Elephants and the Cross-Section of Expected Returns

Nora Laurinaityte[§] Christoph Meinerding[†] Christian Schlag[‡] Julian Thimme[*]

This version: January 15, 2018

Abstract

The population growth of captive Asian elephants explains the cross-section of expected returns of size-value sorted portfolios with a cross-sectional R^2 of 93% and a t-statistic of 4.0 for the market price of risk. Obviously, this factor is economically meaningless. Standard GMM cross-sectional asset pricing tests can generate such spurious explanatory power for factor models when the weight on certain moment conditions is set inappropriately. In fact, by shifting these weights, any desired level of cross-sectional fit can be attained at the price of not matching the factor means. We run placebo tests with factors that by construction do not explain the cross-section of expected returns and obtain spuriously high cross-sectional R^2's. Finally, we document some examples of factor models proposed in the literature that suffer from this problem.

Keywords: Asset pricing, cross-section of expected returns, spurious factors

JEL: G00, G12, C21, C13

[§] Research Center SAFE, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt, Germany.
E-mail: laurinaityte@safe.uni-frankfurt.de

[†] Deutsche Bundesbank, Research Centre, Wilhelm-Epstein-Straße 14, 60431 Frankfurt, Germany.
E-mail: christoph.meinerding@bundesbank.de

[‡] Goethe University Frankfurt and Research Center SAFE, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt, Germany. E-mail: schlag@finance.uni-frankfurt.de

[*] Goethe University Frankfurt, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt, Germany.
E-mail: julian.thimme@hof.uni-frankfurt.de

The authors would like to thank Campbell Harvey, Stephan Jank, Emanuel Moench, Jantje Soenksen, and seminar participants at Goethe University Frankfurt and Deutsche Bundesbank for valuable comments and suggestions. No animals were harmed during the making of this paper. This work represents the authors' personal opinions and does not necessarily reflect the views of the Deutsche Bundesbank or its staff.

1 Introduction

The amount of factor models that have been proposed to explain the cross-section of expected stock returns is enormous and the literature keeps on growing. Some attempts have been undertaken to domesticate this zoo of factors.[1] However, one factor seems to have been overlooked by researchers so far. Figure 1 shows the cross-sectional fit of a GMM estimation of a new one-factor model featuring the log growth rate of the number of captive Asian elephants as the only factor. Using a quarterly postwar sample of the usual 25 Fama-French size and book-to-market sorted portfolios as test assets, the cross-sectional R^2 is 92.7%, and the estimate for the market price of risk is statistically significantly different from zero with a t-statistic of 4.0.

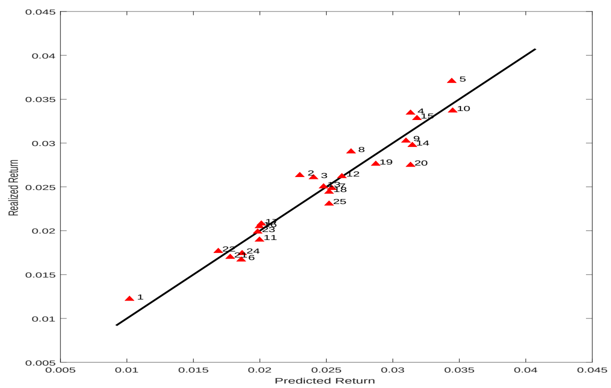

Figure 1: The figure depicts realized average excess returns of the 25 Fama-French size- and book-to-market-sorted portfolios, plotted against their model-implied expected excess returns. The latter are determined via a one-stage GMM procedure as outlined in Section 5. The explanatory factor is the log growth rate in the total number of captive Asian elephants living in zoos worldwide. The data have been downloaded from http://www.asianelephant.net/database.htm. In the GMM estimation we use the identity matrix as the weighting matrix for the moment conditions.

Does this mean that elephant population growth should be considered as the Holy Grail of empirical asset pricing? Although a creative economist might come up with a story explaining the finding in Figure 1, we regard it as economically meaningless. We argue in this paper that, in their search for the "right" factor model for the cross-section of expected returns, researchers may (unknowingly) propose models with spuriously strong explanatory power. We show that

[1]For a selection of the most recent approaches see, e.g., the papers by Harvey et al. (2016), Kogan and Tian (2014), Freyberger et al. (2017), or Harvey and Liu (2017).

the standard GMM cross-sectional asset pricing test is prone to this error when the weighting matrix is set inappropriately.

There are several ways to test models for the cross-section of asset returns. Fama and MacBeth (1973) have established the standard two-pass regression approach. Cochrane (2005) discusses how this two-stage approach can be translated into a GMM test in which both betas and market prices of risk are estimated jointly. A third way is to use the Euler equation directly as a moment condition, which is sometimes labeled the “sdf approach”. Finally, a related and also popular procedure relies on GMM as an estimation tool and uses the moment condition $0 = E\left[R_i^e - R_i^e(F - \mu_F)\lambda\right]$, which is based on the covariance between excess returns R^e and a candidate factor F. In this procedure the factor mean μ_F is typically estimated by adding a further moment condition of the form $0 = E[F - \mu_F]$. In this paper, we mainly focus on the latter approach. We show that too small a weight on the moment condition for the factor mean can lead to an imprecise estimate of μ_F in favor of an improved cross-sectional fit, i.e., a smaller pricing error.

In the elephant example above, we use the identity matrix for weighting the different moment conditions, i.e., pricing errors for the test assets and deviations from the sample mean of the factor are equally relevant in the optimization. Figure 2 shows the cross-sectional fit from an alternative estimation, in which we leave all ingredients but the weighting matrix unchanged. Now, the weight on the moment condition that identifies the factor mean is high to ensure that the estimate is close to the sample average of the factor. The resulting cross-sectional R^2 is negative and the market price of risk estimate is insignificant.

What is happening in the estimation when the weights are chosen inappropriately? To see this, assume a one-factor pricing model, where the factor is uncorrelated with all test asset returns. Assume further that the mean of the pricing factor is unknown, and is treated as a free parameter in the GMM procedure. In this setting it is easy to prove that, in the extreme case where one only targets the standard cross-sectional moment conditions in the GMM procedure and puts zero weight on the condition pinning down the factor mean, the following results hold: (i) betas with respect to the factor are fixed multiples of average excess returns, and (ii) the average pricing error of each asset is identically equal to zero, which, translated into the language of Fama-MacBeth regressions, means that (iii) the cross-sectional R^2 in the second-pass regression is 1. Stated differently, it is possible to generate misleading evidence in favor of a (spurious) factor model by allowing the mean of the pricing factor to be misestimated. These theoretical statements hold only in the knife-edge case of a zero correlation between the factor and the test asset excess returns and a zero weight on the moment condition targeting the

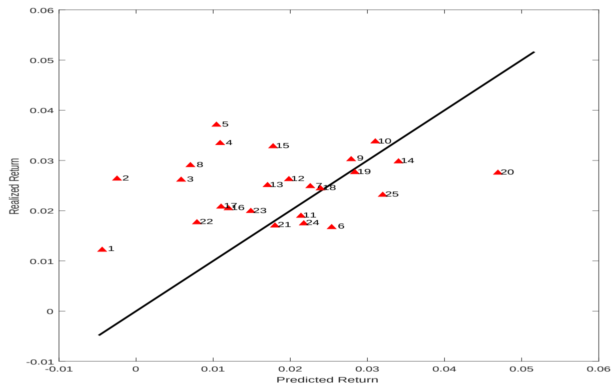

Figure 2: The figure depicts realized average excess returns of the 25 Fama-French size- and book-to-market-sorted portfolios, plotted against their model-implied expected excess returns. The latter are determined via a one-stage GMM procedure as outlined in Section 5. The explanatory factor is the log growth rate in the total number of captive Asian elephants living in zoos worldwide. The data have been downloaded from `http://www.asianelephant.net/database.htm`. In the GMM estimation the weighting matrix assigns a weight of 1 to the pricing errors and a weight of 250,000 to the moment condition $F - \mu_F = 0$ that identifies the factor mean.

factor mean. However, we show in a range of applications that spuriously high cross-sectional R^2's and statistically significant market prices of risk estimates can still be obtained even when these knife-edge assumptions are weakened, as is also documented in Figure 1.

