The Multifactor Risk-Return Tradeoff ☆

Fahiz Baba-Yara* Martijn Boons† Rik Frehen‡
Indiana University Nova SBE Tilburg University

Abstract

Previous work finds mixed evidence for the market risk-return tradeoff (RRT). We instead study the multifactor RRT and show it is strongly positive when accounting for factor covariances. Our multifactor risk model (i) reveals that covariances drive more variation in factor returns than variances and (ii) predicts multiple factors at least as well as benchmark models designed for single-factor prediction, both in- and out-of-sample. Consistent with a positive RRT, conditional multifactor alphas for many anomalies are similar to unconditional alphas. These results are largely overlooked in recent work on factor predictability and volatility-timing and matter for practitioners using conditional moments in asset allocation.

Keywords: Multifactor Models, Unconditional vs. Conditional Alpha, Factor Return Predictability, Variance and Covariance Risk, Principal Component Factors

*Indiana University, Kelley School of Business, 1309 E 10th St, Bloomington Indiana, USA. E-mail: `fababa@iu.edu` Web: `www.babayara.com`

†Corresponding author. Nova School of Business and Economics, Rua Holanda 1, 2775-405 Carcavelos, Portugal. E-mail: `martijn.boons@novasbe.pt` Web: `https://sites.google.com/site/martijnboons/`

‡Tilburg University. E-mail: `r.g.p.frehen@tilburguniversity.edu` Web:

☆ We are grateful to Victor DeMiguel, Elise Gourier (discussant), Amit Goyal, Paul Huebner, Andreas Neuhierl (discussant), Tobias Sichert (discussant), Rossen Valkanov, seminar participants at Nova School of Business and Economics as well as participants at the Safe Asset Pricing Workshop, Sorbonne Alliance Asset Pricing Workshop, and the Wabash River Finance Conference.

Relative to a vast literature on the market risk-return tradeoff, this tradeoff is severely understudied in multifactor models. Do long-short factor risk premia increase in factor variances and what is the role of factor covariances? Answering these questions is important because real-world portfolios load strongly on multiple factors (think of thematic mutual funds and ETFs or portfolio tilts more generally) and many portfolio management decisions therefore explicitly rely on conditional multifactor moments.

A related question is whether conditional multifactor models can explain long-short characteristic-sorted portfolio returns or anomalies. The difference between an asset's unconditional and conditional alpha is due to comovement between the asset's betas and multifactor risk premia and risk. Lewellen and Nagel (2006) find that the conditional CAPM fails to explain anomalies just like the unconditional CAPM.[1] Given that multifactor betas jointly vary a lot more over time than market betas and multifactor models explain anomalies better than the CAPM unconditionally (see, e.g., Fama and French, 1996, 2016; Hou et al., 2015), one might expect a multifactor model to perform even better conditionally.[2] However, we show that that it does not: conditional multifactor alphas are largely indistinguishable from unconditional alphas. In this paper, we flesh out the singular reason why: a strong and positive multifactor risk-return tradeoff that starkly contrasts previous work that has failed

[1] Using state-of-the-art methods for estimating conditional market risk and studying a broad global sample and a larger set of characteristic-sorted portfolios, Gormsen and Jensen (2024) find that the performance of the conditional CAPM improves slightly. Their conditional market risk factor explains on average 18% of unconditional CAPM alphas.

[2] Many other efforts have been undertaken recently to "tame the factor zoo." See, for instance, Harvey et al. (2016); McLean and Pontiff (2016); Green et al. (2017); Hou et al. (2015, 2020); Jensen et al. (2023); van Binsbergen et al. (2022).

to agree on the sign of the market risk-return tradeoff.[3]

We start by studying monthly long-short decile portfolios sorted on the 153 characteristics collected in Jensen et al. (2023) from 1970 to 2023. While we initially focus on the Fama and French (2015, FF5) five-factor model, which is a staple in both academic and practitioner literature on multifactor models, we also analyze RP-PCA factors extracted from the large set of characteristic-sorted portfolios following Lettau and Pelger (2020a,b). We document a new fact: the difference between unconditional and conditional FF5 alphas, $\alpha^u - \alpha^c$, is generally small, just like in the CAPM. In a regression of α^c on α^u, the coefficient (intercept) is indistinguishable from one (zero). Here, α^u is the intercept from the usual full sample regression and $\alpha^c = \frac{1}{T}\sum_t (R_{t+1} - \widehat{\beta}_t' F_{t+1})$, where the conditional beta of the portfolio is aggregated up from rolling stock-level exposures.

To understand this result, we derive unconditional alphas under the null of a conditional multifactor model.[4] We show that $\alpha^u - \alpha^c$ is due to comovement between an asset's betas and (i) factor risk premia (factor timing) and (ii) factor variances or covariances (factor (co-)variance timing). This leaves three explanations for the fact that $\alpha^u - \alpha^c$ is small in the multifactor setting. First, betas may not vary much over time. While the (lack of) time-variation in market betas is key in Lewellen and Nagel (2006), we find that multifactor betas jointly vary over three times as much. Second, time-variation in multifactor betas may be uncorrelated to factor risk

[3] This disagreement arises from using different proxies for conditional market variance. See, for instance, French et al. (1987); Campbell (1987); Nelson (1991); Campbell and Hentschel (1992); Glosten et al. (1993); Harvey (2001); Ghysels et al. (2005); Moreira and Muir (2017); DeMiguel et al. (2024).

[4] Our derivation closely follows Jagannathan and Wang (1996); Lewellen and Nagel (2006); Boguth et al. (2011), who study a single factor setting.

premia or risk, either because betas are noisy or because the process driving risk exposures is disconnected from the factors' risk-return tradeoff. Because we estimate large factor timing and factor (co-)variance timing terms for many characteristics (the standard deviation across characteristics of their contribution to α^u is about 30 bps per month), this explanation is similarly implausible. This leaves only one explanation, which is that the factor timing and factor (co-)variance timing terms largely cancel each other out. This explanation implies that multifactor risk premia and risk are positively related. In the remainder of the paper, we impose the structure of linear multifactor models to formalize this risk-return tradeoff and estimate it, both in- and out-of-sample.

The main challenge in estimating the multifactor risk-return tradeoff is that the model implies a pooled specification:

$$E_t(F_{t+1}) = Var_t(F_{t+1})\Gamma, \tag{1}$$

where $E_t(F_{t+1})$ are conditional risk premia, $Var_t(F_{t+1})$ is the conditional variance-covariance matrix of the factors, and Γ is a vector of risk-prices. Thus, covariance of a factor with any of the other factors in the model contributes to time-variation in both of these factors' risk premia and the respective contributions are determined by the risk-prices of the two factors.[5] Under the null of the multifactor model, the usual regression in the risk-return tradeoff literature, that is, of factor return on

[5]We are not arguing that other predictors, perhaps motivated by specific economic theories, could not find additional sources of predictability. In fact, any variable that predicts future factor risk and return with the same sign contributes to a positive risk-return tradeoff and thus helps to understand why $\alpha^u - \alpha^c$ is generally small. In turn, any variable that predicts return and risk with opposite sign would make it more surprising that $\alpha^u - \alpha^c$ is small.

conditional factor variance, suffers from omitted variables bias. Empirically, we find that this bias is large, both for the market and long-short factors.

Throughout, we follow extant literature and proxy for conditional factor risk, $Var_t(F_{t+1})$, using realized factor variances and covariances estimated over a six month window of daily returns. We confirm that this proxy strongly predicts future risk. Moreover, we find that covariances are more predictable than variances, which is important because each covariance contributes to two risk premia in the multifactor setting. A positive risk-return tradeoff then implies that our risk-proxies also predict factor risk premia with a positive sign (i.e., the risk-prices in the vector $\Gamma > 0$). That is indeed what we find in our pooled regressions.

We estimate a market price of risk of about two to three (depending on the specification, but always with $t > 3$), which is consistent with the usual interpretation of this parameter as measuring the representative agent's risk aversion. Of the other four risk-prices, three are large and positive (i.e., SMB at 4.9, t =2.2; RMW at 7.7, t=2.5; and, CMA at 5.6, t = 1.3), while one is negative though small and insignificant (HML at -2.4, $t = -1.0$). Our pooled estimation reveals at least two other interesting facts. First, partial R^2s range from 0.7% for SMB to 2.1% for RMW. Thus, our model captures sizeable variation in expected returns for all factors even at the monthly frequency. Second, even though they are largely ignored in previous work, covariances have a large contribution to expected return variation. For the market, covariances contribute about half of the total variation in expected return. For SMB, RMW, and CMA, covariances contribute more than variances (with contributions ranging from 65% to 113%). We show that our risk-price estimates are largely insensitive to

(i) the window over which we estimate factor risk, (ii) eliminating outliers, and (iii) the inclusion of additional test assets, such as industry portfolios (as advocated in, e.g., Lewellen et al., 2010). More importantly, we show that our conclusions extend beyond the FF5 model. We also find a robust and positive multifactor risk-return tradeoff for a variety of RP-PCA factors. In fact, we find that the risk-return tradeoff is stronger for RP-PCA factors, consistent with the idea that these factors better capture the true variation in priced risks and therefore risk premia than FF5 (see also Lettau and Pelger, 2020b).

Although the parameters in our pooled regression for FF5 are more unstable over shorter historical windows, we find that this instability is easily managed in real-time using off-the-shelf methods. Both a forecast combination approach (see, e.g., Rapach et al., 2010) and a Ridge penalty yield positive out-of-sample R^2s of 0.50% or more at the monthly horizon for all factors except SMB. The out-of-sample R^2 for an optimal combination of the factors is even larger at over 1.1%. Given that we are using the same model and parameters to predict the returns of all factors, the performance of our model is remarkable. For instance, our out-of-sample R^2 for the market is similar to recent work by Kelly and Pruitt (2013) and Martin (2017). These authors analyze models that are specifically designed to predict market returns, but we require our model to predict long-short factor returns as well. While we show that investors can increase portfolio returns by trading on information in realized factor risks, we show that such trading does not increase (mean-variance) utility. We conclude that the multifactor risk-return tradeoff is not only strong and positive, but risk is compensated in return in such a way that reasonably risk-averse investors cannot benefit from trading on it.

Our results relate to at least four strands of the literature.

The risk-return tradeoff: We are not the first to argue that the simple regression of market returns on conditional variance suffers from omitted variables bias (see, e.g., Scruggs, 1998; Guo and Whitelaw, 2006; Bali and Engle, 2010; Rossi and Timmermann, 2015). These papers study the Intertemporal CAPM of Merton (1973) and find a positive market price of risk when controlling for conditional covariance with a variety of state variables. We differ from these papers by accounting for the market's conditional covariance with tradable factors, rather then imposing these covariances to be functions of variables that are known to predict returns, like bond yields or the price-dividend ratio. We show that omitting the market's covariance with tradable factors also leads to a large downward bias in the market price of risk. Our evidence of a positive risk-return tradeoff for both the market and other factors is much stronger than that in Barroso and Maio (2024), because the multifactor specification studied by these authors ignores factor covariances (see Gormsen and Jensen, 2024, for a similar specification). We conclude that covariances are key in the identification of the multifactor risk-return tradeoff.

Volatility timing: Moreira and Muir (2017) show that factor returns in month $t+1$ tend to load negatively on realized variance in month t, which gives rise to single factor volatility-timing strategies that generate positive unconditional alpha. DeMiguel et al. (2024) study a multifactor volatility-timing strategy that conditions on one-month realized market volatility.[6,7] We differ by focusing on factor risk estimated over longer windows and by appropriately accounting for factor covariances. We show that our "low frequency" risk measures strongly predict future risk, even controlling for one month risk. This result implies that our measures are relevant instruments for the risk-return tradeoff. One month risk measures are relevant instruments as well, and we show in robustness checks that these also generate a multifactor risk-return tradeoff that is, if anything, positive.

Our results have two further implications. First, some of the seemingly attractive return of volatility-managed strategies may simply be due to priced covariance risks. For instance, we estimate that about half of the CAPM alpha of a volatility-managed market portfolio is due to covariance risks. Second, DeMiguel et al. (2024) argue that prices of risk for many factors are negatively correlated to realized market volatility. If risk-prices are negatively correlated to amounts of risk, this would mute the risk-return tradeoff over time. However, we find no evidence that the FF5 or RP-PCA risk-prices vary over time in this way. If anything, our pooled risk-price estimates are increasing in market volatility, such that the risk-return tradeoff is stronger in times of market stress.

Return predictability: A large literature studies time-series predictability of aggregate stock market returns. A number of recent papers have also analyzed the predictability of long-short factors. Among others, Cohen et al. (2003); Asness et al.

[6] Cederburg et al. (2020) argue that these strategies fail out-of-sample, while Barroso and Detzel (2021) argue they do not survive transaction costs. The multifactor strategy of DeMiguel et al. (2024) survives both, however.

[7] Note that this choice of conditioning variable falls outside the general specification of the multifactor model in Eq. (1). The model suggests that it is covariance between the market and other factors that predicts factor returns, not market variance or volatility.

(2017); Baba Yara et al. (2021) use the value spread to predict value returns in different markets. Haddad et al. (2020) extract principal components from the valuation ratios of a large set of portfolios. Other predictor variables include: the issuer-purchaser spread (Greenwood and Hanson, 2012), sentiment (Baker and Wurgler, 2006; Stambaugh et al., 2012), and flows (Akbas et al., 2015). Many of these predictors are either common to all assets (such as sentiment) or asset-specific (such as a valuation ratio). Our model based on arguably the most fundamental economic concept – the risk-return tradeoff – is common to all factors, but the predictors are specific to each factor. Our model accounts for the important role of factor correlations, which is often ignored in previous literature (Haddad et al., 2020, is an exception). While controlling for conditional market variance is common whenever a new predictor of market returns is proposed in the literature, our results suggest that future work on predictability should control for factor risk more broadly.

Interpreting factors: It is largely irrelevant for our paper whether the FF5 or RP-PCA factors capture economic risks that households desire to hedge or whether they reflect common components of sentiment-driven demand. In the latter case, fully rational risk averse arbitrageurs will still require compensation to trade aggressively against this demand because it exposes them to systematic factor risk (Kozak et al., 2018). We show that the expected return compensation households or arbitrageurs require is time-varying and commensurate to factor risk. This conditional implication is common to all empirical linear multifactor models and a moment that theoretical models on the origins of cross-sectional factors may want to match.[8] Further, our

[8]For instance, Bordalo et al. (2024a) and Bordalo et al. (2024b) (see, also, Bordalo et al., 2019) broadly argue that overreaction to earnings news and subsequent surprises (rather than expected

results contribute to the evidence in Daniel et al. (2020) and Herskovic et al. (2019). These papers show that exposure to factors can be hedged at virtually no cost on average. We instead argue that hedging factor risk is costly, because this risk is priced conditionally. However, the fact that factor returns are commensurate to factor risk effectively caps the factors' Sharpe ratios and therefore also the cost of hedging.

1. Data and motivation

In this section, we compare unconditional and conditional alphas for the 153 characteristics collected in Jensen et al. (2023).[9] All characteristics are signed such that the unconditional CAPM alpha of a value-weighted long-short decile portfolio is positive. The multifactor model we consider here is the Fama and French (2015) five-factor model (FF5), because it features prominently in academic and practitioner literature on factors.

The unconditional alpha is commonly estimated by regressing long-short portfolio returns on a set of factors using monthly observations over a long historical sample:

$$R_{L-S,t+1} = \alpha^u + \beta' F_{t+1} + \epsilon_{t+1}. \tag{2}$$

compensation for risk) explain a large share of average factor returns. As such, these papers are silent on the risk-return tradeoff. If high factor risk coincides with low signal-to-noise ratios, standard filtering logic would suggest that agents shrink more in times of high risk. Thus, a challenge for these models could be to generate more overreaction in times of high factor risk.

[9] See their Appendix J or `https://jkpfactors.com/stock-char` for details.

The conditional alpha is commonly defined as:

$$\alpha^c = \frac{1}{T}\sum_t (R_{L-S,t+1} - \widehat{\beta}_t' F_{t+1}), \tag{3}$$

where $\widehat{\beta}_t'$ is a conditional exposure. We follow a large body of work and estimate conditional betas over a one-year window of daily returns that ends at time t. We estimate these betas at the stock-level and then aggregate them up to the portfolio-level using time t value-weights. As noted in Boguth et al. (2011), these "lagged component betas" are preferable over "lagged portfolio betas," because the portfolio composition varies over time and strongly so for some characteristics.[10]

We first present the main stylized facts about unconditional versus conditional multifactor alphas. We then link these alphas to the multifactor risk-return tradeoff that is the ultimate focus of our paper.

1.1 Unconditional versus conditional alphas: Evidence

To benchmark our results for the FF5, we start our analysis with the CAPM and present a scatter plot of conditional versus unconditional alphas for all long-short characteristic-sorted portfolios in Panel A of Figure 1. First, 79.1% of the unconditional alphas are significant at the 10% level ($t > 1.65$). As noted already in Jensen et al. (2023): a large share of anomalies is replicable. More importantly, we find

[10] We show that our main conclusion on the alpha difference is robust to estimating betas differently in Figures IA.2 - IA.4. Here, we consider lagged component betas estimated over either a six-month window of daily returns or a five-year window of monthly returns. These alternative betas provide a slightly different tradeoff between noise and timeliness. We also consider realized portfolio betas estimated using daily portfolio returns in month $t + 1$, as in Lewellen and Nagel (2006). These realized betas suffer from an overconditioning bias, however (Boguth et al., 2011).

that conditional CAPM alphas are similar: 79.7% is significant and the correlation between the two alphas is 96.0%. The fact that the anomalies scatter closely around the 45-degree line confirms that the conditional CAPM does not explain anomalies better than the unconditional model, as shown already in Lewellen and Nagel (2006) for portfolios sorted on Size, Book-to-Market and Momentum.

In Panel B, we present the evidence for FF5. Perhaps unsurprisingly, we find that the number of significant unconditional alphas is lower at 60%. More surprisingly, we again find similar alphas in the conditional specification. 64% of the conditional alphas is significant and the correlation between conditional and unconditional alphas is large at 90%. Note that the correlation in Panel B is lower by construction. We have to estimate firms' exposures to a larger set of factors, which increases noise. Even so, unconditional and conditional alphas are very similar in the multifactor model, just like they are in the CAPM.

1.2 Unconditional versus conditional alphas: Theory

It is well understood theoretically what drives the difference between conditional and unconditional alphas in a single factor model (see, e.g., Jagannathan and Wang, 1996; Lewellen and Nagel, 2006; Boguth et al., 2011). We extend the analysis in these papers to the multifactor setting. Like them, we start from the null of a conditional model specified as follows:

$$E_t(R_{t+1}) = \beta_t' E_t(F_{t+1}). \tag{4}$$

Under this null, α^u is the bias from estimating the model unconditionally, while it holds conditionally. Alternatively, we could include in Equation (4) a conditional alpha α_t. If α_t is uncorrelated to factor risk premia, $E_t(F_{t+1})$, the alpha bias derived below is defined analogously as $\alpha^u - E(\alpha_t)$.

