

Esben Hedegaard and Robert Hodrick
**Measuring the Risk-Return Tradeoff with
Time-Varying Conditional Covariances**

Measuring the Risk-Return Tradeoff with Time-Varying Conditional Covariances

Esben Hedegaard*

Robert J. Hodrick[†]

W.P. Carey School of Business

Graduate School of Business

Arizona State University

Columbia University

June 8, 2014

Abstract

We examine the prediction of Merton's intertemporal CAPM that time varying risk premiums arise from the conditional covariances of returns on assets with the return on the market and other state variables. We find a positive and significant price of risk for the covariance with the market return that is driven by the time series variation in the conditional covariances, and the risk-premium on the market remains positive and significant after controlling for additional state variables. Our method estimates the risk-return tradeoff in the ICAPM using multiple portfolios as test assets.

We thank Andrew Ang, Turan Bali, Jules van Binsbergen, Oliver Boguth, and Xiaoyan Zhang for useful comments and discussions. This research was supported by a grant from the Network for Study on Pensions, Aging, and Retirement to the Columbia Business School.

*esben.hedegaard@asu.edu

[†]rh169@columbia.edu

The paper develops a multistep estimation procedure to investigate versions of Merton's (1973) intertemporal capital asset pricing model (ICAPM). The first step estimates the conditional covariances of asset returns with the risk factors using Engle's (2002) dynamic conditional correlation (DCC) model, and the second step estimates the prices of risks in a panel regression. Simulations demonstrate that the estimation method produces unbiased estimates in situations with up to 100 test assets. The results are robust to using different portfolios as test assets as well as to including state variables in addition to the market return. We invariably find that the conditional covariance of an asset's return with the market return is a priced risk. When including additional state variables in the analysis we find that the covariance of asset returns with the bond yield does not have a significant price of risk, but we find that the covariance of returns with the return on the HML of Fama and French (1992) does have a significant price of risk. Surprisingly, the significant results on the risk-return tradeoff relative to using only the market return as a test asset are not caused by an increase in statistical power from using the cross section of returns, but rather from a change in the point estimates. The results provide support for Merton's (1973) ICAPM and show that a risk-return tradeoff exists in the time series.

Merton's (1973) ICAPM states that the conditional expected return on any asset is proportional to its conditional covariance with the return on the market, as well as to its conditional covariances with state variables that describe future investment opportunities. The coefficient linking the expected return on an asset to the conditional covariance of the asset's return with the return on the market is the relative risk aversion of a representative investor, and the prices of the risks induced by the conditional covariances of returns with state variables are determined, as Merton (1973) states, by investors' desires to hedge against unfavorable shifts in the investment opportunity set, that is, by what Cochrane (2013) calls investors' state-variable aversion. An immediate implication of the ICAPM is that the conditional expected return on the market is proportional to the conditional variance of

the market, and to the covariances of the market return with innovations in state variables. Many existing studies focus on the implication for the expected return on the market and test the ICAPM by estimating the relation between the return on the market and the conditional variance of the market. These studies find mixed evidence on the relation between risk and expected return.¹

A time-series relation between conditional measures of risk and conditional expected returns should be present both for the market return, as well as for individual assets or portfolios of assets. The expected return of any asset should vary through time proportionally to the asset's conditional covariances with the market return and with state variables, and the same coefficients should describe this relation for all assets.

When we use the market return as the only test asset, we find only a marginally significant estimate of the risk-return tradeoff, which confirms the findings in previous studies. However, when we include additional test assets as well as the market, we consistently find positive, statistically significant estimates of the risk-return tradeoff. These estimates indicate strong evidence of a time-series relation between the conditional covariances of the returns on portfolios with the market return and the conditional expected returns on the portfolios. The results are generally robust to using different test assets, as well as to the inclusion of additional state variables.

We demonstrate that our multi-asset estimator is a weighted average of the individual estimates one would obtain by estimating the risk-return relation using individual portfolios. When estimating the risk-return relation using the Fama and French (1992) 25 size and book-to-market portfolios one-by-one, we find that most of the estimates for the four smallest size quintiles are positive and significant, and most are larger than the estimate based on the market return only. Thus, the result that the multi-asset estimate is positive and significant

¹Lettau and Ludvigson (2010) provide a review of the literature on the risk-return tradeoff. See Hedegaard and Hodrick (2014) for recent estimates.

is not simply driven by small stocks, but it arises from the fact that there is a positive risk-return relation for a variety of portfolios.

In the ICAPM, investors not only care about their level of wealth, but also about the distribution of future returns. As a result, the risk of an asset is measured not only by the covariance of the asset's return with the return on wealth, but also by the asset's covariances with state variables containing information about future investment opportunities. Merton (1973) suggests that the risk free interest rate is one such state variable to which investors would desire to hedge against its unanticipated adverse changes. Scruggs (1998) was the first to include the return on a long-term Treasury bond as an additional state variable noting that in his sample its omission induces a negative bias in the estimate of the price of risk associated with the return on the market. When including this second state variable, he finds a positive and significant price of risk for the return on the market, and a negative and significant price of risk for the return on the long-term bond. While we are able to replicate his parameter estimates, we find larger standard errors in his sample implying that the estimates are less significant. Furthermore, including the next 15 years of data results in estimates that are much closer to zero and no longer at all significant, a finding supportive of the conclusions in Scruggs and Glabadanidis (2003). Nevertheless, the conditional covariance of the return on an asset with the return on the market remains a significantly priced risk factor.²

In factor models, often the factors are thought of as mimicking portfolios for unobserved state variables. For instance, Fama and French (1992) suggest that the HML and SMB portfolios are mimicking portfolios for underlying common risk factors in returns. Despite the success of factor models in explaining the cross-sectional variation in returns in unconditional

²Guo and Whitelaw (2006) also stress that estimating the risk-return tradeoff of the market return with its conditional variance requires controlling for the effects of additional state variables. However, they do not attempt to measure the conditional covariances of returns with their state variables. They merely argue that the conditional covariances are proportional to a linear prediction based on the state variables.

frameworks, there has so far been little evidence that these factors are helpful in explaining the time-series variation in expected returns. We include the Fama-French factors as state variables and examine their ability to explain the time-series variation in returns. We find that controlling for the conditional covariation with the market return drives out the size-effect. This confirms the results in Lewellen and Nagel (2006), who use a different estimation procedure and arrive at the same conclusion. On the other hand, the HML portfolio carries a positive and significant price of risk in the time series. This provides support for the idea that the HML portfolio acts as a mimicking portfolio for some underlying state variables.

Two studies of the ICAPM that find strongly supportive time series results are Bali (2008) and Bali and Engle (2010). Both of these papers test the ICAPM using multiple assets, but unfortunately they report incorrect severely understated standard errors. After estimating the conditional covariances, Bali (2008) and Bali and Engle (2010) use a traditional Seemingly Unrelated Regression (SUR) approach to estimate the coefficient of relative risk aversion. As we discuss in Section 2.1, the SUR estimation minimizes the pricing errors on transformed portfolios which may not have any economic relevance. More importantly, Bali (2008) and Bali and Engle (2010) employ the usual SUR standard errors that do not account for the conditional heteroskedasticity of the error terms and that consequently severely understate the true standard errors and overstate the t -statistics. When we develop heteroskedasticity consistent standard errors, the resulting t -statistics are on average four times smaller than the ones reported in Bali and Engle (2010) causing the parameter estimates generally to become statistically insignificant. We demonstrate that while there is no evidence of a risk-return relation in daily data, which is the investment horizon used in Bali and Engle (2010), there is strong evidence of a risk-return tradeoff for longer horizons of one to three months. We demonstrate these horizon effects using the overlapping data inference approach of Hedegaard and Hodrick (2014).

The organization of the paper is as follows. Section 1 lays out the asset pricing frame-

work. Section 2 presents our estimation method which first estimates DCC models and then uses one of several different methods to estimate the model subject to the cross-equation restrictions of the theory. Section 3 describes the data. Section 4 discusses the results of the estimation when the return on the market is the only state variable. Section 5 includes additional state variables such as the yield on long-term bonds and the Fama-French factors. Section 6 relates our findings to the existing literature. In particular, we replicate the results in Bali and Engle (2010), and we provide both corrected t -statistics as well as estimates using the methodology of Section 2. We also update the results in Scruggs (1998). Section 7 concludes.

1 The Asset Pricing Model

Merton (1973) derives the ICAPM by analyzing the optimal portfolio choice of an individual investor, aggregating across investors, and imposing equilibrium. He derives an expression that the conditional expected return on any asset i must satisfy

$$E_t(R_{i,t+1}) - R_{F,t} = \gamma_M \text{cov}_t(R_{i,t+1}, R_{M,t+1}) + \sum_{k=2}^K \gamma_k \text{cov}_t(R_{i,t+1}, S_{k,t+1}), \quad i = 1, \dots, N \quad (1)$$

where $R_{F,t}$ is the risk free return, $R_{M,t+1}$ is the return on the market, S_2, \dots, S_K denote $K - 1$ state variables that contain information about future investment opportunities, and cov_t denotes the conditional covariance operator. Merton shows that the coefficient γ_M is the relative risk aversion of a representative agent, and similarly, the γ_k coefficients are weighted averages across investors of their state-variable aversions. Thus, the sign and magnitude of the γ_k 's are determined by the curvature of the investors' value functions.³ Aggregating

³We treat the prices of risk, the γ_k 's, as constants. Maio and Santa-Clara (2012) argue incorrectly that in this case one can condition down such that equation (3) holds with the unconditional expectation on the left-hand side with unconditional covariances on the right-hand side. This is incorrect because the unconditional expectation of a conditional covariance equals the unconditional covariance minus the

across assets implies that the conditional expected return on the market is given by

$$E_t(R_{M,t+1}) - R_{F,t} = \gamma_M V_t(R_{M,t+1}) + \sum_{k=2}^K \gamma_k \text{cov}_t(R_{M,t+1}, S_{k,t+1}). \quad (2)$$

where $V_t(R_{M,t+1})$ is the conditional variance of the market return. If there are no state variables (or, if the sensitivity of the representative investor's value function with respect to the state variables is small, or the variation in the state variables is very small), the model simplifies to the conditional CAPM in which $E_t(R_{M,t+1}) - R_{F,t} = \gamma_M V_t(R_{M,t+1})$.

As noted in the introduction, many existing studies of the ICAPM focus on equation (2), whereas we include multiple test-assets while also including the market return and focus on equation (1). We simplify notation by letting the first state variable be the return on the market, such that $S_{1,t+1} = R_{M,t+1}$ and $\gamma_1 = \gamma_M$, in which case we can write equation (1) as

$$E_t(R_{i,t+1}) - R_{F,t} = \sum_{k=1}^K \gamma_k \text{cov}_t(R_{i,t+1}, S_{k,t+1}), \quad i = 1, \dots, N. \quad (3)$$

To represent the model in econometric form, compile the prices of risks into the vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_K)'$; let $\mathbf{S}_{t+1} = (S_{1,t+1}, \dots, S_{K,t+1})'$ be the vector of state variables; let $\mathbf{R}_{t+1} = (R_{1,t+1} - R_{F,t}, \dots, R_{N,t+1} - R_{F,t})'$ be the vector of excess returns on the test assets; and let $\mathbf{Z}_{t+1} = (\mathbf{S}'_{t+1}, \mathbf{R}'_{t+1})'$ be the $(K + N)$ -dimensional vector of state variables and returns. Because some state variables may also be returns, we interpret the notation flexibly so as not to double count these returns, and we make the appropriate adjustments to the dimension of \mathbf{Z}_{t+1} . Let \mathbf{H}_t denote the conditional covariance matrix of \mathbf{Z}_{t+1} .⁴ We use the notation $\mathbf{H}_{NK,t}$ to denote the lower left $(N \times K)$ -block of \mathbf{H}_t containing the conditional covariances

unconditional covariance of the conditional means. Since the conditional means vary through time with the conditional variance, equation (3) does not condition down. This point was first made by Grant (1977) and was expanded upon by Jagannathan and Wang (1996), Lewellen and Nagel (2006), and Boguth, Carlson, Fisher, and Simutin (2011).

⁴We denote conditional moments of random variables that will be realized at time $t + 1$ with a subscript t to indicate when they enter the information set.

of \mathbf{R}_{t+1} with the state variables \mathbf{S}_{t+1} . We can now write the asset pricing model as

$$\mathbf{R}_{t+1} = \boldsymbol{\mu} + \mathbf{H}_{NK,t}\boldsymbol{\gamma} + \boldsymbol{\varepsilon}_{t+1}, \quad E_t(\boldsymbol{\varepsilon}_{t+1}) = \mathbf{0}_N, V_t(\boldsymbol{\varepsilon}_{t+1}) = \mathbf{H}_{NN,t} \quad (4)$$

where we have introduced a vector of asset specific constants, $\boldsymbol{\mu}$, and the vector of innovations in returns, $\boldsymbol{\varepsilon}_{t+1}$. The innovations have mean zero and covariance matrix $\mathbf{H}_{NN,t}$ which is the lower right $(N \times N)$ -block of \mathbf{H}_t . We do not assume conditional normality.

We want to test the time series prediction that higher systematic risk at a point in time as given by larger conditional covariances requires the conditional risk premium to be higher. We test this time series prediction by examining constrained regressions of returns on risks, which are measured as the covariances of the returns with the market return and with additional state variables.

