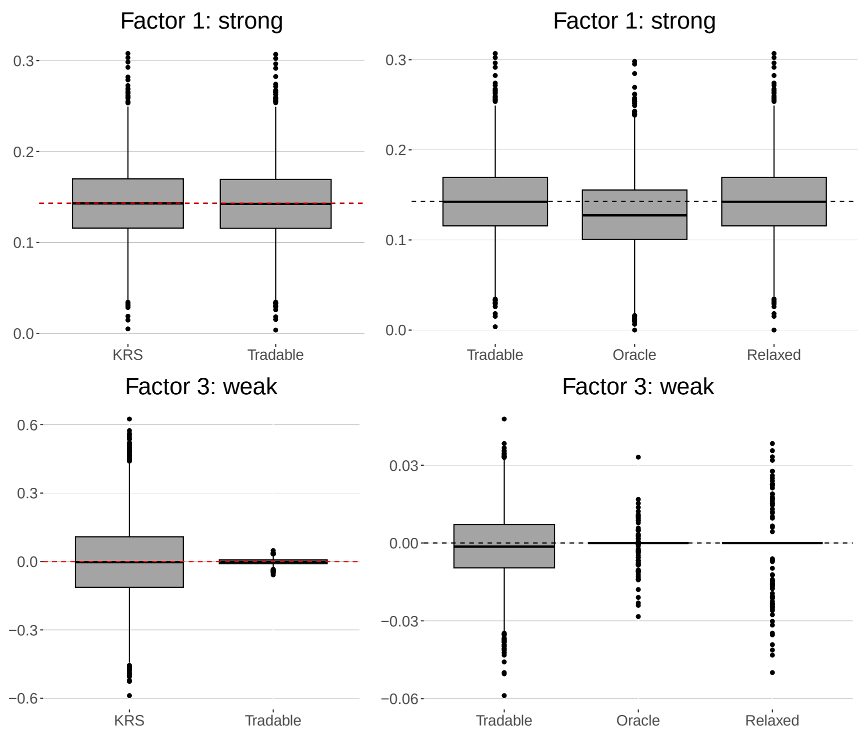

Figure 14: **Estimators' simulated distribution, Model 3 (correctly specified)**: the figure presents the marginal simulated distributions of KRS (left panels), tradable (left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators of the strong (upper panels)and the $1/\sqrt{T}-$weak (lower panels) factors of the correctly specified version of simulation Model 3. We omit the figure of the $1/T^{3/4}-$weak factor as it is qualitatively identical to the one of the $1/\sqrt{T}-$weak factor. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

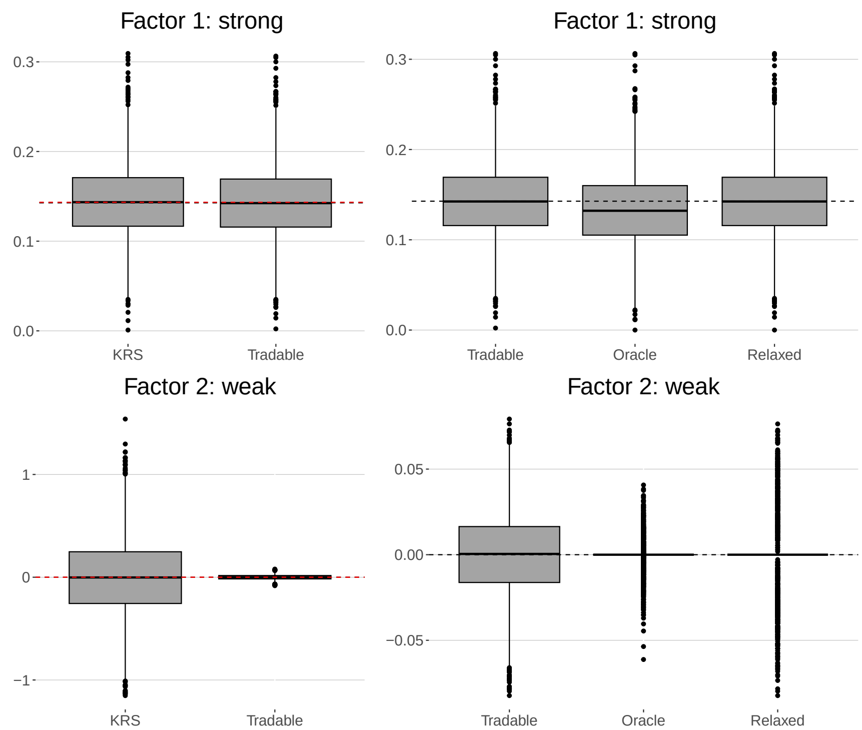

Figure 15: **Estimators' simulated distribution, Model 3 (misspecified)**: the figure presents the marginal simulated distributions of KRS (left panels), tradable (left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators of the strong (upper panels), the $1/T^{3/4}-$weak (central panels) and the $1/\sqrt{T}-$weak (lower panels) factors of the misspecified version of simulation Model 3. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

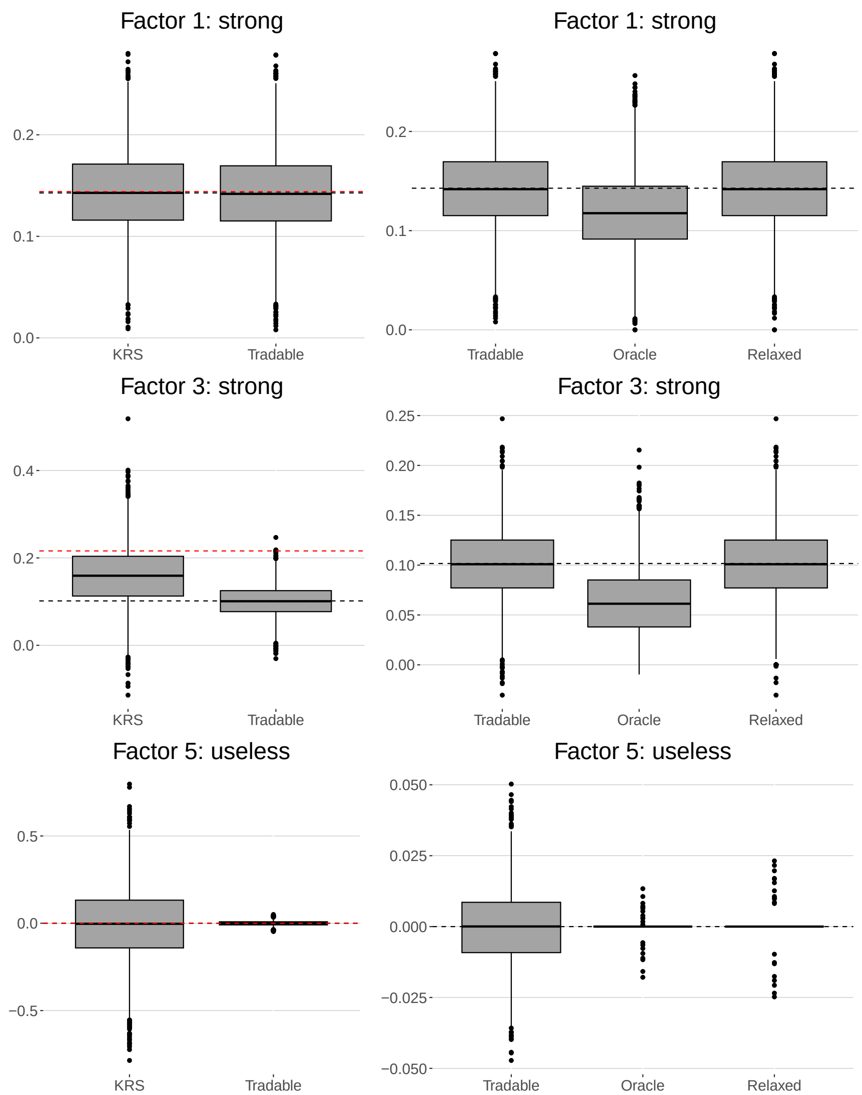

Figure 16: **Estimators' simulated distribution, Model 4 (correctly specified)**: the figure presents the marginal simulated distributions of KRS (left panels), tradable (left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators of the first strong (upper panels), the third strong (central panels) and the first useless (lower panels) factors of the correctly specified version of simulation Model 4. We omit the figure of the other useless factors and of the weak factor as they are qualitatively identical to the one of the first useless factor. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