In Appendix A, we derive the following expression:

$$\text{Alpha bias} = \underbrace{E(\beta_t' F_{t+1})\left[1 + \mu_F' \Sigma_F^{-1} \mu_F\right]}_{\text{Factor Timing (I)}} - \underbrace{E(\beta_t' F_{t+1} F_{t+1}') \Sigma_F^{-1} \mu_F}_{\text{Factor (Co-) Variance Timing (I)}}. \tag{5}$$

The first component is due to factor timing, i.e., comovement between betas and factor returns. The second component is due to factor (co-) variance timing, i.e., comovement between betas and squared factor returns or cross-products of factor returns.[11]

We can rewrite this decomposition into slightly more familiar terms involving conditional expectations of factor risk premia and risk:

$$\text{Alpha bias} = \underbrace{E(\beta_t' \mu_{F,t})\left[1 + \mu_F' \Sigma_F^{-1} \mu_F\right]}_{\text{Factor Timing (II)}} - \underbrace{E\left(\beta_t' \left[\mu_{F,t} \mu_{F,t}'\right]\right) \Sigma_F^{-1} \mu_F}_{\text{Small}} - \underbrace{E\left(\beta_t' \Sigma_{F,t}\right) \Sigma_F^{-1} \mu_F}_{\text{Factor (Co-)Variance Timing (II)}}. \tag{6}$$

While in this expression the alpha bias consists of three terms, the second term is small. Therefore, we simply refer to the first and third terms as factor timing and factor (co-)variance timing, respectively.

Let us now discuss the intuition that applies equally to the representations in

[11]Note that we broadly refer to the terms in the expectations as capturing "comovement." While the unconditional factor risk premia μ_F also contribute to the factor timing term, their contribution to the factor (co-)variance timing term is tiny (because $\mu_F' \mu_F$ is small compared to (cross-) products of demeaned factor returns). In other words, if the factor (co-)variance timing term is large in the data, this can only be due to covariance between betas and (cross-) products of factor returns and that is why we emphasize comovement in the following.

Eqs. (5) and (6). The factor timing terms are large when the asset's betas covary positively with factor risk premia. The intuition is that the unconditional model ignores the covariance between betas and factor risk premia, which thus leaves some of the asset's average return unexplained. Factor timing has a positive impact on alpha bias. To understand the factor (co-)variance timing terms, note that these are multiplied by the weights of the unconditional tangency portfolio of the factors. Let us assume all these weights are positive, consistent with the fact that the factors in the FF5 were added to the CAPM because they generate positive alpha. Then, the variance timing terms impact the bias exactly like in a single factor model: if betas are large when factor risk is high and factor returns are large in absolute magnitude, the unconditional beta will overestimate the true expected conditional beta. This fact leads to a negative alpha bias. The covariance timing terms work similarly, so let us provide some intuition from a simple example. Suppose that all but the first conditional beta are zero, and the first beta is positive and time-varying. If this beta comoves with the conditional covariance between the first and one of the other factors, the unconditional regression will overestimate beta with respect to this other factor and thus bias the unconditional alpha downward.

To start, note that both timing terms in Eq. (5) will be small when multifactor betas do not vary much over time. Lewellen and Nagel (2006) argue that this is the reason why unconditional and conditional CAPM alphas are so similar: market betas are too stable. To address this possibility, we present summary statistics for conditional CAPM and FF5 betas in Table I. We see that the average (across 153 characteristics) of the time-series standard deviation of conditional market betas equals 0.28 in the CAPM. While conditional market betas vary a little less in the FF5, at 0.18, the other betas in the FF5 vary a lot more, with the average standard deviation ranging from 0.30 for SMB to 0.45 for RMW. Also, both the minima and maxima (across the 153 characteristics) of the standard deviations of these other betas tend to be large relative to market beta.[12] Overall, conditional betas vary much more over time in the FF5 than in the CAPM, which thus leaves two possible explanations for small alpha differences.

One possibility is that the time-variation in assets' factor betas does not correlate strongly enough with factor risk premia and risk, in which case both terms on the right hand side of Eq. (5) are small. The other possibility is that both timing terms are large, but they roughly cancel each other out. In Panel A of Figure 2, we present the difference between the unconditional and conditional alpha and the alpha bias estimated as the difference between the two timing terms from Eq. (5) for all characteristics.[13] As expected, the two are highly correlated ($corr = 0.89$).[14] In Panel B, we present the two timing terms separately and we see that each of these terms is economically large for a sizeable subset of the 153 characteristics. Across characteristics, the standard deviation of the factor timing ((co-)variance

[12]One may be concerned that some of the time-variation in betas with respect to different factors offsets each other for a given long-short characteristic portfolio. To address this concern, we also report the standard deviation of the sum of betas. This standard deviation is more than three times as large in the FF5 than in the CAPM at 0.95 versus 0.28.

[13]The advantage of using Eq. (5) to analyze these possibilities is that we will only need to estimate betas. Eq. (6) in addition requires estimates of conditional factor risk premia. Because (co-)variances are relatively easy to estimate using simple methods like those that we use in the next section, it does not matter much empirically whether you estimate Eq. (5) or Eq. (6) to analyze factor (co-) variance timing. Indeed, using our estimates of conditional factor risk, the correlation between the factor (co-)variances timing terms I and II is 0.89.

[14]The correlation is not exactly equal to one because of estimation error: we do not observe conditional betas, but we have to estimate them.

timing) term equals 3% (3.5%), respectively. To put this in perspective, note that the average return for four of the FF5 factors over our sample period is less than 4% (SMB, HML, RMW, and CMA). Thus, in isolation, factor timing and factor (co)-variance timing can have an economically large contribution to the alpha bias. However, these two terms are strongly positively correlated across characteristics ($corr = 0.81$), for which reason they tend to cancel each other out.

Taking stock, we have seen that there is large variation in multifactor betas (Table I), these betas comove with both factor premia and risk (Figure 2), but the unconditional-conditional alpha difference is small (Figure 1). Together, these facts suggest that the conditional factor premia, $\mu_{F,t}$, and conditional factor risks, $\Sigma_{F,t}$, that feature in Eq. (6) are positively related in a multivariate sense. In the next section, we impose the structure of linear multifactor models to formalize this risk-return tradeoff and estimate it, for FF5 in Section 2 and for alternative factors extracted from the large cross-section of 153 portfolios in Section 3.

2. The multifactor risk return tradeoff

Do investors require higher factor returns when factor risk – broadly defined – is higher? Although any variable that predicts factor risk and returns with the same sign contributes to a positive risk-return tradeoff, we rely on the economics of the multifactor model to propose a specific set of predictors to answer this question.

A factor model represents a SDF that is linear in the factors F_{t+1} and can be written as $M_{t+1} = \gamma_0 - (1 + R_{F,t})^{-1} F'_{t+1}\Gamma$. The fundamental asset pricing equation states that for any excess return R^e_{t+1}, we have $E_t(M_{t+1}R^e_{t+1}) = 0$, which implies that

$$E_t(R^e_{t+1}) = Cov_t(R^e_{t+1}, F'_{t+1})\Gamma, \tag{7}$$

where we take expectations conditional on the investor's time t information and we have used that the model also prices the risk-free asset, such that $E_t(M_{t+1}) = (1 + R_{F,t})^{-1}$. Since the factors in the FF5 are traded excess returns, we also have:

$$E_t(F_{t+1}) = Var_t(F_{t+1})\Gamma. \tag{8}$$

This equation states that conditional factor risk premia are determined by the conditional factor risks collected in the variance-covariance matrix of the factors, and this relation is governed by the prices of risk in Γ.

Because estimating time-variation in risk-prices is difficult, we follow most previous work and focus on constant prices of risk in our main analysis. We present results for time-varying risk-prices in Section 4. Assuming that Γ is constant does not imply that factor risk premia $E_t(F_{t+1})$ are constant over time, however. Factor risk premia will vary over time with factor variances and covariances. Thus, we study how much time-variation in expected factor returns we can capture through time-varying risk alone. We believe this is an interesting question given the scarcity of evidence on (i) the risk-return tradeoff in multifactor models and (ii) the time-series predictability of long-short factor returns more generally.

In sum, for the FF5, we have:

$$\begin{pmatrix} E_t(MKT_{t+1}) \\ E_t(SMB_{t+1}) \\ E_t(HML_{t+1}) \\ E_t(RMW_{t+1}) \\ E_t(CMA_{t+1}) \end{pmatrix} = \begin{bmatrix} \sigma^2_{M,t} & \sigma_{M,S,t} & \sigma_{M,H,t} & \sigma_{M,R,t} & \sigma_{M,C,t} \\ \sigma_{M,S,t} & \sigma^2_{S,t} & \sigma_{S,H,t} & \sigma_{S,R,t} & \sigma_{S,C,t} \\ \sigma_{M,H,t} & \sigma_{S,H,t} & \sigma^2_{H,t} & \sigma_{H,R,t} & \sigma_{H,C,t} \\ \sigma_{M,R,t} & \sigma_{S,R,t} & \sigma_{H,R,t} & \sigma^2_{R,t} & \sigma_{R,C,t} \\ \sigma_{M,C,t} & \sigma_{S,C,t} & \sigma_{H,C,t} & \sigma_{R,C,t} & \sigma^2_{C,t} \end{bmatrix} \begin{pmatrix} \gamma_{MKT} \\ \gamma_{SMB} \\ \gamma_{HML} \\ \gamma_{RMW} \\ \gamma_{CMA} \end{pmatrix}. \tag{9}$$

The key challenge is that Eq. (9) represents a pooled specification. Covariance of a factor with any of the other factors contributes to time-variation in both of these factors' risk premia and the contributions are determined by the risk-prices on the two factors. Consequently, risk-prices cannot be estimated in univariate regressions of a factor's return on the conditional variance of that factor and they will be inefficiently estimated when adding the covariance with all the other factors to this univariate regression.

Under the null of FF5, the regression of market returns on conditional market variance – the standard in previous literature on the risk-return tradeoff – suffers from omitted variables bias. This bias (see Appendix B) equals:

$$\begin{aligned} \widehat{\gamma_{MKT}} - \gamma_{MKT} = \gamma_{SMB} \frac{Cov\left(\sigma_{M,S,t}; \sigma^2_{M,t}\right)}{Var\left(\sigma^2_{M,t}\right)} + \gamma_{HML} \frac{Cov\left(\sigma_{M,H,t}; \sigma^2_{M,t}\right)}{Var\left(\sigma^2_{MKT,t}\right)} \\ + \gamma_{RMW} \frac{Cov\left(\sigma_{M,R,t}; \sigma^2_{M,t}\right)}{Var\left(\sigma^2_{M,t}\right)} + \gamma_{CMA} \frac{Cov\left(\sigma_{M,C,t}; \sigma^2_{M,t}\right)}{Var\left(\sigma^2_{M,t}\right)}. \end{aligned} \tag{10}$$

We thus have that the standard regression underestimates the market price of risk, when factor risks covary negatively with market variance. We find that this is the relevant case empirically.

As discussed in more detail below, we proxy for the conditional variances and covariances of the factors using realized variances and covariances. Since these proxies predict future risk with a positive sign, a positive risk-return tradeoff means that these proxies are also positively related to factor returns, and therefore that each $\gamma_{\bullet} > 0$.[15] As we discuss in more detail in Appendix C, positive risk-prices are also consistent with Fama and French's preferred interpretation of their model as an application of Merton's (1973) ICAPM. In the ICAPM, the price of market risk (γ_{MKT}) captures the representative agent's relative risk aversion. The risk-prices on the long-short factors are determined by investors hedging demands for state variable risk. Under this interpretation, increasing factor (co-)variances indicate increasing state variable risk, and this would lead to increasing factor risk premia.

Finally, note that one could alternatively use the unconditional moment $E(M_{t+1}F_{t+1}) = 0$ to get an estimate of the risk-prices $\Gamma^u = Var(F_{t+1})^{-1}E(F_{t+1})$ for which the SDF prices the factors unconditionally. We believe this approach is unattractive for two reasons. First, this approach is less efficient, because it ignores all time-variation in risk premia and risk. Second, this approach is more sensitive to mispricing: If a constant part of average factor returns is due to mispricing, we can account for it by including factor fixed effects in our pooled regressions.[16]

We proceed as follows. In Section 2.1, we introduce and test our proxy for con-

[15] If we estimate these risk-prices to be small or even negative, this would mean that our chosen proxy for factor risk does not help to explain why differences between unconditional and conditional alphas in the multifactor model are so small.

[16] Our estimate of Γ^u in the data is $[5.1, 2.7, -2.2, 9.1, 13.4]$ for MKT, SMB, HML, RMW and CMA. As we will show below, these estimates are similar to what we find in our pooled regressions, with the exception of CMA.

ditional factor risk. In Section 2.2, we estimate the risk-prices in Γ in an in-sample test of the multifactor risk-return tradeoff. Finally, we ask in Section 2.3 whether our proxy for conditional factor risk also helps to predict factor returns in an out-of-sample test.

2.1 Does past factor risk predict future factor risk?

We follow common practice in the literature (see, e.g., Andersen and Bollerslev, 1998; Ghysels et al., 2005; Bollerslev et al., 2018) and estimate realized variances and covariances over rolling windows of daily factor returns to proxy for conditional factor risk. Our main proxy is based on a six month window of daily factor returns. We show that the six month window provides us with a useful and robust estimate of factor risk when considering both variances and covariances of all the factors. We will also discuss the information content of our lower frequency risk measures relative to factor risk measured over a shorter one-month window.

To study the information content of our proxy for conditional factor risk, we first run univariate regressions of realized factor variances ($\sigma^2_{k,t+1}$, based on daily returns in month $t+1$ with $\sigma^2_{k,t}$ referring to its one month lag) on the six month risk measures (denoted $\sigma^2_{k,t,6}$, using all days in months $t-5:t$):

$$\sigma^2_{k,t+1} = b_0 + b_1 \sigma^2_{k,t,6} + \epsilon_{k,t+1} \text{ for } k = [MKT, SMB, HML, RMW, CMA]. \quad (11)$$

We report the estimates in Panel A of Table II. Standard errors are Newey and West (1987) with lag length 6. In short, we see that past variance strongly and significantly predicts future variance with a coefficient ranging from 0.50 (SMB) to 0.84 (HML).

We next run a pooled regression of realized factor risks (denoted $\sigma_{k,j,t+1}$) on the six month risk measures:

$$\sigma_{k,j,t+1} = b_{0,k,j} + b_1 \sigma_{k,j,t,6} + \epsilon_{k,j,t+1}. \tag{12}$$

To start, we include all 15 unique terms in the 5×5 variance-covariance matrix of the factors.[17] Next, we run this pooled regression separately for the five variances ($k = j$) as well as the 10 covariances ($k \neq j$). $b_{0,k,j}$ are fixed effects. The coefficient b_1 captures our pooled estimate of the link between past and future risk. t-statistics are based on Driscoll and Kraay (1998) standard errors that are robust to both serial and cross-sectional correlation (with Hansen and Hodrick (1980) weights up to 6 lags).

The six month risk measures strongly predict both variances and covariances. The pooled loading equals 0.62, which is economically large and statistically significant ($t = 6.11$). Consistent with Frazzini and Pedersen (2014), we find that the persistence of covariances is larger than that of variances: the pooled coefficient estimates equal 0.76 ($t = 9.87$) and 0.58 ($t = 5.08$), respectively. Importantly, these findings are robust to changing the window length to 3, 9 or 12 months as well as when using an exponential weighting scheme (see Table IA.V). In sum, realized factor risk strongly predicts future factor risk.

Note that our chosen window length differs from the one month window commonly

[17] This approach is conservative. The multifactor model suggests that each covariance contributes to the risk premia of two factors, while each variance only contributes to the risk premium of one factor. Thus, economically it is sensible to include each covariance twice in Eq. (12). Because 6-month risk measures generally predict covariance better than variance, this would only strengthen our results.

studied in the volatility-timing literature (see, e.g., Moreira and Muir, 2017; Ferson and Mo, 2016; DeMiguel et al., 2024). These timing strategies take advantage of the weak and, if anything, negative relation between last month's market variance and market returns. If (i) this negative relation extends to other factors and covariances and (ii) our six month risk measures contain largely the same information as one month risk measures, then our six month risk measures are unlikely to generate a positive risk-return tradeoff. To address the second part of this concern, we now test whether our risk measures help to predict future variances and covariances controlling for one month risk. To address the first part, we analyze whether the negative relation between one month market variance and market returns extends to the variance of other factors as well as to covariances in the following subsection.

In columns 4 to 7 of Table II, we see that the coefficients on our six month risk measures are typically halved controlling for one month risk. However, the six month measures remain positive and significant in all models (both univariate and pooled). We thus conclude that our risk measures contain significant independent information about future risk that is not picked up by the short-term proxy and this conclusion is robust (see Table IA.V). The one month risk measure is also significant in most specifications, but – with the exception of the univariate test for market variance – the increase in R^2 from adding the one month measure is not overwhelming. In particular, when jointly predicting all covariances, it is our six-month measure that is relatively important with a coefficient of 0.46 relative to 0.31 for the one-month measure (and an R^2 that increases from 0.33 to 0.38 when we add the one-month measure). We thus extend evidence reported in Adrian and Rosenberg (2008) and

Bollerslev et al. (2018) to the setting of a multifactor asset pricing model: tradable factor (co-)variances contain both low and high frequency components. Because the shortest one month risk proxies do not contain all relevant information, we caution against using these directly to proxy for conditional risk.

We conclude that six-month risk measures are useful for testing the multifactor risk-return tradeoff. Previous literature largely ignores longer-term components of multifactor risk, because it almost exclusively studies the market in isolation using relatively short-term risk measures. Adrian and Rosenberg (2008) show that exposure to innovations in both high and low frequency components of market volatility is priced in cross-sections of portfolios. We next ask whether time-variation in these components of multifactor risk is compensated in factor risk premia.

2.2 Does past factor risk predict future factor return?

Our main interest is in pooled regressions of factor returns on the conditional variance and covariance estimates from the previous subsection:

$$F_{t+1} = A_F + \Sigma_{F,t}\Gamma + \nu_{t+1}. \tag{13}$$

where $\Sigma_{F,t}$ is the 5×5 matrix of six month variances and covariances, Γ is a vector containing the five estimated risk-prices and A_F is a vector of factor-specific intercepts. As discussed above, we include factor fixed effects to address the concern that part of the return of empirical factors may be due to mispricing rather than risk. That said, we show below that our results are robust to not including these fixed effects. Including fixed effects is also attractive because it ensures our model matches average factor returns in-sample, such that these tests zoom in on time-variation in factor returns alone. We report the risk-price estimates in Panel B of Table III with t-statistics using either Driscoll and Kraay (1998) asymptotic standard errors (with 6 lags) or standard errors from a stationary bootstrap (with average block length equal to 6 months).Before discussing these pooled estimates, however, let us discuss some benchmark estimates from time-series regressions. We report these estimates in Panel A with t-statistics based on Newey and West (1987) standard errors.

We start with the standard regression of market returns on past market variance. We find a positive, but insignificant coefficient of 0.98. Although this number is somewhat low as an estimate of the representative agent's coefficient of risk aversion, it is at least positive in contrast to a large body of previous work. The main reason for this difference is that we use variance estimated over a six month window. Indeed, if we use one month variance, the estimate is negative at -1 (t=-1.5, unreported). When we run the univariate regression of returns on six month past variance for the other factors, we find positive risk price estimates as well. These estimates range from an insignificant 0.5 for HML to (marginally) significant estimates of over 4 for SMB, RMW, and CMA. Overall, time-series regressions of factor returns on six month variances provide weak evidence of a positive multifactor risk-return tradeoff.