2 Estimation

2.1 Estimating the Model

In this section we first discuss alternative ways to estimate the system (4) under the assumption that the conditional covariance matrices are known. We then turn to the estimation of the conditional covariance matrix process in Section 2.3 below. We use a sequential estimation approach in which we first estimate the conditional covariances, and then, taking these as given, we estimate the fundamental parameters of the model. We demonstrate below through simulations that the estimates of the fundamental parameters in the second step are unbiased and that their standard errors are correctly estimated, even though the conditional covariances are estimated in a first step and then treated as known.⁵

⁵Cochrane (2005) credits Shanken (1992) with deriving correct standard errors in the i.i.d. environment of unconditional beta models. In discussing the need for such corrections in an analogous GMM setting, Cochrane (2005, p. 258) notes, “In my experience so far with this method, the correction for the fact that

To simplify notation, define

$$\mathbf{X}_t = \left[\mathbf{I}_N \mid \mathbf{H}_{NK,t} \right] \quad (5)$$

and $\boldsymbol{\delta} \equiv (\boldsymbol{\mu}', \boldsymbol{\gamma}')$. Then, the system (4) can be written as

$$\mathbf{R}_{t+1} = \mathbf{X}_t \boldsymbol{\delta} + \boldsymbol{\varepsilon}_{t+1}. \quad (6)$$

There are several ways of estimating the parameters of this system of equations.

Each of the methods we consider can be thought of as minimizing the sum of quadratic forms in the pricing errors

$$\sum_{t=1}^T (\mathbf{R}_{t+1} - \mathbf{X}_t \boldsymbol{\delta})' \mathbf{W}_t^{-1} (\mathbf{R}_{t+1} - \mathbf{X}_t \boldsymbol{\delta}) \quad (7)$$

for some weighting matrices, \mathbf{W}_t , which are in the time- t information set. By setting $\mathbf{W}_t = \mathbf{I}_N$, we obtain the constrained OLS estimator.

We explore three other choices of \mathbf{W}_t . First, if we set $\mathbf{W}_t = \mathbf{H}_{NN,t}$, we obtain the conditional generalized least squares (CGLS) estimator, which is equivalent to maximum likelihood under the assumption of conditional normality for this part of the model. This estimation method is efficient, but it is likely less robust than OLS. Second, we scale the time- t observations by their conditional standard deviations, but we do *not* account for the correlations of the innovations. That is, we set $\mathbf{W}_t = \text{diag}(h_{1,t}, \dots, h_{N,t})$, where $h_{i,t}$ is the conditional variance of $R_{i,t+1}$. The advantage of this approach is that it scales the time- t observations by their conditional standard deviations to avoid the results being driven by large returns in periods with high volatility. At the same time, by not scaling the system

Ef is estimated is very small in practice, so that little damage is done in ignoring it (as is the case with the Shanken correction).” Newey and McFadden (1994) discuss how to derive corrections to the standard errors resulting from first stage estimation, but such analysis is quite difficult in our application.

by the inverse of the full conditional covariance matrix, we maintain a robust estimation procedure. Finally, we estimate the parameters of the system using SUR, which corresponds to $\mathbf{W}_t = \mathbf{\Sigma}$, where $\mathbf{\Sigma}$ is the estimated unconditional covariance matrix of the innovations, but we do not use the traditional SUR standard errors.

In each case, the first order conditions of the minimization are considered to be orthogonality conditions of Hansen’s (1982) GMM method, and we construct the standard errors appropriately with GMM allowing for the conditional heteroskedasticity and cross-correlation of the residuals. We derive explicit expressions for $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\gamma}}$ as well as the GMM estimates of their standard errors in Appendix A.

When reporting our results below, we primarily focus on the constrained OLS estimation. We examine all estimation procedures in more detail using simulations, and we discuss the results of the alternative estimation procedures in sections 4.4 and 4.5.

2.2 On the Importance of Including a Constant

Before discussing the estimation of the conditional covariance matrices, we briefly discuss the specification of the model. Scruggs (1998) and Nyberg (2012) estimate models such as equation (4) with and without constant terms in the conditional means of the asset returns finding that the estimated prices of risks in their models are much more significant when the models are estimated without constants. Lanne and Saikkonen (2006) explicitly advocate estimating asset pricing models without constants to impose the null hypothesis. Scruggs (1998, p. 589) cautions that estimating without a constant can lead to “misleading estimates” of the prices of risks because the market excess return is positive on average and the variance of the market return is positive by construction.

In Hedegaard and Hodrick (2014) we demonstrate the importance of including a constant in the maximum likelihood estimation of the single factor conditional CAPM when the only asset in the estimation is the market return. In that framework, if a constant is not included

in the conditional mean, the estimate of γ_M is merely the average return on the market divided by the average conditional variance of the market, which does not examine the time series prediction of the conditional CAPM that the conditional expected return on the market moves simultaneously with the conditional variance of the market return.

Appendix A demonstrates an analogous point in the context of multivariate conditional asset pricing models. Including constant terms in the conditional means of the returns is necessary to allow for a test of the time series prediction of the model that the conditional expected returns vary contemporaneously with the conditional covariances of the returns with the risk factors. If asset-specific constants are not included in the estimation, the constrained OLS estimator of $\hat{\gamma}$ is

$$\hat{\gamma} = \left(\frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{H}_{NK,t} \right)^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{R}_{t+1}. \quad (8)$$

Notice that this estimator of γ does not test the prediction of the theory that the conditional expected returns move with the conditional covariances of the returns with the risk factors. Equation (8) is the multivariate analogue to the findings of Hedegaard and Hodrick (2014) for the single factor conditional CAPM. Consider the case of the conditional CAPM, in which the market is the only state variable. Since the covariances of the test-assets with the market are generally positive, and average returns on the test-assets are positive, the estimator mechanically produces a positive and significant estimate of γ_M .

On the other hand, when asset-specific constants are included, the constrained OLS estimator of γ is

$$\hat{\gamma} = \left(\frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} (\mathbf{H}_{NK,t} - \overline{\mathbf{H}_{NK,t}}) \right)^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} (\mathbf{R}_{t+1} - \overline{\mathbf{R}_{t+1}}) \right), \quad (9)$$

where $\overline{\mathbf{X}_t}$ denotes the time series average of the matrix process \mathbf{X}_t . Consider again the

case of the conditional CAPM. Now, since the returns are demeaned in the second term, the estimator of γ_M is only positive when the sample covariance between the conditional covariance with the market and future returns on the test-assets is positive. In this sense, the significance of $\hat{\gamma}$ arises purely from the time series predictability of returns. These results are, of course, merely a reinterpretation of the standard results for the intercept and slope in a linear time series regression.

2.3 Multivariate GARCH Models

The challenges in conditional multivariate asset pricing are estimating the conditional covariances between the returns and the risk factors and linking these covariances to the conditional means of the asset returns. Bollerslev, Engle, and Wooldridge (1988) were the first to use a multivariate generalized autoregressive conditional heteroskedasticity (GARCH) model to estimate conditional second moments and link them simultaneously to conditional first moments.⁶ One can think of this approach as adding simultaneous estimation of the conditional second moments to the GMM system of equations of the previous section. Unfortunately, this simultaneous estimation approach is computationally difficult and is tractable only for a small number of returns. It quickly becomes intractable when the cross section of returns is expanded.

To examine the time series predictions of conditional asset pricing models for multiple assets, one must estimate the conditional covariance matrices, \mathbf{H}_t , of both the state variables and returns. Here we explore the dynamic conditional correlation (DCC) model of Engle (2002).⁷ Under the null hypothesis in which the returns are not predictable but have constant

⁶Garch-in-Mean models with multiple assets have also been used by Ng (1991) and DeSantis and Gerard (1997, 1998). Ng (1991) uses decile portfolios ranked on either size or beta to test the conditional CAPM. DeSantis and Gerard (1997, 1998) use currency portfolios as additional sources of risks that must also be simultaneously priced.

⁷We also examined the constant conditional correlation (CCC) model of Bollerslev (1990), which is used by Ng (1991). Because the results are similar, we only report the results of the DCC models. Other approaches, such as mixed data sampling (MIDAS) could also be used. See Hedegaard and Hodrick (2014)

means with conditionally heteroskedastic innovations, we can write the mean equations and the traditional GARCH (1,1) models for each return as follows (similarly for the K state variables):

$$R_{i,t+1} = \alpha_0^i + \varepsilon_{i,t+1}, \quad \varepsilon_{i,t+1} \sim N(0, h_{i,t}) \quad (10)$$

$$h_{i,t+1} = \omega_i + \alpha_i \varepsilon_{i,t+1}^2 + \delta_i h_{i,t}. \quad (11)$$

As in Engle (2002), we model \mathbf{H}_t as

$$\mathbf{H}_t = \mathbf{D}_t \mathbf{P}_t \mathbf{D}_t \quad (12)$$

where $\mathbf{D}_t = \text{diag}(h_{1t}^{1/2}, \dots, h_{Nt}^{1/2})$ and $\mathbf{P}_t = [\rho_{ij,t}]$ represent positive definite correlation matrices with $\rho_{ii,t} = 1$ on the diagonal. We assume that

$$\rho_{ij,t} = \frac{q_{ij,t}}{\sqrt{q_{ii,t}q_{jj,t}}}, \quad (13)$$

and that the matrix process $\mathbf{Q}_t = [q_{ij,t}]$ follows

$$\mathbf{Q}_{t+1} = \mathbf{S}(1 - a_1 - a_2) + a_1(\mathbf{u}_{t+1}\mathbf{u}'_{t+1}) + a_2\mathbf{Q}_t \quad (14)$$

where \mathbf{S} is the unconditional correlation matrix, the standardized innovations are $\mathbf{u}_{t+1} = \mathbf{D}_t^{-1}\boldsymbol{\varepsilon}_{t+1}$, and a_1 and a_2 are scalar parameters that describe the persistence of the correlations. We use maximum likelihood to estimate the parameters of the DCC model and the conditional covariances as described in the Appendix.

Of course, under the alternative hypothesis, the conditional means of the returns depend on the conditional covariances of the returns with the risk factors, and the likelihood function

for a comparison of GARCH and MIDAS approaches to estimating the risk-return tradeoff in the single equation case when the market return is the only priced asset.

does not separate into GARCH and correlation parts as demonstrated in the Appendix. Estimating under this hypothesis leads to the curse of dimensionality. Rather than doing everything simultaneously, we perform a four-step estimation as follows:

- i. Estimate GARCH(1,1) models for the conditional variances of the individual asset returns, including the market, as well as the state variables. For each process, we estimate the model in equations (10)-(11).
- ii. Estimate the DCC model for the correlation matrix process for all assets and state variables, using the standardized residuals from step one.
- iii. Compute the conditional covariance matrix process using the conditional variances from step one and the conditional correlation matrix from step two.
- iv. Estimate the risk-return parameters in the constrained system (4).

We use simulations to demonstrate that the multistep estimation produces well-behaved point estimates and t -statistics. Intuitively, if the time variation in expected returns is small relative to the innovation variance, the error in estimating the conditional covariances is small when the conditional mean is assumed to be constant.

2.4 ODIN Estimation

In addition to the basic parameter estimates and standard errors described above that are derived for non-overlapping monthly data, for the actual estimation we also report parameter estimates and standard errors calculated with overlapping data inference (ODIN) as in Hedegaard and Hodrick (2014). The ODIN methodology allows the econometrician to use all of the available daily data, even when the forecast horizon is set to an arbitrarily longer interval. Theoretical models are essentially silent on the period of time over which

the predictions of the model are thought to hold, and Merton (1980) suggested that a one month interval is “not unreasonable.”

Given that the forecast horizon is monthly, though, there is no reason why the returns must be sampled as calendar months. Indeed, with a monthly forecast interval, the first sampling day could be chosen to be any of the (on average) 22 trading days in the month. Consider a forecast interval F , where F equals 5 for weekly data, 22 for monthly data, or 66 for quarterly data. The ODIN parameter estimate is derived by maximizing the average of the F non-overlapping MLE equations for each of the F possible sampling days. As a result, the ODIN estimation uses all of the available daily data, for any given forecast horizon.

In this application, we first estimate the DCC model using ODIN. For, say, a 22 day forecast interval, we calculate 22 day returns for all assets for the 22 different possible starting dates. To estimate the DCC model, we maximize the sum of the 22 individual likelihoods for the 22 different models, resulting in *one* set of GARCH parameters for each return process, and *one* set of DCC parameters. Using the estimated parameters, we calculate the 22-day-ahead conditional covariance matrix of the returns, for each day.

Finally, we use overlapping constrained regressions to estimate the risk-return tradeoff. We regress 22-day-ahead returns on the conditional covariances, for each day in our sample. The standard errors are then corrected for the overlap in the forecasts, as in Hansen and Hodrick’s (1980) application of Hansen’s (1982) GMM, but allowing for conditionally heteroskedastic errors in the estimator of the spectral density matrix. The simulations in Hedegaard and Hodrick (2014) suggest that the standard errors from this strategy are well behaved in sample sizes such as ours and are generally smaller than the standard errors from any particular non-overlapping sample. Using ODIN also forces the econometrician to consider all F parameter estimates than arise from different sampling dates, and as in Hedegaard and Hodrick (2014), we demonstrate in Section 4 that the non-overlapping estimates vary greatly with changes in the sampling start date.

3 Data

We begin by estimating the conditional CAPM in which case the market return is the only risk factor. For the market portfolio, we use the value-weighted return on the CRSP market portfolio. When estimating the model, we always include the market return as well as one of the following five additional data sets: *i)* six size and book-to-market sorted portfolios, *ii)* size sorted decile portfolios, *iii)* book-to-market sorted decile portfolios, *iv)* 25 size and book-to-market sorted portfolios, *v)* 100 size and book-to-market sorted portfolios, and *vi)* 10 market beta sorted portfolios. We obtain the beta-sorted portfolios from CRSP and the remaining data from Ken French's web site. In all cases, we use excess returns that are constructed as the difference of the raw returns and the one-month risk-free rate. Including the market return does not induce a redundant set of orthogonality conditions, which would induce singularity into the covariance matrix, because the market is value weighted, and the weights on the portfolios vary over time.

