Online Appendix to
One Vol to Rule Them All:
Common Volatility Dynamics in Factor Returns

August 11, 2022

This not-for-publication document contains additional material and results referenced in the main paper.

This Online Appendix provides additional discussion and results associated with the paper. Section A.1 discusses measurement error associated with realized (co)-variances and whether associated biases could explain our evidence of CFV. Section A.2 considers commonality in correlations across different factor sets. Section A.3 describes how we construct statistical factors that are analyzed in the main paper. Section A.4 provides details concerning our construction of mimicking portfolios for macroeconomic factors. Section A.5 discusses an alternative approach to forecasting factor volatilities using a reduced rank regression approach. Section A.6 considers impulse response functions for shocks to the vector autoregressive index model introduced in Section A.5. Finally, Section A.7 further discusses tests involving CFV and aggregate measures of financial and operating leverage.

A.1. Measurement error in volatility proxies

In the main paper, we measure the volatility of various factors and anomaly portfolios via realized (co)-variances that are computed as the sum of squared daily returns over lower frequency intervals (months, quarters, or years). Here we discuss measurement error associated with these proxies and evaluate the possible argument that corresponding measurement error could spuriously drive our main results concerning common factor volatility (CFV). We first discuss the nature of the potential problem. Then we discuss results under alternative approaches that mitigate or avoid the problem.

A.1.1. Measurement error and biases in realized variances and co-variances

The literature on volatility measurement using high-frequency data assumes that asset values follow a continuous time jump-diffusion process. To avoid technical details that are inessential to the key measurement issues, we focus on a setting without jumps. In this case, it is commonly assumed that the d-dimensional log price process evolves according to a (multivariate) diffusion process:

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s, \tag{A.1}$$

where b_s is a d-dimensional progressively measureable drift process, σ is a $d \times d$-dimensional progressively measurable stochastic volatility process, and W_s is a d-dimensional standard Brownian motion.[1] Some additional technical assumptions are imposed on the process, see, e.g., Aït-Sahalia and Jacod (2014). Define the $d \times d$ integrated covariance of X_t as:

$$IC_t = \int_0^t \Sigma_s ds, \tag{A.2}$$

where $\Sigma_s = \sigma_s \sigma_s'$ is the $d \times d$ 'spot-covariance' process associated with X_t.

Under ideal conditions, high-frequency financial data permit highly precise measurement of the integrated covariance IC_t. Suppose, for example, that the process X_t is observed on a regular sampling grid with time interval Δ_n between observations. Define the realized covariance as the $d \times d$ matrix

$$RC(\Delta_n)_t = \sum_{j=1}^{\lfloor t/\Delta_n \rfloor} \Delta X_j \Delta X_j', \tag{A.3}$$

where $\Delta X_j = X_j - X_{j-1}$, i.e., the $d \times 1$ return vector during the j-th intraperiod interval. In

[1] It is possible to consider more general jump-diffusion processes as well, by adding a (finite variation) jump component J_t to the price process in Eq. (A.1). Truncated alternatives to the realized (co)-variation measures discussed here allow for the separation of the jump versus diffusive components of return variation as the sample frequency increases. However, our application primarily involves daily returns and we do not aim to separate the jump versus diffusive component of return variation in the paper. Consequently, we omit jumps from this discussion for tractability.

this setting, the realized covariance is a consistent estimator of the integrated covariance, i.e.. $RC(\Delta_n)_t \xrightarrow{p} IC_t$ as $\Delta_n \to 0$. Heuristically, this result implies that given sufficiently finely sampled returns, the integrated covariance is effectively observed. Barndorff-Nielsen and Shephard (2004) provide an asymptotic distribution theory associated with the convergence result here, which has subsequently been extended to even more general processes.