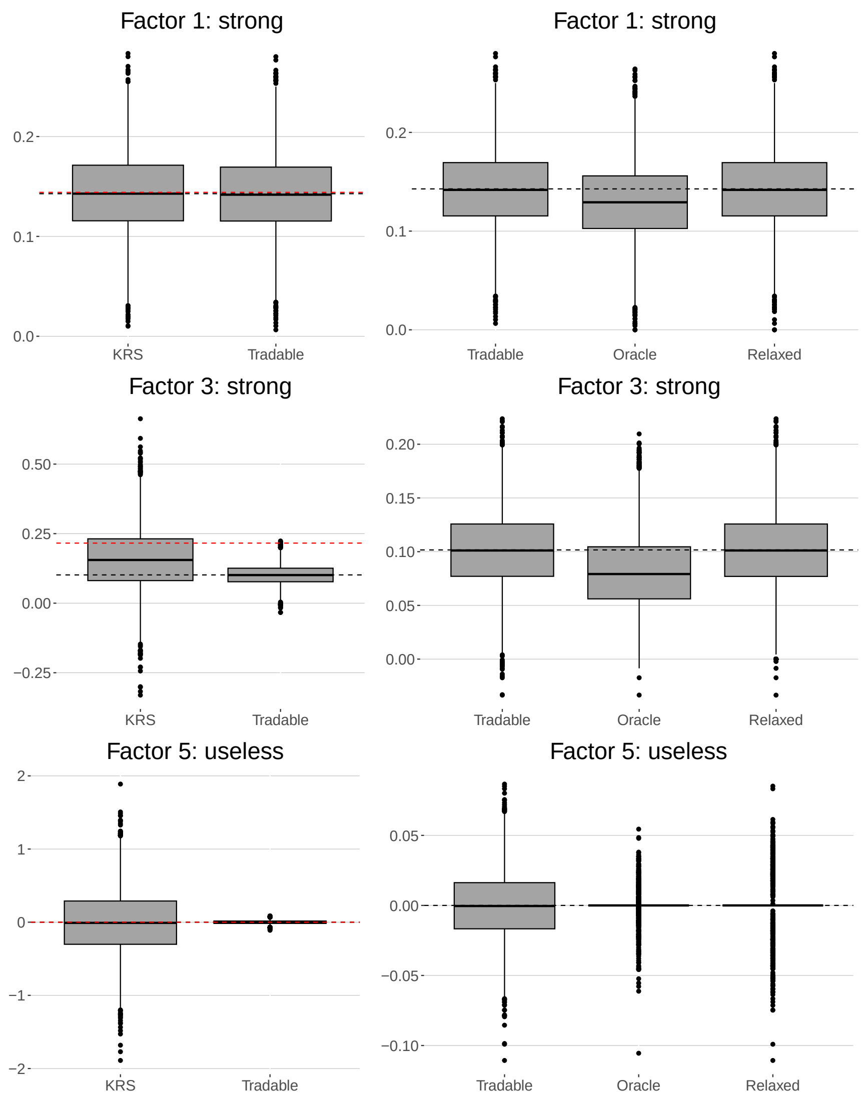

Figure 17: **Estimators' simulated distribution, Model 4 (correctly specified)**: the figure presents the marginal simulated distributions of KRS (left panels), tradable (left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators of the first strong (upper panels), the third strong (central panels) and the first useless (lower panels) factors of the correctly specified version of simulation Model 4. We omit the figure of the other useless factors and of the weak factor as they are qualitatively identical to the one of the first useless factor. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

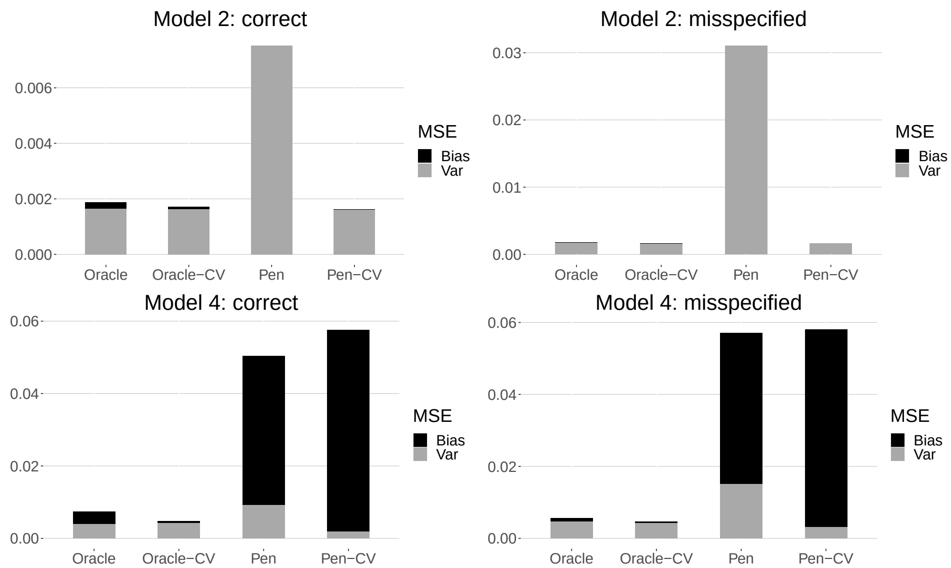

Figure 18: **Estimators' MSE decomposition, Models 2 and 4**: MSE bias-variance decomposition of the Pen-FM (Bryzgalova (2015)) and our Oracle factor risk premia estimators, for both correctly specified (left panel) and misspecified (right panel) versions of simulation Models 2 (upper panels) and 4 (lower panels). Pen and Oracle (Pen-CV and Oracle-CV) are tuned using Generalized Cross Validation (5-fold Cross Validation).

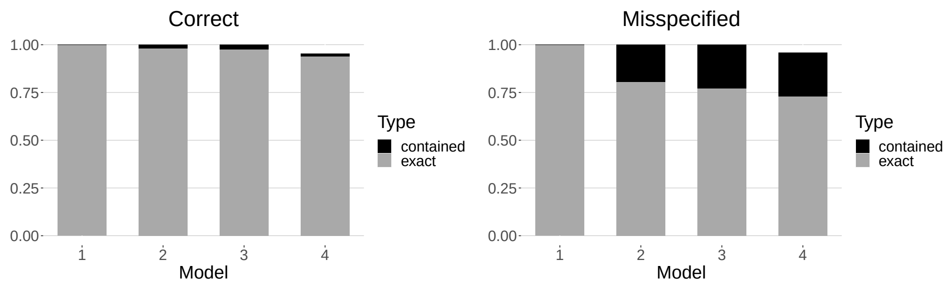

Figure 19: **Oracle factor risk premia selection properties, Models 1–4**: the figure presents the "exact" and "contained" variable selection properties of our Oracle factor risk premia estimator for both correctly specified (left panel) and misspecified (right panel) versions of simulation Models 1–4. The "exact" ("contained") variable selection property is computed as the frequency over the simulation runs that the active set of the population tradable factor risk premia coefficient equals (is contained in) the estimated active set of our Oracle risk premia estimator.

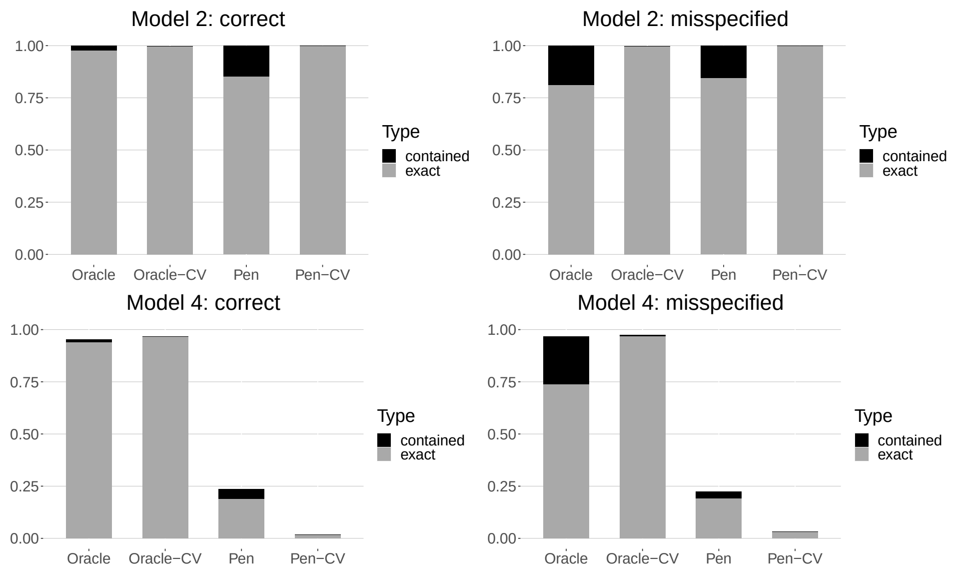

Figure 20: **Factor risk premia selection comparison, Models 2 and 4**: the figure presents the "exact" and "contained" variable selection properties of the Pen-FM (Bryzgalova (2015)) and our Oracle factor risk premia estimators, for both correctly specified (left panel) and misspecified (right panel) versions of simulation Models 2 (upper panels) and 4 (lower panels). The "exact" ("contained") variable selection property is computed as the fraction over the simulation runs that the active set of the KRS (tradable) population factor risk premia coefficient equals (is contained in) the estimated active set of the Pen-FM (our Oracle) estimator. Pen and Oracle (Pen-CV and Oracle-CV) are tuned using Generalized Cross Validation (5-fold Cross Validation).