We next add the covariances of the market with the other factors to the time-series regression:

$$R_{M,t+1} = a_M + \sigma^2_{M,t}\gamma_M + \sigma_{M,S,t}\gamma_S + \sigma_{M,H,t}\gamma_H + \sigma_{M,R,t}\gamma_R + \sigma_{M,C,t}\gamma_C + \nu_{M,t+1}. \quad (14)$$

Specification (14) estimates a market price of risk that is positive and significant at 2.97, which previous literature argues is a sensible estimate of risk aversion. This result confirms that the standard estimate of the risk-return tradeoff based on market variance alone is negatively biased. In Appendix Table IA.I, we decompose the total bias of -2 (= 0.98 – 2.97) into the contribution of covariance with each of the long-short factors. To this end, we combine the estimated conditional variances and covariances with the risk-price estimates from Eq. (14). In short, we find that all long-short factors contribute negatively to the bias. The largest contribution comes from the profitability factor RMW (-0.8), because conditional market variance is strongly negatively correlated with the covariance between the market and RMW ($corr = -0.42$) and the price of risk for RMW is large and positive at 10.05. HML and CMA also have a large contribution (at -0.75 and -0.4, respectively), because market variance is positively (negatively) correlated to the covariance between the market and HML (CMA) and the estimated risk-price for HML (CMA) is negative (positive). These findings underscore the importance of studying the risk-return relation from a multifactor perspective: without controlling for the comovement between market variance and other factor risks, it is difficult to find a positive risk-return tradeoff. That said, one may be worried that only the prices of risk for the market and RMW are significant in this specification. Therefore, let us now consider the more efficient pooled estimates.

In Panel B, we see that the pooled estimates of the risk-prices are quite similar to those from Specification (14). The market price of risk is sensible and significant at 2.2. We see that two additional factors have significant prices of risk, that is, SMB and RMW at 4.86 and 7.71. To see why these risk prices tend to be much larger in magnitude than the market, note that the factor returns are roughly the same order of magnitude for all factors, while the variance of the market is 2 to 4 times the variance of the other factors. Thus, to explain the same level of return, the risk prices on the other factors need to be larger. Interpreting these results through the lens of the ICAPM, we conclude that investors demand significantly higher returns for exposure to the combination of state variable risks that SMB and RMW represent. The risk price on CMA is economically large ($\gamma_{CMA} = 5.60$), but insignificant. The risk price on HML is negative, but relatively small and insignificant. This finding suggests that covariance with HML is not priced conditionally. This finding extends Fama and French (2015), who show that HML does not provide an unconditional alpha relative to the remaining factors either.[18] In conclusion, when we control for factor covariances as the multifactor model prescribes, we find a mostly positive risk-return tradeoff that is strong for the market, SMB, RMW, and – to a lesser extent – CMA.

To further grasp the economic content of our pooled model, let us also discuss the R^2. The pooled R^2 is 0.012, but this R^2 is hard to compare to the usual time-series predictive regression R^2 in the literature. Therefore, we estimate the partial R^2 for each factor:

$$\text{Partial } R^2 = 1 - \frac{Var(F_{k,t+1} - P(F_{k,t+1}))}{Var(F_{k,t+1})}, \tag{15}$$

[18] Over our sample period, the unconditional alpha of HML with respect to the four factors MKT, SMB, RMW, and CMA is -0.11% ($t = -1.1$). The alphas of the other long-short factors (in each case relative to the remaining four factors in the FF5) are instead positive and equal 0.2% ($t = 1.67$) for SMB, 0.27% ($t = 4.20$) for CMA and 0.4% ($t = 4.60$) for RMW.

where $P(F_{k,t+1})$ is the predicted factor return. We see that these R^2s equal 0.7% for SMB, 0.8% for HML, 1.6% for CMA, 1.9% for MKT, and 2.1% for RMW. To see that these numbers are economically large, just note that many prominent predictors of market returns studied in previous literature (see, e.g., Welch and Goyal, 2008) generate smaller in-sample R^2s. This fact is remarkable given that we are using a single set of risk-prices for five different factor return series.

We have already discussed that covariances change insights about the risk-return tradeoff in a profound way. However, we have not answered yet how much covariances contribute to time-variation in factor risk premia. To this end, we calculate the relative contribution to expected return variation of each factor's variance and each factor's covariances with the other factors. Panel B of Table III shows that half of the time-variation in the market risk premium is due to the covariance terms. The contribution of covariances is large also for a number of other factors, such as SMB (92%) and CMA (65%). In fact, covariances drive virtually all of the time-variation in the expected return of the RMW factor. In short, the omission of covariances from previous literature has severely biased insights about the time-variation in long-short factor premia.

To see the practical relevance of covariances, we analyze the returns to a volatility-managed market portfolio in Appendix D. In short, we show that our estimated multifactor risk-return tradeoff implies that about half of the CAPM alpha of this strategy is due to covariance of the market with the other factors. In other words, a volatility-managed market portfolio carries large factor risk. Our results thus support the idea that exploiting the alpha of such strategies may be difficult in practice, as concluded also for different reasons in Cederburg et al. (2020) and Barroso and Detzel (2021).

In sum, our conditional risk proxies (based on six month realized variances and covariances) reveal a strong and positive multifactor risk-return trade off, which helps to explain why differences between unconditional and conditional multifactor alphas are so small in Section 1. We discuss in detail in Appendix E that this conclusion is robust in important dimensions. First, we show that our risk-price estimates are largely insensitive to removing the factor-specific fixed effects, which alleviates concerns about Stambaugh bias. The out-of-sample tests of the next section further alleviate this concern. Next, we vary the window over which we estimate conditional factor risk or the horizon over which we predict factor returns. For instance, we show that the risk-return tradeoff is noticeably stronger when risk is measured over longer windows, but it is weakly positive using one-month risk-measures as well. Finally, we show that the estimated risk-prices are largely insensitive to removing outliers as well as adding a variety of test portfolios (following the suggestion in Lewellen et al., 2010).

2.3 Out-of-sample factor return predictability

We next study whether variances and covariances help to predict factor returns out-of-sample. This analysis is interesting from various perspectives. First, given that variances and covariances can exhibit large sudden changes, one may be worried about the stability of our pooled estimates. Second, empirical evidence on time-series predictability of factors beyond the market is relatively scarce, and even less is known about predictive performance in out-of-sample tests. Third, we will tie our hands to models that link risk to return as prescribed by the multifactor theory and therefore will always use the same model to predict returns on all FF5 factors, rather than constructing an optimal model for each factor.

To assess performance, we study the out-of-sample (OOS) R^2, defined as:

$$1 - \frac{\sum_t (F_{k,t+1} - P(F_{k,t+1}|t))}{\sum_t (F_{k,t+1} - \overline{F_{k,t}})}. \tag{16}$$

Here, $P(F_{k,t+1}|t)$ is the predicted factor return for month $t+1$ derived from a model with parameters estimated using data available at time t. $\overline{F_{k,t}}$ is the historical mean of the factor up to and including time t. Throughout, we impose the economic restriction that factor risk premia are positive, i.e., $P(F_{k,t+1}|t) > 0$, and we cap them at the 99th percentile assuming a normal distribution for the factor risk premium based on the historical mean and standard error (estimated using factor return data up to and including t). Imposing such economic restrictions is a staple in out-of-sample forecasting in the literature since Campbell and Thompson (2008).[19] We also study the aggregate fit across factors by calculating an OOS R^2 for an optimal multifactor portfolio, which places more weight on the factors that are more relevant for investors. At each time t, we calculate the optimal weights in the five factors, w_t, based on the historical monthly variance-covariance matrix of the factors as well as the historical average return. We rescale w_t at each point in time to sum to

[19] Table IA.VI reports largely similar results if we floor the factor risk premia at the 1st percentile. The non-negativity constraint is standard for the market and for the long-short factors it is consistent with theory that rationalizes why sorts on these characteristics should capture a positive risk premium.

one. This benchmark represents the optimal portfolio for an investor that ignores the information in current variances and covariances. The OOS R^2 is analogously calculated as:

$$1 - \frac{\sum_t (F_{t+1} - P(F_{t+1}|t))' w_t}{\sum_t (F_{t+1} - \overline{F_t})' w_t}. \tag{17}$$

To start, we simply estimate the pooled specification of Eq. (13) over expanding windows. Note that this specification nests the historical mean, because it includes factor fixed effects. The first window size is 25 years, such that the first one month-ahead prediction is made at the end of December 1994.

We see in the first row of Table IV that the OOS R^2 for this method is positive but small for three factors (MKT, HML, and RMW) and strongly negative for two factors (SMB and CMA, at -1.1% and -0.5%, respectively). For the optimal portfolio, the OOS R^2 is positive as well, though not large at 0.14%. In short, simply applying our main specification out-of-sample does not work well. We believe the reason for the lack of outperformance would have been obvious to an investor in real-time.

To see this, we plot the risk-price estimates in each expanding window in Panel A of Figure 3. We immediately see large instability in the risk-prices. Up to the early 2000s, the risk-prices for HML and CMA hover around values of +50 and -50, respectively, which is an order of magnitude larger than what we find over the full sample and unreasonable economically. Indeed, over the first expanding window of 25 years, the mean-to-variance ratio of factor returns ranges from about 2 for the market to 12 for CMA. Thus, either the investor believes that the contribution of covariances to risk-price estimates is huge or, perhaps more likely, he/she understands that estimating these risk-prices with relatively little data is a bad idea. Since HML and CMA returns are highly correlated over the first 25 year window ($corr = 0.77$), the three terms $Var_t(HML_{t+1})$, $Var_t(CMA_{t+1})$, and $Cov_t(HML_{t+1}, CMA_{t+1})$ are also highly correlated. Consequently, it is difficult to separate the contributions of these three terms to HML and CMA returns through only two risk-prices, γ_{HML} and γ_{CMA}. Fortunately, recent approaches straightforwardly deal with such parameter instability.

We consider two approaches, both adhering to our pooled setup. Our first approach is a forecast combination following Rapach et al. (2010). The motivation is that averaging forecasts from different models reduces uncertainty and parameter instability inherent in forecasts from any individual model. The basic idea behind our applications is as follows. Consider our pooled regression to estimate risk-prices:

$$F_{t+1} = A_F + \Sigma_{F,t}\Gamma + \nu_{t+1}. \tag{18}$$

where $\Sigma_{F,t}$ is a 5×5 matrix and Γ a 5×1 vector. This regression uses information from all five factor returns to estimate each risk-price in Γ, because the return of the kth factor depends on the variance of the kth factor as well as covariance with the other factors. Our first forecast combination (Combination I) averages predicted factor risk premia from five models, where each model relies on a pooled regression to estimate one of the risk-prices in Γ. Intuitively, we thus use each column of $\Sigma_{F,t}$ in separate pooled regressions. Our second forecast combination (Combination II) considers all $\sum_{k=1}^{5}\binom{5}{k}$ =31 possible combinations of risk-prices to be estimated in separate pooled regressions. We then average the 31 forecasts derived from these estimations.

Our second approach estimates the pooled regression of Eq. (18) while imposing a Ridge penalty. This method is specifically designed to shrink coefficients. For each month in the out-of-sample period, we choose an optimal ridge penalty that maximizes the OOS R^2 using data up to that point in time. We divide the expanding in-sample data into a training and validation set using a 90%-10% split that we randomly choose 200 times. In each of the training samples, we estimate the parameters for the model using a specific ridge penalty.[20] We then use the estimated parameters in each of the validation samples to compute the OOS R^2. The optimal ridge penalty for a given month is the penalty with the highest OOS R^2 averaged over all 200 validation samples.

In short, each of these alternative approaches leads to a large improvement in the out-of-sample performance of our model. For instance, for the market we find positive OOS R^2s ranging from 0.51% for Ridge to 0.65% for Combination II. With the exception of SMB, the OOS R^2s are similarly large for the long-short factors HML, RMW and CMA. For these factors, it is either Combination II or Ridge that works best. For SMB, all methods fail to improve meaningfully on the historical mean. This finding is consistent with our in-sample evidence in the sense that SMB also generated the smallest in-sample R^2 among all five factors. This result means that variances and covariances are not useful to predict SMB returns over time, but it does not mean that including the information from SMB in the model is useless overall. Indeed, covariances with SMB do contribute to variation in the returns of the other factors and our model is quite succesful in predicting returns of these other

[20]We search over a wide range of Ridge penalties that is roughly centered around the penalty that in the first 25-year window generates a sum of absolute risk-prices for the five factors that is equal to 5× the ratio between average market return and its variance in that same historical window. Intuitively, a reasonable prior is that all risk-prices are equal to the market price of risk.

factors out-of-sample.

To understand why these alternative methods significantly improve performance, we plot the time-series of estimated risk-prices from the Ridge setup in Panel B of Figure 3. Noting the change in y-axis range compared to Panel A, it is easy to see that these risk-prices are much smaller and the economic insights align surprisingly well with what we find over the full sample. For instance, the risk-price for RMW is largest over the full sample (without penalty) at 7.7, and it remains largest with the Ridge penalty at a value of 5.5 at the end of the sample. The amount of shrinkage (about one-third at the end of the sample) is similar for other factors.

To appreciate the economic magnitude of these OOS R^2, we compare them to the squared monthly Sharpe ratio of each factor. Campbell and Thompson (2008) show that a mean-variance investor that is timing an individual factor can use our model to increase the average monthly portfolio return by a proportional factor equal to $\frac{\text{OOS } R^2}{\text{Sharpe}^2}$. This improvement equals about 33.3% for MKT and CMA, 50% for RMW, and 75% for HML. The forecast combination and Ridge approaches also work well predicting the optimal joint portfolio with OOS R^2s of 0.85% for Combination I, 1.1% for Combination II, and 1.6% for Ridge. Given a squared Sharpe ratio of about 5% for this portfolio, the proportional increases in average returns for mean-variance timers are also economically large ranging from 17 to 32%.

In conclusion, the variance-covariance matrix of the factors helps to predict factor returns out-of-sample through a consistently positive risk-return tradeoff. This conclusion is interesting from at least two perspectives. First, we use the same model for all factors and we are using the simplest off-the-shelf methods to deal with parameter instability. Second, Welch and Goyal (2008) show that many popular market return predictors generate negative OOS R^2. Nevertheless, our OOS R^2 for the market is similar to two recent predictors that are specifically designed to predict the market, that is, the SVIX from Martin (2017) and the three pass regression procedure of Kelly and Pruitt (2013). That said, we caution against interpreting our evidence as an invitation to time factors. The main insight from our paper is that the multifactor risk-return tradeoff is positive. Hence, any increases in portfolio return achievable through factor timing based on our estimates will be coupled with greater risk. In fact, when we analyze such factor timing strategies in Section F of the Appendix, we find that the increases in risk are commensurate to the increases in return. As a result, Certainty Equivalent improvements from timing are small even for moderately risk averse investors.

3. Principal component factors and the risk-return tradeoff

To analyze whether our results extend beyond the FF5 model, we extract principal component (PC) factors from the large set of long-short characteristic-sorted portfolios studied in Section 1. Recent work argues that the SDF can be suitably approximated using a few PC factors when long-short portfolios do not each represent an independent source of priced risk (e.g., Kozak et al., 2020; Kelly et al., 2019; Haddad et al., 2020; Lettau and Pelger, 2020a,b). Intuitively, the PCs provide us with an alternative approach to account for covariances, which we have already argued are key to identifying a positive multifactor risk-return tradeoff. While standard PC analysis generates factors that are unconditionally uncorrelated, the factors' conditional covariances will vary over time and we study if this time-variation is priced. To this end, we first regress monthly returns of PC factors on their conditional variance as well as conditional covariances with the other PC factors, using a pooled specification analogous to Eq. (13). We also combine the PC factors into their mean-variance optimal combination to test whether the return of the Tangency portfolio is predictable by its conditional variance.[21]

We follow Lettau and Pelger (2020a,b) and extract the PCs using the RP-PCA approach. The RP-PCA objective function jointly minimizes unexplained time-series variation in portfolio returns and cross-sectional variation in pricing errors. Thus, much like the factors of Fama and French, RP-PCA factors capture both time-series and cross-sectional variation in returns. The weight on the cross-section relative to the times-series is controlled by a single parameter y. Following the authors' suggestion, we set $y = 10$ in our main analyses. We report qualitatively and often quantitatively similar results for values ranging from −1 (standard PCA) to 15 in Table IA.VIII. Variances and covariances are estimated as before, using six months of daily returns. We show broadly similar results in Table IA.IX for (co-)variances estimated over a shorter one month window. Throughout, we drop 11 characteristics that have some missing data early in the sample as well as all 8 characteristics from

[21]We focused on the individual risk-prices in FF5, because each risk-price is easily interpretable as the compensation for covariance with a single factor portfolio. The risk-prices for the PC factors are harder to interpret, because each factor is a combination of many long-short portfolios. Constructing an optimal portfolio of the PC factors also accommodates an out-of-sample test. It is difficult to set up an out-of-sample test for the separate PC factors, because their identity varies over time.

the momentum category, leaving us with a set of 134 characteristics. In a robustness check, we show that including the relatively small set of momentum characteristics quantitatively affects some of our results, even though it leaves the main takeaways qualitatively unchanged.

We first analyze the pooled regression for a five-factor model that appends the first four RP-PCA factors from the long-short characteristic-sorted portfolios to the market portfolio. Because we sign each PC factor to have a positive average excess return, positive coefficient estimates in this regression indicate a positive multifactor risk-return trade-off.[22] In Panel A of Table V, we see that the estimated coefficient is positive for all five factors. As in FF5, the market risk premium is economically large at 2.04%, although it is not significant ($t = 1.4$). The risk-prices for the remaining PC factors range from 2.29% to 18.77% and three out of four are marginally significant.

We next note that the PC factors are considerably more predictable than the FF5 factors: the R^2 of this pooled regression equals 2.7%, which is substantially above the 1.2% we found in Table III. The partial R^2s are also quite large, especially for the most volatile factors at values of 1.9%, 3.7%, and 2.6% for PC1 (=MKT), PC2, and PC3, respectively. This finding suggests that either the conditional (co-)variances of the PC factors better capture the true variation in priced risk or the PC factors better capture the true variation in risk premia that drive the returns of characteristic-sorted portfolios, or a combination of both. To shed light on this question, we regress in Panel B the returns of the FF5 factors on their covariances with the five PC factors. We thus estimate risk-prices for covariance with the PC

[22]We scale the factors such that the first PC of long-short portfolios has a volatility of 3.25% per month, similar to the FF5 factor with the largest volatility (SMB).

factors as in Panel A, but now using the FF5 factors as test assets. First, we see that the estimated risk-prices are all positive and similar in magnitude to Panel A, which confirms that the tradeoff between PC factor risk and return is positive. Second, the R^2 in this pooled regression is about 50% larger than what it is in Table III at 1.8%. The partial R^2s show that this increase is largely due to an improved fit for HML, RMW, and CMA. We can interpret these increases in R^2 as reflecting that covariances with PC factors better capture the true variation in priced risks (than FF5 (co-)variances). Under this interpretation, the further increases in R^2 when we use the PC factors as test assets in Panel A reflect that the PC factor returns better capture the true variation in risk premia (than FF5 returns).