Our paper makes an important contribution to the empirical asset pricing literature in that it points out a fundamental flaw in cross-sectional asset pricing tests that may lead researchers to think that a model has explanatory power when it actually does not. However, we want to emphasize that the goal of our paper is not to provide a complete econometric analysis of the distributional properties of Fama-MacBeth or GMM parameter estimates. In fact, the main theoretical result presented in this paper is easy to grasp, very straightforward to prove and potentially nested in existing papers on the econometrics of cross-sectional asset pricing tests.

Instead, our analysis addresses the deeply rooted economic issue of whether the asset pricing performance of certain factor models potentially comes at the cost of falling short on matching fundamentals. The moment condition that, when disregarded, drives up the explanatory power of the model is simply the one related to the mean of the pricing factor. We show that there is a trade-off between estimating the correct mean of the pricing factor and matching the cross-section of expected returns. More precisely, we document that, by shifting the weights in the GMM accordingly, any desired level of cross-sectional fit is possible if one is willing to sacrifice matching basic fundamentals. Instead of being primarily an econometric exercise, our analysis is thus more directed towards the greater issue to what extent a calibrated asset pricing model should match both financial and fundamental macroeconomic dynamics.

Our paper comprises a short theoretical motivation and a large part with applications. In the theoretical part, we prove the main result stated above: Assume that a factor is uncorrelated with all test asset excess returns. If then, in a GMM estimation, the mean of the pricing factor is chosen exogenously without imposing additional moment conditions (like, e.g., the mean being equal to the sample average), then the standard set of cross-sectional moment conditions will deliver a market price of risk estimate such that the pricing error is zero for all test assets. A notable special case is the situation when the sample mean of the factor is nonzero, but the econometrician assumes it to be zero (for instance, if the econometrician "forgets" to demean the factor, but treats it as having mean zero). We show that in this situation the misspecified GMM translates into a very simple, misspecified Fama-MacBeth regression. More precisely, we prove that, when a factor has zero explanatory power for the cross-section of expected returns, a cross-sectional R^2 of 1 is obtained if the intercept in the first stage of the Fama-MacBeth regression is omitted.

The second part of the paper deals with applications to demonstrate the actual severeness of the matter. We perform four exercises. First, we run a placebo test in a controlled environment with simulated data. We simulate i.i.d. normal test asset returns, i.e., from a model without factors. We then make up a pricing factor that is by construction uncorrelated with all test asset returns in the sample. We perform the misspecified one-stage GMM procedure, where we vary the relative weight on the moment condition that identifies the factor mean from 0 to 1. We show that any desired level of cross-sectional fit can be achieved and that the spuriously good pricing performance can even show up for relative weights close to 1.

In practice it is impossible to find a factor that has exactly zero correlation with any test asset return in sample. As a second exercise, we therefore run tests with the elephant factor described above. We choose this factor for two reasons: i) it has no economic meaning, and ii) it is almost uncorrelated with the excess returns on the 25 Fama-French size- and book-to-market portfolios (the largest absolute correlation across the 25 portfolios is 0.075). We show that the results obtained in the controlled simulation environment are almost perfectly replicated in this example with real data.

As a third exercise, we repeat the placebo simulation study, where we now simulate returns from a one-factor model and run the empirical test with a two-factor model (the true factor and an additional useless factor). We do this in order to document that the problem is not resolved by the presence of possibly true pricing factors. Varying the relative weight of the moment conditions identifying the factor means, we find again that any cross-sectional R^2 between the theoretically true R^2 (which is larger than 0 in this case) and 1 can be achieved. Moreover, again, for low relative weights on the factor mean conditions, the estimate for the market price of risk of the second (useless) pricing factor is statistically different from 0.

Finally, we apply our theoretical findings to a well-known example from the empirical consumption-based asset pricing literature. More precisely, we document that durable and non-durable consumption growth, which have been proposed as factors by Yogo (2006), do not have explanatory power for the cross-section of expected stock returns when the weight on the additional moment condition in the GMM is sufficiently large.[2] Other papers that estimate factor models as described above and are thus prone to the spuriousness documented in our paper are Dhume (2010), Darrat et al. (2011), Maio and Santa-Clara (2012), Maio (2013), Lioui and Maio (2014), Da et al. (2016), and Chen and Lu (2017). None of these papers mention the tradeoff that we discuss.

Although the issue that we raise in our paper is very general and, in theory, applies to any factor model, it is particularly relevant in practice for models related to the consumption CAPM, as, for instance, Parker and Julliard (2005) already point out in a footnote. The reason is that macroeconomic variables like consumption often tend to have little correlation with excess returns. For instance, the *cay* variable proposed by Lettau and Ludvigson (2001) is used as an example in the paper by Kleibergen and Zhan (2015) about the asymptotics of weak factors. Similarly in spirit, Savov (2011), who estimates both betas and market prices of risk jointly using the moment conditions proposed by Cochrane (2005), writes: "Fixing the betas to their OLS estimates is important because doing so does not allow the unwanted flexibility of manipulating betas to obtain a better cross-sectional fit" (p.194). Our paper concretizes this "flexibility" and shows that the trade-off between estimating "correct" betas and fitting the cross-section translates into a trade-off between cross-sectional fit and fit to macro fundamentals. Our findings thus imply that matching the cross-section of expected returns with consumption-based models remains a major challenge in asset pricing.

Our paper is linked to several strands of the literature. First of all, there are papers dealing with the econometric details behind our main theoretical result, like e.g. Kan and Zhang (1999b), Kan and Zhang (1999a), Kleibergen (2009), Kleibergen and Zhan (2015), Burnside (2016), Gospodinov et al. (2014) or Lewellen et al. (2010). Moreover, in the late 1990s a debate

[2]The paper of Yogo (2006) has recently also been criticized by Borri and Ragusa (2017).

emerged whether the classical Fama-MacBeth test or the more advanced GMM/SDF method provides more reliable estimates of market prices of risk (see Kan and Zhou (1999), Cochrane (2001), Jagannathan and Wang (2002)). However, the emphasis of our paper is not so much on the statistical properties of Fama-MacBeth or GMM estimates, but rather on their economic significance, i.e., on quantifying the underlying tradeoff between matching asset price data and fundamental dynamics. The analysis in Bryzgalova (2016) also nests the finding that treating the factor mean as a free parameter can produce spurious results, but this issue is not at the core of her discussion. Instead, she argues that the (potentially low) unconditional correlation between factors and excess returns should be taken into account and proposes an ad-hoc solution by adding a penalty term for low correlation to the GMM objective function. We make the more general argument that the spuriousness is rooted in not matching macroeconomic fundamentals properly and can be overcome by just putting more weight on the respective moment conditions. Given the choice of our examples, our paper is also linked to the large literature that tries to evaluate the performance of consumption-based asset pricing models for the cross-section of expected returns. Major advances in this literature have been made recently by, e.g., Savov (2011), Ferson et al. (2013), Boguth and Kuehn (2013), and Kroencke (2017).

2 Cross-Sectional Regressions with GMM

Assume that there are n test assets with excess returns $R^e_{i,t}$ $(i = 1, \dots, n)$ and a single candidate pricing factor F. The standard moment conditions for a cross-sectional GMM estimation of this one-factor model are the following (for $i = 1, \dots, n$):[3]

$$\begin{aligned} & E[R_i^e] = Cov[R_i^e, F]\lambda \\ \Leftrightarrow \quad & E[R_i^e] = E[R_i^e F]\lambda - E[R_i^e]E[F]\lambda \\ \Leftrightarrow \quad & 0 = E\left[R_i^e - R_i^e(F - E[F])\lambda\right]. \end{aligned} \tag{1}$$

Here λ denotes the market price of F-risk, scaled by the variance of F (i.e. $\lambda = \frac{MPR_F}{Var[F]}$) and is the parameter to be estimated. Using the sample average of the factor, $\bar{F}$, as a proxy for $E[F]$ the moment condition would become

$$0 \;=\; E\left[R_i^e - R_i^e(F - \bar{F})\lambda\right].$$

[3] See, e.g., Cochrane (2005)

Suppose that the pricing factor F is uncorrelated with all test asset excess returns, i.e. $Cov(R_i^e, F) = 0$ for all i. Then λ cannot be identified from these moment conditions alone since either Equation (1) holds for any value of λ (if all expected excess returns are zero) or it holds for no value of λ (if at least one of the expected excess returns is nonzero).