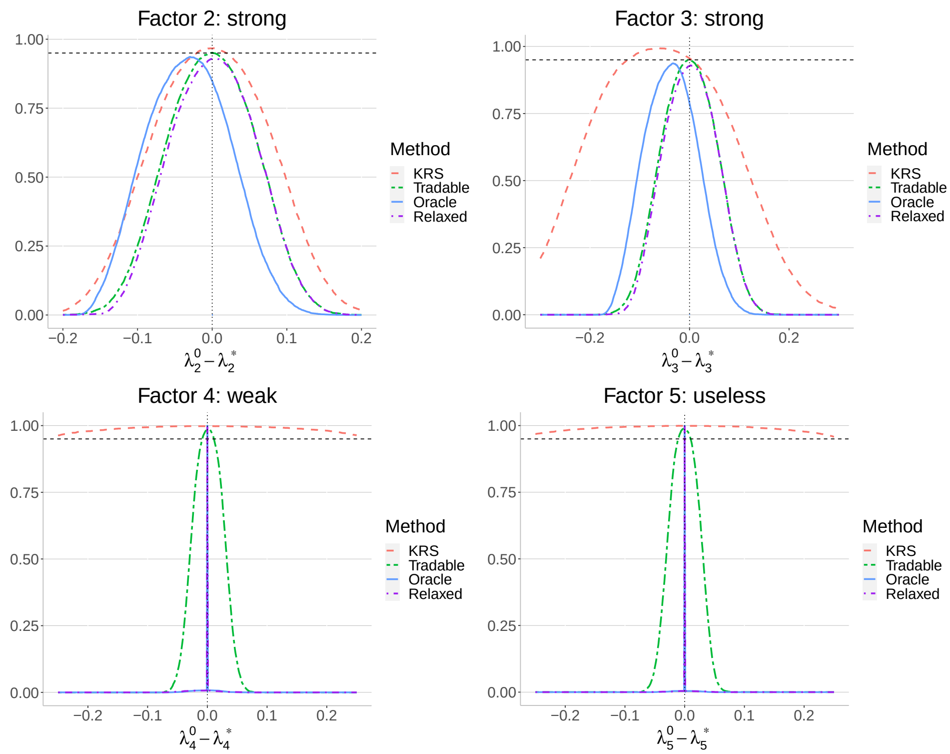

Figure 21: **Inclusion frequencies of estimators' 95% confidence intervals, Model 4 (correctly specified)**: the figure presents the frequencies with which various values of factor's j risk premium λ_j^* around the population coefficient λ_j^0 are contained in the corresponding KRS (red dashed line), tradable (green dashed-dotted line), Oracle (solid blue line) and Relaxed (dashed-dotted purple line) estimators' 95% confidence interval across the simulation runs, for the first strong factor (upper left panel, $j = 1$), the second strong factor (upper right panel, $j = 2$), the $1/\sqrt{T}-$weak factor (lower left panel, $j = 4$) and the first useless factor (lower right panel, $j = 5$) of the correctly specified version of simulation Model 4. To facilitate the comparison, we display these frequences as a function of the deviation from the estimators' corresponding population risk premia coefficient λ_j^0, which is the population KRS risk premia coefficient for the KRS estimator, and the population tradable risk premia coefficient for the tradable, Oracle and Relaxed estimators. The horizontal dashed black line represents the nominal level of the confidence intervals, namely 95%.

8.4 Additional Monte Carlo simulation evidence

We report the simulation evidence for the two-factor model in Gospodinov et al. (2014), in which returns on the test assets and factors are simulated in a similar manner as in Section 8.3, but with moments of the factor that is neither weak nor useless calibrated to the market factor and moments of the useless factor given by a zero mean and a unit variance. As shown in Figures 22–24 below, the analysis and conclusions regarding the MSEs and finite sample distributions of our sample, Oracle and Relaxed estimators of tradable factor risk premia, in relation to the misspecification-robust estimators of Kan et al. (2013) of misspecification-robust factor risk premia, remain unchanged. Similarly, Figure 25 displays the high degree of factor selection accuracy achieved by our Oracle estimator in this simulation setting.

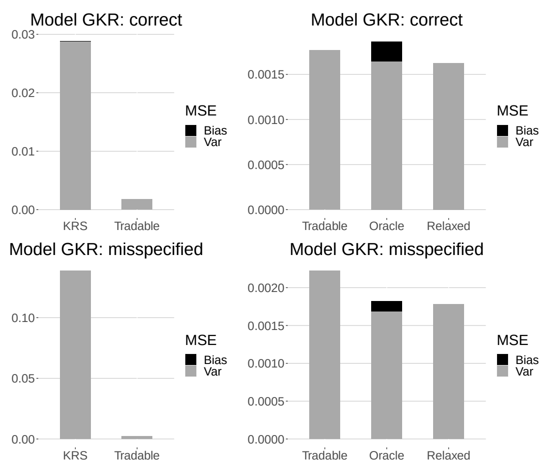

Figure 22: **Estimators' MSE decomposition, Model GKR**: the figure presents the MSE bias-variance decomposition of KRS (left panels), tradable (both left and right panels), Oracle and Relaxed (right panels) factor risk premia estimators in simulation Model GKR, for both correctly specified (upper panels) and misspecified (lower panels) version.

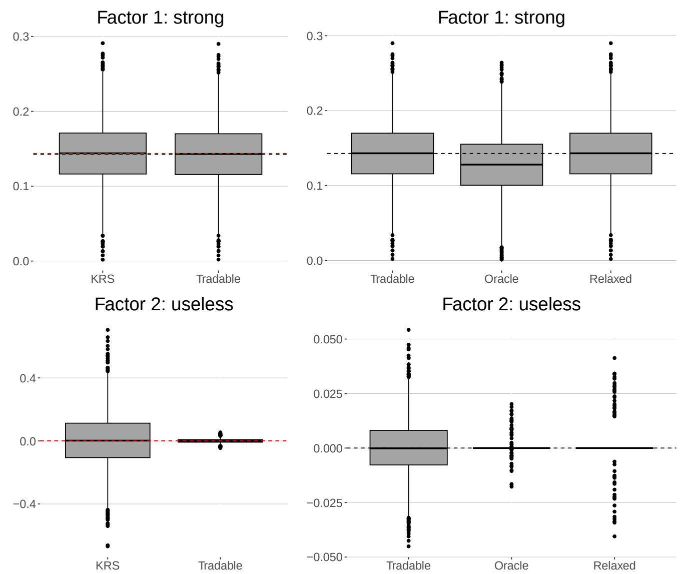

Figure 23: **Estimators' empirical distribution, Model GKR (correctly specified)**: the figure presents the empirical distribution of KRS (left panels), tradable (left panels), Oracle and Relaxed (right panels) factor risk premia estimators of the strong (upper panel) and the two useless (lower two panels) factors of the correctly specified version of simulation Model GKR. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

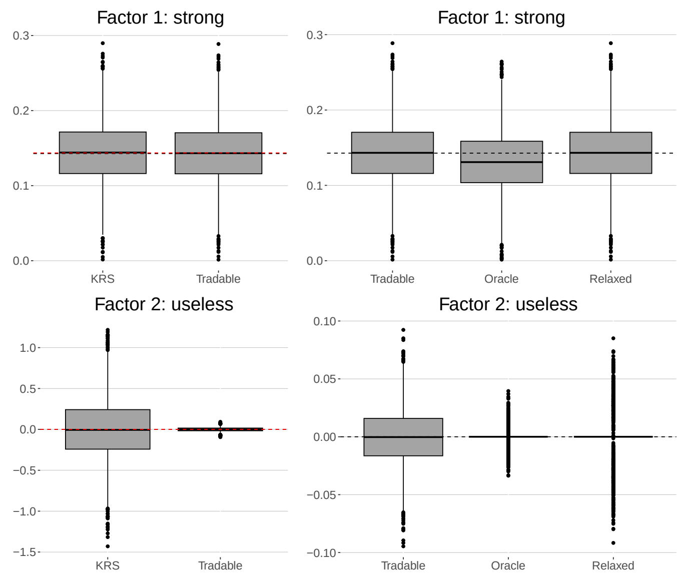

Figure 24: **Estimators' empirical distribution, Model GKR (misspecified)**: the figure presents the empirical distribution of KRS (left panels), tradable (left panels), Oracle and Relaxed (right panels) factor risk premia estimators of the strong (upper panel) and the two useless (lower two panels) factors of the misspecified version of simulation Model GKR. The population tradable and KRS factor risk premia are represented by the horizontal dashed black and red lines, respectively.