Our analysis in the main paper effectively treats the integrated covariance as truly 'realized,' i.e., we implicitly assume $RC(\Delta_n)_t = IC_t$. Given this implicit assumption, we focus on the properties of IC_t for factor sets and in particular on measuring common dynamics in the diagonal elements of IC_t and their square roots (the realized volatilities for the factors). To more clearly fix ideas, suppose that the object of interest is a real-valued function of the time series of integrated covariances for a set of factors or assets, denoted $g(IC_1, .., IC_Q)$, where $1, .., Q$ indexes, e.g., quarters. For example, the function of interest might be the proportion of variation explained by the first principal component of the realized volatility series for the underlying factors or assets. In the main paper, we estimate this quantity of interest by replacing the true unknown integrated covariances with their empirical counterparts, i.e., we consider $g(RC(\Delta_n)_1, ..., RC(\Delta_n)_Q)$. Although this 'plug-in' approach seems natural, in general the estimator $g(IC_1, .., IC_Q)$ suffers from small sample bias that in some cases can be economically significant (see, e.g., Vetter (2015) and Kalnina and Tewou (2019)).

The key concern with respect to our analysis involves whether small sample bias has a first-order effect upon the measured strength of commonality in factor volatility series. In other words, the concern is whether there could be a sufficiently large upward bias in the estimated first eigenvalue associated with a PCA decomposition of realized volatility series so as to drive apparent strong commonality among factor volatility series when in fact the true commonality is much weaker. One can raise similar small sample bias concerns regarding other related empirical exercises in the paper, such as loadings regressions of factor volatility series on a (high-frequency) proxy for common factor volatility.

Before proceeding, we emphasize that the concern described above is distinct from a different concern oriented around the role of frictions in the observed price process, i.e., market microstructure noise. In the presence of noise, the econometrician does not observe the true price process X_t but rather observes $Y_t = X_t + U_t$, where U_t is a process that captures measurement error driven by various forms of microstructure effects. The precise effects of microstructure noise on the standard realized (co)-variance depend on the particular assumptions imposed on the noise process. One leading case assumes Gaussian white noise, i.e., a noise distortion that is $N(0, \phi^2)$ independent of the sampling frequency. In this case, as the sampling frequency becomes arbitrarily fine, the standard realized variance no longer converges in probability to the integrated variance but instead explodes and reflects more the variance of the contamination noise than the true integrated variance (see, e.g., Zhang, Mykland, and Aït-Sahalia (2005)). From the perspective of our analysis, if the severity of microstructure noise varies over time in common among a set of factors or anomaly portfolios, then realized volatility measures that contain bias due to microstructure noise might co-move due to this bias rather than true commonality in the underlying portfolio variances.

A.1.2. Robustness of key results

In this section, we discuss various robustness checks that indicate that our main result of CFV is unlikely to be driven by biases associated with realized (co)-variance quantities. First, concerning microstructure noise, we note that microstructure noise is most severe at high intraday sampling frequencies. Most of our results rely on daily returns, which offer less precise estimates of volatility and related quantities under ideal conditions relative to intradaily data, but are less affected by microstructure-driven biases. Still, because factor and anomaly portfolios potentially contain positions in small, relatively illiquid stocks, there remains a concern in this direction. However, we address this concern in the main paper by looking at an alternative set of daily factor returns based on the mid-point between the daily bid and ask closing prices recorded by CRSP (see Table 2). We obtain similar evidence of CFV using these alternative realized variances, and therefore we conclude that market microstructure effects are unlikely to be the ultimate source of CFV that we document in the paper.

We now turn to the issue of whether CFV could be driven by small sample biases associated with computing functions of realized covariance matrices such as eigenvalues or correlations. We conduct two different robustness checks. The first check continues to measure volatility as the sum of squared daily (market-adjusted) returns, but measures volatility at an *annual* rather than quarterly frequency. Moving to an annual measurement frequency reduces the noise associated with realized (co)-variances due to the fact that the corresponding estimates now utilize essentially four times the intra-period data relative to the quarterly case.[2] If it were the case that our evidence of CFV at the quarterly frequency was largely driven by biases associated with estimation noise in volatility proxies, then we would expect to observe materially different results upon moving to the annual frequency. However, this is not the case. We continue to find strong evidence of CFV when measuring factor volatility at the annual frequency. The first principal component of annual factor volatility series remains about 60% for the same factor sets reported in Table 1 of the paper.