Now suppose that we change the GMM procedure slightly. The expectation of the pricing factor, $E[F]$, is generally unknown. Of course, $\bar{F} \equiv E_T[F]$ is an unbiased estimate for it.[4] Let us pretend not to have any information about $E[F]$ and replace the expectation by a new choice variable μ_F. Let us set μ_F to a value that is not exactly equal to the true mean $E[F]$. Then the moment conditions

$$0 = E\left[R_i^e - R_i^e(F - \mu_F)\lambda\right] \tag{2}$$

allow us to pin down a unique λ if $E[R_i^e] \neq 0$ for at least one i:

$$\begin{aligned} 0 &= E\left[R_i^e - R_i^e(F - \mu_F)\lambda\right] \\ \Leftrightarrow E[R_i^e] &= E[R_i^e F]\lambda - E[R_i^e]\mu_F\lambda \\ \Leftrightarrow E[R_i^e] &= E[R_i^e]E[F]\lambda - E[R_i^e]\mu_F\lambda \\ \Leftrightarrow \lambda &= (E[F] - \mu_F)^{-1} \end{aligned}$$

Here we make use of the fact that $E[R_i^e]E[F] = E[R_i^e F]$, since $Cov(R_i^e, F) = 0$ by assumption.

Using the sample equivalent of Equation (2) as a moment condition and using $\mu_F \neq \bar{F}$, i.e., assuming a factor mean that is different from the sample average, will deliver the estimate

$$\widehat{\lambda} = (\bar{F} - \mu_F)^{-1} \tag{3}$$

We have thus proven the following theorem.

Theorem 1 *Assume there are n test assets with excess returns $R_{i,t}^e$ and $E_T[R_i^e] \neq 0$ for at least one i. There is a pricing factor F that is uncorrelated in sample with all test asset excess returns, i.e. $Cov_T(R_i^e, F) = 0$ for all i. Assume μ_F is an arbitrary number different from the*

[4]Throughout the paper, we use the notation of Hansen (1982) in which a subscript T denotes the sample equivalent of a given moment.

sample average $\bar{F}$. Then the following set of sample moment conditions

$$g_T(\lambda) = E_T \begin{bmatrix} R_1^e - R_1^e(F - \mu_F)\lambda \\ \vdots \\ R_n^e - R_n^e(F - \mu_F)\lambda \end{bmatrix}$$

in a GMM estimation delivers the unique solution $\widehat{\lambda} = (\bar{F} - \mu_F)^{-1}$ which satisfies $g_T(\widehat{\lambda}) = 0$.

Note that this theorem holds irrespective of the sample size or of how the moment conditions are weighted. It also implies that all test assets are priced perfectly, so that the cross-sectional R^2 is equal to 1. Summing up, in the theoretical case that $Cov_T(R_i^e, F) = 0$ for all i, any choice of $\mu_F \neq \bar{F}$ allows to explain the cross-section of expected returns perfectly, i.e. there exists a $\widehat{\lambda}$ for which the average pricing errors are equal to zero for all test assets.

3 Fama-MacBeth regressions

To better understand Theorem 1, it is instructive to formulate its implications for Fama-MacBeth two-stage regressions. The first stage involves time series regressions of excess returns on the factor to estimate the betas for all assets. In the second stage, average excess returns are regressed on these betas to estimate the market price of risk from the cross-section of average returns.

The beta of an asset with respect to the factor F is given as

$$\beta_i^F = \frac{Cov[R_i^e, F]}{Var[F]} = \frac{E[R_i^e(F - E[F])]}{Var[F]} = \frac{E[R_i^e(F - E[F])]}{E[(F - E[F])^2]}$$

which is estimated via the sample analogues as

$$\widehat{\beta}_i^F = \frac{E_T[R_i^e(F - \bar{F})]}{E_T[(F - \bar{F})^2]}.$$

An econometrician who sets the mean of the factor to μ_F instead of estimating it from the data will obtain the following beta:

$$\widehat{\beta}_i^F = \frac{E_T[R_i^e(F - \mu_F)]}{E_T[(F - \mu_F)^2]} = \frac{Cov_T[R_i^e, F - \mu_F] + E_T[R_i^e]E_T[(F - \mu_F)]}{E_T[(F - \mu_F)^2]}.$$

Assuming again that the sample covariance $Cov_T[R_i^e, F - \mu_F]$ is equal to zero for all i, this implies

$$\widehat{\beta}_i^F = \frac{E_T[F - \mu_F]}{E_T[(F - \mu_F)^2]} \, E_T[R_i^e].$$

The beta measured by the econometrician is thus equal to the asset's average excess return, multiplied by a constant which does not depend on i.

Given this, the second-pass cross-sectional regression

$$E_T[R_i^e] \;\; = \;\; \ell_0 + \ell_1 \hat{\beta}_i^F + \varepsilon_i$$

can be rewritten as

$$E_T[R_i^e] \;\; = \;\; \ell_0 + \ell_1 \frac{E_T[F - \mu_F]}{E_T[(F - \mu_F)^2]} E_T[R_i^e] + \varepsilon_i.$$

This regression obviously delivers $\ell_0 = 0$ and $\ell_1 = \frac{E_T[(F-\mu_F)^2]}{E_T[F-\mu_F]}$ with a cross-sectional R^2 of 1, i.e. a perfect explanatory power of the (misspecified) model. The coefficient ℓ_1 is related to the GMM estimate via $\widehat{\lambda} = \frac{\ell_1}{Var_T[F]}$. Since $E_T[(F - \mu_F)^2]$ is the (incorrect) estimate of the variance of F in this context, we again obtain the result $\widehat{\lambda} = (\bar{F} - \mu_F)^{-1}$.

To sum up, we have shown the following: Suppose a pricing factor is completely unrelated to returns in the time series. If an econometrician then falsely believes that the mean of the pricing factor is $\mu_F \neq \bar{F}$ and consequently applies this incorrect mean in a standard GMM estimation or in a Fama-MacBeth two-stage regression, then the estimated betas are just multiples of the average excess returns. The cross-sectional R^2 in a second-pass Fama-MacBeth regression then automatically equals 1, which is equivalent to zero average pricing errors for all test assets in GMM.

A special case of this setup would be to set $\mu_F = 0$, while the sample mean $E_T[F]$ is nonzero. Then the falsely estimated beta would be

$$\widehat{\beta}_i^F = \frac{E_T[R_i^e(F - \mu_F)]}{E_T[(F - \mu_F)^2]} = \frac{Cov_T[R_i^e, F - \mu_F] + E_T[R_i^e]E_T[(F - \mu_F)]}{E_T[(F - \mu_F)^2]} = \frac{E_T[F]}{E_T[F^2]} \, E_T[R_i^e],$$

and this is just the coefficient from a first-pass regression of R_i^e on F without intercept.[5] Again,

[5] In a similar vein, for $\mu_F \neq 0$, the estimated beta would equal the coefficient from a first-pass regression of R_i^e on F with an exogenously fixed nonzero intercept. For brevity, we do not discuss this case here any further.

the cross-sectional R^2 from the second-pass regression would be 1. So, in the special case $\mu_F = 0$ Theorem 1 is equivalent to the following theorem:

Theorem 2 *Assume there are n test assets with excess returns $R^e_{i,t}$. There is a pricing factor F uncorrelated with all test asset excess returns, i.e., $Cov(R^e_i, F) = 0$ for all i. Assume that the first pass regression in a standard Fama-MacBeth two stage regression is performed without intercept. Then the cross-sectional R^2 in the second stage is equal to 1.*

We have shown that the misspecified GMM discussed above is equivalent to a Fama-MacBeth two-stage regression without intercept in the first stage. If the factor that is tested is uncorrelated with all test asset returns, then this procedure will generate a seemingly perfect, but only spurious explanatory power for the factor model. Note that Theorems 1 and 2 hold regardless of the *structure* of the true data-generating process. In particular, there does not even have to be a hidden factor structure in returns that we want to uncover. This is an important distinction between our paper and econometric papers like Lewellen et al. (2010) or Kleibergen and Zhan (2015). We make no assumption on the data-generating process, other than returns being uncorrelated with the factor.