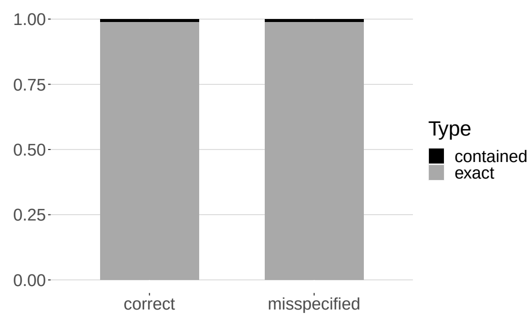

Figure 25: **Proximal factor risk premia selection properties, Model GKR**: the figure presents the "exact" and "contained" variable selection properties of our Oracle factor risk premia estimator for both correctly specified and misspecified simulation Model GKR. The "exact" ("contained") variable selection property is computed as the fraction over the simulation runs that the active set of the tradable population factor risk premia coefficient equals (is contained in) the estimated active set of our proximal risk premia estimator.

8.5 Additional empirical evidence

This section complements the empirical analysis outlined in Section 6 of the primary text. In Subsection 8.5.1, we assess model identification, post-screening factor dimension and factor selection across randomized models for the same sets of tests assets as in the main text but with initial factor dimension ranging from 1 to 4. Conversely, Subsection 8.5.2 is dedicated to scrutinizing the robustness of our empirical findings across different sets of test assets.

8.5.1 Evidence for low-dimensional randomized factor models

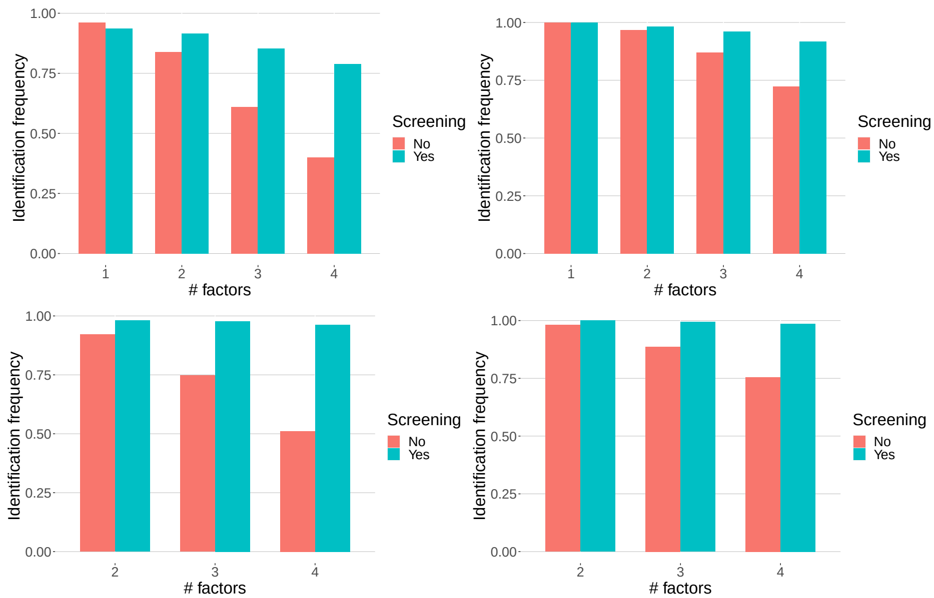

Figure 26: **Model identification frequencies**: Frequency of null hypothesis (6.1) rejections under a significance level $\alpha = 0.05$ by the Chen and Fang (2019) test, across randomized factor models including 1–4 initial factors. The red (blue) bars indicate model identification frequencies before (after) having applied our Oracle factor selection. The upper (lower) panels report the model identification frequencies for selections with no (the market as) ubiquitous factor across models. The left (right) panels report the results for test assets comprising the 25 size/book-to-market (the 25 size/book-to-market and the 17 industry) portfolios. For models having 1 (2) randomized factors, we consider all 52 (1326) possible model combinations. For models including more than 2 randomized factors, we consider 10'000 random factor combinations.

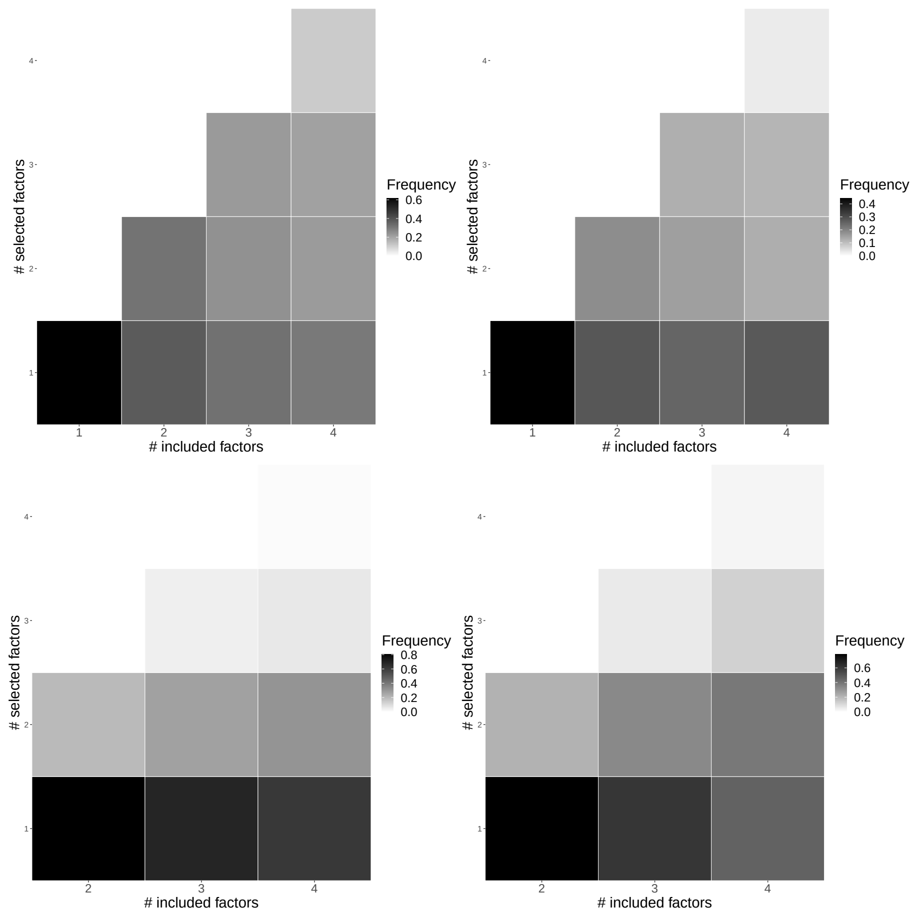

Figure 27: **Post-selection model size**: Frequency of post-selection model dimensions, i.e., number of factors selected by our Oracle factor selection, across randomized factor models including 1–4 initial factors. The red (blue) bars indicate model identification frequencies before (after) having applied our Oracle proximal factor selection. The upper (lower) panels report the model identification frequencies for selections with no (the market as) ubiquitous factor across models. The left (right) panels report the results for test assets comprising the 25 size/book-to-market (the 25 size/book-to-market and the 17 industry) portfolios.