In Panel C, we report the results from the time-series regression of the optimal combined portfolio of the five PC factors on its own conditional variance. In this test, all variance and covariance risk is aggregated into a single pricing factor. For the case of 5 PCs, the coefficient estimate in this time-series regression is positive and significant at 17.5 (t = 2.01). For comparison, if we regress the mean-variance optimal combination of the five factors in FF5 on its own conditional variance, we estimate a similar coefficient of 15.0 (t = 3.24). These coefficient estimates are so large, because the optimal portfolio generates relatively high returns per unit of risk. In Panel C, we also vary the number of PCs from three to seven. While the coefficient estimate is positive in all cases, it is only large and significant when we consider at least 4 PC factors (i.e., the MKT plus the first three PCs of long-short portfolios). In Appendix G, we discuss additional evidence showing that the risk-return tradeoff is positive even in out-of-sample tests that use no forward-looking information in the construction of (i) the RP-PCA factors and (ii) the optimal combination of these factors. We conclude that the risk-return trade-off is robustly positive for PC factors extracted from a large set of 134 characteristic-sorted portfolios. Moreover, as discussed already in Section 2, the key to this finding is to account for covariances between more than a few factors.

We report in Table IA.X the evidence when we add 8 characteristics from the momentum category. In Panel A, we see that the pooled evidence for a positive multifactor risk-return trade-off weakens. For instance, the risk-price on PC3 turns negative. The fact that the risk-return tradeoff is weaker when accounting for momentum-type strategies is consistent with Barroso and Santa-Clara (2015), who show that momentum returns are negatively correlated with conditional variance. That said, when combining the PCs into a single pricing factor (Panel B), we see that this single factor loads positively on its own conditional variance. In short, there is some evidence in favor of a positive multifactor risk-return tradeoff even when accounting for momentum characteristics.

4. Time-variation in risk-prices

We have so far assumed that prices of risk are constant over time. Since measures of risk aversion, hedge demands, and the risk-return tradeoff more generally may vary over time, we now relax this assumption. With risk-prices that are linear in some state variable z_t, the main equation of interest becomes:

$$E_t(F_{t+1}) = \Sigma_{F,t,6}\Gamma_t \text{ where } \Gamma_t = \Gamma_0 + \Gamma_1 z_t. \tag{19}$$

To bring this equation to the data, the key choice is z_t. In a recent contribution, (DeMiguel et al., 2024, DMU) argue that the ratio of factor returns to factor variances – a measure of risk-prices that ignores factor covariances – is decreasing in one-month realized market volatility, $\sigma_{MKT,t}$. However, basic economic intuition would suggest the opposite: if risk aversion and hedge demands are increasing in times of market stress, risk-prices should increase rather than decrease in $\sigma_{MKT,t}$.

In this section, we follow DMU and set z_t equal to demeaned inverse one month realized market volatility, denoted $\sigma^{-1,\star}_{MKT,t}$, and revisit their evidence. We ask if our positive multifactor risk-return tradeoff varies with market volatility, while accounting for the role of factor covariances exactly like we have done so far in this paper. Because six month market variance is correlated to this conditioning variable, we also present results for a dummy specification that makes the time-variation easier to interpret. We define a dummy I_σ that is equal to one when $\sigma^{-1}_{MKT,t}$ is above its 70th percentile. We thus run the following pooled regressions:

$$F_{t+1} = A_F + \Sigma_{F,t,6}(\Gamma_0 + \Gamma_1 \times \sigma^{-1,\star}_{MKT,t}) + \nu_{t+1} \tag{20}$$

$$F_{t+1} = A_F + \Sigma_{F,t,6}(\Gamma_0 \times (1 - I_\sigma) + \Gamma_1 \times I_\sigma) + \nu_{t+1}, \tag{21}$$

with each regression estimating a total of fifteen parameters. Table VI below shows coefficient estimates from this specification for the FF5 model, but we find similar results for the RP-PCA factors studied in the previous section.

In the continuous specification, we see that all risk-price estimates in Γ_0 are positive, but insignificant. In turn, four out of five estimates in Γ_1 are negative, but only one is significant. In the dummy specification, which is easier to interpret because it suffers less from correlation between the variables being interacted, we see that the estimates are consistent in sign, but now also significant. Of the five unconditional risk-price estimates in Γ_0, four are economically large and positive (MKT, SMB, RMW and CMA) and three of these estimates are significant. For the same four factors, the estimates in Γ_1 are negative and significant, which means that these risk-prices are lower when inverse realized market volatility is high. For instance, for the MKT (RMW), the unconditional risk-price is significant at 2.37 (8.41) and this risk-price is lower by a significant -3.66 (-14.46) when inverse realized market volatility is above its 70th percentile. In other words, when realized market volatility is high, risk-prices are higher for four out of five factors.

In Table IA.VII, we present various robustness checks that provide similar insights. First, we run the pooled regressions without factor-specific intercepts, as in DMU. Second, we estimate the variance-covariance matrix of the factors using only one month of daily data to be consistent with (i) the definition of the interaction variable and (ii) Section IA.23 of DMU. Thus, in contrast to DMU, we conclude that the risk-return tradeoff estimated in Section 2 is, if anything, stronger in periods of high market volatility. This conclusion aligns well with basic economic intuition. To understand why our results are so different, note that DMU study a specification with nine factors. Consequently, they have to estimate three times as many (co-

)variance terms than we do and because the one-month risk measures are relatively noisy, they resort to the shrinkage estimator of Ledoit and Wolf (2004). Our main setup using only five factors and six-month (co-)variances is less noisy and therefore we can stick to the simplest possible realized (co-)variance estimates. Indeed, we show in Table IA.VII that our results are identical when we use shrinkage.

5. Conclusion

This paper provides new insights into the multifactor risk-return tradeoff and its implications for asset pricing, broadly through two results. First, we show that unconditional and conditional FF5 alphas are similar for a large set of long-short anomaly portfolios. Previous work shows that this alpha difference is small also in the CAPM and we extend the theoretical analyses in these papers to argue that the lack of alpha difference in the FF5 model suggests there is a positive multifactor risk-return tradeoff. Our second set of results reveals this tradeoff in the data. To this end, we run pooled regressions of factor returns on factor variances and covariances. Our risk-price estimates are strongly positive and robust across many alternative specifications. For instance, we find a positive multifactor risk-return trade-off not only for FF5, but also for alternative models that extract principal component factors from a large set of characteristic-sorted portfolios. Also, we find that factor (co-)variances are at least as useful in predicting factor returns out-of-sample as prominent alternative predictors in the literature, which is remarkable given that our pooled model fixes parameters across factors. We conclude that previous estimates of the risk-return tradeoff in the literature are biased, because they ignore the important role of factor covariances. When properly accounting for covariances, multifactor risk generates time-variation in both market and long-short factor risk premia in the way that standard models predict.

These results are important practically. The risk-return tradeoff of factors is a key input in portfolio management, because real-world portfolios load strongly on these factors through portfolio tilts. Our results suggests that if these tilts are time-varying, such as through factor rotation or volatility- and value-timing, the portfolio is exposed to covariance risks that so far have gone largely unnoticed. Finally, we leave open the question why the risk-price for one factor is larger than another. Even though our results do not depend on whether factors capture economic or (merely) statistical risks, studying whether and how different investors' demands for exposure to characteristics interact with factor risk is an interesting avenue for future work.

References

Adrian, T., Rosenberg, J., 2008. Stock returns and volatility: Pricing the short-run and long-run components of market risk. The journal of Finance 63, 2997–3030.

Akbas, F., Armstrong, W. J., Sorescu, S., Subrahmanyam, A., 2015. Smart money, dumb money, and capital market anomalies. Journal of Financial Economics 118, 355–382.

Andersen, T. G., Bollerslev, T., 1998. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. International economic review pp. 885–905.

Arnott, R. D., Kalesnik, V., Linnainmaa, J. T., 2023. Factor momentum. The Review of Financial Studies 36, 3034–3070.

Asness, C. S., Liew, J. M., Pedersen, L. H., Thapar, A. K., 2017. Deep value. Available at SSRN 3076181 .

Baba Yara, F., Boons, M., Tamoni, A., 2021. Value return predictability across asset classes and commonalities in risk premia. Review of Finance 25, 449–484.

Baba-Yara, F., Boons, M., Tamoni, A., 2024. Persistent and transitory components of firm characteristics: Implications for asset pricing. Journal of Financial Economics 154, 103808.

Baker, M., Wurgler, J., 2006. Investor sentiment and the cross-section of stock returns. The Journal of Finance 61, 1645–1680.

Bali, T. G., Engle, R. F., 2010. The intertemporal capital asset pricing model with dynamic conditional correlations. Journal of Monetary Economics 57, 377–390.

Barroso, P., Detzel, A., 2021. Do limits to arbitrage explain the benefits of volatility-managed portfolios? Journal of Financial Economics 140, 744–767.

Barroso, P., Maio, P., 2024. The risk-return trade-off among equity factors. Journal of Empirical Finance p. 101518.

Barroso, P., Santa-Clara, P., 2015. Momentum has its moments. Journal of Financial Economics 116, 111–120.

Boguth, O., Carlson, M., Fisher, A., Simutin, M., 2011. Conditional risk and performance evaluation: Volatility timing, overconditioning, and new estimates of momentum alphas. Journal of Financial Economics 102, 363–389.

Bollerslev, T., Hood, B., Huss, J., Pedersen, L. H., 2018. Risk everywhere: Modeling and managing volatility. The Review of Financial Studies 31, 2729–2773.

Bordalo, P., Gennaioli, N., La Porta, R., Shleifer, A., 2024a. Finance without exotic risk. Tech. rep., National Bureau of Economic Research.

Bordalo, P., Gennaioli, N., Porta, R. L., Shleifer, A., 2019. Diagnostic expectations and stock returns. The Journal of Finance 74, 2839–2874.

Bordalo, P., Gennaioli, N., Porta, R. L., Shleifer, A., 2024b. Belief overreaction and stock market puzzles. Journal of Political Economy 132, 1450–1484.

Campbell, J. Y., 1987. Stock returns and the term structure. Journal of financial economics 18, 373–399.

Campbell, J. Y., Hentschel, L., 1992. No news is good news: An asymmetric model of changing volatility in stock returns. Journal of financial Economics 31, 281–318.

Campbell, J. Y., Thompson, S. B., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? The Review of Financial Studies 21, 1509–1531.

Campbell, J. Y., Vuolteenaho, T., 2004. Bad beta, good beta. American Economic Review 94, 1249–1275.

Cederburg, S., O'Doherty, M. S., Wang, F., Yan, X. S., 2020. On the performance of volatility-managed portfolios. Journal of financial Economics 138, 95–117.

Cohen, R. B., Polk, C., Vuolteenaho, T., 2003. The value spread. The Journal of Finance 58, 609–641.

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

DeMiguel, V., Martin-Utrera, A., Uppal, R., 2024. A multifactor perspective on volatility-managed portfolios. Journal of Finance forthcoming .

Driscoll, J. C., Kraay, A. C., 1998. Consistent covariance matrix estimation with spatially dependent panel data. Review of economics and statistics 80, 549–560.

Fama, E. F., French, K. R., 1996. Multifactor explanations of asset pricing anomalies. The Journal of Finance 51, 55–84.

Fama, E. F., French, K. R., 2015. A five-factor asset pricing model. Journal of Financial Economics 116, 1–22.

Fama, E. F., French, K. R., 2016. Dissecting anomalies with a five-factor model. The Review of Financial Studies 29, 69–103.

Ferson, W., Mo, H., 2016. Performance measurement with selectivity, market and volatility timing. Journal of Financial Economics 121, 93–110.

Frazzini, A., Pedersen, L. H., 2014. Betting against beta. Journal of Financial Economics 111, 1–25.

French, K. R., Schwert, G. W., Stambaugh, R. F., 1987. Expected stock returns and volatility. Journal of financial Economics 19, 3–29.

Ghysels, E., Santa-Clara, P., Valkanov, R., 2005. There is a risk-return trade-off after all. Journal of financial economics 76, 509–548.

Glosten, L. R., Jagannathan, R., Runkle, D. E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. The journal of finance 48, 1779–1801.

Gormsen, N. J., Jensen, C. S., 2024. Conditional risk. Journal of Financial Economics 162, 103933.

Green, J., Hand, J. R., Zhang, X. F., 2017. The characteristics that provide independent information about average us monthly stock returns. The Review of Financial Studies 30, 4389–4436.

Greenwood, R., Hanson, S. G., 2012. Share issuance and factor timing. The Journal of Finance 67, 761–798.

Guo, H., Whitelaw, R. F., 2006. Uncovering the risk–return relation in the stock market. The Journal of Finance 61, 1433–1463.

Haddad, V., Kozak, S., Santosh, S., 2020. Factor timing. The Review of Financial Studies 33, 1980–2018.

Hansen, L. P., Hodrick, R. J., 1980. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. Journal of political economy 88, 829–853.

Harvey, C. R., 2001. The specification of conditional expectations. Journal of Empirical Finance 8, 573–637.

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

Herskovic, B., Moreira, A., Muir, T., 2019. Hedging risk factors. Working Paper .

Hjalmarsson, E., 2008. The stambaugh bias in panel predictive regressions. Finance Research Letters 5, 47–58.

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

Hou, K., Xue, C., Zhang, L., 2020. Replicating anomalies. The Review of Financial Studies 33, 2019–2133.

Jagannathan, R., Wang, Z., 1996. The conditional capm and the cross-section of expected returns. The Journal of finance 51, 3–53.

Jensen, T. I., Kelly, B., Pedersen, L. H., 2023. Is there a replication crisis in finance? The Journal of Finance 78, 2465–2518.

Kelly, B., Pruitt, S., 2013. Market expectations in the cross-section of present values. The Journal of Finance 68, 1721–1756.

Kelly, B. T., Pruitt, S., Su, Y., 2019. Characteristics are covariances: A unified model of risk and return. Journal of Financial Economics 134, 501–524.

Keloharju, M., Linnainmaa, J. T., Nyberg, P., 2021. Long-term discount rates do not vary across firms. Journal of Financial Economics 141, 946–967.

Kozak, S., Nagel, S., Santosh, S., 2018. Interpreting factor models. The Journal of Finance 73, 1183–1223.

Kozak, S., Nagel, S., Santosh, S., 2020. Shrinking the cross-section. Journal of Financial Economics 135, 271–292.

Ledoit, O., Wolf, M., 2004. A well-conditioned estimator for large-dimensional covariance matrices. Journal of multivariate analysis 88, 365–411.

Lettau, M., Pelger, M., 2020a. Estimating latent asset-pricing factors. Journal of Econometrics 218, 1–31.

Lettau, M., Pelger, M., 2020b. Factors that fit the time series and cross-section of stock returns. The Review of Financial Studies 33, 2274–2325.

Lewellen, J., Nagel, S., 2006. The conditional capm does not explain asset-pricing anomalies. Journal of Financial Economics 82, 289–314.

Lewellen, J., Nagel, S., Shanken, J., 2010. A skeptical appraisal of asset pricing tests. Journal of Financial economics 96, 175–194.

Martin, I., 2017. What is the expected return on the market? The Quarterly Journal of Economics 132, 367–433.

McLean, R. D., Pontiff, J., 2016. Does academic research destroy stock return predictability? The Journal of Finance 71, 5–32.

Merton, R. C., 1973. An intertemporal capital asset pricing model. Econometrica: Journal of the Econometric Society pp. 867–887.

Moreira, A., Muir, T., 2017. Volatility-managed portfolios. The Journal of Finance 72, 1611–1644.

Nelson, D. B., 1991. Conditional heteroskedasticity in asset returns: A new approach. Econometrica: Journal of the econometric society pp. 347–370.

Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55(3), 703–708.

Petkova, R., Zhang, L., 2005. Is value riskier than growth? Journal of Financial Economics 78, 187–202.

Polk, C., Thompson, S., Vuolteenaho, T., 2006. Cross-sectional forecasts of the equity premium. Journal of Financial Economics 81, 101–141.

Rapach, D. E., Strauss, J. K., Zhou, G., 2010. Out-of-sample equity premium prediction: Combination forecasts and links to the real economy. The Review of Financial Studies 23, 821–862.

Rossi, A. G., Timmermann, A., 2015. Modeling covariance risk in merton's icapm. The Review of Financial Studies 28, 1428–1461.

Scruggs, J. T., 1998. Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach. The Journal of Finance 53, 575–603.

Stambaugh, R. F., Yu, J., Yuan, Y., 2012. The short of it: Investor sentiment and anomalies. Journal of financial economics 104, 288–302.

van Binsbergen, J. H., Ma, L., Schwert, M., 2022. The factor multiverse: The role of interest rates in factor return measurement and discovery. Available at SSRN 4231626 .

Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. The Review of Financial Studies 21, 1455–1508.

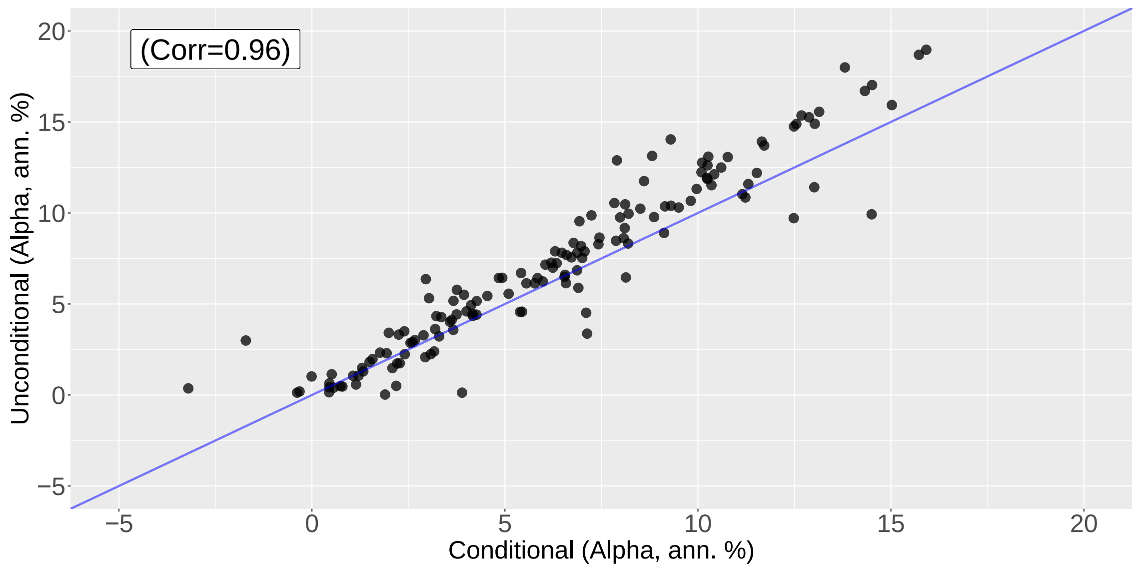

(A) CAPM alphas

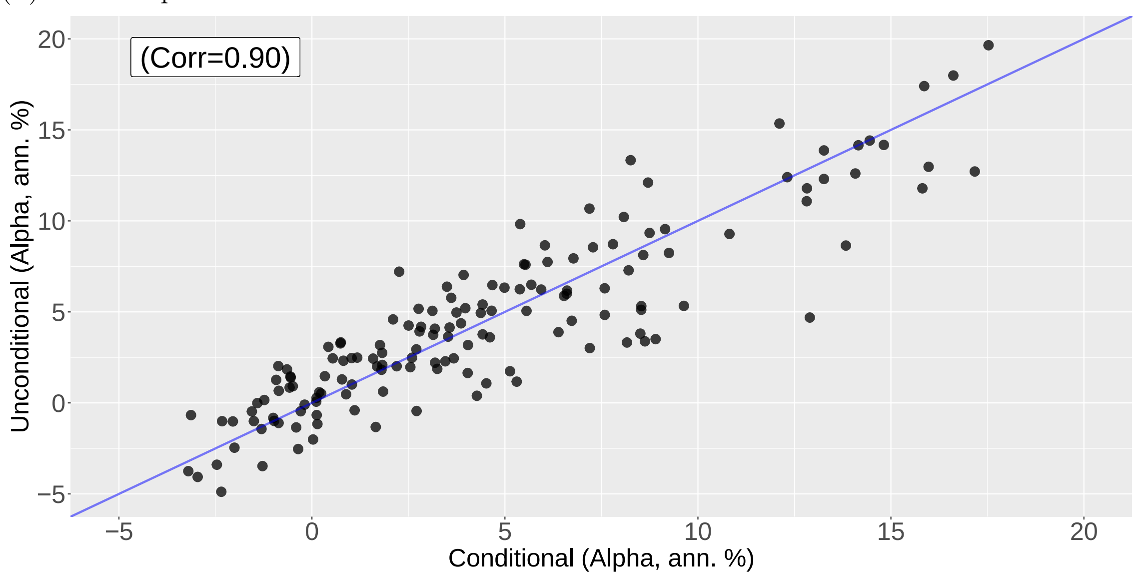

(B) FF5 alphas

FIGURE 1: **Unconditional vs Conditional Alphas**
This figure presents scatterplots of unconditional and conditional alphas for 153 long-short characteristic-sorted portfolios. Panel A shows CAPM alphas, while Panel B shows Fama-French five-factor (FF5) alphas. Unconditional alphas are estimated using a full sample regression of monthly long-short returns on the factors: $R_{L-S,t+1} = \alpha^u + \beta' F_{t+1} + \epsilon_{t+1}$. Conditional alphas are estimated as $\alpha^c = \frac{1}{T}\sum_t (R_{L-S,t+1} - \widehat{\beta}_t' F_{t+1})$, where $\widehat{\beta}_t'$ is a lagged component beta estimated using 12 month windows of daily firm-level returns. The sample period is from January 1970 to December 2023.