In applications, the factor mean is typically not just set to a value that is different from the sample average of the factor. Instead, researchers often add the moment condition $F - \mu_F$ to estimate the factor mean. We show in the following how the weight on this moment condition, relative to the weights on the pricing error moment conditions, may have an impact on the inference.

4 Simulation study: one factor

We start the second part of the paper, which contains applications of the theory presented above, by performing a simulation exercise. We do so in order to study the impact of weighting different moment conditions in a clean environment. In particular, throughout this first "placebo test" we maintain the assumption that the candidate factor is uncorrelated with the excess returns of all test assets, even though this extreme case is certainly impossible to find in reality.

We assume the data generating process $R^e_{i,t} = \alpha_i + \sigma\varepsilon_{i,t}$ for 25 test assets, where the alphas are randomly assigned and vary between 0.3 and 2.7 percentage points per quarter. The return volatility is set to 8 percentage points quarterly for all assets and the $\varepsilon_{i,t}$ are i.i.d. normally distributed. We simulate only one sample with a sample size of 240, which corresponds to a post-war dataset with quarterly data. The realizations of the useless factor F_t are also drawn from a normal distribution with a mean of 2 percentage points and a standard deviation of 8 percentage points. To make sure that the factor is uncorrelated with all excess returns even in the simulated sample, we then make the factor orthogonal to all 25 test assets. Afterwards we scale and shift the factor so that its mean and standard deviation are exactly equal to 2 and 8 percentage points in the finite sample, in order to ease the interpretation of the following numerical results.

We use the simulated sample to estimate the market price of risk λ and the factor mean μ_F using the moment conditions

$$g_t(\lambda, \mu_F) = \begin{bmatrix} R^e_{1,t} - R^e_{1,t}(F_t - \mu_F)\lambda \\ \vdots \\ R^e_{n,t} - R^e_{n,t}(F_t - \mu_F)\lambda \\ F_t - \mu_F \end{bmatrix} \tag{4}$$

To weight the corresponding sample moment conditions g_T we use a prespecified weighting matrix of the form

$$W = \begin{pmatrix} 1 & \dots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & 1 & 0 \\ 0 & \dots & 0 & x \end{pmatrix} = \text{diag}(1, \dots, 1, x).$$

x denotes the weight of the last moment condition, which pins down the factor mean of the useless factor. In the following, we vary x between 0 and $+\infty$ and study the result of the cross-sectional regression in terms of R^2's and point estimates.[6]

Note that the GMM will deliver two parameters, λ and μ_F, whereas the theorems derived in the previous sections all assume that μ_F is chosen exogenously. The smaller the weight x, the further away from the true sample mean $\bar{F}$ will the estimate μ_F be. Given, μ_F, the point estimate for λ is then pinned down by the relationship $\widehat{\lambda} = (\bar{F} - \mu_F)^{-1}$. We constrain the estimate $\widehat{\lambda}$ to be between $-B$ and B. Without such a constraint for $\widehat{\lambda}$ the algorithm will set μ_F arbitrarily close to $\bar{F}$ and then let $\widehat{\lambda}$ explode in order to match asset pricing moments in

[6] Of course, the GMM weighting matrix has always been an issue of large debates. Cochrane (2005) discusses advantages and disadvantages of using a prespecified weighting matrix or to iteratively find the "optimal" weighting matrix (see Hansen (1982)) in an asset pricing context. Since we want to keep our analysis as simple as possible, we perform single-stage GMM estimations only.

the extreme case of zero correlation between the factor and the excess returns. We set $B = 50$, but any other number would deliver the same qualitative results.

The upper graph in Figure 3 shows the cross-sectional R^2 as a function of the relative weight of the last moment condition, which is given by $x/\text{trace}(W) = x/(x + 25)$. A relative weight of 0 corresponds to $x = 0$, while a relative weight of 1 corresponds to the limiting case $x = +\infty$. We find that for a relative weight of 0, i.e. if we assign zero weight on estimating the factor mean correctly, the cross-sectional R^2 is equal to 1. Increasing the relative weight leads to a decrease in R^2. The true R^2 of 0 is matched as the relative weight approaches 1.

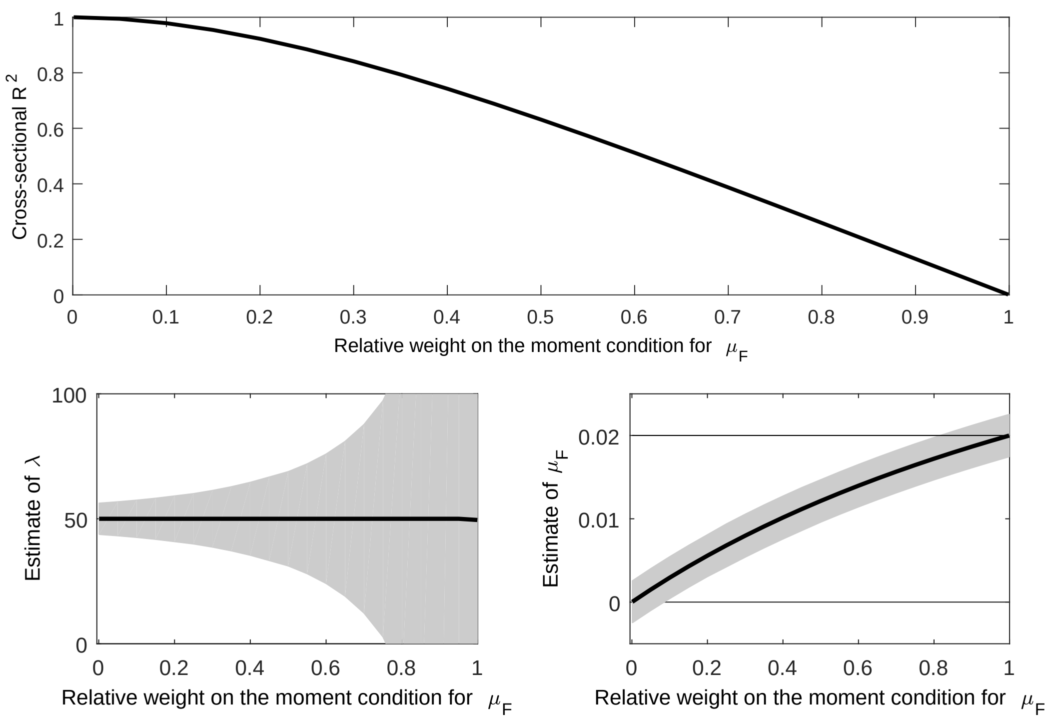

Figure 3: The figure depicts R^2 and point estimates of coefficients as functions of the relative weight on the moment condition that identifies μ_F. The relative weight is weight/(#test assets+weight). Test asset returns and the factor are simulated as outlined in Section 4.

The lower two graphs show point estimates of λ in the left and μ_F in the right plot. The point estimate of λ is equal to the upper bound B for all values of the relative weight by construction as explained above. Economically more interesting are the shaded areas, which indicate confidence bounds produced by the estimation algorithm. These are obtained by adding and subtracting 1.96 standard errors to the point estimates. It is standard in the asset pricing literature to consider such confidence bounds in order to evaluate whether an estimate is statistically significantly different from zero. In our case, we would conclude that the useless factor is useful in explaining cross-sectional variation in expected returns as long as the relative weight is below 0.77.