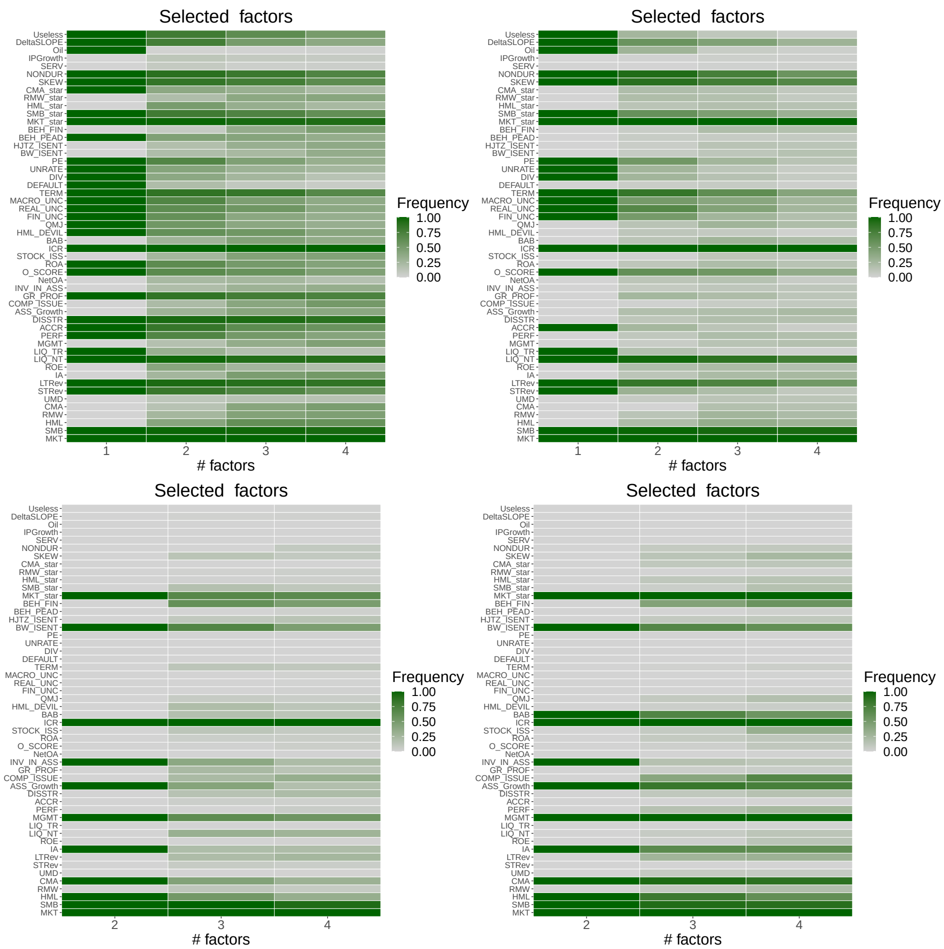

Figure 28: **Factor selection frequencies**: Selection frequencies of individual factors determined through our Oracle factor screening procedure, observed across randomized factor models containing the specific factor along with 0–3 additional factors. The upper (lower) panels report the frequencies of the post-selection model dimensions for selections with no (the market as) ubiquitous factor across models. The left (right) panels report the results for test assets comprising the 25 size/book-to-market (the 25 size/book-to-market and the 17 industry) portfolios.

8.5.2 Robustness to the choice of test assets

To assess the robustness of the empirical results obtained in Section 6 to the choice of the test assets, we perform the same analysis for (i) the

25 size/book-to-market and 25 operating profitability/investment portfolios, and (ii) the 25 size/book-to-market and a varying number of principal components extracted from 17 industry portfolios and 310 double-sorted portfolios. This latter large cross-section of sorted portfolios is constructed from various combinations of asset characteristics, including size, book-to-market, operating profitability, investment, net issuance, beta, variance, accruals, short-term reversal, long-term reversal, and momentum. Overall, the analysis performed with these sets of test assets conforms remarkably to the conclusions obtained in Section 6. Importantly, Figure 34 shows that, even when applied to these different sets of test assets, our Oracle screening procedure detects the same low dimensional set of factors that give rise to well-identified models and have a nonzero tradable risk premium. These factors include the market, size (SMB), market with a hedged unpriced component (MKT_star), intermediary capital ratio (ICR), long-term behavioral factor (BEH_FIN) and the management mispricing factor (MGMT).

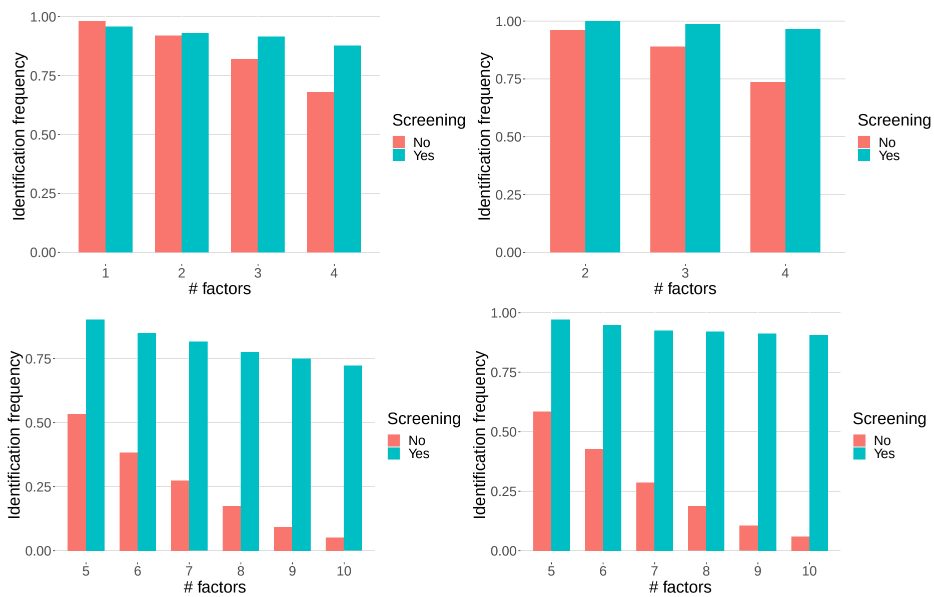

Figure 29: **Model identification frequencies**: Frequencies over various combinations of factor models having 1–10 initial number of factors that null hypothesis (6.1) is rejected, with a size $\alpha = 0.05$, by the Chen and Fang (2019) test. The red (blue) bars indicate the model identification frequencies before (after) having applied our proximal factor selection. The left (right) panels report the model identification frequencies for a model with no (the market as) common factor in the various models. The results are obtained with test assets comprising the 25 size/book-to-market and the 25 operating profitability/investment portfolios.

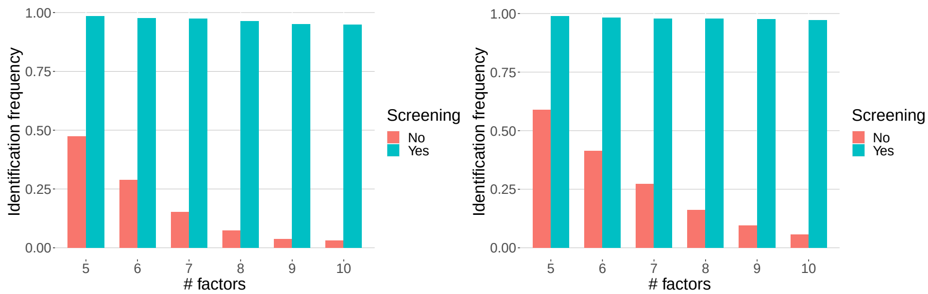

Figure 30: **Model identification frequencies**: Frequencies over various combinations of factor models having 5–10 initial number of factors that null hypothesis (6.1) is rejected, with a size $\alpha = 0.05$, by the Chen and Fang (2019) test. The red (blue) bars indicate the model identification frequencies before (after) having applied our proximal factor selection. The results are obtained by always including the market factor in the initial model, and the test assets used comprise the 25 size/book-to-market and 5 PCs (left panel) or 8 PCs (right panel) extracted from 17 industry and 310 double-sorted portfolios.