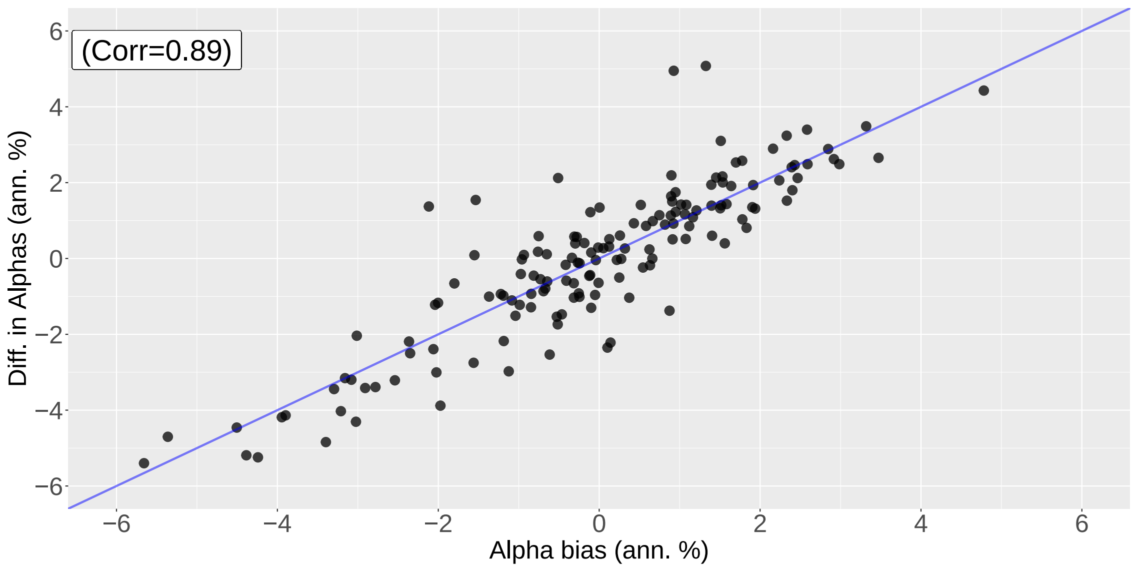

(A) Alpha bias

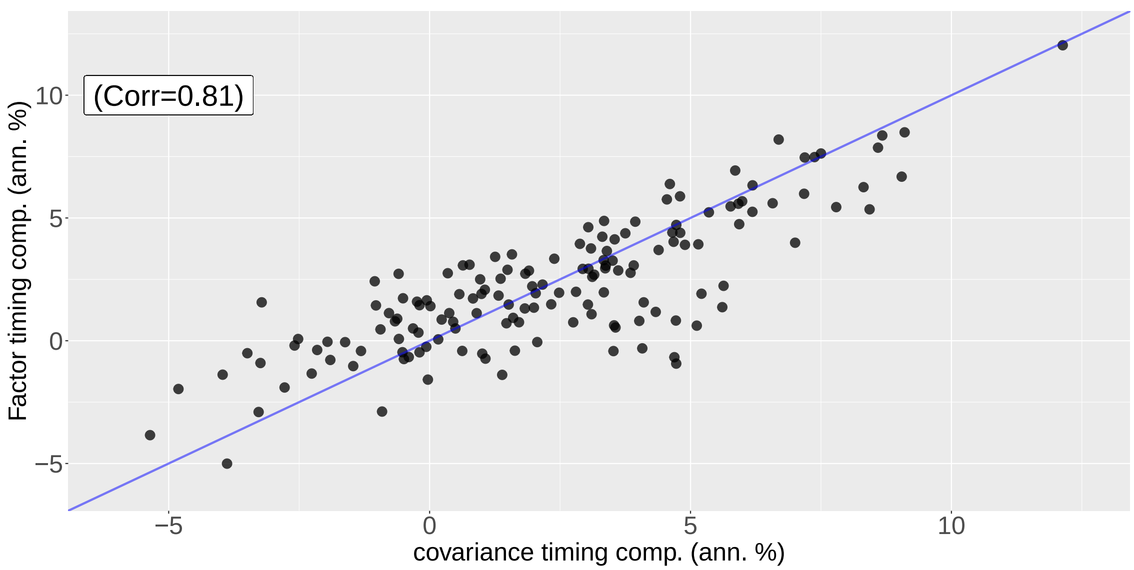

(B) Factor versus (Co-)Variance Timing

FIGURE 2: **Alpha Bias Decomposition**
This figure decomposes the alpha bias in the Fama-French five-factor model for 153 well-known anomalies. Panel A plots this alpha bias estimated as the difference between the unconditional and conditional alphas against the approximation in (IA.5), that is, the difference between the factor timing and (co-)variance timing terms. Panel B plots these timing terms separately for each anomaly. The sample period is January 1970 to December 2023.

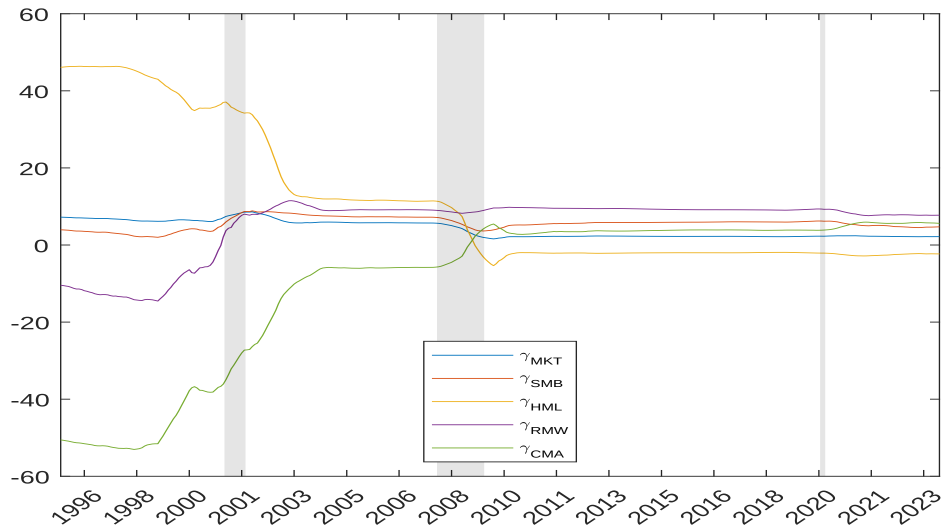

(A) Expanding window estimates of risk-prices (No penalty)

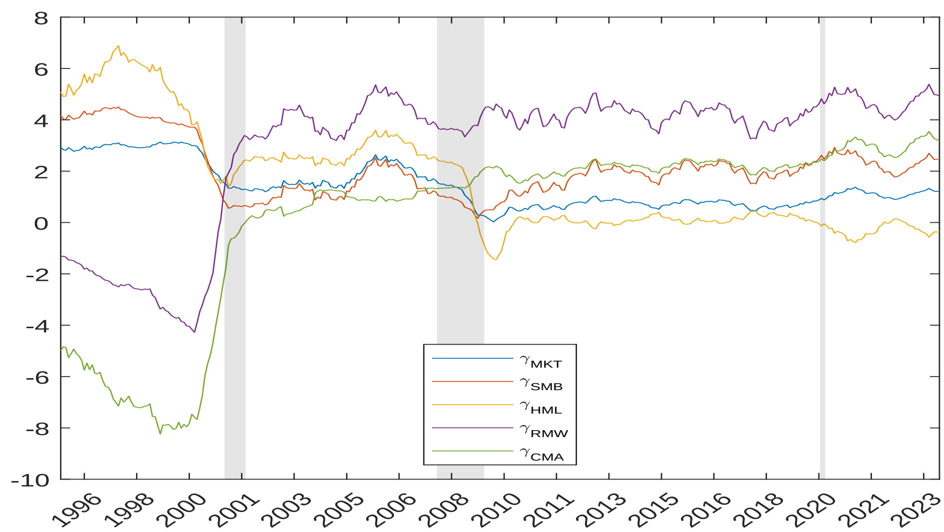

(B) Expanding window estimates of risk-prices (Optimal ridge penalty)

Figure 3: **Risk-prices over time**
This figure plots risk-price estimates (smoothed, 12 month average) over expanding windows that are used to make out-of-sample predictions for returns ranging January 1995 to December 2023. Panel A runs our pooled regression without any adjustment over each window. Panel B uses a Ridge penalty.

Table I: **Summary Statistics for Lagged Component Betas**

This table reports summary statistics for lagged component betas estimated using the Capital Asset Pricing Model (CAPM) and the Fama-French Five-Factor Model (FF5). Betas are estimated using daily firm-level returns over the previous 12 months and then aggregated up to the portfolio-level. Panel A presents statistics for the CAPM market factor. Panel B presents statistics for the FF5 factors: market (MKT), size (SMB), value (HML), profitability (RMW), and investment (CMA). We report the cross-sectional average, minimum, and maximum of the time-series means (μ) and standard deviations (σ) across all characteristics for each factor. The "Sum" row in Panel B reports the standard deviation of the sum of factor betas.

	Average β_t	**Min**	**Max**	**Standard Deviation β_t**	**Min**	**Max**
			Panel A: CAPM			
MKT	-0.20	-1.56	0.67	0.28	0.10	0.56
			Panel B: FF5			
MKT	-0.11	-1.14	0.52	0.18	0.07	0.36
SMB	-0.15	-1.08	0.61	0.30	0.09	0.57
HML	0.20	-0.63	1.33	0.42	0.18	0.82
RMW	0.28	-0.69	1.15	0.45	0.17	1.02
CMA	0.12	-0.34	0.80	0.44	0.19	0.88
Sum	0.34			0.95		

TABLE II: **Predicting variances and covariances**

In this table we run predictive regressions of realized variances and covariances - estimated using daily factor returns in month $t+1$ - on their lagged values estimated using 1 month and 6 month windows of daily data (up to t). We run predictive time series regressions for the variance of the market, SMB, HML, RMW and CMA factors in the first five pairs of rows. In the first three columns of Panel A we use predictors estimated over a six month window and in the last four columns predictors estimated over one month and six month windows. In Panel B, we estimate pooled panel regressions for all variance and covariance terms between each of the risk factors in our model. We include fixed effects for each element in the variance-covariance matrix that we consider. In the first pair of rows we consider all elements, in the second pair of rows only the diagonal terms and in the final set only the off-diagonal elements. We report corresponding t-statistics below the parameter estimates.

	Panel A: Time series				
	$\sigma^2_{k,t,6}$	R^2	$\sigma^2_{k,t,6}$	$\sigma^2_{k,t}$	R^2
$\sigma^2_{MKT,t+1}$	0.53	0.11	0.18	0.38	0.21
	3.93		2.76	2.44	
$\sigma^2_{SMB,t+1}$	0.50	0.09	0.31	0.21	0.12
	4.64		3.62	1.65	
$\sigma^2_{HML,t+1}$	0.84	0.48	0.36	0.50	0.57
	8.80		2.18	2.95	
$\sigma^2_{RMW,t+1}$	0.79	0.42	0.35	0.48	0.52
	8.42		3.68	4.79	
$\sigma^2_{CMA,t+1}$	0.82	0.47	0.49	0.36	0.52
	5.59		3.84	4.73	
	Panel B: Pooled panel				
	$\sigma_{k,j,t,6}$	R^2	$\sigma_{k,j,t,6}$	$\sigma_{kj,t}$	R^2
All ($\sigma_{k,j,t+1} \quad \forall k,j$)	0.62	0.17	0.27	0.37	0.26
	6.11		4.26	3.10	
Variances ($\sigma^2_{k,t+1} \quad \forall k$)	0.58	0.14	0.23	0.38	0.24
	5.08		3.39	2.69	
Covariances ($\sigma_{k,j,t+1} \quad \forall k \neq j$)	0.76	0.33	0.46	0.31	0.38
	9.87		4.89	4.71	

Table III: **Risk-return tradeoff**

In this table we report price of risk estimates. In Panel A, we estimate a time series regression of next month's factor returns on variances (and covariances) estimated using daily realized returns in the past six months. In the first two rows we focus on the market as risk factor and regress market returns on conditional market variance. In the next rows we run the analogous test for the other four Fama-French factors. In the last rows of Panel A, we predict market returns controlling for covariances with the other four risk factors. In Panel B, we regress next period's factor returns jointly on realized variances and covariances in a pooled panel regression. We report the five price of risk estimates with corresponding Driscoll-Kraay and bootstrapped t-statistics below and the pooled R^2 in the last column. In Panel B, we further report the partial (time-series) R^2 for each factor as well as the fraction of total variation in each factor's predicted return that is due to variance versus covariances.

	Panel A: Time series					
	γ_{MKT}	γ_{SMB}	γ_{HML}	γ_{RMW}	γ_{CMA}	R^2
MKT_{t+1} on $\sigma^2_{MKT,t,6}$	0.98					0.003
	1.04					
$SMBt+1$ on $\sigma^2_{SMB,t,6}$		4.23				0.006
		2.47				
HML_{t+1} on $\sigma^2_{HML,t,6}$			0.51			-0.001
			0.30			
RMW_{t+1} on $\sigma^2_{RMW,t,6}$				5.65		0.017
				2.84		
CMA_{t+1} on $\sigma^2_{CMA,t,6}$					5.00	0.008
					1.65	
MKT_{t+1} on $\sigma_{MKT,j,t,6}$ $\forall j$	2.97	7.33	-4.20	10.05	5.56	0.012
	3.00	1.61	-0.91	2.22	0.99	
	Panel B: Pooled panel					
F_{t+1} on $\Sigma_{F,t,6}$	2.20	4.86	-2.38	7.71	5.60	0.012
t_{DK}	3.61	2.22	-1.01	2.47	1.35	
t_{Boot}	2.83	1.97	-0.91	2.11	1.16	
	MKT	SMB	HML	RMW	CMA	
Partial R2	0.019	0.007	0.008	0.021	0.016	
% Own Variance	0.50	0.08	1.19	-0.13	0.35	
% Covariance w. Other Factors	0.50	0.92	-0.19	1.13	0.65	

TABLE IV: **Out-of-Sample R^2s**

In this table we present the out-of-sample R^2s for four approaches. This R^2 is calculated as described in the text over the out-of-sample period from January 1995 to December 2023, for both individual factors and optimal portfolio with conditional weights determined by the average return and variances and covariances estimated over each expanding window. All models rely on estimating pooled regressions (similar to Eq. (18) over expanding windows), such that return forecasts for all five factors are based on five risk-price estimates. We constrain the forecast for each factor to be non-negative and smaller than the factor's average return plus 2.33 standard errors (where both the average and standard error are estimated over the expanding window).

	MKT	SMB	HML	RMW	CMA	Optimal portfolio
Simple	0.0002	-0.0108	0.0004	0.0011	-0.0053	0.0014
Combination I	0.0054	0.0001	0.0048	0.0050	0.0033	0.0085
Combination II	0.0065	-0.0035	0.0067	0.0072	0.0034	0.0109
Ridge	0.0051	-0.0056	0.0081	0.0071	0.0060	0.0159

Table V: **The Risk-Return Tradeoff for Principal Component Factors**
We extract RP-PCA factors from the large set of characteristic-sorted portfolios studied in Section 1. In Panel A, we consider a five-factor model including the market and the first four RP-PCs extracted from long-short portfolios (using weighting parameter $y = 10$). We report coefficient estimates from the pooled regression of factor returns on factor variances and co-variances, analogous to Eq. (13). In Panel B, we run a pooled regression of FF5 factor returns on their covariances with the five PC factors (thus estimating the same five risk-prices as in Panel A). In both panels, we also report the pooled R^2 as well as the partial R^2 (defined as in Eq. (15)) for each of the factors included on the left-hand side (PC factors in Panel A and FF5 factors in Panel B). In Panel C, we report the coefficient estimate and R^2 from a time-series regression of the returns of the mean-variance optimal combination of RP-PC factors on its own conditional variance. In this case, we consider models with three up to seven factors.

Panel A: PC factor returns on PC conditional (co-)variances (Pooled)						
	PC1=MKT	PC2	PC3	PC4	PC5	R^2
Estimate	2.04	2.29	8.01	18.77	7.45	0.029
t-stat	(1.40)	(1.73)	(1.99)	(1.45)	(1.74)	
	PC1=MKT	PC2	PC3	PC4	PC5	
Partial R^2	0.019	0.037	0.026	0.001	0.016	

Panel B: FF5 returns on conditional covariances with PC factors (Pooled)						
	PC1=MKT	PC2	PC3	PC4	PC5	R^2
Estimate	2.37	3.13	8.05	13.29	10.20	0.018
t-stat	(1.48)	(2.04)	(2.19)	(1.00)	(2.12)	
	MKT	SMB	HML	RMW	CMA	
Partial R^2	0.017	0.003	0.015	0.026	0.021	

Panel C: Optimal combination of K PCs on conditional variance (Time-series)					
K =	3	4	5	6	7
Estimate	1.51	18.54	17.50	19.09	20.79
t-stat	(0.95)	(2.93)	(2.01)	(2.10)	(1.94)
R^2	0.000	0.037	0.016	0.016	0.016

Table VI: **Time-Variation in the Multifactor Risk-Return Tradeoff**
This table reports coefficient estimates from the regressions in Eqs. (20) and (21) that ask whether the risk-return tradeoff varies over time as a function of (demeaned) inverse realized market volatility. t-statistics use Driscoll and Kraay (1998) standard errors with 6 lags.