The lower right plot shows point estimates of μ_F. With a small relative weight, the estimated factor mean is $\bar{F} - \frac{1}{B}$, the solution of $B = (\bar{F} - \mu_F)^{-1}$, which is equal to zero in our case. The confidence band does not include the sample mean of 0.02. Putting more weight on the last moment condition brings the point estimate of μ_F closer to the sample mean $\bar{F} = 0.02$, which is marked by the thin line. With a relative weight of 1, the estimate of μ_F is equal to $\bar{F} = 0.02$ and, at the same time, the confidence band around the estimate of λ is huge, which is in line with the true structure of expected returns. λ should not be identified given that all sample covariances of returns with the factor are zero.[7]

To sum up, our analysis shows that there is a tradeoff between the two objectives of having a high explanatory power for the cross-section of expected returns and of matching fundamentals, i.e., estimating the factor mean precisely. When testing a factor that does not explain anything, one can "pump up" the R^2 at the price of not matching fundamentals. Our analysis further documents that the problem not only exists for very low weights for the moment condition targeting the factor mean. Even for relatively large weights, i.e., when the estimation of μ_F is treated as rather important, the misspecified GMM test delivers spurious evidence in favor of the factor model.

5 Real data: elephant population growth

To demonstrate that the documented spurious evidence does not only occur in a controlled environment with simulated returns, we repeat the exercise from Section 4 using actual data.

[7] Note that, for a given choice of the weighting matrix, the objective function resulting from the moment condition (4) typically has two local minima, one where $\mu_F > \bar{F}$ and one where $\mu_F < \bar{F}$. Depending on which local minimum the numerical minimization will run into, the estimated $\widehat{\lambda}$ will either be positive or negative. In the clean environment with zero correlation discussed in this subsection, these two local minima will both be global minima with identical function values. In the empirical cases discussed below, only one of these two minima is the global minimum. However, numerically, this implies that small changes in the weighting matrix can make the estimated $\widehat{\lambda}$ switch sign and turn from a large positive number to a large negative number or vice versa. In our implementation, we try to avoid this numerical problem by hand, for instance by trying out different starting values for the minimization algorithm. In particular, in the figures reported in this paper, we make sure that the $\widehat{\lambda}$'s never switch sign (in the simulation exercises) or that we indeed find the global minimum (in the empirical cases).

As test assets we use 25 size and book-to-market sorted portfolios from Ken French's webpage. To form excess returns, the 3-month Treasury bill rate is subtracted from quarterly portfolio returns. We use a quarterly sample from 1952Q2 to 2014Q2.

As a factor, we use the quarterly log growth rate in the number of captive Asian elephants that live in zoos worldwide.[8] We choose this factor because it is basically uncorrelated with the returns on our test assets, with correlation coefficients ranging from -0.006 to 0.075. The average quarterly growth rate is 0.47 percentage points with a standard deviation of 0.71 percentage points. To avoid possible relations between economic conditions in North America and the endowment of zoos in this region, we exclude all U.S. and Canadian zoos from our sample.

Figure 4 shows the same type of pictures as Figure 3. The interpretation is also similar. The cross-sectional R^2 is equal to one if we set the relative weight of the last moment condition to zero. The point estimate of μ_F is far away from the true value 0.0047. Indeed, the estimate is significantly negative. If one were to interpret this point estimate literally, one would conclude that the number of captive elephants in zoos should have been declining heavily over the past 60 years, whereas it has in fact been increasing. The estimates of λ are significantly positive.

Increasing the relative weight of the moment condition related to the factor mean lets the R^2 drop to even negative values.[9] Using a large relative weight close to 1 also leads to a point estimate of μ_F which is close to the sample average and to an insignificant estimate of λ.

Figures 1 and 2 discussed in the introduction depict the explanatory power of elephant growth exposures for the cross-section of size-value sorted portfolios graphically, for $x = 1$ (which corresponds to a relative weight of $\frac{1}{26}$) and $x = 250,000$ (which corresponds to a relative weight of 0.9999). The detailed analysis above reveals that the elephant factor could be ruled out for relative weights above 0.9. Any weight below 0.9 produces spurious evidence in favor of the elephant factor.

6 Simulation study: multiple factors

So far, our analysis has focused on one-factor models. In practice, models often feature multiple factors. Next, we document that our results carry over to multi-factor models and are also valid

[8] The data has been downloaded from http://www.asianelephant.net/database.htm. We thank the creators of this website, Jonas Livet and Torsten Jahn, for making these data publicly available.

[9] This may well happen if no constant term is included in a regression.

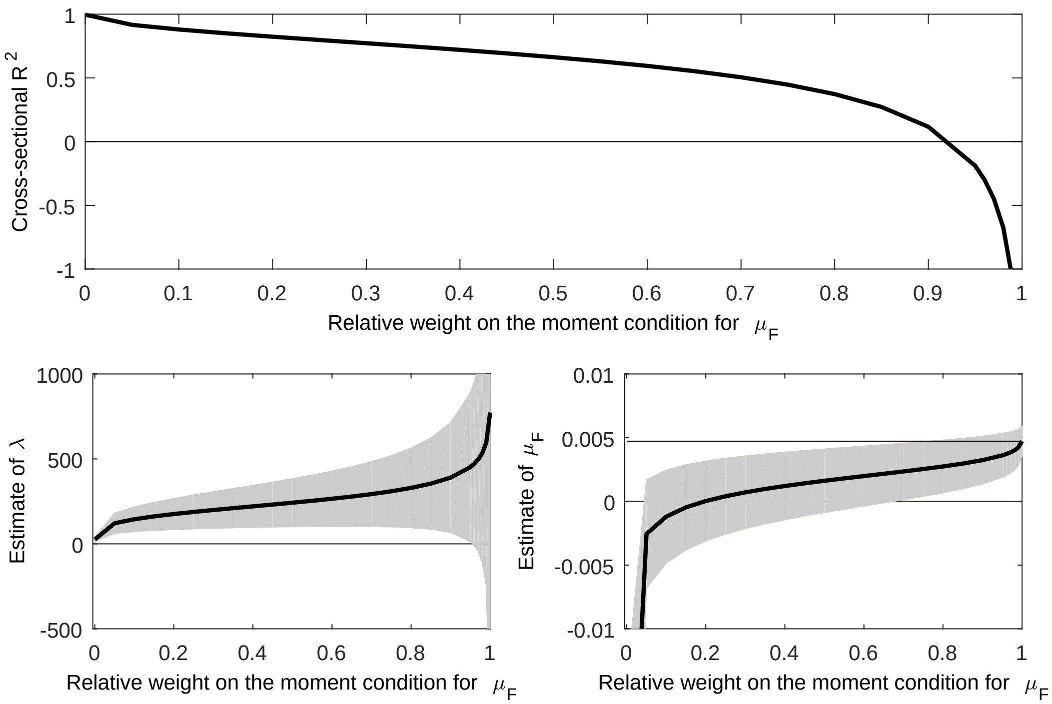

Figure 4: The figure depicts R^2 and point estimates of coefficients as functions of the relative weight on the moment condition that identifies μ_F. The relative weight is weight/(#test assets+weight). As outlined in Section 5, we use the 25 size and book-to-market sorted portfolios as test assets and the log growth rate of captive Asian elephants as explanatory factor.

in the presence of additional factors which do in fact have explanatory power. We repeat the simulation exercise from Section 4, but with different data-generating processes.

We now assume $R^e_{i,t} = \alpha_i + \beta_i^{(1)} F_{1,t} + \varepsilon_{i,t}$ for 25 test assets with different randomly assigned α_i's ranging from 0.3 to 2.7 percentage points, as in Section 4. By construction, the true beta of portfolio i ($i = 1, \ldots, 25$) with respect to the useful factor is $0.5 + \frac{i-1}{24}$, i.e. betas range between 0.5 and 1.5. The noise terms ε are i.i.d. Gaussian with mean 0 and a standard deviation of 8 percentage points quarterly. As before, we simulate one sample with 60 years of quarterly data.

We construct a second factor $F_{2,t}$ that is orthogonal to all the test portfolio excess returns and to the factor F_1 in the finite sample. Both F_1 and F_2 have a mean of 2 and a standard deviation of 4 percentage points in quarterly terms. As before, we make sure that these conditions exactly hold in the sampled time series.