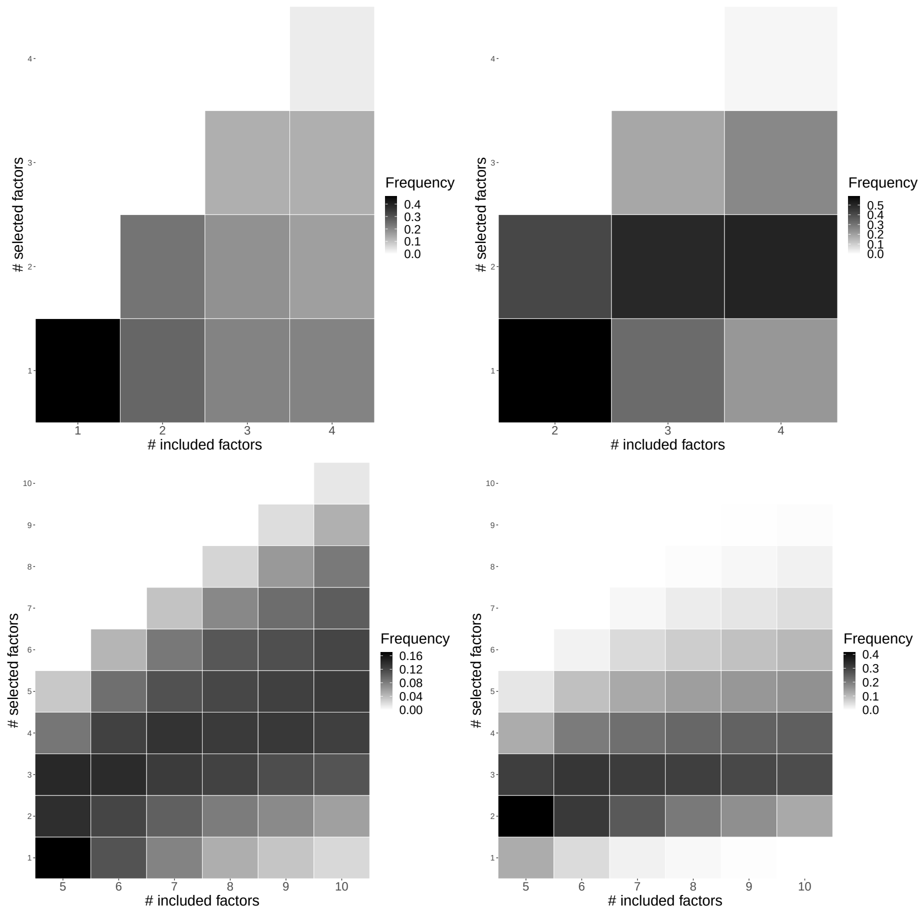

Figure 31: **Post-selection model size**: Frequencies of the post-selection model size, i.e., number of factors selected by our proximal factor selection, over the various combinations of factor models having 1–10 initial number of factors. The left (right) panels report the frequencies of the post-selection model size for a model with no (the market as) common factor in the various models. The results are obtained with test assets comprising the 25 size/book-to-market and the 25 operating profitability/investment portfolios.

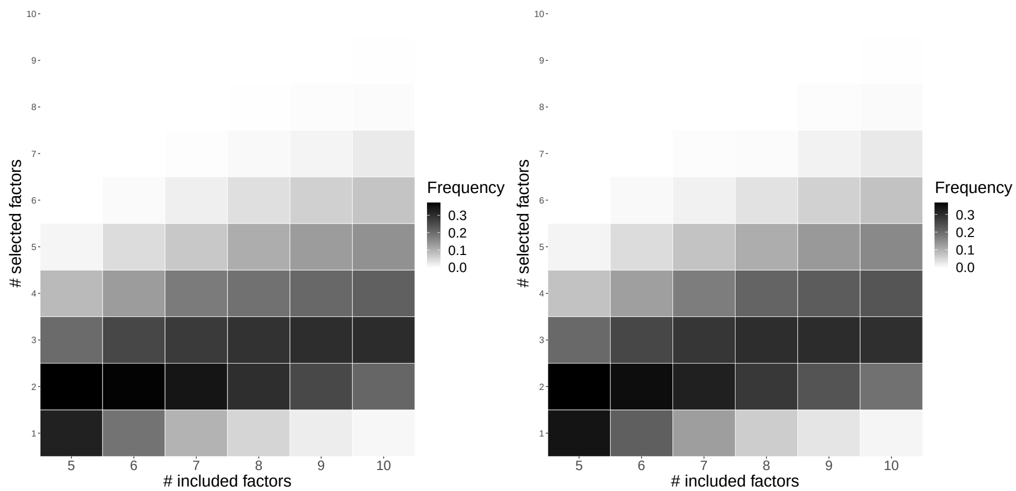

Figure 32: **Post-selection model size**: Frequencies of the post-selection model size, i.e., number of factors selected by our proximal factor selection, over the various combinations of factor models having 1–10 initial number of factors. The results are obtained by always including the market factor in the initial model, and the test assets used comprise the 25 size/book-to-market and 5 PCs (left panel) or 8 PCs (right panel) extracted from 17 industry and 310 double-sorted portfolios.

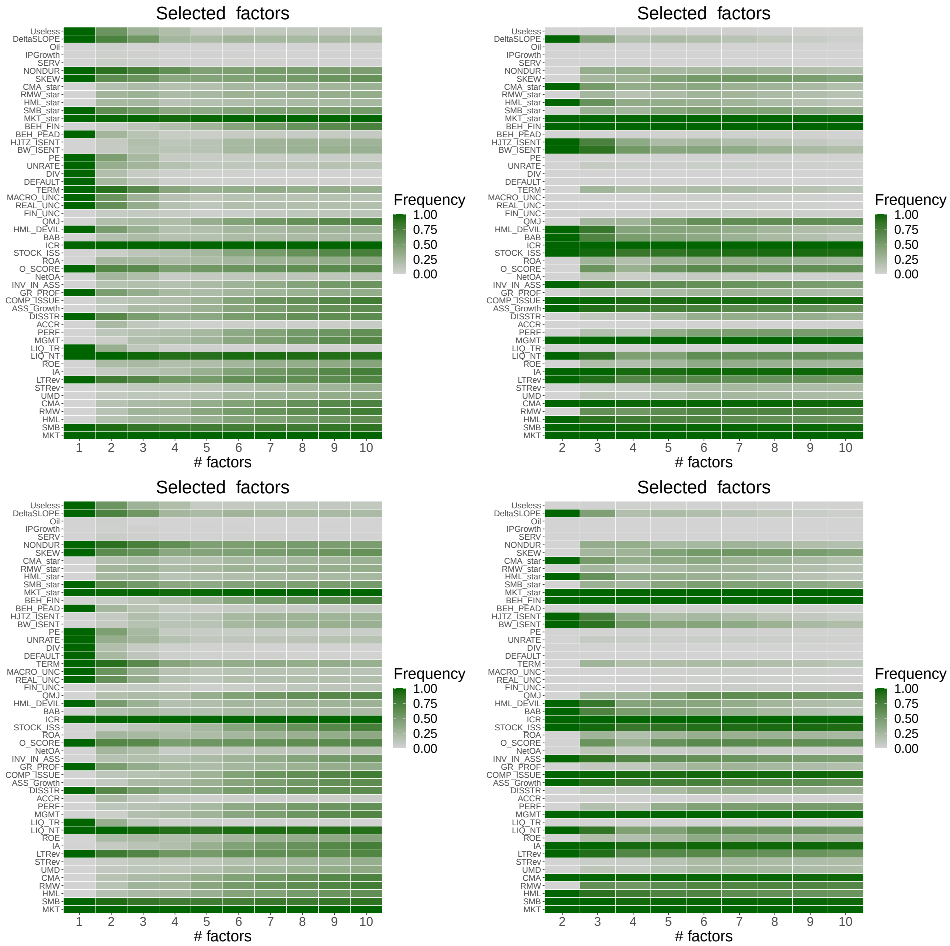

Figure 33: **Factor selection frequencies**: Selection frequencies of individual factors determined through our Oracle factor screening procedure, observed across randomized factor models containing the specific factor along with 4–9 additional factors. The left (right) panel reports the factor selection frequencies for a model with no (the market as) common factor in the various models. These results are obtained with test assets comprising the 25 size/book-to-market and the 25 operating profitability/investment portfolios.