For our first multi-factor model, we use the yield on the 10-year Treasury bond from the Federal Reserve Bank of St. Louis as an additional state variable. Since monthly observations on the yield are available from February 1954 to December 2012 we limit all of our samples to these 707 monthly observations, and in our simulations we consequently use this same number of observations. For daily data, the 10-year Treasury yields are available from January 1962 to December 2012.

The next multi-factor model is the conditional Fama and French (1992) model in which we examine whether the HML and SMB portfolios are mimicking portfolios for underlying common risk factors in returns that can explain the time-series variation in returns. The HML and SMB series are also taken from Ken French's web site. Bali and Engle (2010) also include conditional covariances with a number of other potential risk factors like the change in the fed funds rate, the term premium, the default spread, and stock market volatility, and

we discuss their data sources in Section 6.1.

4 The Market as the only State-Variable

Table 1 presents the estimation results with monthly returns, reporting results both for the standard non-overlapping estimation from calendar months and for the ODIN estimation based on constrained OLS estimation of the system of equations. We report asymptotic and bootstrap standard errors and p -values, where the bootstrapped values are based on 1,000 simulations.

The estimation with the market return as the only test asset produces a positive point estimate of the risk-return tradeoff of 3.61 with a standard error of 2.14 and a p -value of .09 indicating only marginal significance of the parameter estimate. As noted in the introduction, previous papers have generally reported relatively insignificant estimates of γ_M using the market as the only test asset. One of our main results is that when adding additional test assets, the parameter estimates of γ_M become more significant as the p -values fall to between .01 and .04 (except when using the 10 book-to-market portfolios in which case the p -value is .08).⁸

[Table 1 about here.]

These estimates use the traditional, non-overlapping, calendar monthly returns. However, the estimates are very sensitive to the day of the month on which returns are sampled. If we instead use a fixed 22-day forecast interval and consider the various estimates that are obtained from the 22 different possible starting dates using the return on the market as the only test-asset, we find that the point estimates of γ_M vary from 0.97 to 3.71. Thus, the reported estimate of 3.61 in Table 1 is in the upper range of the different monthly estimates.

⁸This discussion uses covariances based on a DCC model estimated to all assets. We also estimate bivariate DCC models of each test-asset and the market, which allows the mean-reversion dynamics to depend on the test-asset. This only leads to minor differences as the point estimates of γ change by less than 0.05.

When we use the 25 size and book-to-market sorted portfolios as test assets, the various non-overlapping estimates of γ_M range from 1.54 to 4.64. The calendar month estimate of 4.74 reported in Table 1 falls outside this range since all months do not have the same number of trading days, such that the standard monthly returns do not correspond to a fixed 22 day forecast interval for any starting date. In general, the estimates based on the usual non-overlapping calendar monthly sampling are at the upper range of the different estimates obtained from varying the sampling start date. Due to the sensitivity of the estimates to the sampling start date, we also consider the ODIN estimation, in which we maximize the sum of the 22 likelihood functions associated with the 22 different possible starting dates associated with the 22 different possible starting dates associated with the 22 different possible starting dates to derive the parameters of the DCC model and the conditional covariances and then minimize the sum of the 22 constrained OLS objective functions to estimate the price of market risk.

The ODIN estimates are all positive, but they are less significant than their non-overlapping counterparts. The standard errors for the ODIN model are always smaller than the standard errors for the non-overlapping estimations, illustrating the increased precision that the ODIN estimation offers. However, the point estimates are also smaller in magnitude because ODIN average across all the starting dates, and it is the differences across the starting dates that results in less significant estimates. The ODIN estimates of γ_M are significant at the 10% level for 5 portfolios, while the estimate of γ_M is significant at the 1% level for the 10 beta-sorted portfolios. Overall, the evidence in favor of a risk-return tradeoff at the 1-month horizon remains relatively weak, but as we demonstrate below, it becomes very strong at the 2-month to 3-month horizon.

4.1 Cross-Sectional Implications

We also examine the cross-sectional implications of the model by testing the hypothesis that all intercepts jointly equal zero, using a Wald test. The model is generally rejected when including multiple test assets.

Nevertheless, although the model is rejected, the conditional CAPM explains a large part of the time-series variation in returns. To demonstrate the partial success of the model, we focus on the market portfolio and the 25 size and book-to-market sorted portfolios as test assets. The top plot in Figure 1 shows the average returns on each of the 26 test assets along with two standard deviation intervals. These are the pricing errors under the null hypothesis of $\gamma_M = 0$, that is, the pricing errors that would arise if there were no risk-return relationship and the estimated value of γ_M were zero. The standard deviations are based on average returns, and the standard deviations would be larger if γ_M were estimated. The usual patterns in average returns across size and book-to-market are clearly visible—on average, value stocks outperform growth stocks, and small stocks outperform large stocks.

The bottom plot in Figure 1 shows the constants from the estimated model with their two standard deviation intervals. Although the usual patterns across book-to-market are still clearly visible, the conditional CAPM explains the average returns across portfolios very well. The conditional model overadjusts the expected returns for low book-to-market stocks as the constant terms are negative.

Note that the constants are parameters of the model that are not directly minimized by the estimation. As in any time series regression, the constants pick up the average of the dependent variable minus the slope coefficient times the average of the regressor. Nothing in the model or the estimation procedure makes the constants close to zero. If there were no risk-return tradeoff in the time series, the estimated value of γ_M would be zero and the constants would equal the average returns in the top plot. In particular, even if there were a cross-sectional relation between average returns on portfolios and the portfolios' unconditional

covariance with the market, but no time-series relation, the estimated γ_M would be zero, and the constants would pick up the average returns on the different portfolios. The estimate of γ_M will only be non-zero if a time series relation exists between the future returns on the portfolios and their conditional covariances with the market. In this sense, the conditional CAPM is a success.

While the conditional CAPM is useful in explaining the average time series variation in returns, it clearly fails to explain the time-series and cross-sectional variation jointly. As noted above, the bottom plot in Figure 1 shows the usual patterns across book-to-market sorted portfolios. This finding is similar to the results in Lewellen and Nagel (2006), who find that the conditional CAPM captures the variation in size-portfolios, but that it does not explain the variation in book-to-market sorted portfolios. We thus agree with Lewellen and Nagel (2006) that the conditional CAPM does not fully explain asset pricing anomalies.

[Figure 1 about here.]

4.2 Alternative Horizons

To explore if the risk-return relationship of the conditional CAPM is present at alternative horizons, we vary the forecast interval from 1 to 60 days and estimate the model based on the 25 size and book-to-market sorted portfolios.⁹ Figure 2 shows the point estimates of γ_M as a function of the forecast interval, along with a band indicating plus and minus two standard deviations. The estimated risk-return tradeoff increases as a function of the forecast interval, up to about a 2 month horizon. It is insignificant at horizons shorter than 1 month, at which point it is significant at the 10% level, and becomes significant at the 1% level at a 2 month forecast horizon. The ODIN estimation methodology ensures that the estimations are robust even for long forecast intervals. With quarterly observations we only

⁹The results are similar for the other data sets.

have 236 non-overlapping intervals, but using overlapping data results in stable parameter estimates. When estimating the ODIN model in Table 1 using two month returns instead of one month returns, four specifications are significant at the 1% level and the remaining three specifications are significant at the 5% level. Hence, while the evidence in favor of a risk-return tradeoff at the one-month horizon is relatively weak, it is very strong at longer horizons.

There are several reasons why the risk-return tradeoff may not be detectable at shorter horizons. First, aspects of the trading process induced by market microstructure frictions, non-synchronous portfolio investment decisions, and individual stock illiquidity that are outside the theory dominate the autocorrelations in short-horizon returns. More importantly, when more volatile trading environments arise, theory predicts that stock returns are expected to be contemporaneously negatively correlated with the increase in volatility because prices must fall to provide an increase in expected returns, as in Campbell and Hentschel (1992). If the adjustment of expected returns to news that increases the conditional variance is not precisely contemporaneously correlated with the increase in the conditional variance because of market illiquidity or the non-synchronous trading of investors, using a short horizon for testing the conditional risk-return tradeoff may find a negative relation as volatility increases and prices fall slightly later.¹⁰

[Figure 2 about here.]

¹⁰Müller, Durand, and Maller (2011) discuss estimation of the risk-return relation at the daily horizon, with the market as the only test-asset. They argue that in this case one must explicitly account for weekend effects which they include in a COGARCH framework. When ignoring weekend effects, they find a positive and significant risk-return relation, but adjusting for weekend effects the significance disappears, and they conclude that using daily data does not provide support for Merton's ICAPM. Bali and Engle (2010) use daily data, and we discuss their results below.

4.3 The OLS Panel Data Estimator

Appendix A.2 shows that the constrained OLS estimator of γ_M is a weighted average of the individual estimates, $\gamma_{i,M}$ obtained by OLS estimation using a single portfolio. Specifically,

$$\hat{\gamma}_M = \sum_{i=1}^N \frac{\hat{V}(\hat{\sigma}_{iM,t}) \hat{\gamma}_{i,M}}{\sum_{k=1}^N \hat{V}(\hat{\sigma}_{kM,t})}. \quad (15)$$

where $\hat{\gamma}_{i,M} = \frac{\text{cov}(\hat{\sigma}_{iM,t}, R_{i,t+1})}{\hat{V}(\hat{\sigma}_{iM,t})}$ is the estimate of γ_M based on portfolio i only, and \hat{V} , cov , and $\hat{\sigma}$ denote sample moments. The weights are determined by the sample variances of the conditional covariances of each asset return with the return on the market, and the estimate associated with the asset with the highest variation in the conditional covariance with the market receives the highest weight.

[Figure 3 about here.]

The top panel of Figure 3 presents the individual $\hat{\gamma}_{i,M}$'s from the one asset case for each of the 26 test assets. The first asset is the market portfolio, and the other assets are the 25 size and book-to-market sorted portfolios ordered within size quintiles from low to high book-to-market beginning with the smallest quintile. The circles indicate the point estimates and the vertical lines indicate the two standard error bounds. One sees that while estimating with just the market portfolio implies an insignificant risk-return tradeoff, most of the $\hat{\gamma}_{i,M}$'s for the four smallest size quintiles are larger than the point estimate from the market only, and the vast majority are more than two standard errors from zero. The horizontal line is the weighted average of the individual estimates, which is the constrained OLS estimate, and this is significantly different from zero.

The bottom panel in Figure 3 presents the weights that the 26 individual estimates receive in the constrained estimation. Within each size quintile, the estimate of γ_M based on growth stocks receives the highest weight. Interestingly, these are the point estimates of γ_M with the

lowest values. Despite this, since most of the individual estimates of γ_M are larger than the estimate based on the market only, the constrained estimate from the system of equations is significant. Importantly, one sees that the significance of the constrained estimate is not driven purely by small stocks. An individually significant risk-return relation exists for all portfolios except for those in the largest quintile.

estimator from the panel data regression will be more efficient than the estimator using the market alone.

4.4 GMM Estimations of the Panel Data System

As Cochrane (2005) notes, using a weighting matrix in GMM can be interpreted as forming portfolios of assets, and GMM estimation can be interpreted as choosing the parameters to minimize the pricing errors on these statistical portfolios. The counterpart of this point in our framework is found by taking the Cholesky decomposition of $\mathbf{W}_t^{-1} = \mathbf{C}_t \mathbf{C}'_t$, in which case one sees that the GMM estimation method minimizes the sum of squared pricing errors on the N portfolios

$$\mathbf{C}'_t(\mathbf{R}_{t+1} - \mathbf{X}_t \boldsymbol{\delta}). \tag{16}$$

The estimation method with $\mathbf{W}_t = \text{diag}(h_{1,t}, \dots, h_{N,t})$ simply scales the observations by their conditional standard deviations and then minimizes the squared errors of the resulting standardized pricing errors. The SUR and CGLS methods both exploit the correlation structure of the assets to extract statistically orthogonal components with the lowest variance. Because the different portfolios used as test assets are highly correlated, $\mathbf{H}_{NN,t}$ and $\boldsymbol{\Sigma}$ are close to singular, and the resulting priced portfolios consequently have extreme long/short weights on the original test assets. Although CGLS and SUR may asymptotically improve the efficiency of the estimator, the weighting matrix may in practice put a lot of weight on portfolios with little economic relevance. We analyze the different methods using simulations

below, but we first discuss the results of the alternative estimations.

Our first finding is that standardizing the time- t observations by their conditional standard deviations (but not using correlations) does not change the estimation results in any economically or statistically significant way. This finding is important because it demonstrates that the previous results are not spuriously driven by large returns in periods with high volatility.

Next, we turn to the SUR and CGLS estimation procedures which both use the correlation structure of the test assets to standardize the system (4). Table 2 shows the results for the CGLS estimation (the results for the SUR estimation are similar). The estimates of γ_M are systematically lower than for the constrained OLS estimation, and they are now all insignificantly different from 0 at traditional marginal levels of significance. To better understand the SUR and CGLS estimations, we examine the Cholesky decomposition of Σ^{-1} for the SUR estimation with the market and the 10 book-to-market sorted portfolios, and we find that the ratio of the largest absolute weight to the smallest absolute weight for the 11 rows is generally above 100. Using the market and the 25 size- and book-to-market sorted portfolios the ratios of absolute weights are again above 100, and for several moment conditions the ratio exceeds 1000. As a result, the SUR procedure chooses the parameters of the model to minimize the squared pricing errors on highly levered portfolios that are not economically meaningful. We further examine the statistical properties of the SUR and CGLS estimators below using simulations.

[Table 2 about here.]