We then perform a one-step GMM estimation of λ_1, λ_2, $\mu_{F,1}$, and $\mu_{F,2}$ using the moment conditions

$$g_t(\lambda_1, \lambda_2, \mu_{F,1}, \mu_{F,2}) = \begin{bmatrix} R^e_{1,t} - R^e_{1,t}(F_{1,t} - \mu_{F,1})\lambda_1 - R^e_{1,t}(F_{2,t} - \mu_{F,2})\lambda_2 \\ \vdots \\ R^e_{n,t} - R^e_{n,t}(F_{1,t} - \mu_{F,1})\lambda_1 - R^e_{n,t}(F_{2,t} - \mu_{F,2})\lambda_2 \\ F_{1,t} - \mu_{F,1} \\ F_{2,t} - \mu_{F,2} \end{bmatrix}.$$

We again assign various weights to the latter two moment conditions using a prespecified weighting matrix

$$\begin{pmatrix} 1 & \dots & 0 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & \dots & 1 & 0 & 0 \\ 0 & \dots & 0 & x & 0 \\ 0 & \dots & 0 & 0 & x \end{pmatrix} = \mathrm{diag}(1, \dots, 1, x, x)$$

We vary x from 0 to $+\infty$ and normalize x to a $[0, 1]$-scale by considering $x/(25 + x)$ as in Section 4. We produce the same pictures as above, but for the two-factor model, i.e. we plot $\widehat{\mu}_{F,1}$, $\widehat{\mu}_{F,2}$ and the two market prices of risk estimates $\widehat{\lambda}_1$ and $\widehat{\lambda}_2$ with confidence bands. Just as in Section 4, we constrain $\widehat{\lambda}_2$ to be between -50 and 50.

Figure 5 depicts the results. The upper picture again shows the cross-sectional R^2 as a function of the relative weight assigned to the moment condition targeting the factor means. For a relative weight of zero, the cross-sectional R^2 is again equal to 1, which (falsely) indicates that the two-factor model delivers a perfect cross-sectional fit. When the relative weight approaches one, the R^2 converges to the true R^2 which is, by construction, 0.385.

The lower pictures again show point estimates for the market prices of risk and factor means. The market price of risk of the (useful) first factor is correctly estimated to be significantly positive. Only for very small weights on the factor means, one may conclude that it is insignificant. At the same time, with small weights on the factors means, the market price of risk of the (useless) second factor is estimated to be significantly different from zero. This shows that the choice of inappropriate weights has the potential to bias the inference in favor of a useless factor at the expense of an actually useful factor.

Turning to the estimates of μ_F, we find that especially the estimate of the sample mean of the useless factor $\mu_{F,2}$ is biased if the estimation puts too little emphasis on the factor means.

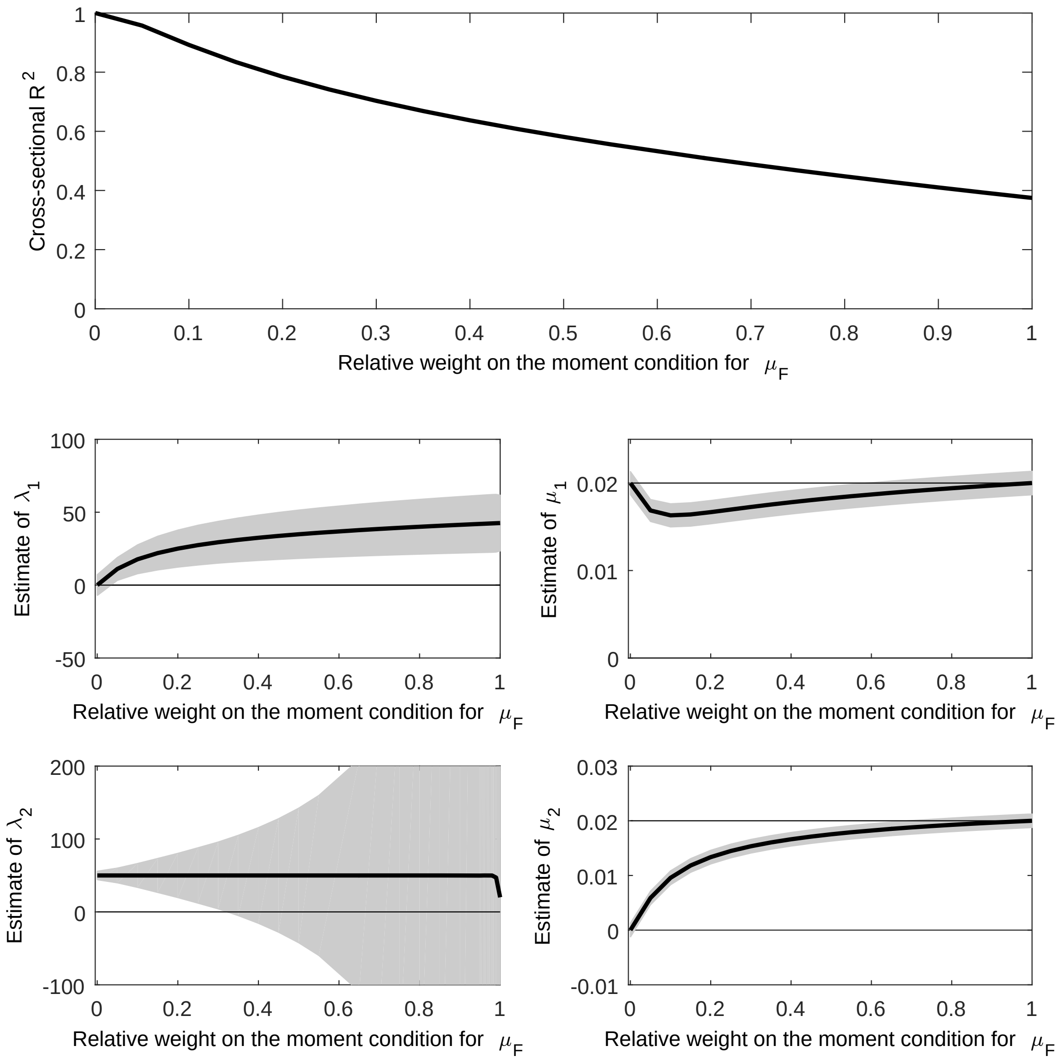

Figure 5: The figure depicts R^2 and point estimates of coefficients as functions of the relative weight on the moment conditions for $\mu_{F,1}$ and $\mu_{F,2}$. The relative weight is weight/(#test assets+weight). Test asset returns and factors are simulated as outlined in Section 6. F_1 denotes the useful factor and F_2 denotes the useless factor.

As described in Section 2, we have the relation $\widehat{\lambda}_2 = (\bar{F}_2 - \hat{\mu}_{F,2})^{-1}$. One takeaway from this is that the estimates of the factor means should always be reported in papers that perform GMM estimations of cross-sectional relations. Whenever the estimate of the mean of a tested factor is far from its sample mean and the market price of risk is equal to (or close to) the inverse of the difference between sample mean and estimated mean, the inference in favor of this factor is very likely spurious.

7 Empirical evaluation of the durable consumption model

The final section is devoted to an actual example from the asset pricing literature where success in pricing the cross-section is claimed based on the GMM presented in this paper. The example we choose is the durable consumption factor which has been established by Yogo (2006). He argues in favor of a three-factor model with durable and nondurable consumption growth and the return on the market index as factors. Although the issue that we raise in our paper is very general and, in theory, is valid for any factor model, it is particularly relevant in practice for models related to the consumption CAPM, since macroeconomic variables like consumption often tend to have little correlation with excess returns. For instance, the correlation between durable consumption growth and the excess returns of the 25 Fama-French size and book-to-market-sorted portfolios ranges from -0.139 to -0.064 on the sample described below. As a comparison, the same returns exhibit correlations with the stock market factor between 0.746 and 0.941.