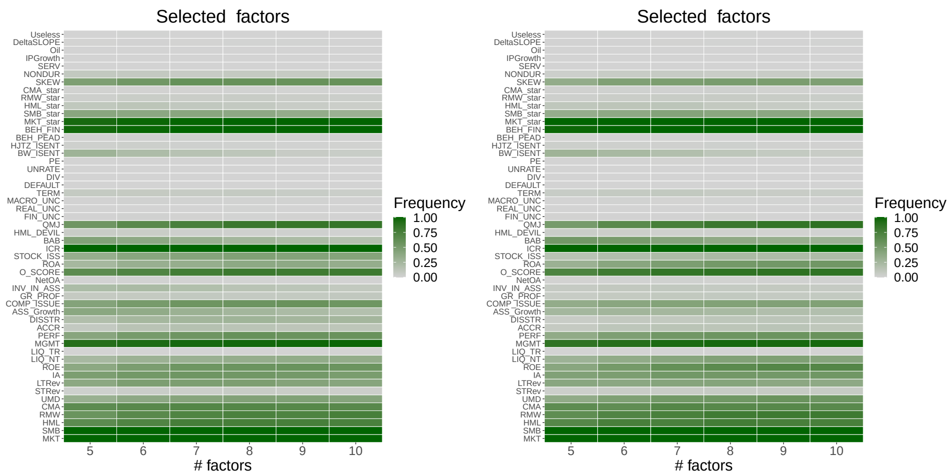

Figure 34: **Factor selection frequencies**: Selection frequencies of individual factors determined through our Oracle factor screening procedure, observed across randomized factor models containing the specific factor along with 4–9 additional factors. The results are obtained by always including the market factor in the initial model, and the test assets used comprise the 25 size/book-to-market and 5 PCs (left panel) or 8 PCs (right panel) extracted from 17 industry and 310 double-sorted portfolios.

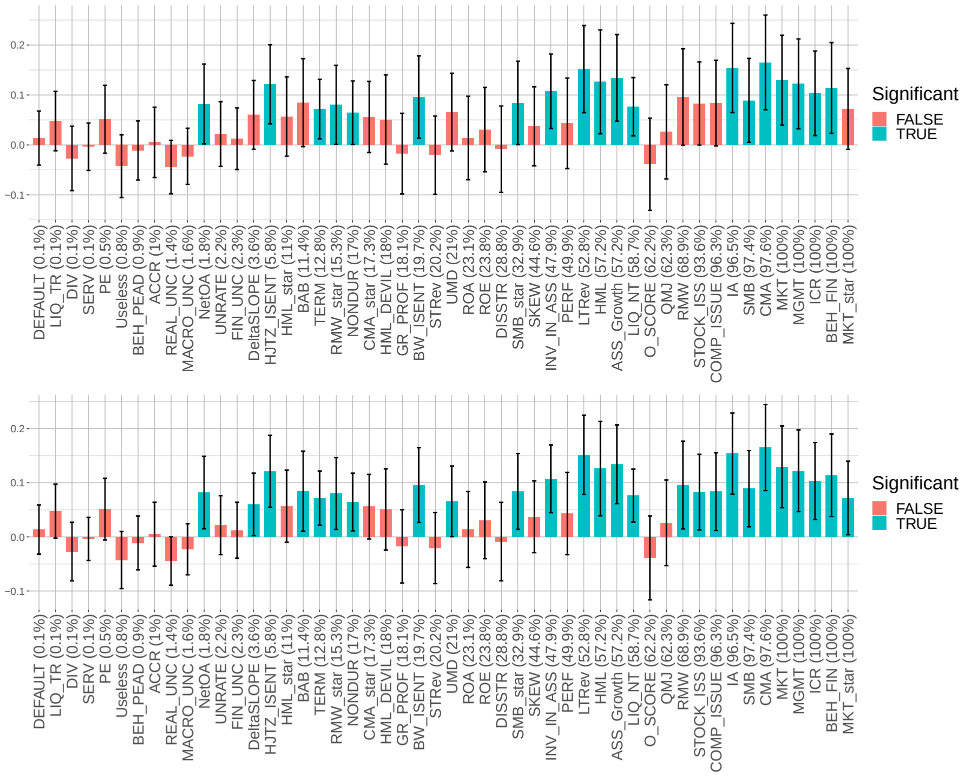

Figure 35: **Tradable factor risk premia**: Tradable factor risk premia and corresponding confidence intervals at approximate level 95% (90%) in the upper (lower) panel. Factors are ordered by their selection frequency (conditional on the factor being present in the initial model) in randomized 10-factor models always including the market, reported in parenthesis, and only factors with positive selection frequency are retained. The results are obtained with test assets comprising the 25 size/book-to-market and the 25 operating profitability/investment portfolios.

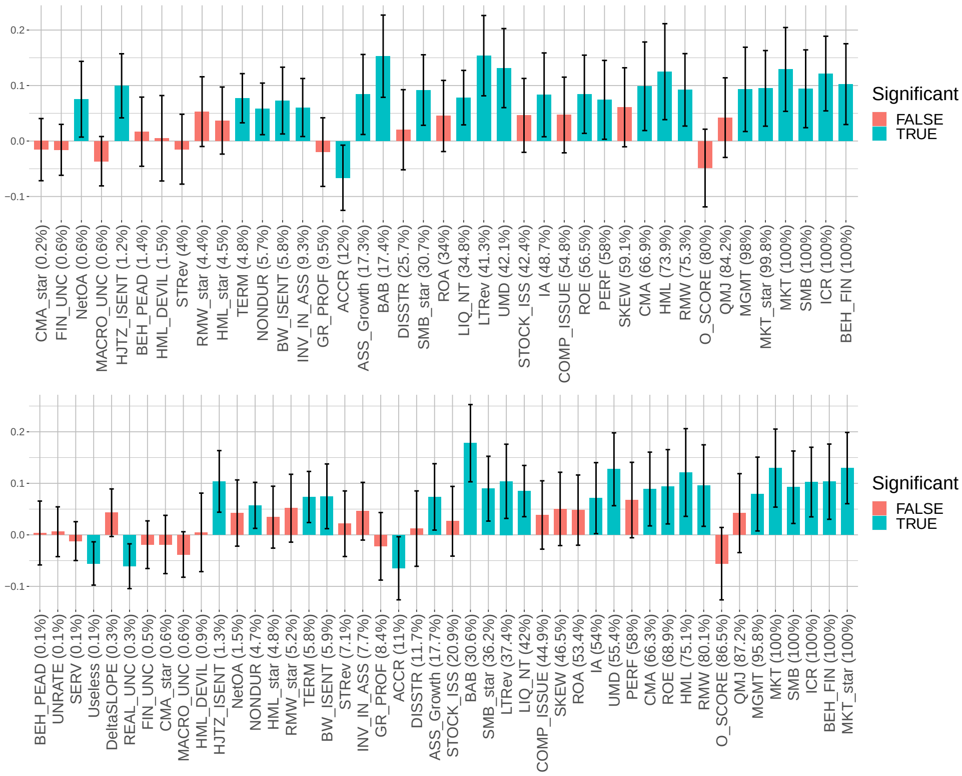

Figure 36: **Tradable factor risk premia**: Tradable factor risk premia and corresponding confidence intervals at approximate level 90%. Factors are ordered by their selection frequency (conditional on the factor being present in the initial model) in randomized 10-factor models always including the market, reported in parenthesis, and only factors with positive selection frequency are retained. The results are obtained with test assets comprising the 25 size/book-to-market and 5 PCs (upper panel) or 8 PCs (lower panel) extracted from 17 industry and 310 double-sorted portfolios.