4.5 Simulations

As seen in Table 1 with the actual data, the standard error of γ_M generally does not decrease as we add more test assets. To understand why, in this section we analyze the estimation

procedure in more detail using simulations. Our goals are to determine the degree of bias that is introduced by the multistep estimation, to determine the accuracy of the standard errors that do not account for the first stage estimation, and to examine if there is an increase in power from using additional returns. Moreover, we compare the different ways of estimating the constrained system of equations.

For each of the data sets described above, we estimate a DCC model as described in steps one and two above. We then assume a value of γ_M and simulate from the GARCH-in-Mean model of equation (3). We perform 1,000 simulations, and for each one we draw the innovations randomly with replacement from the original standardized residuals. For each of the simulated data sets we repeat the four step estimation procedure above to estimate the price of risk, γ_M . Finally, we test the null hypothesis that $\gamma_M = 0$. We focus first on the constrained OLS estimation of the system, and we then turn to the alternative estimations below.

Figure 4 presents the mean of γ_M in Panel A on the left side. The estimates only have small biases, which are larger for higher values of γ_M . The online appendix presents QQ-plots of the t -statistics under the null hypothesis $\gamma_M = 0$ indicating that they are very well behaved (nevertheless, we also use the empirical distributions of the t -statistics under the null for inference).

Panel B on the right side presents plots of the power of the test of the null hypothesis $\gamma_M = 0$ against the alternative hypotheses $\gamma_M = 2, 4$ or 6 , for the different data sets. Note that the empirical size of the test is exactly 5%, since we use the empirical distribution of the t -statistic under the null. Somewhat surprisingly, though, we see essentially no increase in power from including additional test assets.

[Figure 4 about here.]

To understand in more detail why the power of the test does not increase as more test

assets are included, we perform additional simulations. Instead of using the correlation structure from the actual test portfolios, we consider scenarios in which the assets have a lower or higher correlation. When the assets display a low correlation with each other, adding more test assets clearly improves the power to reject the null. However, as we increase the correlation between assets, the power increase gradually disappears. Intuitively, when the correlation between test assets is very high, the returns move closely together and the additional assets fail to increase the precision of the estimation in these sample sizes.

We also analyze the SUR and CGLS methods using simulations. Whereas the constrained OLS method produces estimates of γ_M that were nearly unbiased for all simulations, both the SUR and CGLS methods result in a downward bias for γ_M . This bias is around -0.5 when the true value is $\gamma_M = 0$ and increases to -2 when $\gamma_M = 6$. The distribution of the t -statistics also deteriorates when using 25 and 100 test assets. Interestingly, the power of the test of $\gamma_M = 0$ against a given alternative value of γ_M increases as more test assets are included. The SUR and CGLS methods both have the ability, at least for simulated data, to improve inference when using highly correlated portfolios. The reason is, as described above, that the two methods are able to extract orthogonal portfolios from the original test assets.

We thus face a tradeoff between constrained OLS, which prices economically interesting portfolios and produces unbiased estimates of γ_M with larger standard errors, and SUR and CGLS which price economically uninteresting portfolios, and produce severely biased, but more efficient estimates. We have chosen to focus mainly on the constrained OLS estimator, since it is unbiased and directly minimizes the pricing errors on the original test assets. Ultimately, in the data, adding additional test assets allows us to reject the null of no risk-return relation using constrained OLS, not because of an increase in power, but because of an increase in the point estimate of γ_M . As noted above, this increase is driven by all stocks except those the largest quintile (all but the largest decile when using the 100 size-

and book-to-market sorted portfolios).

5 The ICAPM with Additional State-Variables

5.1 The Market and the Long-Term Yield as State Variables

We now add the yield on a long-term bond as a second risk factor. Merton (1973) notes that "...there exists at least one element of the opportunity set which is directly observable: namely, the interest rate, and it is definitely changing stochastically over time." This makes the interest rate a prime candidate for a state variable. In Merton's model, the yields on all bonds are perfectly correlated because they are driven by a one-factor model with a stochastic instantaneous risk-free rate. When yields are not perfectly correlated, we must decide on a specific maturity. We use the yield on the 10-year Treasury bond as the state variable, since indexed bonds are not available for a sufficiently long time period, and we first-difference the yields when estimating the GARCH model in equations (10) and (11).

Since we use stock portfolios as test assets, the price of risk for the covariance with the 10-year Treasury yield is only identified if the correlations between the portfolio returns and innovations in the yield are non-zero. The conditional correlation of the return on the market and the change in the 10-year Treasury yield is negative between 1960 and 2001, after which the correlation becomes positive. The magnitude is sizable and varies from -0.45 to 0.4 .¹¹

A decrease in the bond yield is a deterioration in investment opportunities, since a long-term investor will earn a lower return on long-term bonds (which are risk-free to the long-term investor, ignoring inflation). Assets that do well when the yield falls (such as long-term bonds) have a negative covariance with changes in the yield. Such assets provide

¹¹When we estimate the DCC GARCH model on the different data sets, the estimates of the DCC GARCH parameters change. Consequently, the precise estimate of the conditional correlation between the market and the innovations in the Treasury yield are different for the different data sets. However, the differences are small.

a hedge against deteriorations in investment opportunities and should command a lower return. Thus, we would expect a positive price of risk for γ_F .

Table 3 presents the results for both the non-overlapping estimations and the ODIN estimations. The non-overlapping estimations use monthly data from 1954 to 2012, while the ODIN estimations use daily data from 1962 to 2012 since daily observations on 10-year Treasury yields are only available after 1962. The estimated values of γ_M for the standard non-overlapping estimations are similar to the results for the univariate model in Table 1. The estimates of γ_M vary from 3.47 using the market as the only test asset to 6.71 using the 10 beta-sorted portfolios. The estimate of γ_M is insignificant when only the market is used as a test asset, but it is significant at the 5% level in 3 estimations and at the 1% when using the 10 beta-sorted portfolios as test assets. The estimates of γ_F are all negative and insignificant, as they are smaller than their standard deviations.¹²

For the ODIN estimation, the estimates of γ_M are similar to the estimates from the non-overlapping data, and the standard errors are always smaller despite the shorter sample period, again illustrating the increased precision of the ODIN estimation. As a result, the estimates of the risk-return tradeoff are now significant at the 10% level or better for all portfolios, at the 5% level for 4 portfolios, and at the 1% level using the 10 beta-sorted portfolios. As in the non-overlapping estimations, the ODIN estimates of γ_F are all negative and insignificant.

[Table 3 about here.]

When varying the forecast intervals from one day to three months, we observe the same patterns as for the conditional CAPM. For short forecast horizons, the estimate of γ_M is

¹²As above, the covariances are based on a DCC model estimated on all test-assets and state variables. We also estimate bivariate DCC models (for the test-asset and the market return, as well as the test-asset and innovations in the yield), which allows the DCC parameters to depend on the test-asset, and in particular allows for different mean-reversion speeds for the correlation of a test-asset with the market return and the correlation of a test-asset with innovations in the yield. The results are similar.

small and insignificant, but it becomes significant at the 5% at a two-week horizon and is significant at the 1% level for a two-month horizon. The estimate of γ_F is always negative and never significant.

Scruggs (1998) shows that if the market is the only test asset, and the true model is a two-factor model, the bias arising from omitting the yield factor is

$$\widehat{\gamma}_M - \gamma_M = \gamma_F \frac{\text{cov}(\sigma_{M,t}^2, \sigma_{MF,t})}{V(\sigma_{M,t}^2)}. \quad (17)$$

Scruggs reports that in his study, this bias is around -6.7 . He reports an insignificant estimate of γ_M in a model with only the market as a state variable, but a positive and significant estimate when including the risk-free rate as an additional state variable. Thus, he concludes that the differences in the estimated parameters can be ascribed to an omitted-variable bias.

In our sample, on the contrary, the bias is much smaller, since the correlation between stock returns and bond return has changed over time. For our estimation with multiple test assets, we analyze the potential bias using simulations. By simulating from a two-factor model estimated on the actual data sets, and then estimating a one-factor model in which we omit the innovations in the yield, we find that the bias in γ_M is between 0.1 and 0.2, depending on the data set. That is, the bias has the opposite sign of the bias reported by Scruggs (1998) and is much smaller in magnitude. As seen in Table 1 and 3, the estimates of γ_M are actually smaller in the two-factor model, which again is contrary to the findings by Scruggs. We report an update of Scruggs's results below.

5.2 Other State-Variables

Fama and French (1992) argue that the market return, and the excess returns on portfolios of small minus big stocks, $R_{SMB,t+1}$, and of high minus low book-to-market stocks, $R_{HML,t+1}$,

are risk factors that explain the average returns on various portfolios. They run OLS regressions of portfolio returns on the three risk factors, and they interpret the significance of the intercepts in these regressions as measures of the inability of the three risk factors to explain the average returns. Grant (1977) was the first to argue that such a methodology encounters problems when the conditional means of returns and the conditional exposures of the returns to the risk factors vary over time. Jagannathan and Wang (1996), Lewellen and Nagel (2006), and Boguth et al (2011) subsequently expand on this issue. The fundamental problem is that the intercepts in the regressions are not expected to be zero if the conditional model is true. Hence, they do not represent pricing errors of the model.

We are interested in whether $R_{SMB,t+1}$ and $R_{HML,t+1}$ are conditional risk factors or state variables in the Merton (1973) context. We therefore estimate the conditional covariances of the test asset returns with $R_{SMB,t+1}$ and $R_{HML,t+1}$ using the DCC model as above and then perform constrained OLS on the following model:

$$\begin{aligned}
 R_{i,t+1} = & \mu_i + \gamma_M \text{cov}_t(R_{i,t+1}, R_{M,t+1}) + \gamma_{SMB} \text{cov}_t(R_{i,t+1}, R_{SMB,t+1}) \\
 & + \gamma_{HML} \text{cov}_t(R_{i,t+1}, R_{HML,t+1}) + \varepsilon_{i,t+1}
 \end{aligned} \tag{18}$$

Table 4 presents the results. The sample period for the standard estimation in the top panel is monthly data from February 1954 to December 2012.

[Table 4 about here.]

For the standard specification we consistently find that $R_{M,t+1}$ and $R_{HML,t+1}$ are priced risk factors. The estimates of γ_M are larger than in Table 1 and retain their significance. The conditional covariances of returns with HML also produce in a significant price of risk in the standard estimation with estimates of γ_{HML} ranging from 6.22 to 17.34 depending on the test assets. Consistent with the results in Lewellen and Nagel (2006), we do not find a

significant price of risk for covariances with $R_{SMB,t+1}$. Any size effect appears to be driven out by the time variation in the conditional covariances.

The ODIN estimation uses daily data from January 1, 1962 to December 31, 2012. Inference about the significance of γ_M is similar to Table 1, but the statistical significance of $R_{HML,t+1}$ is now much reduced. The coefficient estimates of γ_{SMB} are generally smaller than their standard errors.

6 Relation to the Literature

The papers in the literature most closely related to ours are Bali (2008), Bali and Engle (2010) and Scruggs (1998). This section demonstrates that Bali (2008) and Bali and Engle (2010) both severely understate the standard errors of their parameter estimates leading to false statistical significance. The results in Scruggs (1998), while also having impressive test statistics, are found to be sample specific as extending the sample for another 15 years results in insignificant prices of risks. This is consistent with the findings of Scruggs and Glabadanidis (2003) who report reduced significance of the original specification of Scruggs (1998) when adding more data.

6.1 The Bali-Engle (2010) Approach: Estimating the Risk-Return Tradeoff with Seemingly Unrelated Regression

Bali (2008) and Bali and Engle (2010) use daily data and the SUR procedure described above to estimate the prices of risks for various state variables. Bali (2008) estimates the conditional covariances using a bivariate GARCH model, and Bali and Engle (2010) use bivariate DCC GARCH models to estimate the conditional covariances. After estimating the conditional covariances, both papers estimate the risk-return tradeoff using a traditional SUR method. Both papers report highly significant parameter estimates of the risk-return

tradeoff. Since their estimation methods are similar, we focus on Bali and Engle (2010). For the conditional CAPM, which is their simplest model, they do the following:

- i. Estimate AR(1) models for the conditional means of the individual asset returns and the market return, while simultaneously estimating GARCH (1,1) models for the conditional variances of the residuals. The model estimated is therefore

$$R_{i,t+1} = \alpha_0^i + \alpha_1^i R_{i,t} + \varepsilon_{i,t+1}, \quad \varepsilon_{i,t+1} \sim N(0, h_{i,t}) \quad (19)$$

$$h_{i,t+1} = \omega_i + \alpha_i \varepsilon_{i,t+1}^2 + \delta_i h_{i,t}. \quad (20)$$

- ii. Estimate the conditional covariances for each test asset and the market portfolio using the standardized residuals from step one and separate, bivariate DCC models
- iii. Use Zellner’s (1962) Seemingly Unrelated Regression (SUR) to estimate the common price of risk on the conditional covariances of the individual returns with the market return in a panel of returns that excludes the market return.

In describing their methodology, Bali and Engle (2010, p. 382) state, “The panel estimation method with SUR takes into account heteroskedasticity and autocorrelation as well as contemporaneous cross-correlation of error terms.” In stark contrast to Bali and Engle, we do not find significant estimates of the risk-return tradeoff using a daily forecast horizon. We therefore conjectured that the standard errors in Bali and Engle (2010) are incorrect.¹³ Because basic SUR estimation imposes an assumption of conditional homoskedasticity in a highly conditionally heteroskedastic model, the Bali and Engle (2010) standard errors are too small, which severely biases the statistical evidence in favor of the significance of the asset pricing models.

¹³Bali confirmed in an email correspondence that the Bali and Engle (2010) panel estimation was done in EVIEWS, which produces traditional SUR standard errors, as we duplicate below.