We repeat the analysis from Section 4 with the same data Yogo (2006) uses in his paper.[10] The test asset returns are quarterly excess returns of the 25 size and book-to-market sorted portfolios from Kenneth French's webpage. The factors are growth rates of aggregate nondurable and services consumption and of durable consumption, both obtained from NIPA, and the excess return on the value-weighted CRSP stock market portfolio. The sample covers the period from 1951Q1 to 2001Q4. We employ the same moment conditions as outlined in the previous sections. The only difference between our analysis and the one of Yogo (2006) is the prespecified weighting matrix. To make the results from this section comparable to the ones from the previous sections and for tractability, we use the weighting matrix

$$\begin{pmatrix} 1 & \dots & 0 & 0 & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & \dots & 1 & 0 & 0 & 0 \\ 0 & \dots & 0 & x & 0 & 0 \\ 0 & \dots & 0 & 0 & x & 0 \\ 0 & \dots & 0 & 0 & 0 & x \end{pmatrix} = \mathrm{diag}(1, \dots, 1, x, x, x).$$

[10]The data can be downloaded from https://sites.google.com/site/motohiroyogo/.

Figure 6 depicts the results from this exercise. The upper plot again shows the cross-sectional R^2 as a function of the relative weight of the latter three moment conditions that identify the means of the three factors. As before, the R^2 is equal to one if the weight on the factor means is zero. The other extreme, i.e. setting the relative weight of the factor means to a very high value, thus forcing the estimated factor means to be equal to the sample averages of the factors, results in a cross-sectional R^2 of only 1.18%.

The six plots below show point estimates and confidence bands of the market prices of risk (left column of graphs) and the factor means (right column of graphs) for (from top to bottom) the growth rate of nondurable and services consumption, the growth rate of durable consumption, and the stock market index. We find that the market prices of risk of the non-durable consumption factor and of the market portfolio are at best marginally significant for all values of the relative weight. The durable consumption factor shows a pattern that is similar to those found for the useless factors in the previous sections. The estimate of the market price of risk is significantly different from zero for all but the highest values of the relative weight on the moment conditions that identify the factor means. Looking at the estimates of the factor means, we find that they are close to the sample average only when a high weight is put on the latter three moment conditions. The exception is the market return for which the factor mean estimate is always literally identical to the time series average.

Economically, we can interpret our findings as follows. The return on the CRSP index is highly correlated with the returns on the test assets in the time series. Thus, the estimated betas are meaningful. Still, the cross-sectional R^2 is close to zero for a high weight on the factor mean estimation, because, as is well-known, market betas can hardly explain any cross-sectional variation in size-value sorted portfolios (as pointed out by Fama and French (1992)). Stated differently, the variation in market betas does not explain the variation in average returns.

The two consumption factors are only weakly correlated with returns on the test assets. To decrease the pricing errors, the optimization algorithm chooses factor means that are different from the sample averages. In case of low relative weights on the factor mean condition, we find a global minimum of the objective function in which $\hat{\mu}_1 < \bar{F}_1$ and $\hat{\mu}_2 > \bar{F}_2$. This causes the estimates of the market prices of risk to be positive for the nondurable consumption factor and negative for the durable consumption factor.

We, however, also find many local minima that yield similarly low values of the objective function, but different signs for the market prices of risk. For rather low relative weights on the latter three moment conditions, using positive starting values for the market prices of risk in the minimization routine delivers a local minimum in which the market price of risk estimate

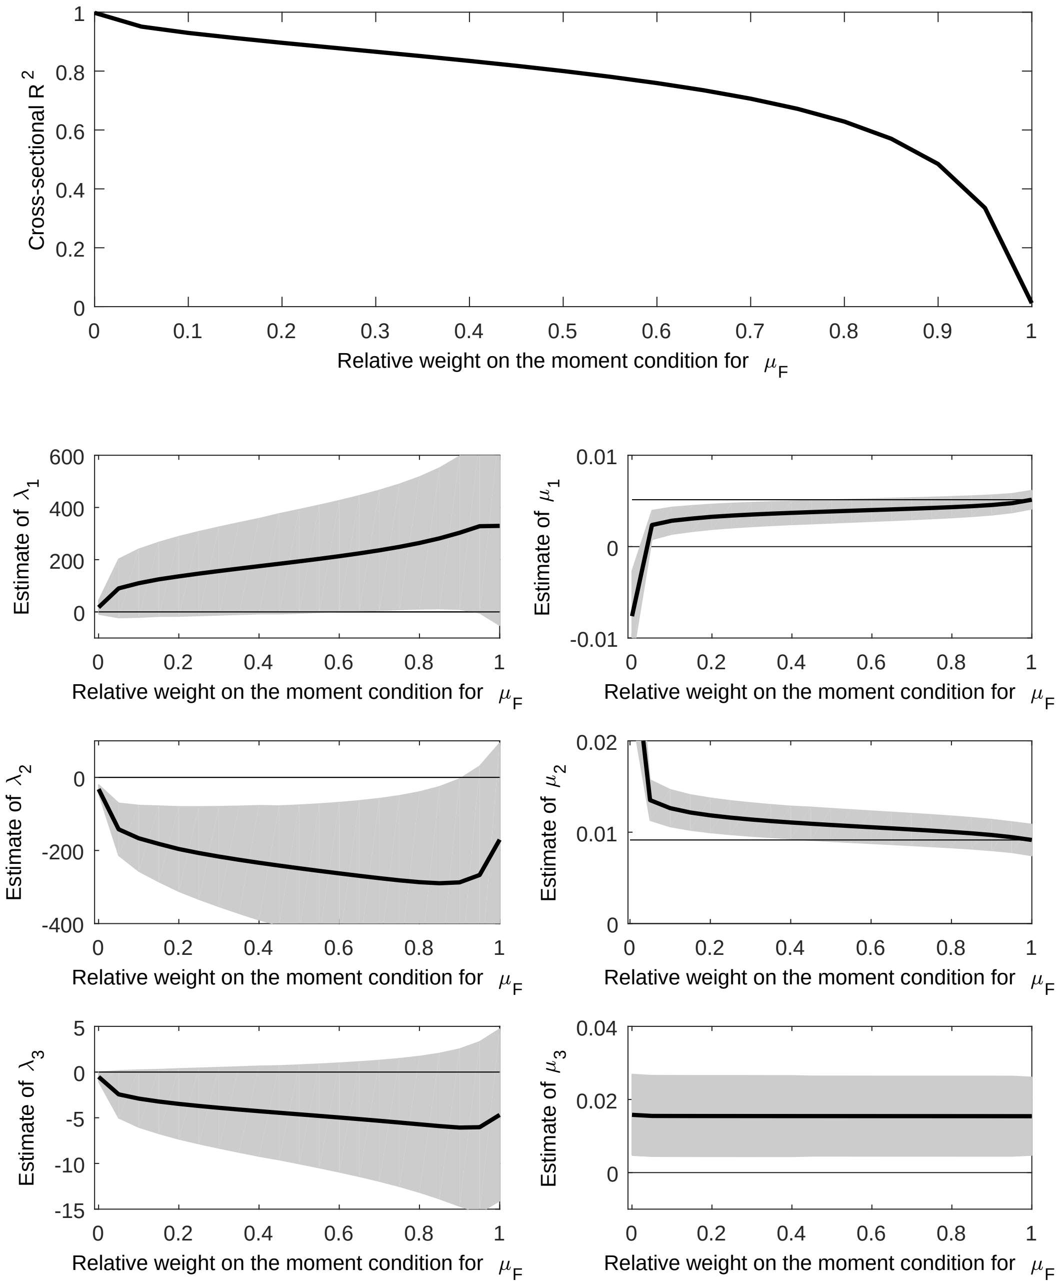

Figure 6: The figure depicts R^2 and point estimates of coefficients as functions of the relative weight on the moment conditions that identify the $\mu_{F,j}$. The relative weight is weight/(#test assets+weight). Test asset returns and factors have been downloaded from Motohiro Yogo's webpage, as outlined in Section 7. F_1 (F_2) denotes the log growth rate of aggregate nondurables and services (durables) consumption, F_3 denotes the return on the CRSP value-weighted stock return index.

of the durable factor is large and positive. In particular, using a relative weight of 0.07 gives values that are close to the ones reported in Yogo (2006): $R^2 = 93.5\%$ (compared to $R^2 = 93.5\%$ in the original paper), $\hat{\lambda}_2 = 178.52$ (170.57), $\hat{\mu}_2 = 0.0041$ (0.0028)[11], and the estimates of the other two market prices of risk are insignificant as in Yogo's paper.