	Γ_0 MKT	SMB	HML	RMW	CMA	Γ_1 MKT	SMB	HML	RMW	CMA
	Demeaned inverse realized market volatility									
Estimate	0.69	0.87	1.90	4.32	2.31	-0.10	-0.31	0.28	-0.23	-0.21
t-stat	(0.64)	(0.28)	(0.52)	(0.86)	(0.33)	(-1.61)	(-2.22)	(1.59)	(-0.82)	(-0.49)
	Demeaned inverse realized market volatility dummy									
Estimate	2.37	5.77	-2.43	8.41	5.68	-3.66	-10.11	7.87	-14.46	-21.09
t-stat	(3.94)	(2.83)	(-1.04)	(2.80)	(1.38)	(-1.91)	(-2.29)	(1.89)	(-2.09)	(-2.15)

Online Appendix

A. Alpha bias

Taking unconditional expectations of Eq. (4) and applying the law of iterated expectations leads to:

$$E(R_{t+1}) = E(\alpha_t) + E(\beta_t' F_{t+1}). \tag{IA.1}$$

If we run an unconditional regression, the unconditional alpha and betas will equal:

$$\alpha^u = E(R_{t+1}) - \beta^u \mu_F \tag{IA.2}$$

$$\beta^u = Cov(R_{t+1}, F_{t+1}') \Sigma_F^{-1}, \tag{IA.3}$$

where $\mu_F = E(F_{t+1})$ and $\Sigma_F = Var(F_{t+1})$. Simple substitution then leads to

$$\beta^u = [E(\beta_t' F_{t+1} F_{t+1}') - E(\beta_t' F_{t+1}) \mu_F'] \Sigma_F^{-1}. \tag{IA.4}$$

Substituting Eqs. (IA.1) and (IA.4) in Eq. (IA.2) provides an expression for the alpha bias:

$$\alpha^u = \underbrace{E(\beta_t' F_{t+1}) \left[1 + \mu_F' \Sigma_F^{-1} \mu_F\right]}_{\text{Factor Timing (I)}} - \underbrace{E(\beta_t' F_{t+1} F_{t+1}') \Sigma_F^{-1} \mu_F}_{\text{Factor (Co-) Variance Timing (I)}}. \tag{IA.5}$$

We can rewrite this decomposition into slightly more familiar terms involving conditional expectations of factor risk premia and risk. Using the law of iterated expectations, we have

$$E(\beta_t' F_{t+1}) = E(\beta_t' \mu_{F,t}) \text{ and} \tag{IA.6}$$

$$E(\beta_t' F_{t+1} F_{t+1}') = E(\beta_t' \Sigma_{F,t}) + E\left(\beta_t' \left[\mu_{F,t} \mu_{F,t}'\right]\right), \tag{IA.7}$$

where $\mu_{F,t} = E_t(F_{t+1})$ and $\Sigma_{F,t} = Var_t(F_{t+1})$. This yields the following expression for the alpha bias:

$$\alpha^u = \underbrace{E(\beta_t' \mu_{F,t}) \left[1 + \mu_F' \Sigma_F^{-1} \mu_F\right]}_{\text{Factor Timing (II)}} - \underbrace{E\left(\beta_t' \left[\mu_{F,t} \mu_{F,t}'\right]\right) \Sigma_F^{-1} \mu_F}_{\text{Small}} - \underbrace{E\left(\beta_t' \Sigma_{F,t}\right) \Sigma_F^{-1} \mu_F}_{\text{Factor (Co-)Variance Timing (II)}} . \tag{IA.8}$$

Eq. (IA.8) is analogous to the single factor representations in Eq. (4) of Lewellen and Nagel (2006) and Eq. (2.4) in Boguth et al. (2011).

B. Omitted variable bias in the price of risk

In this appendix we show that single factor identifications of the market price of risk are negatively biased under the null of the FF5 model, because conditional market variance is negatively correlated with the conditional covariance of the market with various of the long-short factors. For the sake of argument, we assume in this appendix that conditional market variance is observed without error. IID measurement error would yield additional errors-in-variables bias in the estimated market price of risk.

Consider the usual regression of market returns on conditional market variance:

$$MKT_{t+1} = \text{constant} + \gamma_{MKT}\sigma^2_{MKT,t} + \epsilon_{t+1}. \tag{IA.9}$$

Under the null of the CAPM (i.e., an SDF with the market as only factor and a constant risk-price), the price of risk estimate will equal:

$$\widehat{\gamma_{MKT}} = \frac{Cov\left(MKT_{t+1}, \sigma^2_{MKT,t}\right)}{Var\left(\sigma^2_{MKT,t}\right)} = \frac{Cov\left(E_t(MKT_{t+1}), \sigma^2_{MKT,t}\right)}{Var\left(\sigma^2_{MKT,t}\right)} = \gamma_{MKT} \tag{IA.10}$$

However, under the null of the multifactor model, expected market returns follow the first row in Eq. (9):

$$E_t(MKT_{t+1}) = \gamma_{MKT}\sigma^2_{MKT,t} + \gamma_{SMB}\sigma_{MKT,SMB,t} + ... + \gamma_{CMA}\sigma_{MKT,CMA,t}. \tag{IA.11}$$

This means that the single factor price of risk estimate will have the following bias:

$$\widehat{\gamma_{MKT}} - \gamma_{MKT} = \\ = \gamma_{SMB} \frac{Cov\left(\sigma_{MKT,SMB,t}; \sigma^2_{MKT,t}\right)}{Var\left(\sigma^2_{MKT,t}\right)} + \gamma_{HML} \frac{Cov\left(\sigma_{MKT,HML,t}; \sigma^2_{MKT,t}\right)}{Var\left(\sigma^2_{MKT,t}\right)} \\ + \gamma_{RMW} \frac{Cov\left(\sigma_{MKT,RMW,t}; \sigma^2_{MKT,t}\right)}{Var\left(\sigma^2_{MKT,t}\right)} + \gamma_{CMA} \frac{Cov\left(\sigma_{MKT,CMA,t}; \sigma^2_{MKT,t}\right)}{Var\left(\sigma^2_{MKT,t}\right)}. \tag{IA.12}$$

Our empirical estimates indicate that the bias in the single factor setting is large. To see this, we report the magnitude of each of the terms in Table IA.I below. To this end, we use our our estimates of the risk-prices from Eq. (14) together with our empirical estimates of the conditional variances and covariances of the factors.[23] Overall, we estimate the bias to be almost -2, which indicates that the single factor identification underestimates by a large amount the market price of risk under the null of the multifactor model.

[23]Our results are similar when we use the risk-price estimates from the pooled specification. We prefer the time-series regression here, because the bias exactly estimates the difference between the market price of risk in the regression of market returns on market variance versus the regression of market returns on market variance and the four covariances.

Table IA.I: **Estimates of the bias in Eq.** (IA.12)

This table presents our estimate of the bias in a single factor identification of the MKT price of risk under the null of a multifactor model. The total bias is the sum of the contribution of the four factors.

	MKT	SMB	HML	RMW	CMA
Risk-price γ	2.97	7.33	-4.20	10.05	5.56
$\frac{Cov(\sigma^2_{MKT,t},\sigma_{MKT,F,t})}{Var(\sigma^2_{MKT,t})}$	1	-0.012	0.180	-0.078	-0.067
$Corr(\sigma^2_{MKT,t},\sigma_{MKT,F,t})$	1	-0.06	0.47	-0.42	-0.37
Contribution to bias		-0.088	-0.756	-0.784	-0.373
Total bias	-2.00				

C. ICAPM

In this appendix, we relate Eq. (19) to the ICAPM of Merton (1973). The aggregate demand for risky assets in the ICAPM can be written as (see his Eq. (23)):

$$w_{m,t} = \gamma_m^{-1} \Sigma_{rr,t}^{-1} \mu_{r,t} - \sum_{s=1}^{K-1} q_s \Sigma_{rr,t}^{-1} \Sigma_{rs,t}, \tag{IA.13}$$

where γ_m is the continuous time equivalent of the representative agent's risk aversion, formally denoted $-\frac{J_W}{J_{WW}W} > 0$; $\Sigma_{rr,t}$ is the $N \times N$ covariance matrix of the risky assets with expected excess return $\mu_{r,t}$; q_s represents exposure to each of $s = 1 : K - 1$ state variables that drives changes in future consumption-investment opportunities, denoted $-\frac{J_{Ws}}{J_{WW}W}$ in Merton (J_{Ws} is the derivative of the indirect marginal utility of wealth with respect to the state variable s); and, $\Sigma_{rs,t}$ is a vector with the covariances of each of the N assets with state variable s. For the sake of comparison to the rest of our paper, we write the exposures γ_m and q_s as constants, but in general they could also vary over time. An asset's optimal portfolio weight is thus determined by its standard Markowitz demand and the degree to which the asset hedges against changes in investment opportunities. $J_{Ws} > 0 (< 0)$ indicates that an increase in s is bad (good) news that increases (decreases) marginal value of wealth (J_W). Hence, investor wants to invest more (less) in assets that covary positively with s and these assets will therefore have lower (higher) expected returns in equilibrium. To be precise, the equilibrium expected return relation is:

$$\mu_{r,t} = \gamma_m \Sigma_{rm,t} - \sum_{s=1}^{K-1} \gamma_m q_s \Sigma_{rs,t}, \tag{IA.14}$$

such that covariance with both the market and the state variables is priced.

To understand the link to the portfolios that feature in empirical factor models, let us re-scale the hedging weights to sum to one for each risk factor $s = 1 : K - 1$:

$$H_{s,t} = (\iota' \Sigma_{rr,t}^{-1} \Sigma_{rs,t})^{-1} \Sigma_{rr,t}^{-1} \Sigma_{rs,t} \tag{IA.15}$$

$$z_{s,t} = (\iota' \Sigma_{rr,t}^{-1} \Sigma_{rs,t}) q_s, \tag{IA.16}$$

so that $H_{s,t}$ is the hedge portfolio for risk s. We can then re-write Eq. (IA.13) as follows:

$$w_{m,t} = \gamma_m^{-1} \Sigma_{rr,t}^{-1} \mu_{r,t} - H_t z_t,$$

where H is $N \times K - 1$ and contains each of the $H_{s,t}$ terms in Eq. (IA.15) and z_t a $K - 1$-vector with the $z_{s,t}$ terms from Eq. (IA.16). Simple re-writing leads to

$$\mu_{r,t} = \Sigma_{rm,t} \gamma_m + \Sigma_{rH,t} z_t \gamma_m. \tag{IA.17}$$

Note that Eq. (IA.17) holds for any asset and if we substitute the market and hedge portfolios we find:

$$\begin{pmatrix} \mu_{m,t} \\ \mu_{h_1,t} \\ \ldots \\ \mu_{h_{K-1},t} \end{pmatrix} = \begin{pmatrix} \sigma_{m,t}^2, \sigma_{m,h_1,t}, \ldots, \sigma_{m,h_{K-1},t} \\ \sigma_{m,h_1,t}, \sigma_{h_1,t}^2, \ldots, \sigma_{h_1,h_{K-1},t} \\ \ldots \\ \sigma_{h_{K-1},m,t}, \sigma_{h_{K-1},h_1,t}, \ldots, \sigma_{h_{K-1},t}^2 \end{pmatrix} \begin{pmatrix} \gamma_m \\ z_{1,t}\gamma_m \\ \ldots \\ z_{K-1,t}\gamma_m \end{pmatrix} \tag{IA.18}$$

where the first term on the right hand side represents the variance-covariance matrix of the market and hedge portfolios and these are multiplied by a vector of risk prices. For the market, the risk-price γ_m measures risk aversion. For the state variable hedge portfolios, the risk-price $z_{s,t}\gamma_m$ is a scaled version of the exposure to risk s, q_s. Assuming the scaling does not vary over time, these risk-prices are also constant. Note that this equation will hold for any linear combination of the hedge portfolios. This is exactly the interpretation that Fama and French (2015) advocate for their model: the factors in FF5 represent diversified portfolios that provide different combinations of exposures to the unknown state variables that govern time-variation in consumption-investment opportunities in Merton's ICAPM. Under that interpretation, the risk-prices we estimate empirically are linear combinations of the scaled exposures to each risk s and represent the compensation that investors require for exposure to a particular combination of state variables.

D. Volatility-timing

Another way to see the importance of the covariance terms is by estimating their contribution to the attractiveness of a volatility-managed strategy that invests a time-varying amount w_t in the market. The return of this strategy equals:

$$MKT_{t+1}^{VT} = w_t MKT_{t+1}, \text{ where} \tag{IA.19}$$

$$w_t = \frac{c}{\widehat{\sigma^2_{MKT,t})}}, \tag{IA.20}$$

and we set the scalar c such that the beta in a regression of the timed strategy to the original strategy equals one. As noted in Moreira and Muir (2017), the improvement in Sharpe ratio from having access to the timed strategy equals: $\sqrt{\text{SR}^2_{Old} + AR^2} - \text{SR}_{Old}$, where the appraisal ratio, $AR = \frac{\alpha}{\sigma_\epsilon}$. We have considered realized variance estimated over both a 1 and 6 month window to define w_t. Since we find that the improvement is virtually zero when using 6 month realized variance, we focus on the 1 month measure in the following.[24] We will also show results using 1 month realized volatility, rather than variance.

We find that the Sharpe ratio improvement (in annualized terms) for the strategy in Eq. (IA.20) equals 0.045 over our sample from 1970 to 2023. The alpha in the regression of MKT^{VT}_{t+1} on MKT_{t+1} equals 3.2% ($t = 1.5$). To understand where this alpha comes from, we use that the beta of MKT^{VT}_{t+1} on MKT_{t+1} is normalized to 1. The alpha is defined as:

$$\alpha = E(w_t MKT_{t+1}) - E(MKT_{t+1}) = E((w_t - 1)MKT_{t+1}) = E((w_t - 1)E_t(MKT_{t+1})). \tag{IA.21}$$

We can plug in our pooled regression estimates to approximate this alpha as:

$$\widehat{\alpha} = E((w_t - 1)(A_{MKT} + \sigma^2_{MKT,t}\gamma_{MKT} + \sigma_{MKT,SMB,t}\gamma_{SMB} + \ldots + \sigma_{MKT,CMA,t}\gamma_{CMA})), \tag{IA.22}$$

[24] The reason is that the coefficient in the regression of market returns on 6 month realized variance is positive (0.54, t-stat=0.50), rather than negative as it is for 1 month realized variance (-1, t-stat=-1.74). As a result, volatility-timing using the 6 month measure increases (decreases) risk exactly when expected returns are high (low).

such that we split the conditionally expected market return in three components: the intercept A_{MKT}, the part due to market variance $\sigma^2_{MKT,t}\gamma_{MKT}$ and the remainder due to the market's covariance with the other factors. This is only an approximation because the residual market return from the pooled regression may be correlated to conditional market variance (as well as the covariances) as a result of the orthogonality restrictions being imposed in the pool. For the market return, the approximation error is quite small though and we show below that the results are similar if we estimate $E_t(MKT_{t+1}))$ using a time-series regression.

We calculate the average of the product of each of these three terms with $(w_t - 1)$ in the data and report them in the first row of Table IA.II below. First, we find that the contribution of the intercept term to α is large at +2.6%. The reason is that MKT^{VT}_{t+1} invest on average more than 1 unit in the market portfolio.[25] Second, the contribution of conditional market variance is negative at -0.6%. The reason is that 6 month variance predicts returns with a positive sign. Since 6 month variance is positively correlated with 1 month variance, we have that the volatility-managed strategy increases exposure to the market exactly when expected market returns are lower. Third, the contribution of the covariance with the other four factors is large and positive at 1.4%. Combining, we have that about 43% (=1.4%/3.2%) of the alpha of a volatility-managed market portfolio is due covariance with the other factors. In the third row of estimates in the table, we see that this contribution increases to 58% when we estimate $E_t(MKT_{t+1})$ using a time-series regression of the

[25]Note that $1 = \frac{Cov(w_t MKT_{t+1}, MKT_{t+1})}{Var(MKT_{t+1})}$ and also $Cov(w_t MKT_{t+1}, MKT_{t+1}) = E(w_t)Var(MKT_{t+1}) + E(MKT_{t+1})Cov(w_t MKT_{t+1}, MKT_{t+1})$. Since our pooled regression indicates that past variance predicts returns with a positive sign, we have that $Cov(w_t MKT_{t+1}, MKT_{t+1}) < 0$. Hence, $E(w_t)$ must be >1.

market on its own variance and the covariances with the other factors. In the second and fourth row, we see similar effects when we time using 1 month volatility rather than variance. We conclude that an important reason why a volatility-managed market strategy seems attractive is that it increases market exposure exactly when the market covaries more with other factors, and this covariance is compensated in average returns. This analysis shows that the omitted variables bias discussed in Section 2 is practically relevant.

TABLE IA.II: **Decomposing the Performance of Volatility-Timing**
We decompose the alpha in a regression of the returns of a timed market strategy, where the time-varying weight in the market is defined as in Eq. (IA.20), on the market. We consider both 1 month variance and volatility to define the weight. We decompose the alpha in three parts following Eqs. (IA.21) and (IA.22): the contribution of the intercept in the predictive regression for the market, the contribution of market variance and the contribution of the covariances between the market and the other factors. These contributions are based on either the pooled or time-series regression estimates of the factor risk-prices (see Panels A and B of Table III).

	Alpha (MKT^{VT}_{t+1} on MKT_{t+1})	Contribution			Fraction
		Intercept	Market variance	Covariances	(Covariances/Alpha)
		Pooled regression			
Variance	3.234	2.611	-0.553	1.389	0.43
Volatility	3.057	2.445	-0.703	1.426	0.47
		Time-series regression			
Variance	3.234	1.625	-0.753	1.865	0.58
Volatility	3.057	1.522	-0.957	1.898	0.62

E. Robustness checks for pooled risk-return trade-off

We report in Table IA.III results for a large variety of robustness checks. First, we show that our risk-price estimates are largely insensitive to removing the factor-specific fixed effects. This result is important, because pooled specifications without fixed effects are less subject to concerns about Stambaugh bias (Hjalmarsson, 2008).[26] Second, we ask how sensitive our results are to the window length used to estimate conditional risk. For all window lengths ≥ 3 as well as exponential weighting, we find virtually identical risk prices (in terms of magnitude and significance) as in our main specification using six month risk measures. Only if we estimate past variances and covariances over the shortest one month window do our results weaken somewhat. In that case, the market price of risk is insignificant at 0.64 (t = 0.96), for instance. This result is in line with findings from the volatility-timing literature, but in stark contrast with all our other specifications. The short-term market risk-return tradeoff also stands out relative to the other factors. Indeed, although the risk-prices for SMB and RMW drop as well when we estimate risk over a one month window, they remain economically large and marginally significant at 2.61 (t = 2.23) and 5.30 (t = 2.66). Overall, we conclude that conditional factor risk is positively related to factor risk premia, regardless of the length of the window over which we estimate risk. Given

[26]This concern is further alleviated noting that monthly factor returns are contemporaneously largely uncorrelated to innovations in variances and covariances. Perhaps unsurprisingly, the largest (in absolute value) time-series correlation between any factor return and an innovation in a variance or covariance is -0.27 for the market return and market variance. All remaining correlations are substantially smaller.

that this risk-return tradeoff is noticeably stronger when risk is measured over longer windows than one month, we also conclude that long-short anomaly portfolios' betas comove more with low frequency variation in factor risk (and thus factor risk premia). Indeed, such comovement is necessary to generate large factor timing and factor (co-) variance timing terms that cancel each other out. This conclusion is consistent with the idea that short-term factor risk is more transitory than betas (see Keloharju et al., 2021; Baba-Yara et al., 2024, for related discussions on the transitory and persistent components of firm characteristics).