Bali and Engle (2010) use daily data on the Dow 30 stocks from July 10, 1986 to June 30, 2009 for 5,795 observations, as well as daily data on value-weighted decile portfolios from sorts based on size, book-to-market, momentum, and industry portfolios from January 3, 1972 to June 30, 2009 for 9,462 observations. We downloaded the daily returns on the size, book-to-market, momentum, and industry portfolios from Kenneth French’s online data library, the same source used by Bali and Engle (2010).¹⁴

[Table 5 about here.]

Panel A of Table 5 provides the original estimates and standard errors from Bali and Engle (2010), who find quite reasonable estimates of $\hat{\gamma}_M$ that range from 1.85 for the industry portfolios to 3.32 for the momentum portfolios. All their estimates seem highly significant when using the incorrect standard errors. Panel B of Table 5 reports our attempt to replicate their results and provides corrected standard errors for the SUR estimation. Differences between our replication and the original estimation no doubt mostly arise from the first stage DCC models given the variety of choices that must be made regarding the non-linear estimation. Despite these differences, our point estimates are reasonably close to the original results, except for the size-sorted portfolios.¹⁵ Note, though, that the correct standard errors are much larger than the incorrectly calculated ones. The estimates of γ_M now range from 1.74 for the industry portfolios to 3.84 for the Dow stocks (apart from the size-sorted portfolios), and the estimates are no longer significantly different from zero. Panel C reports the results from our constrained OLS methodology. Here, the different weighting of the pricing errors results in lower point estimates of γ_M , and all estimates are insignificant.

¹⁴We were unable to download the daily returns on the investment-to-assets IA and return-on-assets ROA portfolios that Bali and Engle (2010) obtained from Long Chen’s and LuZhang’s online data library because this data library no longer exists.

¹⁵We find it puzzling that our results for the size-sorted portfolios are so different from the results in Bali and Engle (2010). We use the same program for all of the replications, and we downloaded the data from the same source.

We also explore the Bali and Engle (2010) SUR methodology with simulations. When the simulations match the parameters from the size decile portfolios with daily data for the range used in the estimation, we find that the average, across the 1,000 simulations, of the ratio of the standard deviations for the standard SUR estimation to its heteroskedastic counterpart calculated on the same data is 0.23 indicating that the t -statistics are on average inflated by a factor of 4.3 in these samples. Indeed, as shown in Panel B of Table 5, for the actual data, we find that the Bali and Engle (2010) t -statistics are generally more than four times larger than the correctly calculated ones. This is enough bias to render the estimates insignificantly different from zero.

Bali and Engle (2010) extend their tests to include additional state variables, such as the default premium, the term spread, the federal funds rate, and the implied volatility of the S&P500 index. They conclude that conditional covariances with the default premium, the term spread, and the implied volatility are also sources of risk in addition to conditional covariance with the market return. However, when we include the conditional covariances of the first three variables in addition to the conditional covariance with the market return as risk factors, as in Bali and Engle's (2010) Table 4, in our constrained OLS analysis, we find that the significance found by Bali and Engle (2010) is due to the incorrectly calculated standard errors from the unconditional SUR approach as none of the correctly calculated t -statistics is larger than 1.12. Similarly, including the conditional covariance with implied volatility along with the conditional covariance with the market return, as in Bali and Engle's (2010) Table 6, produces insignificant prices of risk as the absolute value of the largest t -statistic is 1.4. These results are available in an online Appendix.

6.2 Scruggs (1998)

Scruggs (1998) incorporates innovations to the risk-free rate as a state variable, allowing the expected return on the market to depend on its covariance with innovations in the return on

long-term bonds. He uses a constant-correlation EGARCH model, and models the market risk premium as a linear function of the conditional market variance and the conditional market covariance with the return on a long-term government bond. The full model is¹⁶

$$r_{M,t} = \gamma_0 + \gamma_M \sigma_{M,t}^2 + \gamma_F \sigma_{MF,t} + \varepsilon_{M,t} \quad (21)$$

$$r_{F,t} = \mu_F + \varepsilon_{F,t} \quad (22)$$

$$\log(\sigma_{M,t}^2) = \omega_M + \alpha_M g(\psi_{M,t-1}) + \delta_M \log(\sigma_{M,t-1}^2), \quad \psi_{M,t-1} = \varepsilon_{M,t-1} / \sigma_{M,t-1} \quad (23)$$

$$\log(\sigma_{F,t}^2) = \omega_F + \alpha_F g(\psi_{F,t-1}) + \delta_F \log(\sigma_{F,t-1}^2), \quad \psi_{F,t-1} = \varepsilon_{F,t-1} / \sigma_{F,t-1} \quad (24)$$

$$\sigma_{MF,t} = \rho_{MF} \sigma_{M,t} \sigma_{F,t} \quad (25)$$

$$g(\psi_{M,t-1}) = \left(|\psi_{M,t-1}| - \sqrt{2/\pi} \right) + \theta_M \psi_{M,t-1} \quad (26)$$

$$g(\psi_{F,t-1}) = \left(|\psi_{F,t-1}| - \sqrt{2/\pi} \right) + \theta_F \psi_{F,t-1} \quad (27)$$

[Table 6 about here.]

The estimates in Scruggs (1998) are reported in the second column of Table 6. Scruggs uses the excess returns on the Ibbotson Associates long-term government bond total return index from March 1950 to December 1994. Although we use the return on the Fama Bond Portfolios for 61-120 months from CRSP over the period January 1952 to December 1994, we are able to closely replicate his results, as shown in the third column of Table 6.¹⁷ Note that we now use the return on long term bonds as opposed to the change in yield, as we did above. Thus, we now expect γ_F to be negative. Our point estimates of γ_M and γ_F are of the same magnitude as the estimates reported by Scruggs, however, we find that they are

¹⁶Scruggs (1998) also considers specifications where the volatility processes are functions of the short-term risk-free rate. His findings for these specifications are similar to the finding for the model reported here, which corresponds to his model 3a.

¹⁷However, for his specifications 3b(i) and 3b(iii), we obtain opposite signs of both γ_M and γ_F .

not significant.¹⁸ When updating the sample to 2012, the estimates of γ_M and γ_F are now close to zero and insignificant.

In fact, for the original sample period considered by Scruggs, the likelihood function has multiple local minima. Scruggs reports the global minimum, but the likelihood function also has a local minimum for which γ_M and γ_F are close to zero and insignificant. When updating the sample, this local minimum becomes the global minimum, and the other local minimum disappears. Also note that in our sample, the estimate of the unconditional correlation between the return on the market and the bond return is close to zero and not statistically significant. As a result, the price of risk on the covariance with the bond return is poorly identified, with a standard error which is several orders of magnitude larger than the point estimate. This contrasts with the original findings in Scruggs (1998), who finds a positive and significant correlation between the return on the market and the bond return in his sample.

[Table 7 about here.]

[Table 8 about here.]

7 Conclusion

This paper explores strategies for estimating the prices of risks in Merton's (1973) ICAPM. Our preferred constrained OLS method is tractable, scales to allow multiple test assets, and produces unbiased estimates and standard errors in our sample sizes. When including

¹⁸The many parameters and the non-linearities of the model make it difficult to precisely calculate standard errors. We take care in evaluating the Hessian of the likelihood function. In particular, the Hessian returned from the MatLab optimization routine is not reliable, as it is a so-called 'pseudo-Hessian' constructed with the purpose of choosing sensible step-sizes, not to be a high-precision estimate of the second derivatives. Instead, we use the DERIVEST suite by D'Errico (2011), an adaptive numerical differentiation toolbox that provides high-precision first-order and second-order derivatives. When basing the t -statistics on the Hessian matrix returned by the MatLab optimization routine, they increase from 1.0 to 1.5 for γ_M and from -0.73 to -1.1 for γ_F .

additional test assets to the conditional CAPM, we find a positive and significant coefficient of relative risk aversion, which controls the intertemporal risk-return tradeoff. The estimates are robust across different test assets, and they remain significant when we include additional state variables. Whereas the previous literature has mostly found insignificant estimates of the risk-return tradeoff when using only the market return, including test assets in addition to the market return consistently produces positive and significant estimates.

The increased significance of the risk-return tradeoff does not arise from an increase in statistical power from including more test assets, as might have been expected. Using simulations, we demonstrate that the high correlation in stock returns mitigates the effective increase in power from adding assets. Instead, the increased significance arises from higher point estimates. When using the market as the only test asset, the point estimate of relative risk aversion is around 3 with a standard error of 2, which does not allow us to reject the null of no intertemporal risk-return tradeoff. However, when including additional test assets, the point estimates increase to above 4, still with standard errors around 2, and the null of no risk-return relation is now strongly rejected.

We also analyze more efficient estimation methods such as seemingly unrelated regressions and conditional generalized least squares. Although these methods are able to extract more power from the highly correlated test assets, this comes at a great cost. The SUR and CGLS methods effectively minimize the pricing errors on orthogonalized portfolios, which are linear combinations of the original test assets. Due to the correlation structure in the original test assets, the orthogonalized portfolios have extreme long/short positions and are not economically meaningful. Moreover, both the SUR and CGLS estimators are severely biased in our sample sizes. We also demonstrate that the impressive test statistics in Bali (2008) and Bali and Engle (2010) are shown to result from the use of incorrect standard errors from the traditional SUR estimation which are four times too small in this heteroskedastic environment.

8 References

Bali, Turan G., 2008, The intertemporal relation between expected returns and risk, *Journal of Financial Economics* 87,101-131.

Bali, Turan G., and Robert F. Engle, 2010, The intertemporal capital asset pricing model with dynamic conditional correlations, *Journal of Monetary Economics* 57, 377-390.

Bali, Turan G., Robert F. Engle, and Yi Tang, 2012, Dynamic conditional beta is alive and well in the cross-section of daily stock returns, manuscript, Georgetown University.

Boguth, Oliver, Murray Carlson, Adlai Fisher, and Mikhail Simutin, 2011, Conditional risk and performance evaluation: Volatility timing, overconditioning, and new estimates of momentum alphas, *Journal of Financial Economics* 102,363-389.

Bollerslev, Tim, 1990, Modelling the coherence in short-run nominal exchange rates: A multivariate generalized arch model, *Review of Economics and Statistics* 72, 498-505.

Bollerslev, Tim, Robert F. Engle, and Jeffrey M. Wooldridge, 1988, A capital asset pricing model with time-varying covariances, *Journal of Political Economy* 96, 116-131.

Campbell, John Y., 1996, Understanding risk and return, *Journal of Political Economy* 104, 298-345.

Campbell, John Y., Stefano Giglio, Christopher Polk, and Robert Turley, 2012. An intertemporal CAPM with stochastic volatility, manuscript, Harvard University.

Campbell, John Y. and Ludger Hentschel, 1992. No news is good news: An asymmetric model of changing volatility. *Journal of Financial Economics* 31, 281-318.

Cochrane, John, 2005, *Asset Pricing: Revised Edition*, Princeton, NJ: Princeton University Press.

Cochrane, John, 2013, A mean-variance benchmark for intertemporal portfolio theory. Forthcoming, *Journal of Finance*.

DeSantis, Giorgio, and Bruno Gerard, 1997, International asset pricing and portfolio

diversification with time-varying risk, *Journal of Finance* 52, 1881-1912.

DeSantis, Giorgio, and Bruno Gerard, 1998, How big is the premium for currency risk? *Journal of Financial Economics* 49, 375-412.

Engle, Robert F., 2002, Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models, *Journal of Business and Economic Statistics* 20, 339-350.

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427-465.

Person, Wayne E., and Campbell R. Harvey, 1991, The variation of economic risk premiums, *Journal of Political Economy* 99, 385-415.

Grant, Dwight, 1977, Portfolio performance and the "cost" of timing decisions, *Journal of Finance* 32, 837-846.

Guo, Hui, and Robert F. Whitelaw, 2006. Uncovering the risk-return relation in the stock market, *Journal of Finance* 61, 1433-1463.

Hansen, Lars P., 1982. Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029-1054.

Hansen, Lars P., and Robert J. Hodrick, 1980. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy* 88, 829-853.

Harvey, Campbell R., 1989, Time-varying conditional covariances in tests of asset pricing models, *Journal of Financial Economics* 24, 289-317.

Hedegaard, Esben, and Robert J. Hodrick, 2014, Estimating the conditional CAPM with overlapping data inference, NBER Working Paper No. 19969.

Jagannathan, Ravi, and Zhengyu Wang, 1996, The conditional capm and the cross-section of expected returns, *Journal of Finance* 51, 3-53.

Lanne, Markku, and Pentti Saikkonen, 2006, Why is it so difficult to uncover the risk-

return tradeoff in stock returns? *Economics Letters* 92, 118-125.

Lettau, Martin, and Sydney C. Ludvigson, 2010. Measuring and modeling variation in the risk-return tradeoff, Chapter 11 in Yacine Ait-Sahalia and Lars P. Hansen, eds., *Handbook of Financial Econometrics: Volume 1 - Tools and Techniques*, Elsevier: Amsterdam, The Netherlands.

Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289–314.

Lundblad, Christian, 2007, The risk return trade off in the long run: 1836–2003, *Journal of Financial Economics* 85, 123–150.

Maio, Paulo, and Pedro Santa-Clara, 2012, Multifactor models and their consistency with the ICAPM, *Journal of Financial Economics* 106, 586-613..

Merton, Robert C., 1973, An intertemporal asset pricing model, *Econometrica* 41, 867–887.

Müller, Gernot, Robert B. Durand, and Ross A. Maller, 2011. The risk-return tradeoff: A COGARCH analysis of Merton’s hypothesis, *Journal of Empirical Finance* 18, 306–320.

Newey, Whitney K., and Daniel McFadden, 1994, Large sample estimation and hypothesis testing, in Robert F. Engle and Daniel L. McFadden, eds., *Handbook of Econometrics: Volume 4*. Amsterdam, The Netherlands: Elsevier Science B.V.