It is also worth noting that $(\bar{F}_2 - \hat{\mu}_{F,2})^{-1}$, which, based on our theory, approximates the market price of risk of a useless factor, equals $(0.0091 - 0.0028)^{-1} = 158.73$ and is thus close to the actual point estimate $\widehat{\lambda} = 170.57$. Of course, our analytical result in Section 2 only applies when all sample covariances of test asset returns with the factor are exactly zero, but since durable consumption growth is only weakly correlated with the test asset returns, the deviation from the theoretical result is small.

Finally, Yogo (2006) does not use a prespecified weighting matrix, but performs a two-step GMM estimation. The weighting matrix in the first stage is already different from ours, so that we do not know the weights put on the moment conditions for the factor means. Nevertheless, comparing the estimated factor mean of the durable consumption factor with the sample average (and relating their difference to the estimate of λ) suggests that the evidence in favor of the durable consumption model is spurious.

8 Conclusion

When calibrating or estimating equilibrium asset pricing models, researchers are often confronted with the general tradeoff between matching fundamental macroeconomic dynamics and matching asset prices and returns. An ideal model should, of course, work well in both dimensions, but models often fall short on (at least) one of them. In this paper, we argue that such a tradeoff is also present in the large subfield of cross-sectional asset pricing which has brought about an overwhelming number of factor models in the past 30 years.

More precisely, standard GMM cross-sectional asset pricing tests can generate spurious explanatory power for factor models when the weight on certain moment conditions is set inappropriately and the mean of the pricing factor is allowed to be misestimated. In a range of examples we vary the weighting matrix in a single-stage GMM and show that any desired level of cross-sectional fit can be achieved, depending on how the fundamental fit (i.e., matching the factor means) is traded off against the cross-sectional fit (i.e., minimizing the pricing errors). For

[11] This value is not reported in the published paper, but can be found in the supplementary material provided on https://sites.google.com/site/motohiroyogo/.

instance, in terms of the motivating example presented in the introduction, even the population growth of captive Asian elephants can explain the cross-section of expected stock returns as long as too little weight is put on matching mean of this “factor”.

A technical conclusion from our paper is that researchers using the estimation strategy discussed in our paper should not only report the estimated market prices of risk and the cross-sectional R^2, but also the estimated factor means. An economic implication is that the quest for a consumption-based explanation of the value premium and other cross-sectional anomalies continues. Against the backdrop of the current scientific debate that is concerned with taming the “zoo of pricing factors”[12], our paper makes an important contribution to the empirical asset pricing literature in that it points out a fundamental flaw in cross-sectional asset pricing tests that may lead researchers to think that a model has explanatory power when it actually does not. Instead, statements claiming success at the cross-sectional front sometimes have to be taken with a certain dose of skepticism since they may be the result of sacrificing the fit to fundamentals.

[12]See, for instance, Harvey et al. (2016), Kogan and Tian (2014), Freyberger et al. (2017), or Harvey and Liu (2017)

References

Boguth, O. and L. Kuehn (2013): "Consumption Volatility Risk," *Journal of Finance*, 68, 2589–2615.

Borri, N. and G. Ragusa (2017): "Sensitivity, Moment Conditions and the Risk-free Rate in Yogo (2006)," *forthcoming: Critical Finance Review.*

Bryzgalova, S. (2016): "Spurious Factors in Linear Asset Pricing Models," *Working Paper.*

Burnside, C. (2016): "Identification and Inference in Linear Stochastic Discount Factor Models with Excess Returns," *Journal of Financial Econometrics*, 14, 295–330.

Chen, Z. and A. Lu (2017): "Seeing the Unobservable from the Invisible: The Role of CO2 in Measuring Consumption Risk," *forthcoming: Review of Finance.*

Cochrane, J. (2001): "A Rehabilitation of Stochastic Discount Factor Methodology," *Working Paper.*

——— (2005): *Asset Pricing*, New Jersey: Princeton University Press.

Da, Z., W. Yang, and H. Yun (2016): "Household Production and Asset Prices," *Management Science*, 62, 387–409.

Darrat, A., B. Li, and J. Park (2011): "Consumption-based CAPM models: International evidence," *Journal of Banking and Finance*, 35, 2148–2157.

Dhume, D. (2010): "Using Durable Consumption Risk to Explain Commodities Returns," *Working Paper.*

Fama, E. and K. French (1992): "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47, 427–465.

Fama, E. and J. MacBeth (1973): "Risk, return, and equilibrium: Empirical tests," *Journal of Political Economy*, 81, 607–636.

Ferson, W., S. Nallareddy, and B. Xie (2013): "The "out-of-sample" performance of long-run risk models," *Journal of Financial Economics*, 107, 537–556.

Freyberger, J., A. Neuhierl, and M. Weber (2017): "Dissecting Characteristics Nonparametrically," *Working Paper.*

Gospodinov, P., R. Kan, and C. Robotti (2014): "Misspecification-Robust Inference in Linear Asset Pricing Models with Irrelevant Risk Factors," *Review of Financial Studies*, 27, 2139–2170.

Hansen, L. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1084.

Harvey, C. and Y. Liu (2017): "Lucky Factors," *Working Paper.*

HARVEY, C., Y. LIU, AND H. ZHOU (2016): "... and the Cross-Section of Expected Returns," *Review of Financial Studies*, 29, 5–68.

JAGANNATHAN, R. AND Z. WANG (2002): "Empirical Evaluation of Asset Pricing Models: A Comparison of SDF and Beta Methods," *Journal of Finance*, 57, 2337–2367.

KAN, R. AND C. ZHANG (1999a): "GMM Tests of Stochastic Discount Factor Models with Useless Factors," *Journal of Financial Economics*, 54, 103–127.

——— (1999b): "Two-Pass Tests of Asset Pricing Models with Useless Factors," *Journal of Finance*, 54, 203–235.

KAN, R. AND G. ZHOU (1999): "A Critique of the Stochastic Discount Factor Methodology," *Journal of Finance*, 54, 1221–1248.

KLEIBERGEN, F. (2009): "Tests of Risk Premia in Linear Factor Models," *Journal of Econometrics*, 149, 149–173.

KLEIBERGEN, F. AND Z. ZHAN (2015): "Unexplained Factors and their Effects on Second-Pass R-squared's," *Journal of Econometrics*, 189, 101–116.

KOGAN, L. AND M. TIAN (2014): "Firm Characteristics and Empirical Factor Models," *Working Paper*.

KROENCKE, T. (2017): "Asset Pricing without Garbage," *Journal of Finance*, 72, 47–98.

LETTAU, M. AND S. LUDVIGSON (2001): "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia are Time-Varying," *Journal of Political Economy*, 109, 1238–1287.

LEWELLEN, J., S. NAGEL, AND J. SHANKEN (2010): "A Skeptical Appraisal of Asset Pricing Tests," *Journal of Financial Economics*, 96, 175–194.

LIOUI, A. AND P. MAIO (2014): "Interest Rate Risk and the Cross-Section of Stock Returns," *Journal of Financial and Quantitative Analysis*, 49, 483–511.

MAIO, P. (2013): "Intertemporal CAPM with Conditioning Variables," *Management Science*, 59, 122–141.

MAIO, P. AND P. SANTA-CLARA (2012): "Multifactor models and their consistency with the ICAPM," *Journal of Financial Economics*, 106, 586–613.

PARKER, J. AND C. JULLIARD (2005): "Consumption Risk and the Cross Section of Expected Returns," *Journal of Political Economy*, 113, 185–222.

SAVOV, A. (2011): "Asset Pricing with Garbage," *Journal of Finance*, 66, 177–201.

YOGO, M. (2006): "A Consumption-Based Explanation of Expected Stock Returns," *Journal of Finance*, 61, 539–580.