Third, we study factor returns over longer windows than 1 month. We find that the risk-return tradeoff is strongest at the three month window, and weakens thereafter. This finding is consistent with the fact that variances and covariances are not as persistent as prominent predictor variables like the dividend-to-price ratio or bond yields.

Fourth, we ask how sensitive our results are to the most extreme observations in our sample. To this end, we winsorize all risk measures (both variances and covariances) at the 1%- and 5%-level. In short, we see that this has little impact on our results.

Finally, we add test portfolios to our estimation of the risk-return tradeoff. To this end, we compute the covariance of each added test portfolio with the five factors over a 6 month window of daily returns. We then add these test portfolio as independent variables in Eq. (13). We estimate the same five risk prices as before, but now also include a fixed effect for each portfolio. We consider three sets of portfolios: 25 portfolios sorted on size and book-to-market, 30 industry portfolios, and finally the combination of these 55 portfolios with another set of 50 portfolios sorted on size and, respectively, profitability and investment. We find that the estimated prices of risk are similar to when we do not include any additional test assets. In fact, the standard errors on the price of risk estimates do not change much either. This result is consistent with the idea that the factors are dominant sources of time-variation in the returns of these diverse sets of portfolios. Consequently, the time-variation in the returns of the factors is mirrored in the portfolio returns, such that they represent the exact same positive multifactor risk-return tradeoff.

F. Certainty Equivalents

We calculate certainty equivalent returns to assess at which risk aversion level an investor would find it useful to condition on information in current variances and co-variances. The benchmark is an investor that uses only information on the historical average and variances-covariances of the factor returns, consistent with the standard definition of an OOS R^2.

For an investor that desires to time an individual factor, the optimal weight in the factor is given by:

$$w_{1,k,t} = P(F_{k,t+1}|t)/(\gamma Var_t(F_{k,t+1})) \quad \text{(IA.23)}$$

$$w_{2,k,t} = \overline{F_{k,t}}/(\gamma \overline{Var(F_{k,t})}) \quad \text{(IA.24)}$$

where $P(F_{k,t+1}|t)$ is the predicted return for the k-th factor in month $t+1$ based on the variance-covariance matrix observed at the end of t. This matrix is then suitably

Table IA.III: Risk-return tradeoff: robustness

In this table we regress next period's factor returns on realized variances and covariances in a pooled panel regression. We report the five price of risk estimates with corresponding Driscoll-Kraay t-statistics below. In Panel A, we remove the factor fixed effects and either estimate the pooled regression without an intercept or with a common intercept. In Panel B, we consider realized variances and covariances over alternative windows of 1, 3, 9 and 12 months as well as exponentially weighted risk measures (with $\lambda = 0.9945$). In Panel C, we regress cumulative returns over 3, 6 and 12 month windows on the factors on the realized (6 month) variances and covariances. In Panel D, we winsorize the risk measures to reduce the impact of outliers. In Panel E, the pooled regression is expanded to include test portfolios. In this case, we include additional fixed effects for each portfolio, although we are still estimating the same five risk-prices as before.

	γ_{MKT}	γ_{SMB}	γ_{HML}	γ_{RMW}	γ_{CMA}	R^2
			Panel A: Removing factor fixed effects			
F_{t+1} on $\Sigma_{F,t,6}$, no intercept	2.92	4.63	-2.67	7.63	7.85	0.015
	4.79	3.01	-1.68	3.78	2.80	
F_{t+1} on $\Sigma_{F,t,6}$, common intercept	2.47	3.75	-2.30	7.30	6.07	0.015
	4.26	2.26	-1.11	2.75	1.61	
			Panel B: Conditional risk estimated over different windows			
F_{t+1} on $\Sigma_{F,t}$	0.64	2.61	-1.22	5.30	5.05	0.014
	0.96	2.23	-0.84	2.66	1.38	
F_{t+1} on $\Sigma_{F,t,3}$	1.45	3.96	-2.44	6.96	5.62	0.012
	1.99	2.55	-1.13	3.05	1.37	
F_{t+1} on $\Sigma_{F,t,9}$	2.53	3.78	-1.41	9.12	4.03	0.013
	3.52	1.71	-0.66	2.81	0.89	
F_{t+1} on $\Sigma_{F,t,12}$	2.23	3.04	-0.28	8.00	2.51	0.010
	2.68	1.28	-0.13	2.36	0.52	
F_{t+1} on $\Sigma_{F,t,EW}$	2.95	3.94	-1.82	9.15	6.72	0.012
	2.79	1.36	-0.64	2.41	1.15	
			Panel C: Long-term return predictability			
$F_{t+1:t+3}$ on $\Sigma_{F,t,6}$	2.01	3.63	-1.33	7.49	3.90	0.032
	4.07	2.17	-0.91	2.75	1.06	
$F_{t+1:t+6}$ on $\Sigma_{F,t,6}$	1.45	2.11	0.46	6.85	1.29	0.060
	2.80	1.51	0.41	2.61	0.32	
$F_{t+1:t+9}$ on $\Sigma_{F,t,6}$	1.14	1.53	1.09	5.68	0.49	0.074
	2.33	1.32	1.11	2.26	0.12	
$F_{t+1:t+12}$ on $\Sigma_{F,t,6}$	1.04	1.54	0.87	4.85	1.00	0.077
	2.46	1.41	1.08	2.13	0.28	
			Panel D: Winsorizing $\Sigma_{F,t,6}$			
[2.5% – 97.5%]	2.77	5.05	-2.38	9.08	7.89	0.010
	3.23	1.98	-0.75	2.33	1.35	
[0.5% – 99.5%]	2.22	4.79	-2.32	7.61	5.89	0.012
	3.54	2.14	-0.96	2.41	1.38	
			Panel E: Adding test assets			
25 Size vs. Book-to-Market	2.75	6.43	-3.95	7.46	5.84	0.012
	2.87	1.65	-1.11	1.75	1.11	
30 Industry	2.74	6.66	-3.35	6.16	5.24	0.010
	3.85	1.77	-0.91	1.80	1.03	
30 Industry + 3×25 Size vs.	2.21	4.88	-2.39	7.74	5.62	0.012
Book-to-Market, Profitability, and Investment	3.61	2.22	-1.01	2.47	1.35	

combined with the pooled regression coefficients estimated with factor return data up to and including month t. $\overline{F_{k,t+1}}$ is the historical average return of the kth factor up to and including month t. The conditional variance of the factor in the denominator is either based on the last 6 months of daily returns, $Var_t(F_{k,t+1})$, or based on the full history of monthly returns up to time t, $\overline{Var(F_{k,t})}$.

Then we define the portfolio returns:

$$R_{1,k,t+1} = w_{1,t}F_{k,t+1} + R_{F,t} \tag{IA.25}$$

$$R_{2,k,t+1} = w_{2,t}F_{k,t+1} + R_{F,t} \tag{IA.26}$$

where we have used that all factors are excess returns so the portfolio return contains a full investment in the risk-free asset. We then calculate the CE improvement over the full out-of-sample period as:

$$\Delta_{CE} = Average(R_{1,k,t+1} - R_{2,k,t+1}) - \gamma/2(Var(R_{1,k,t+1}) - Var(R_{2,k,t+1})) \tag{IA.27}$$

The CE improvement can be interpreted as the annual fee that an investor would be willing to pay to have access to information about current variances and covariances relative to longer historical averages and variances and covariances. A caveat is that the optimal weights in factors are sometimes very large, which is a common issue in a setting with long-short factors. To see why, just note that the unconditional variance of factors like HML and RMW is about a third of the variance of the market. Therefore, we restrict the conditional weight to be between -1 and +3, which is generous given that leveraging up so much can be quite costly. Indeed, restricting the weights makes conditioning on current information in variances and covarianes relatively less attractive.

We also consider an investor that aims to optimally combine the factors. For such an investor, the optimal weights are defined as:

$$w_{1,t} = \gamma^{-1} Var_t(F_{k,t+1})^{-1} P(F_{t+1}|t) \tag{IA.28}$$

$$w_{2,t} = \gamma^{-1} \overline{Var(F_{k,t})}^{-1} \overline{F_t} \tag{IA.29}$$

Using our out-of-sample forecast combination II to forecast factor returns $P(F_{k,t+1}|t)$, we find in Table IA.IV that the CE improvements from using information in current variances and covariances is typically not large.[27] For an investor with moderate risk aversion of three, the CE improvement is sizeable for MKT and RMW at 3 and 5%, respectively, but it is small for HML and CMA, and even negative for SMB. For both lower and higher risk aversion levels, the CE improvements for MKT and RMW are smaller, while those for HML and/or CMA turn negative. For the optimal combination of the factors we find that the certainty equivalent even deteriorates from timing based on current variances and covariances ($Delta_{CE} \approx -4\%$ for all risk aversion levels). Overall, we conclude that timing based on our evidence is likely unattractive for most investors, consistent with the idea that risk increases roughly commensurate with return.

[27] Our results are virtually identical using instead forecast combination I or forecasts derived from our Ridge specification.

Table IA.IV: **Certainty equivalents**

We report in this table certainty equivalents for mean-variance investors with risk aversion $\gamma = 1, 3, 10$. CE_1 is for the investor that uses information in current variances and covariances to choose his/her investment in each factor or their optimal combination, that is, 6-month conditional variances/covariances are used in the denominator and the combination forecast derived from these variances and covariances is used in the numerator of the optimal weights in Eqs. (IA.23) and IA.28. CE_2 is for the investor that uses historical averages as well as historical variances and covariances to choose his/her investment in each factor or their optimal combination (see Eqs. (IA.24) and IA.29). The improvement Δ_{CE} = is simply the difference between the two certainty equivalents.

	MKT	SMB	HML	RMW	CMA	Optimal portfolio
			Risk aversion $\gamma = 1$			
CE_1	0.187	-0.019	0.007	0.130	0.060	0.311
CE_2	0.154	0.016	0.005	0.116	0.071	0.350
Δ_{CE}	0.032	-0.035	0.002	0.014	-0.011	-0.039
			Risk aversion $\gamma = 3$			
CE_1	0.093	-0.006	-0.005	0.081	0.022	0.111
CE_2	0.066	0.020	-0.008	0.037	0.019	0.149
Δ_{CE}	0.027	-0.026	0.003	0.044	0.003	-0.038
			Risk aversion $\gamma = 10$			
CE_1	0.045	0.014	-0.001	0.040	-0.003	-0.065
CE_2	0.035	0.021	0.013	0.026	0.018	-0.027
Δ_{CE}	0.010	-0.008	-0.014	0.014	-0.021	-0.039

G. Out-of-sample test for RP-PCA factors

We first construct RP-PCA factors over expanding windows. The first expanding window runs until December 1994, as in Section 2.3. In each window ending in month t, we then extract the PCs exactly as described in Section 3, but now using returns of the long-short portfolios from the start of our sample up to month t. Over the same window, we calculate the mean-variance optimal weights w_t in the five PC factors (i.e., the market plus the first four PCs of long-short portfolios). Using these weights, we calculate the month $t+1$ return of the Tangency portfolio, $F'_{RP-PCA,t+1}w_t$. The conditional variance of this portfolio, denoted $Var_t(F'_{RP-PCA,t+1}w_t)$, is estimated as the variance of daily returns of the optimal combination of the PC factors from the beginning of month $t-5$ to the end of month t. We thus collect the return and the conditional variance of the Tangency portfolio from January 1995 to December 2023.

Next, we regress return on variance, but now over expanding windows starting in January 2006. These regressions ask whether the risk-return tradeoff has been stable and positive over time, even when we use no forward-looking information. We report the coefficient from this regression in Figure IA.1. In short, the coefficient has been positive over the whole out-of-sample period and averages to around 15, which is only slightly smaller than in our full sample test (see Table V). We next calculate the out-of-sample R^2, analogous to Eq. (16) in the paper:

$$1 - \frac{\sum_t (F_{k,t+1} - P(F_{k,t+1}|t))}{\sum_t (F_{k,t+1} - \overline{F_{k,t}})}. \tag{IA.30}$$

Here, $P(F_{k,t+1}|t)$ is the predicted return of the optimal combination of the PC factors

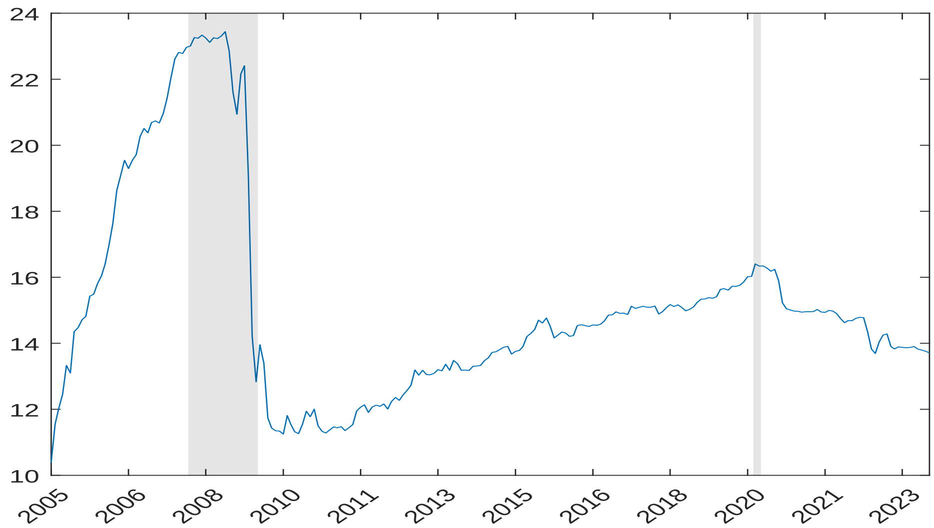

FIGURE IA.1: **Risk-return trade-off of RP-PCA factors (no look-ahead bias)** This figure plots the coefficient γ in an expanding window regression of monthly returns on conditional variance for the mean-variance optimal combination of the RP-PCA factors: $F'_{RP-PCA,s+1}w_s = a + \gamma Var_s(F'_{RP-PCA,s+1}w_s) + \epsilon_{s+1}$ for $s = 1 : t$, where t covers all months from January 2006 onwards. In each window, both the RP-PCA factors as well as the weights to generate their optimal combination are estimated using data observable at the end of month t.

for month $t+1$ derived from the expanding window regression for which the coefficient on conditional variance is reported in Figure IA.1. $\overline{F_{k,t}}$ is the historical mean of the optimal combination up to and including month t. This out-of-sample R^2 equals 3.2%. This fact confirms that the positive multifactor risk-return tradeoff for PC factors is also robust in out-of-sample test, just like it is for the FF5 factors in Section 2.

H. Additional tables and figures

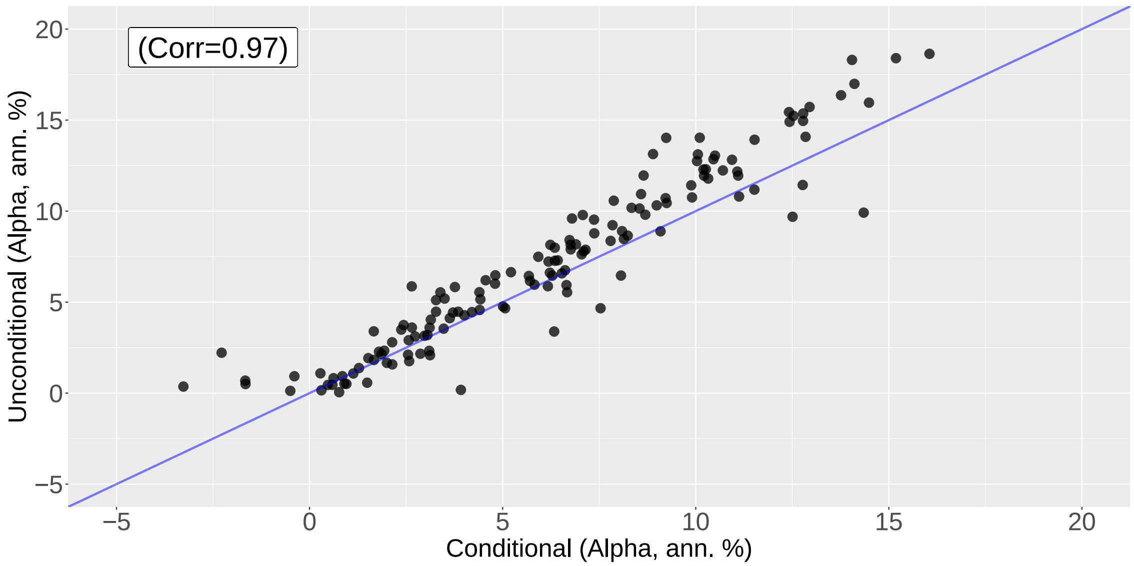

(A) CAPM alphas

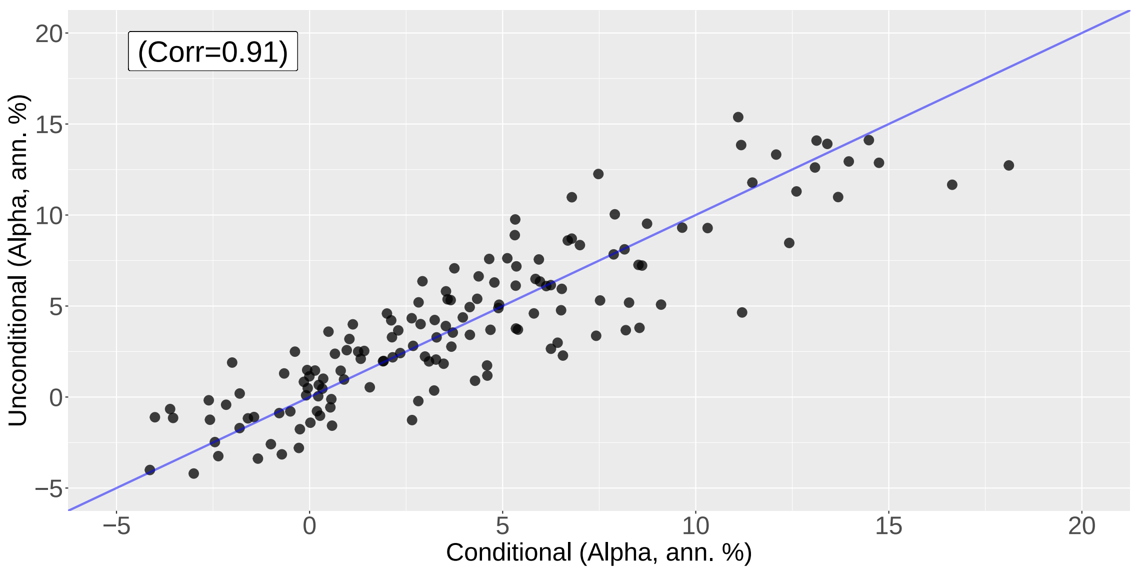

(B) FF5 alphas

FIGURE IA.2: **Unconditional vs Conditional Alphas (Six Month Window)**
This figure is similar to Figure 1 in the main text. The lagged component beta is estimated using a six-month window of daily firm-level returns.