Ng, Lillian, 1991, Tests of the CAPM with time-varying covariances: A multivariate GARCH approach, *Journal of Finance* 46, 1507-1521.

Nyberg, Henri, 2012. Risk-return tradeoff in U.S. stock returns over the business cycle, *Journal of Financial and Quantitative Analysis* 47, 137-158.

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

Scruggs, John T. and Paskalis Glabadanidis, 2003, Risk premia and the dynamic covari-

ance between stock and bond returns, *Journal of Financial and Quantitative Analysis* 38, 295-316.

Shanken, Jay, 1992, On the estimation of beta pricing models, *Review of Financial Studies* 5, 1-34.

Zellner, Arnold, 1962, An efficient method for estimating seemingly unrelated regressions and tests for aggregation bias, *Journal of the American Statistical Association* 57, 348-368.

A Estimation

A.1 Estimating Without a Constant

If asset specific constants are not included in the estimation, equation (7) becomes

$$\sum_{t=1}^T (\mathbf{R}_{t+1} - \mathbf{H}_{NK,t}\boldsymbol{\gamma})' \mathbf{W}_t^{-1} (\mathbf{R}_{t+1} - \mathbf{H}_{NK,t}\boldsymbol{\gamma}) \quad (28)$$

The first order condition with respect to $\boldsymbol{\gamma}$ is

$$\sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} (\mathbf{R}_{t+1} - \mathbf{H}_{NK,t}\boldsymbol{\gamma}) = \mathbf{0}_K \quad (29)$$

where $\mathbf{0}_K$ represents a K -dimensional vector of zeroes. Thus, the solution for $\hat{\boldsymbol{\gamma}}$ is

$$\hat{\boldsymbol{\gamma}} = \left(\frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} \mathbf{H}_{NK,t} \right)^{-1} \frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} \mathbf{R}_{t+1} \quad (30)$$

Notice that this estimator of $\boldsymbol{\gamma}$ does not test the prediction of the theory that the conditional means of returns move with the conditional covariances of the returns with the risk factors. Equation (30) is the multivariate risks analogue to the findings of Hedegaard and Hodrick (2014) for the single factor conditional CAPM.

A.2 Estimating with Constants

When asset-specific constants are included in the estimation, equation (7) becomes

$$\sum_{t=1}^T (\mathbf{R}_{t+1} - \boldsymbol{\mu} - \mathbf{H}_{NK,t}\boldsymbol{\gamma})' \mathbf{W}_t^{-1} (\mathbf{R}_{t+1} - \boldsymbol{\mu} - \mathbf{H}_{NK,t}\boldsymbol{\gamma}) \quad (31)$$

The first order conditions with respect to $\boldsymbol{\mu}$ are

$$\sum_{t=1}^T \mathbf{W}_t^{-1} (\mathbf{R}_{t+1} - \boldsymbol{\mu} - \mathbf{H}_{NK,t}\boldsymbol{\gamma}) = \mathbf{0}_N \quad (32)$$

and the first order conditions with respect to $\boldsymbol{\gamma}$ are

$$\sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} (\mathbf{R}_{t+1} - \boldsymbol{\mu} - \mathbf{H}_{NK,t}\boldsymbol{\gamma}) = \mathbf{0}_K. \quad (33)$$

The solutions are

$$\hat{\boldsymbol{\mu}} = \mathbf{R}_{t+1}^{\mathbf{W}} - \mathbf{H}_{NK,t}^{\mathbf{W}} \hat{\boldsymbol{\gamma}} \quad (34)$$

$$\hat{\boldsymbol{\gamma}} = \left(\frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} (\mathbf{H}_{NK,t} - \mathbf{H}_{NK,t}^{\mathbf{W}}) \right)^{-1} \quad (35)$$

$$\times \left(\frac{1}{T} \sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} (\mathbf{R}_{t+1} - \mathbf{R}_{t+1}^{\mathbf{W}}) \right), \quad (36)$$

where the weighted mean of a matrix process \mathbf{X}_t is defined as $\mathbf{X}_t^{\mathbf{W}} \equiv \left(\sum_{t=1}^T \mathbf{W}_t^{-1} \right)^{-1} \sum_{t=1}^T \mathbf{W}_t^{-1} \mathbf{X}_t$.

If $\mathbf{W}_t \equiv \mathbf{W}$ is constant, the weighted means become sample means, and the solutions are

$$\hat{\boldsymbol{\mu}} = \overline{\mathbf{R}_{t+1}} - \overline{\mathbf{H}_{NK,t}} \hat{\boldsymbol{\gamma}} \quad (37)$$

and

$$\hat{\gamma} = \left(\sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}^{-1} (\mathbf{H}_{NK,t} - \overline{\mathbf{H}_{NK,t}}) \right)^{-1} \left(\sum_{t=1}^T \mathbf{H}'_{NK,t} \mathbf{W}^{-1} (\mathbf{R}_{t+1} - \overline{\mathbf{R}_{t+1}}) \right) \quad (38)$$

Notice that $\hat{\gamma}$ can only be significantly different from zero if weighted sums of the sample covariances of the time t conditional covariances of returns with the risk factors and the future returns at time $t + 1$ are significantly different from zero.

In the special case of constrained OLS, $\mathbf{W} = \mathbf{I}$, and if the only source of risk is the covariance with the market return, equation (38) reduces to

$$\begin{aligned} \hat{\gamma} &= \sum_{t=1}^T \mathbf{H}'_{NM,t} (\mathbf{H}_{NM,t} - \overline{\mathbf{H}_{NM,t}}) \left(\sum_{t=1}^T \mathbf{H}'_{NM,t} (\mathbf{H}_{NM,t} - \overline{\mathbf{H}_{NM,t}}) \right)^{-1} \left(\sum_{t=1}^T \mathbf{H}'_{NM,t} (\mathbf{R}_{t+1} - \overline{\mathbf{R}_{t+1}}) \right) \\ &= \left(\sum_{k=1}^N \hat{V}(\hat{\sigma}_{kM,t}) \right)^{-1} \left(\sum_{i=1}^N \text{c\hat{ov}}(\hat{\sigma}_{iM,t}, R_{i,t+1}) \right) \end{aligned} \quad (39)$$

where $\mathbf{H}_{NM,t}$ is the vector of covariances with market, $\hat{V}(\hat{\sigma}_{kM,t})$ is the sample variance of the conditional covariance, and $\text{c\hat{ov}}(\hat{\sigma}_{iM,t}, R_{i,t+1})$ is the sample covariance of the conditional covariance of asset i with the future return on asset i . Multiplying and dividing within the summation on the last term on the right-hand side by $\hat{V}(\hat{\sigma}_{iM,t})$ and recognizing that individual OLS regressions of $R_{i,t+1}$ on a constant and $\hat{\sigma}_{iM,t}$ would produce estimates of the risk-return tradeoff, $\hat{\gamma}_{i,M} = \frac{\text{c\hat{ov}}(\hat{\sigma}_{iM,t}, R_{i,t+1})}{\hat{V}(\hat{\sigma}_{iM,t})}$, demonstrates that the constrained OLS estimate of γ_M is a weighed average of the individual estimates:

$$\hat{\gamma}_M = \sum_{i=1}^N \frac{\hat{V}(\hat{\sigma}_{iM,t}) \hat{\gamma}_{i,M}}{\sum_{k=1}^N \hat{V}(\hat{\sigma}_{kM,t})}. \quad (40)$$

A.3 GMM Standard Errors

With the notation of Section 2.1, let

$$\mathbf{f}_{t+1} = \mathbf{X}'_t \mathbf{W}_t^{-1} \boldsymbol{\varepsilon}_{t+1}. \quad (41)$$

Then, the GMM orthogonality conditions can be written as $E(\mathbf{f}_{t+1}) = \mathbf{0}_{N+K}$. Denote these sample orthogonality conditions by $\mathbf{g}_T(\boldsymbol{\delta})$ where $\boldsymbol{\delta} = (\boldsymbol{\mu}', \boldsymbol{\gamma}')$ is the $(N + K)$ -dimensional parameter vector. Because the model is just identified, the sample orthogonality conditions can be set to zero exactly.

Hansen's (1982) GMM applied to this case provides the asymptotic distribution of the parameter estimates:

$$\sqrt{T}(\hat{\boldsymbol{\delta}} - \boldsymbol{\delta}) \rightarrow N(\mathbf{0}_{N+K}, \mathbf{d}^{-1} \mathbf{S} \mathbf{d}^{-1}), \quad (42)$$

where $\mathbf{d} \equiv E\left(\frac{\partial \mathbf{f}_{t+1}}{\partial \boldsymbol{\delta}'}\right)$ and $\mathbf{S} \equiv E(\mathbf{f}_{t+1} \mathbf{f}'_{t+1})$. The sample counterparts of \mathbf{d} and \mathbf{S} are $\mathbf{d}_T = \frac{\partial \mathbf{g}_T}{\partial \boldsymbol{\delta}'}$ and $\mathbf{S}_T = \frac{1}{T} \sum_{t=1}^T \mathbf{f}_{t+1} \mathbf{f}'_{t+1}$. For the general case,

$$\mathbf{d}_T = \begin{bmatrix} \frac{1}{T} \sum_{t=1}^T -\mathbf{W}_t^{-1} & \frac{1}{T} \sum_{t=1}^T -\mathbf{W}_t^{-1} \mathbf{H}_{NK,t} \\ \frac{1}{T} \sum_{t=1}^T -\mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} & \frac{1}{T} \sum_{t=1}^T -\mathbf{H}'_{NK,t} \mathbf{W}_t^{-1} \mathbf{H}_{NK,t} \end{bmatrix} \quad (43)$$

A.4 Maximum Likelihood for the DCC Model

The log likelihood function of the DCC model assuming conditional normality with N assets is

$$\begin{aligned} L &= -\frac{1}{2} \sum_{t+1}^T (N \ln(2\pi) + \ln |\mathbf{H}_t| + \boldsymbol{\varepsilon}'_{t+1} \mathbf{H}_t^{-1} \boldsymbol{\varepsilon}_{t+1}) \\ &= -\frac{1}{2} \sum_{t+1}^T (N \ln(2\pi) + 2 \ln |\mathbf{D}_t| + \boldsymbol{\varepsilon}'_{t+1} \mathbf{D}_t^{-1} \mathbf{D}_t^{-1} \boldsymbol{\varepsilon}_{t+1} - \mathbf{u}'_{t+1} \mathbf{u}_{t+1} + \ln |\mathbf{P}_t| + \mathbf{u}'_{t+1} \mathbf{P}_t^{-1} \mathbf{u}_{t+1}) \end{aligned} \quad (44)$$

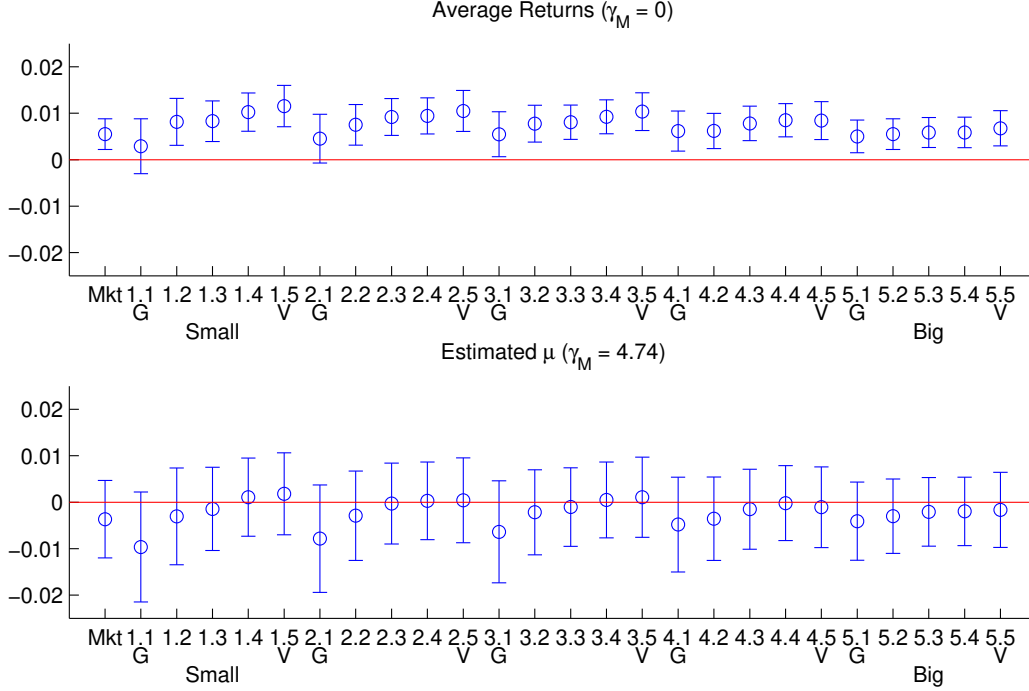
where the term $\boldsymbol{\varepsilon}'_{t+1} \mathbf{D}_t^{-1} \mathbf{D}_t^{-1} \boldsymbol{\varepsilon}_{t+1} - \mathbf{u}'_{t+1} \mathbf{u}_{t+1}$ is zero. As noted by Engle (2002), the likelihood function separates into two parts where the first depends only on the GARCH parameters and the second depends only on the correlation parameters. Let $\boldsymbol{\theta}$ denote the parameters in the \mathbf{D}_t matrices, and let $\boldsymbol{\phi}$ denote the parameters in the \mathbf{P}_t matrices. The volatility part

$$L_V(\boldsymbol{\theta}) = -\frac{1}{2} \sum_{t+1}^T (N \ln(2\pi) + \ln |\mathbf{D}_t| + \boldsymbol{\varepsilon}'_{t+1} \mathbf{D}_t^{-1} \mathbf{D}_t^{-1} \boldsymbol{\varepsilon}_{t+1}) \quad (45)$$

is simply the sum of the individual GARCH likelihoods, which is maximized by separately estimating the individual GARCH processes. Then, conditional on $\hat{\boldsymbol{\theta}}$, the correlation parameters are found by maximizing

$$L_C(\hat{\boldsymbol{\theta}}, \boldsymbol{\phi}) = -\frac{1}{2} \sum_{t+1}^T (-\mathbf{u}'_{t+1} \mathbf{u}_{t+1} + \ln |\mathbf{P}_t| + \mathbf{u}'_{t+1} \mathbf{P}_t^{-1} \mathbf{u}_{t+1}). \quad (46)$$