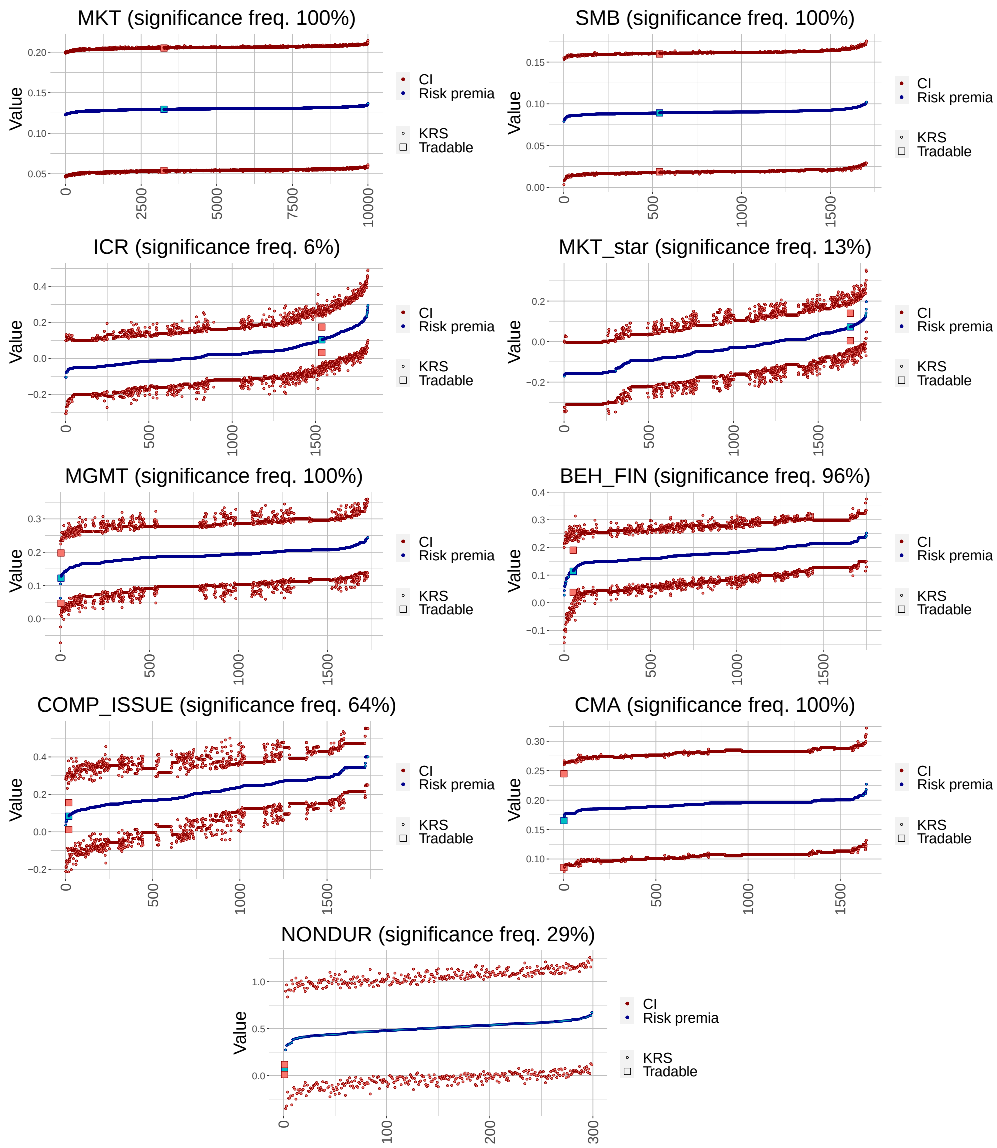

Figure 37: **Factor risk premia and** 95% **confidence intervals**: Misspecification-robust (represented by dots) and tradable (represented by squares) factor risk premia and corresponding 95% confidence intervals of a selection of factors over various random- ized 10-factor models always including the market. The factor significance frequencies (conditional on the factor being present in the initial model) are reported in parenthesis, next to the factor's label. The results are obtained with test assets comprising the 25 size/book-to-market and the 25 operating profitability/investment portfolios.

9 General rank deficiencies in benchmark asset pricing models

General rank deficiencies in asset pricing models arise when the loadings of returns on some of the factors are multicollinear. Surprisingly, this issue, which may at first seem unlikely, can manifest even in seemingly low dimensional benchmark models from the existing literature; see Kleibergen Zhang (2021).

As an illustration, consider the prominent conditional CAPM proposed by Lettau Ludvigson (2001):

$$M_{t+1} = \alpha_t + \boldsymbol{\gamma}_t' \boldsymbol{F}_{t+1}, \tag{26}$$

where the time-varying coefficients are represented as $\alpha_t = \alpha_0 + \alpha_1 cay_t$ and $\boldsymbol{\gamma}_t = \boldsymbol{\gamma}_0 + \boldsymbol{\gamma}_1 cay_t$. Here, cay_t serves as a forecasting variable for excess returns derived from consumption, asset wealth, and labor earnings. The authors incorporate market return or log-consumption growth, and potentially log-labor income (following Jagannathan and Wang (1996)), as risk factors, and estimate and test model (26) via the implied unconditional *scaled multifactor model*:

$$M_{t+1} = \alpha_0 + \alpha_1 cay_t + \boldsymbol{\gamma}_0' \boldsymbol{F}_{t+1} + \boldsymbol{\gamma}_1' \boldsymbol{F}_{t+1} cay_t = \alpha_0 + \boldsymbol{\gamma}' \tilde{\boldsymbol{F}}_{t+1}, \tag{27}$$

where $\boldsymbol{\gamma} := [\alpha_1, \boldsymbol{\gamma}_0', \boldsymbol{\gamma}_1']'$ and $\tilde{\boldsymbol{F}}_{t+1} := [cay_t, \boldsymbol{F}_{t+1}', \boldsymbol{F}_{t+1}' cay_t]'$.

Suppose, without loss of generality, that factors $\boldsymbol{F}_{t+1}$ have zero mean. If cay_t does not predict the cross-moment between test asset excess returns and the market factor, i.e.,

$$\mathbb{E}[\boldsymbol{R}_{t+1} \boldsymbol{F}_{t+1}' | cay_t] = \mathbb{E}[\boldsymbol{R}_{t+1} \boldsymbol{F}_{t+1}'],$$

then we obtain a reduced rank matrix of factor loadings:

$$\begin{aligned}
\boldsymbol{\beta} :=& \mathbb{C}\mathrm{ov}[\boldsymbol{R}_{t+1} \tilde{\boldsymbol{F}}_{t+1}'] \, \mathbb{V}\mathrm{ar}[\tilde{\boldsymbol{F}}_{t+1}]^{-1} \\
=& \mathbb{E}[\boldsymbol{R}_{t+1} \tilde{\boldsymbol{F}}_{t+1}'] \, \mathbb{V}\mathrm{ar}[\tilde{\boldsymbol{F}}_{t+1}']^{-1} \\
=& \Big[\mathbb{E}[\boldsymbol{R}_{t+1} cay_t] \quad \mathbb{E}[\boldsymbol{R}_{t+1} \boldsymbol{F}_{t+1}'] \quad \mathbb{E}[\boldsymbol{R}_{t+1} \boldsymbol{F}_{t+1}'] \mathbb{E}[cay_t] \Big] \, \mathbb{V}\mathrm{ar}[\tilde{\boldsymbol{F}}_{t+1}]^{-1} \; .
\end{aligned}$$

To verify this empirically, we collect quarterly data from 1952 Q1 to 2019 Q3 on the market gross return (proxied by the value weighted S&P 500 index), cay and test asset excess returns on the 25 double-sorted portfolios on size and book-to-market. We find that the Chen Feng 2019 (iterative

Kleibergen Paap 2006) beta rank test with Null hypothesis rank($\boldsymbol{\beta}$) $\leq$ 2 (Null hypotheses rank($\boldsymbol{\beta}$) $=$ $0, 1$ and 2) yield a p-value of 0.484 (0.0001 0.0029 0.9967), thereby indicating a reduced-rank matrix of factor loadings.

Figure 38 presents the risk premia estimates and corresponding standard errors for the Fama-Macbeth two-pass procedure, the misspecification-robust pseudo-risk premia of Kan Robotti Shanken (2013) and the tradable factor risk premia pre- and post-selection of our Oracle estimator. All methodologies estimate a similar significantly nonzero value for the risk premium on the market factor, although the Fama-Macbeth method shows a wider confidence interval. Instead, the evidence for cay and cay multiplied by the market factor varies across methodologies. Specifically, the Fama-Macbeth method yields larger risk premia estimates with confidence intervals that intersect the origin. In contrast, both the Kan-Robotti-Shanken method and the sample tradable factor risk premia estimates are smaller yet still statistically significant. Moreover, our Oracle tradable risk premia estimator sets their risk premia to zero. Since this latter estimator is the only one that is robust to weak identification, we conclude that cay and cay times the market factor are likely weak factors.

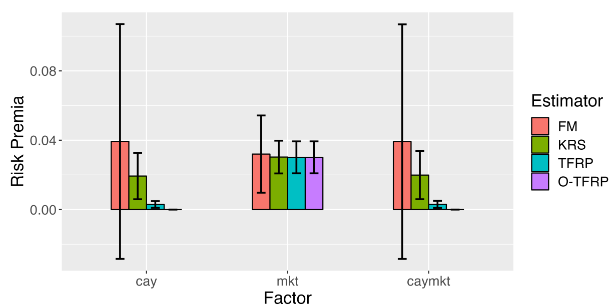

Figure 38: Risk premia and 95% confidence intervals for the scaled multifactor model of Lettau Ludvigson (2001) estimated with Fama-Machbeth (FM) two-pass regression, robust pseudo-risk premia of Kan Robotti Shanken (2013) (KRS), and tradable factor risk premia pre- (TFRP) and post-selection (O-TFRP) of our Oracle estimator.