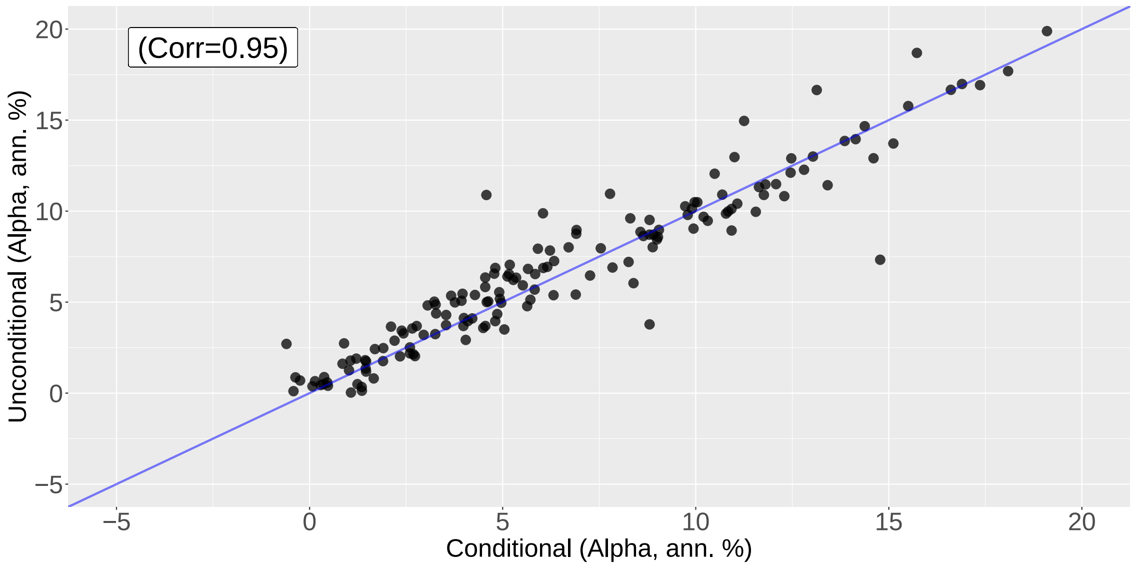

(A) CAPM alphas

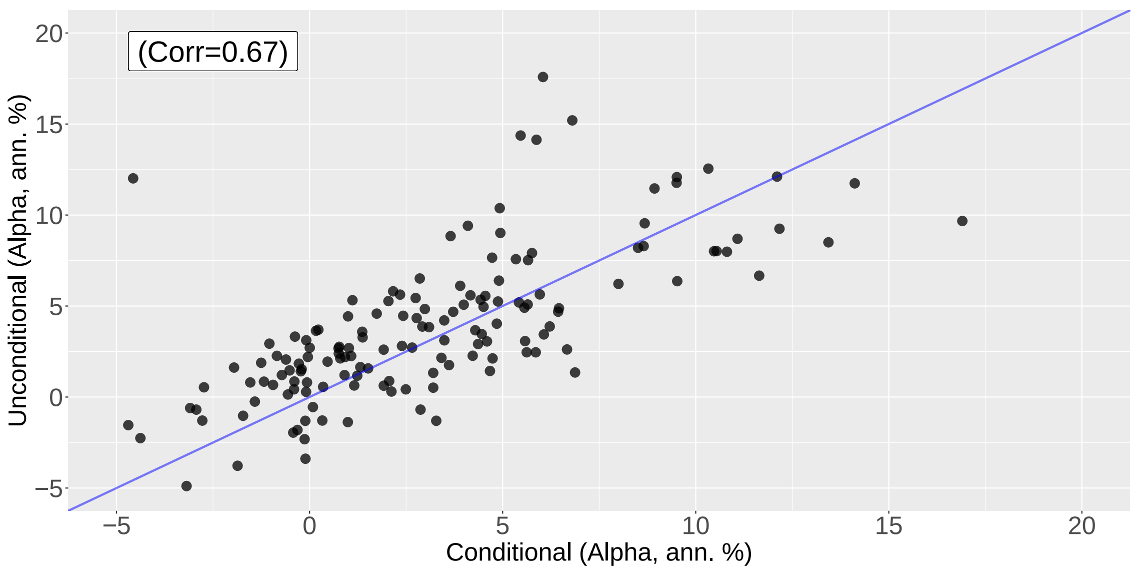

(B) FF5 alphas

FIGURE IA.3: **Unconditional vs Conditional Alphas (Five Year Window)**
This figure is similar to Figure 1 in the main text. The lagged component beta is estimated using a five-year windows of monthly firm-level returns.

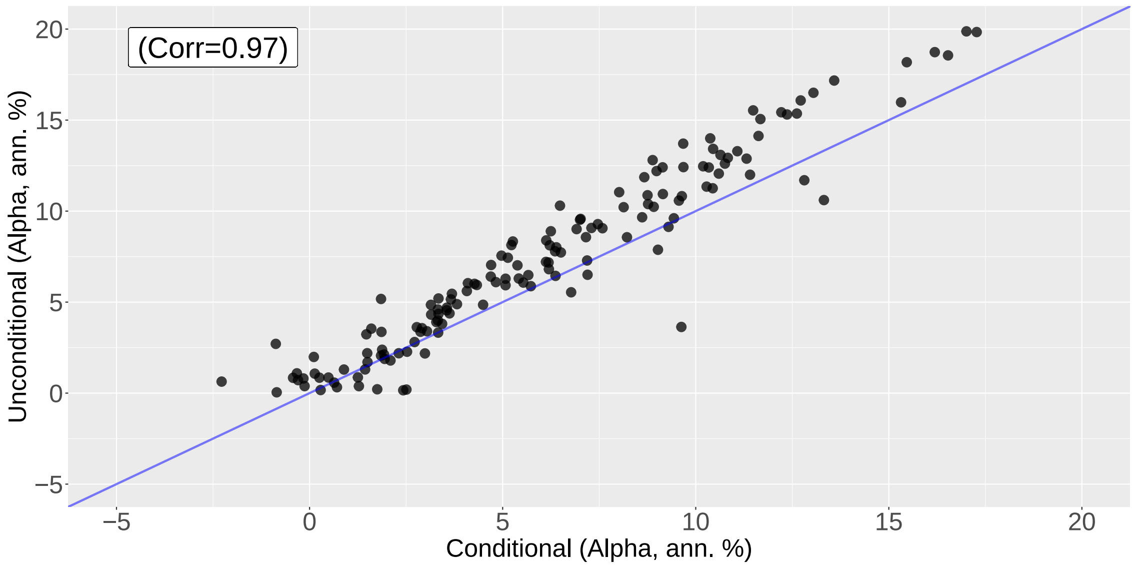

(A) CAPM alphas

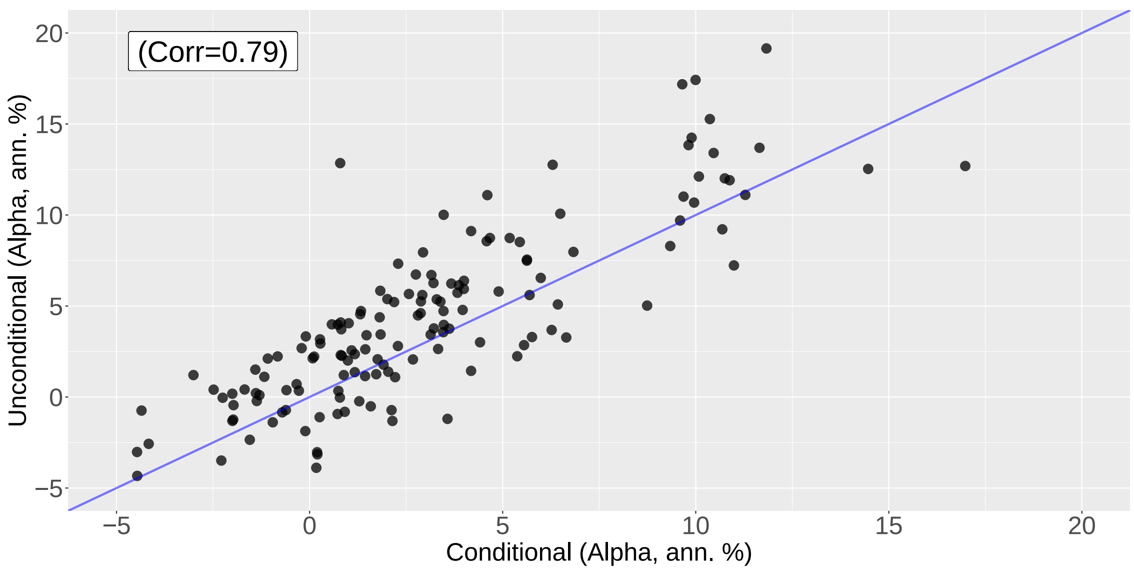

(B) FF5 alphas

FIGURE IA.4: **Unconditional vs Conditional Alphas (contemporaneous)**
This figure is similar to Figure 1 in the main text. The beta component is estimated using a contemporaneous daily firm-level returns.

Table IA.V: **Predicting variances and covariances**

In this table we run predictive regressions of next month's variances and covariances - estimated using one month of daily observations - on their lagged values estimated over alternative windows. We estimate a pooled panel regression for all variance and covariance terms between each of the risk factors in our model. We include fixed effects for each element in the variance-covariance matrix that we consider. In the first pair of rows we consider all elements, in the second pair of rows only the diagonal terms (variances) and in the final set only the off-diagonal elements (covariances). We report corresponding t-statistics below the parameter estimates.

		$\sigma^2_{k,j,t,6}$	R^2	$\sigma_{k,j,t,6}$	$\sigma_{k,j,t}$	R^2
		Panel A: Three-month lookback				
$\sigma_{k,j,t+1}$		0.59	0.22	0.28	0.32	0.26
		6.64		3.91	2.44	
$\sigma^2_{k,t+1}$	$\forall k$	0.55	0.19	0.22	0.34	0.23
		5.30		2.84	2.16	
$\sigma_{k,j,t+1}$	$\forall k \neq j$	0.74	0.38	0.59	0.15	0.39
		16.37		6.44	1.98	
		Panel B: Nine-month lookback				
$\sigma_{k,j,t+1}$		0.62	0.14	0.26	0.40	0.25
		5.01		3.55	3.35	
$\sigma^2_{k,t+1}$	$\forall k$	0.57	0.11	0.22	0.41	0.23
		3.89		3.01	2.90	
$\sigma_{k,j,t+1}$	$\forall k \neq j$	0.76	0.29	0.42	0.37	0.37
		7.63		4.66	6.45	
		Panel C: Twelve-month lookback				
$\sigma_{k,j,t+1}$		0.61	0.12	0.25	0.42	0.25
		4.51		3.30	3.55	
$\sigma^2_{k,t+1}$	$\forall k$	0.56	0.09	0.21	0.42	0.23
		3.36		2.81	3.06	
$\sigma_{k,j,t+1}$	$\forall k \neq j$	0.74	0.25	0.38	0.41	0.36
		6.58		4.53	7.57	
		Panel D: Exponential weighting ($\lambda = 0.9945$)				
$\sigma_{k,j,t+1}$		0.83	0.16	0.36	0.39	0.26
		5.29		3.91	3.15	
$\sigma^2_{k,t+1}$	$\forall k$	0.79	0.13	0.31	0.39	0.23
		4.36		3.17	2.74	
$\sigma_{k,j,t+1}$	$\forall k \neq j$	0.92	0.29	0.51	0.36	0.37
		7.80		4.95	6.31	

Table IA.VI: **Out-of-Sample R^2s with different economic constraint**

This table is similar to Table IV, but now we present the out-of-sample R^2s for three approaches when we constrain the forecasts to be between the average return plus and minus 2.33 standard errors (where both the average and standard error are estimated over the expanding window).

	MKT	SMB	HML	RMW	CMA	Optimal portfolio
Simple	-0.0004	-0.0174	-0.0014	-0.0015	-0.0018	0.0006
Combination I	0.0048	-0.0022	0.0040	0.0031	0.0037	0.0073
Combination II	0.0057	-0.0072	0.0050	0.0046	0.0042	0.0089

TABLE IA.VII: Time-variation in the multifactor risk-return tradeoff: Robustness checks

This table presents robustness checks for the evidence presented in Table VI of the paper. In turn, we present results when we set the factor fixed effects A_F in Eqs. (20) and (21) to zero (Panel A) and when we estimate the variance-covariance matrix of the FF5 factors using 1 month of daily data (with factor fixed effects in Panel B and without factor fixed effects in Panel C). In Panel D, we present results when we shrink the six month rolling variances and covariances using the approach of Ledoit and Wolf (2004). t-statistics use Driscoll and Kraay (1998) standard errors with 6 lags.

	Γ_0 MKT	SMB	HML	RMW	CMA	Γ_1 MKT	SMB	HML	RMW	CMA
Panel A: Six-month variances and covariances without factor fixed effects										
Demeaned inverse realized market volatility										
Estimate	2.67	1.45	0.18	5.59	6.75	-0.04	-0.30	0.21	-0.19	-0.02
t-stat	(2.57)	(0.61)	(0.05)	(1.27)	(1.07)	(-0.55)	(-2.32)	(1.21)	(-0.71)	(-0.05)
Demeaned inverse realized market volatility dummy										
Estimate	3.17	5.79	-2.89	8.56	7.92	-2.15	-10.93	6.10	-13.05	-13.10
t-stat	(5.06)	(3.60)	(-1.38)	(3.22)	(2.07)	(-0.97)	(-2.42)	(1.48)	(-1.96)	(-1.43)
Panel B: One-month variances and covariances										
Demeaned inverse realized market volatility										
Estimate	0.52	-1.41	0.52	-1.75	7.44	-0.01	-0.20	0.07	-0.41	0.18
t-stat	(0.18)	(-0.43)	(0.16)	(-0.32)	(0.91)	(-0.06)	(-1.49)	(0.45)	(-1.45)	(0.40)
Demeaned inverse realized market volatility dummy										
Estimate	0.65	2.78	-1.19	5.44	5.12	-4.01	-7.32	3.34	-11.79	-17.59
t-stat	(0.98)	(2.43)	(-0.83)	(2.74)	(1.41)	(-0.93)	(-1.42)	(0.54)	(-1.46)	(-1.23)
Panel C: One-month variances and covariances without factor fixed effects										
Demeaned inverse realized market volatility										
Estimate	5.05	-0.49	-0.46	1.02	10.91	0.17	-0.16	0.03	-0.31	0.29
t-stat	(3.08)	(-0.21)	(-0.14)	(0.21)	(1.50)	(2.10)	(-1.57)	(0.21)	(-1.19)	(0.71)
Demeaned inverse realized market volatility dummy										
Estimate	1.24	3.15	-1.53	5.71	5.88	5.41	-8.58	3.76	-11.03	-10.00
t-stat	(2.03)	(3.00)	(-1.13)	(3.05)	(1.70)	(1.55)	(-1.83)	(0.66)	(-1.42)	(-0.76)
Panel D: Six-month variances and covariances with Ledoit and Wolf (2004) shrinkage										
Demeaned inverse realized market volatility										
Estimate	0.85	1.28	1.53	5.03	3.37	-0.09	-0.27	0.24	-0.18	-0.13
t-stat	(0.76)	(0.41)	(0.44)	(1.04)	(0.51)	(-1.34)	(-2.01)	(1.52)	(-0.70)	(-0.35)
Demeaned inverse realized market volatility dummy										
Estimate	2.39	5.66	-2.25	8.46	5.53	-3.28	-9.15	6.22	-12.82	-18.01
t-stat	(3.62)	(2.89)	(-1.01)	(2.95)	(1.43)	(-1.64)	(-2.12)	(1.68)	(-1.95)	(-2.14)

Table IA.VIII: **Varying the Lettau and Pelger (2020b) weighting parameter**

This table is identical to Table V of the paper, except that we now consider alternative weighting parameters $y = -1, 5, 15$. Note that $y = -1$ is equivalent to standard PCA, which ignores cross-sectional pricing errors and extracts orthogonal factors from the variance-covariance matrix.

Panel A: PC returns on PC conditional (co-)variances (Pooled)						
	PC1=MKT	PC2	PC3	PC4	PC5	R^2
			$y = -1$			
Estimate	4.07	0.63	6.56	-0.31	14.35	0.03
t-stat	(3.21)	(0.35)	(2.30)	(-0.07)	(1.76)	
			$y = 5$			
Estimate	2.56	0.16	8.29	10.41	19.82	0.03
t-stat	(2.01)	(0.07)	(2.16)	(1.23)	(1.99)	
			$y = 15$			
Estimate	1.95	3.72	9.14	18.68	4.68	0.03
t-stat	(1.30)	(3.19)	(1.95)	(1.43)	(1.26)	

Panel B: Optimal combination of K PCs on conditional variance (Time-series)					
K =	3	4	5	6	7
			$y = -1$		
Estimate	0.90	-0.33	2.20	7.75	9.42
t-stat	(0.60)	(-0.49)	(1.18)	(2.92)	(2.95)
			$y = 5$		
Estimate	1.19	3.65	10.65	16.85	18.73
t-stat	(0.78)	(1.84)	(1.92)	(2.44)	(2.28)
			$y = 15$		
Estimate	1.99	23.22	18.78	19.21	20.69
t-stat	(1.16)	(3.14)	(1.95)	(1.97)	(1.81)

Table IA.IX: **The Risk-Return Tradeoff for Principal Component Factors: 1 Month (Co-)Variances**

This table is identical to Table V in the paper, except that we now use 1 month realized (co-)variances, that is, variances and covariances estimated over a lagged one month window of daily returns rather than the longer 6 month window used in our main tests.

Panel A: PC factor returns on PC conditional (co-)variances (Pooled)						
	PC1=MKT	PC2	PC3	PC4	PC5	R^2
Estimate	0.31	1.63	8.27	20.94	6.15	0.035
t-stat	(0.25)	(1.54)	(2.13)	(1.94)	(1.27)	
	PC1=MKT	PC2	PC3	PC4	PC5	
Partial R^2	0.028	0.043	0.020	-0.010	0.013	

Panel B: FF5 returns on conditional covariances with PC factors (Pooled)						
	PC1=MKT	PC2	PC3	PC4	PC5	R^2
Estimate	1.09	2.62	8.03	17.04	10.10	0.024
t-stat	(0.97)	(2.47)	(2.24)	(1.66)	(2.12)	
	MKT	SMB	HML	RMW	CMA	
Partial R^2	0.028	0.006	0.023	0.018	0.031	

Panel C: Optimal combination of K PCs on conditional variance (Time-series)					
K =	3	4	5	6	7
Estimate	0.77	12.05	12.93	13.65	15.74
t-stat	(0.69)	(1.89)	(1.42)	(1.41)	(1.37)
R^2	0.000	0.025	0.013	0.013	0.014

Table IA.X: **Principal Component Factors Including Momentum Characteristics**
This table is identical to Table V of the paper, except that we now add 8 characteristics from the momentum category to the set of long-short characteristic-sorted portfolios (thus expanding the set of from 134 to 142 characteristics).

Panel A: PC returns on PC conditional (co-)variances (Pooled)						
	PC1=MKT	PC2	PC3	PC4	PC5	R^2
Estimate	1.03	1.53	-4.14	13.96	7.42	0.027
t-stat	(0.58)	(1.07)	(-1.79)	(1.61)	(0.85)	

Panel B: Optimal combination of K PCs on conditional variance (Time-series)					
K =	3	4	5	6	7
Estimate	0.25	6.72	14.68	18.52	24.78
t-stat	(0.14)	(2.65)	(2.79)	(2.42)	(2.08)