Figure 1: Average Returns and Estimated Average Pricing Errors

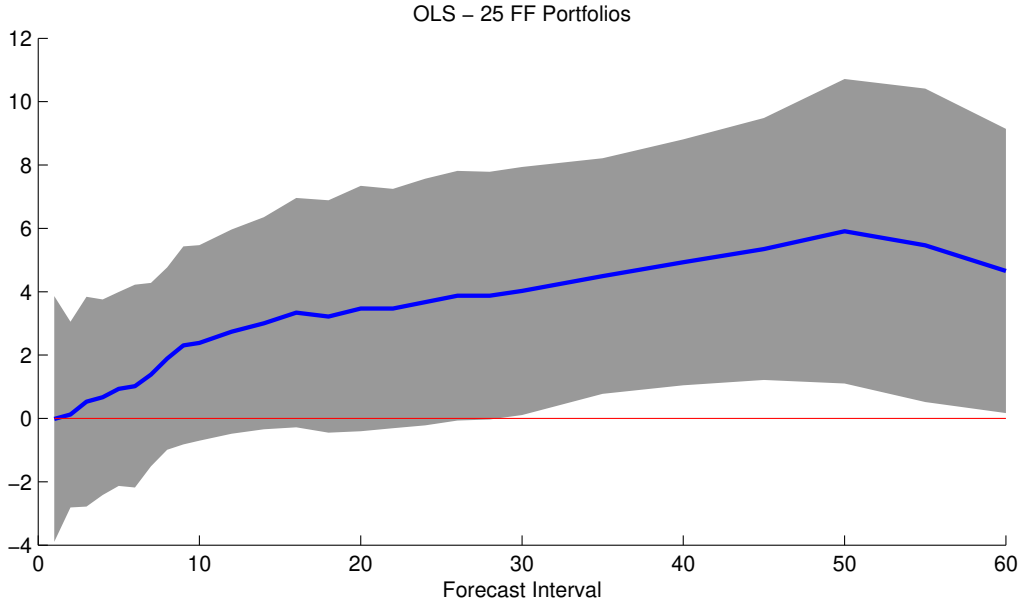


The figure shows estimation results for the intercepts in the ICAPM with the return on the market as the only state variable. The model estimated is

$$\mathbf{R}_{t+1} = \boldsymbol{\mu} + \mathbf{H}_{NK,t}\gamma_M + \boldsymbol{\varepsilon}_{t+1}, \quad \boldsymbol{\varepsilon}_{t+1} \sim D(\mathbf{0}, \mathbf{H}_{NN,t})$$

The system is estimated using OLS, and \mathbf{H}_t is estimated in a first stage using a DCC model. As test assets, we use the market and the 25 size- and book-to-market sorted portfolios. The top plot shows the average monthly returns on each of the 26 test assets along with two standard deviation intervals. These are the pricing errors under the null of $\gamma_M = 0$. The bottom plot shows the constants from the estimated model with their two standard deviation intervals. Although the usual patterns across book-to-market are still clearly visible, the conditional CAPM explains the average return across portfolios very well. Note that the constants are not error terms that are minimized in the estimation, but parameters of the model.

Figure 2: Estimate of γ_M as a function of the forecast interval

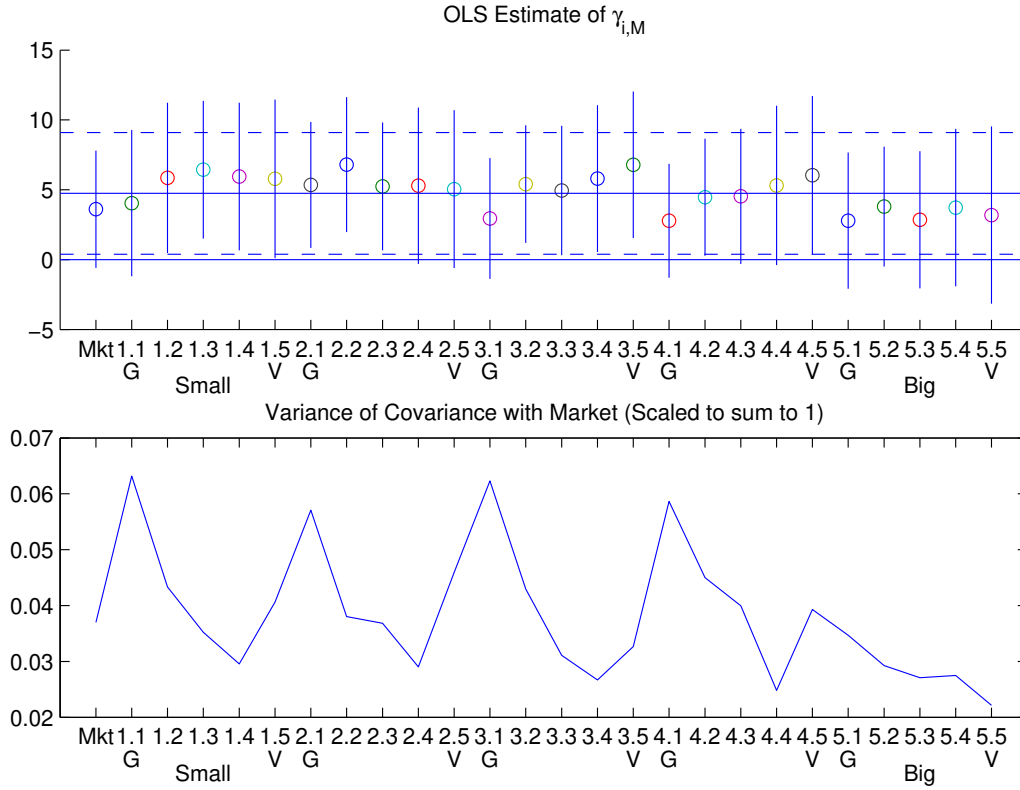


The figure shows estimation results for the risk-return relation in the ICAPM with the return on the market as the only state variable. The model estimated is

$$\mathbf{R}_{t+1} = \boldsymbol{\mu} + \mathbf{H}_{NK,t}\gamma_M + \boldsymbol{\varepsilon}_{t+1}, \quad \boldsymbol{\varepsilon}_{t+1} \sim D(\mathbf{0}, \mathbf{H}_{NN,t})$$

We vary the forecasting interval from one day to 60 days. The system is estimated using OLS, and \mathbf{H}_t is estimated in a first stage using a DCC model. As test assets, we use the market and the 25 size- and book-to-market sorted portfolios. The solid line shows the point estimate of γ_M and the grey area shows plus/minus two standard deviations. The estimation uses daily data from February 1st, 1954, to December 31st, 2012.

Figure 3: Estimates of $\gamma_{i,M}$ and their Weights in the Panel Data Estimation




The top plot shows estimations results for the risk-return relation in the ICAPM based on individual portfolios. The return on the market is the only state variable. The model estimated is

$$R_{i,t+1} = \mu_i + \gamma_{i,M} \text{cov}_t(R_{i,t+1}, R_{M,t+1}) + \varepsilon_{i,t+1}, \quad \varepsilon_{i,t+1} \sim D(0, \sigma_{i,t}^2)$$

The solid horizontal line just below 5 indicates the panel data estimate, and the dashed lines show the corresponding 95% confidence intervals. The panel data estimate is a weighted average of the individual estimates, and the bottom plot shows the individual weight that each estimate receives in the panel data estimator.

Figure 4: Mean of Estimates and Power from Simulations



Panel A shows the mean of the estimates and Panel B shows the power of the test $\gamma_M = 0$ against different alternative values of γ_M , based on simulated data. For each of the seven data sets, we first estimate the parameters of a DCC model to real data. We then simulate from the DCC-GARCH-in-mean model varying γ_M from 0 to 6. For the simulations, we draw innovations randomly with replacement from the original residuals. Using the simulated data, we then estimate a DCC GARCH model and finally we estimate γ_M using the OLS procedure. We perform 1,000 simulations. The left plot shows the mean of the estimates from the 1,000 simulations, for each data set and for each value of γ_M . The estimates are unbiased. The right plot shows the power of the test $\gamma_M = 0$ against different alternative values of γ_M . As the true value of γ_M increases, the power naturally increases. However, the power does not increase as more test assets are added.

Table 1: Estimation Results for the ICAPM with One State-Variable

The table shows estimations results for the risk-return relation with the return on the market as the only state variable. The model estimated is

$$R_{i,t+1} = \mu_i + \gamma_M \text{cov}_t(R_{i,t+1}, R_{M,t+1}) + \varepsilon_{i,t+1}, \quad \varepsilon_{i,t+1} \sim D(0, h_{i,t})$$

The system is estimated using OLS, and \mathbf{H}_t is estimated in a first stage using a DCC model. Robust standard errors are reported in parenthesis. As test assets, we use the portfolios in the first column. For the standard estimation, both asymptotic standard errors and bootstrapped standard errors are shown in parenthesis. For the ODIN estimation, only asymptotic standard errors are reported. We also report the test statistic from a Wald test of the hypothesis that all intercepts are jointly zero, $\hat{\boldsymbol{\mu}}' \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}^{-1} \hat{\boldsymbol{\mu}} \sim \chi^2(N)$, where $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}$ is the estimated covariance matrix for the intercepts and N is the number of test assets. Sample: Monthly observations from February 1954 to December 2012 (707 observations). The ODIN estimation uses daily data from February 1st, 1954, to December 31st, 2012, with a 22 day forecasting interval.

Test-assets	$\hat{\gamma}_M$	$\hat{\sigma}(\hat{\gamma}_M)$	p -value	Wald	p -value	N
Standard Estimation						
Market only	3.61	(2.14)	0.09	0.13	0.71	1
Bootstrap		(2.58)	0.10			
Mkt & 6 Size- and BM	4.46**	(2.16)	0.04	62.72	0.00	7
Bootstrap		(2.28)	0.05			
Mkt & 10 Size	4.77**	(2.19)	0.03	28.39	0.00	11
Bootstrap		(2.39)	0.03			
Mkt & 10 BM	3.91	(2.24)	0.08	18.51	0.07	11
Bootstrap		(2.63)	0.10			
Mkt & 25 Size- and BM	4.74**	(2.22)	0.03	110.99	0.00	26
Bootstrap		(2.64)	0.04			
Mkt & 100 Size- and BM	4.71**	(2.33)	0.04	175.34	0.00	90
Bootstrap		(2.55)	0.04			
Mkt & 10 δ_M sorted	6.86***	(2.20)	0.00	154.01	0.00	11
Bootstrap		(2.65)	0.00			
ODIN Estimation						
Market only	2.65*	(1.59)	0.10	1.51	0.22	1
Mkt & 6 Size- and BM	3.15*	(1.81)	0.08	65.79	0.00	7
Mkt & 10 Size	3.48*	(1.81)	0.05	23.94	0.01	11
Mkt & 10 BM	2.76	(1.86)	0.14	23.91	0.01	11
Mkt & 25 Size- and BM	3.47*	(1.89)	0.07	120.13	0.00	26
Mkt & 100 Size- and BM	3.70*	(2.06)	0.07	207.49	0.00	90
Mkt & 10 δ_M sorted	5.37***	(1.96)	0.01	195.84	0.00	11

***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. For the standard model, the significance is based on bootstrapped p -values.

Table 2: Estimation Results for the CGLS estimation.

The table shows estimations results for the risk-return relation with the return on the market as the only state variable. The model estimated is

$$\mathbf{R}_{t+1} = \boldsymbol{\mu} + \mathbf{H}_{NK,t}\gamma_M + \boldsymbol{\varepsilon}_{t+1}, \quad \boldsymbol{\varepsilon}_{t+1} \sim D(\mathbf{0}, \mathbf{H}_{NN,t})$$

The system is estimated using conditional generalized least squares (CGLS), and \mathbf{H}_t is estimated in a first stage using a DCC model. Both robust asymptotic standard errors and bootstrapped standard errors are shown in parenthesis. As test assets, we use the portfolios in the first column. We also report the test statistic from a Wald test of the hypothesis that all intercepts are jointly zero, $\hat{\boldsymbol{\mu}}' \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}^{-1} \hat{\boldsymbol{\mu}} \sim \chi^2(N)$, where $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}$ is the estimated covariance matrix for the intercepts and N is the number of test assets. Sample: Monthly observations from February 1954 to December 2012 (707 observations). The ODIN estimation uses daily data from February 1st, 1954, to December 31st, 2012, with a 22 day forecasting interval.

Test-assets	$\hat{\gamma}_M$	$\hat{\sigma}(\hat{\gamma}_M)$	p -Value	Wald	p -Wald	N
Market only	2.97	(2.15)	0.17	0.00	0.95	1
Bootstrap		(2.42)	0.12			
Mkt & 6 Size- and BM	-0.43	(1.40)	0.76	81.31	0.00	7
Bootstrap		(1.03)	0.93			
Mkt & 10 Size	2.78***	(1.36)	0.04	18.75	0.07	11
Bootstrap		(1.09)	0.01			
Mkt & 10 BM	0.20	(1.32)	0.88	19.05	0.06	11
Bootstrap		(1.02)	0.49			
Mkt & 25 Size- and BM	-0.15	(1.13)	0.89	156.29	0.00	26
Bootstrap		(0.89)	0.62			
Mkt & 100 Size- and BM	0.31	(0.98)	0.75	246.41	0.00	90
Bootstrap		(0.74)	0.29			
Mkt & 10 δ_M sorted	2.01**	(1.34)	0.13	201.17	0.00	11
Bootstrap		(1.21)	0.02			

***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. For the standard model, the significance is based on bootstrapped p -values.

Table 3: Estimation Results for the ICAPM with the Market and the Long-Term Bond Yield as State-Variables

The table shows estimations results for the ICAPM with the return on the market and the long-term bond yield as state variables. The model estimated is

$$R_{i,t+1} = \mu_i + \gamma_M \text{cov}_t(R_{i,t+1}, R_{M,t+1}) + \gamma_F \text{cov}_t(R_{i,t+1}, \Delta y_{t+1}) + \varepsilon_{i,t+1}, \quad \varepsilon_{i,t+1} \sim D(0, h_{i,t})$$

where Δy_{t+1} is the change in the long-term bond yield. The system is estimated using OLS, and \mathbf{H}_t is estimated in a first stage using a DCC model. Robust standard errors are reported in parenthesis. As test assets, we use the portfolios in the first column. We also report the test statistic from a Wald test of the hypothesis that all intercepts are jointly zero, $\hat{\boldsymbol{\mu}}' \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}^{-1} \hat{\boldsymbol{\mu}} \sim \chi^2(N)$, where $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}$ is the estimated covariance matrix for the intercepts and N is the number of test assets. Sample: Monthly observations from February 1954 to December 2012 (707 observations). The ODIN estimation uses daily data from January 1st, 1962, to December 31st, 2012, with a 22 day forecasting interval.

Test-assets	$\hat{\gamma}_M$	$\hat{\gamma}_F$	Wald	N
Standard Estimation				
Market only	3.47 (2.22)	-23.29 (56.81)	0.18 [0.67]	1
Mkt & 6 Size- and BM	4.34** (2.21)	-24.76 (50.19)	51.12 [0.00]	7
Mkt & 10 Size	4.73** (2.30)	-6.24 (67.49)	29.94 [0.00]	11
Mkt & 10 BM	3.77 (2.32)	-29.77 (64.58)	18.60 [0.07]	11
Mkt & 25 Size- and BM	4.65** (2.36)	-15.16 (87.28)	111.51 [0.00]	26
Mkt & 100 Size- and BM	4.69* (2.70)	-5.83 (163.37)	176.66 [0.00]	90
Mkt & 10 δ_M sorted	6.71*** (2.35)	-27.39 (89.43)	138.00 [0.00]	11
ODIN Estimation				
Market only	3.22** (1.55)	-49.39 (50.30)	1.26 [0.26]	1
Mkt & 6 Size- and BM	3.82** (1.80)	-36.99 (42.40)	22.76 [0.00]	7
Mkt & 10 Size	4.20** (1.85)	-37.08 (64.79)	80.30 [0.00]	11
Mkt & 10 BM	3.35* (1.81)	-54.71 (50.90)	86.08 [0.00]	11
Mkt & 25 Size- and BM	4.08** (1.90)	-60.15 (90.81)	76.70 [0.00]	26
Mkt & 100 Size- and BM	3.85* (2.11)	-98.48 (143.79)	148.27 [0.00]	77
Mkt & 10 δ_M sorted	5.77*** (1.97)	-63.40 (82.07)	157.27 [0.00]	11

***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. For the standard model, the significance is based on bootstrapped p -values.

Table 4: Estimation Results for the ICAPM with the Market, SMB and HML as State-Variables

The table shows estimations results for the ICAPM with the return on the market, SMB and HML as state variables. The model estimated is

$$R_{i,t+1} = \mu_i + \gamma_M \text{cov}_t(R_{i,t+1}, R_{M,t+1}) + \gamma_{SMB} \text{cov}_t(R_{i,t+1}, R_{SMB,t+1}) + \gamma_{HML} \text{cov}_t(R_{i,t+1}, R_{HML,t+1}) + \varepsilon_{i,t+1}, \quad \varepsilon_{i,t+1} \sim D(0, h_{i,t})$$

The system is estimated using OLS, and \mathbf{H}_t is estimated in a first stage using a DCC model. Robust standard errors are reported in parenthesis. As test assets, we use the portfolios in the first column. We also report the test statistic from a Wald test of the hypothesis that all intercepts are jointly zero, $\hat{\boldsymbol{\mu}}' \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}^{-1} \hat{\boldsymbol{\mu}} \sim \chi^2(N)$, where $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}$ is the estimated covariance matrix for the intercepts and N is the number of test assets. Sample: Monthly observations from February 1954 to December 2012 (707 observations). The ODIN estimation uses daily data from January 1st, 1962, to December 31st, 2012, with a 22 day forecasting interval.

Test-assets	$\hat{\gamma}_M$	$\hat{\gamma}_{SMB}$	$\hat{\gamma}_{HML}$	Wald	N
Standard Estimation					
Market only	7.90*** (2.81)	-10.68 (6.86)	11.87** (5.79)	0.38 [0.54]	1
Mkt & 6 Size- and BM	5.48** (2.58)	-0.94 (3.37)	6.22* (3.48)	45.91 [0.00]	7
Mkt & 10 Size	6.76** (2.76)	-1.98 (4.44)	8.38* (5.03)	13.52 [0.26]	11
Mkt & 10 BM	5.52* (2.98)	-4.13 (5.95)	9.55* (5.03)	7.77 [0.73]	11
Mkt & 25 Size- and BM	6.64** (2.88)	-1.59 (4.77)	10.10* (5.73)	93.80 [0.00]	26
Mkt & 100 Size- and BM	7.54** (3.17)	-2.41 (5.10)	17.34*** (6.47)	158.30 [0.00]	90
Mkt & 10 δ_M sorted	9.38*** (3.59)	-3.81 (6.35)	16.55** (7.52)	117.67 [0.00]	11
ODIN Estimation					
Market only	3.81* (2.10)	0.28 (6.71)	5.62 (4.68)	0.25 [0.62]	1
Mkt & 6 Size- and BM	3.59* (2.07)	0.68 (3.92)	2.78 (3.38)	15.07 [0.04]	7
Mkt & 10 Size	4.46** (2.25)	0.19 (4.89)	4.69 (5.18)	70.75 [0.00]	11
Mkt & 10 BM	3.14 (2.28)	1.14 (6.05)	3.25 (4.40)	88.11 [0.00]	11
Mkt & 25 Size- and BM	4.37* (2.29)	-0.08 (5.04)	5.14 (5.65)	64.34 [0.00]	26
Mkt & 100 Size- and BM	5.16** (2.57)	-0.74 (5.76)	9.61 (6.45)	138.77 [0.00]	77
Mkt & 10 δ_M sorted	5.39* (2.85)	2.75 (6.82)	7.80 (6.52)	110.62 [0.00]	11

***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively. For the standard model, the significance is based on bootstrapped p -values.

Table 5: Replication of Bali and Engle (2010) Table 1

The table shows the original results from Table 1 in Bali and Engle (2010) in Panel A, as well as our replications in Panel B and Panel C. The return on the market is the only state variables, and the model estimated is

$$\mathbf{R}_{t+1} = \boldsymbol{\mu} + \mathbf{H}_{NK,t}\gamma_M + \boldsymbol{\varepsilon}_{t+1}, \quad \boldsymbol{\varepsilon}_{t+1} \sim D(\mathbf{0}, \mathbf{H}_{NN,t})$$

The system is estimated using seemingly unrelated regressions (SUR) in Panel A and Panel B, and using OLS in Panel C. \mathbf{H}_t is estimated in a first stage using a DCC model. In Panel B, the SUR standard errors correct for contemporaneous (i.e., cross-equation) correlation and heteroskedasticity, but not for time-series heteroskedasticity. The GMM standard errors correct for both contemporaneous correlation and heteroskedasticity, as well as for time-series heteroskedasticity. As test assets, we use the portfolios in the first column. We also report the test statistic from a Wald test of the hypothesis that all intercepts are jointly zero, $\hat{\boldsymbol{\mu}}' \hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}^{-1} \hat{\boldsymbol{\mu}} \sim \chi^2(N)$, where $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\mu}}$ is the estimated covariance matrix for the intercepts and N is the number of test assets. Sample: Daily observations from January 3, 1972 to June 30, 2009 (9,462 observations). The estimation for Dow 30 stocks is based on daily observations from July 10, 1986 to June 30, 2009 (5,795 observations).

Test-assets	γ_M	Wald	p -value	N
Panel A: Original Results in Bali and Engle (2010) Table 1				
10 Size	1.86*** (0.37)	13.29	0.21	10
10 Book-to-Market	2.05*** (0.39)	15.73	0.11	10
10 Momentum	3.32*** (0.38)	36.71	0.00	10
10 Industry	1.85*** (0.38)	9.53	0.48	10
30 Dow Stocks	2.21*** (0.29)	18.17	0.96	30

Continued...

Table 5 continued...

Test-assets	γ_M	Wald	p -value	N
Panel B: Replication using SUR				
10 Size	-0.18			
SUR Standard err.	(0.39)	9.10	0.52	10
GMM Standard err.	(1.41)	9.25	0.51	
10 Book-to-Market	1.78			
SUR Standard err.	(0.37)	15.74	0.11	10
GMM Standard err.	(1.63)	16.10	0.10	
10 Momentum	3.04*			
SUR Standard err.	(0.37)	34.89	0.00	10
GMM Standard err.	(1.57)	38.25	0.00	
10 Industry	1.74			
SUR Standard err.	(0.37)	9.71	0.47	10
GMM Standard err.	(1.60)	9.66	0.47	
30 Dow Stocks	3.84*			
SUR Standard err.	(0.42)	21.43	0.87	30
GMM Standard err.	(2.19)	22.52	0.83	
Panel C: Replication using OLS				
10 Size	-0.43	8.08	0.62	10
	(2.04)			
10 Book-to-Market	0.23	18.06	0.05	10
	(2.00)			
10 Momentum	1.06	34.57	0.00	10
	(1.93)			
10 Industry	0.62	10.48	0.40	10
	(2.00)			
30 Dow Stocks	1.63	21.25	0.88	30
	(2.18)			

***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Replication and Update of Scruggs (1998)

This table shows the original results from Scruggs (1998) in the second column, as well as our replication in the third column and updated results in the fourth column. The model estimated is

$$\begin{aligned}
 r_{M,t} &= \gamma_0 + \gamma_M \sigma_{M,t}^2 + \gamma_F \sigma_{MF,t} + \varepsilon_{M,t} \\
 r_{F,t} &= \mu_F + \varepsilon_{F,t} \\
 \log(\sigma_{M,t}^2) &= \omega_M + \alpha_M g(\psi_{M,t-1}) + \delta_M \log(\sigma_{M,t-1}^2), \quad \psi_{M,t-1} = \varepsilon_{M,t-1} / \sigma_{M,t-1} \\
 \log(\sigma_{F,t}^2) &= \omega_F + \alpha_F g(\psi_{F,t-1}) + \delta_F \log(\sigma_{F,t-1}^2), \quad \psi_{F,t-1} = \varepsilon_{F,t-1} / \sigma_{F,t-1} \\
 \sigma_{MF,t} &= \rho_{MF} \sigma_{M,t} \sigma_{F,t} \\
 g(\psi_{M,t-1}) &= \left(|\psi_{M,t-1}| - \sqrt{2/\pi} \right) + \theta_M \psi_{M,t-1} \\
 g(\psi_{F,t-1}) &= \left(|\psi_{F,t-1}| - \sqrt{2/\pi} \right) + \theta_F \psi_{F,t-1}
 \end{aligned}$$

Scruggs' original study uses monthly data from March 1950 to December 1994 (538 observations) and the return on long-term government bonds from Ibbotson Associates, whereas we base our replication on data from January 1952 to December 1994 (516 observations) and use the return on the Fama Bond Portfolios for 61-120 months from CRSP. The updated results are based on data from January 1952 to December 2012 (732 observations). Heteroskedasticity robust standard errors are reported in parenthesis.

Parameter	Original	Replication	Update
Panel A: Conditional Mean Equation Parameters			
$\gamma_0 \times 100$	-0.11 (0.55)	-0.07 (0.74)	0.54 (0.92)
γ_M	10.57** (4.48)	7.62 (7.59)	-0.30 (12.96)
γ_F	-49.75** (23.47)	-61.58 (63.87)	0.20 (417.97)
$\mu_F \times 100$	-0.08* (0.04)	-0.03 (0.04)	0.02 (0.04)

Continued...

Table 6 continued...

Panel B: Conditional Second Moment Equation Parameters			
ω_M	-0.49** (0.25)	-0.52 (0.39)	-0.87 (1.50)
α_M	0.19*** (0.05)	0.18*** (0.04)	0.19** (0.09)
δ_M	0.92*** (0.04)	0.92*** (0.06)	0.86*** (0.24)
θ_M	-0.17 (0.24)	-0.27 (0.48)	-0.89 (1.92)
ω_F	-0.18** (0.07)	-0.14 (0.11)	-0.14* (0.08)
α_F	0.38*** (0.05)	0.31*** (0.10)	0.27*** (0.06)
δ_F	0.97*** (0.01)	0.98*** (0.01)	0.98*** (0.01)
θ_F	-0.27** (0.12)	-0.41*** (0.15)	-0.25* (0.14)
ρ_{MF}	0.24*** (0.05)	0.19*** (0.05)	0.05 (0.07)
L	3318.53	2892.40	4050.01
N	538	516	732

***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.