As a second alternative approach, we apply parametric generalized autoregressive conditional heteroskedasticity (GARCH) models of the conditional variance of low frequency

[2] As Kalnina and Tewou (2019) note, "The biases due to preliminary estimation of volatility can be made theoretically negligible when an additional, long-span, asymptotic approximation is used."

(quarterly) factor and anomaly portfolio returns.[3] A simple but appealing benchmark model is the GARCH (1,1) model with constant mean:

$$r_t = \mu + u_t \tag{A.4}$$

$$h_t = \omega + \alpha u_{t-1}^2 + \beta h_{t-1} \tag{A.5}$$

$$u_t \sim N(0, h_t), \tag{A.6}$$

where r_t denotes the excess return on a particular factor or anomaly portfolio, u_t is the unpredictable return component, h_t equals the conditional variance of factor returns, and μ, ω, α, and β are parameters.[4] We estimate the model of Eq. (A.4)–(A.6) for each anomaly or portfolio in a particular set using quarterly returns. Next, we compute the time series of fitted conditional variances $\hat{h_{i,t}}$, where $i = 1, \ldots, N$ denotes different portfolios in the set. We then extract a common component from the fitted conditional volatility series ($\sqrt{\hat{h_t}}$) using several alternative approaches:

- As the cross-sectional average of the fitted volatility series
- As the first principal component extracted from the fitted volatility series
- As the first principal component extracted from the standardized fitted volatility series

Key results are as follows. First, the first PC extracted from the fitted volatility series explains a significant fraction of variation in the data. As a concrete example, the first PC explains almost 60% of the variation in the GARCH-based volatility series for these

[3] A third possible approach would involve following Pelger (2020) (2017 working paper version), who uses implied volatilities from option data to mitigate the measurement problems associated with realized quadratic covariations over short intervals. However, Pelger (2020) examines volatilities for (total) returns on individual stocks. In contrast, we focus on long-short factor and anomaly portfolios. Unfortunately, traded options associated with these factor portfolios are not generally available, and therefore we do not pursue this approach.

[4] More explicitly, $h_t = E_{t-1}[r_t - E_{t-1}(r_t)]$.

portfolios. This is, of course, roughly analogous to results we obtain using the 'realized variance' approach in the main paper. Second, the common volatility component extracted using this alternative GARCH-based approach is similar to the CFV proxy obtained using the realized variance approach in the main paper. Figure A.1 illustrates this result for the KNS anomaly portfolios. The contrasts the standardized quarterly CFV measure based on the realized variance approach used in the main paper with the analog based on the alternative GARCH-based approach. Note that the GARCH model is for the conditional variance, i.e., the ex ante expectation of volatility conditional on available information. Thus, we filter the realized variance-based CFV proxy via a ARMA(1,1) filter to obtain the analogous conditional expectation. These are the series plotted in Figure A.1. The time series correlation of the two series is approximately 0.71. We explored several variations on the GARCH modeling approach including, for example, allowing for some persistence in quarterly factor returns by including lagged returns in the mean equation. These variations produce qualitatively similar results.

The low-frequency GARCH-based approach described above has some weaknesses. First, the approach is parametric and requires an explicit specification of both the conditional mean and volatility function for returns. Thus, there is a concern regarding misspecification that is absent from the nonparametric high-frequency approach to volatility measurement. In addition, model parameters must be estimated from data, i.e., the fitted conditional volatility series do not equal the true volatility series even if the model is correctly specified due to parameter estimation error. But this parameter estimation error is of a rather different nature than the high-frequency estimation error associated with realized variance series constructed from daily returns. Consequently, the fact that we obtain similar core insights regarding the presence and nature of CFV using the parametric GARCH method is quite reassuring.

A.2. Commonality in correlations?

We now consider commonality in the dynamics of correlations among factor and anomaly portfolios. Table A.2 provides summary statistics regarding quarterly correlation measures for raw and market-adjusted factor and anomaly portfolios. For each portfolio within a specified set, we compute time series of quarterly realized correlations as the sample pairwise correlations of daily returns within the corresponding quarter. As an example, if the factor set consists of five factors, we obtain 10 quarterly time series, each reflecting the dynamics of a particular pairwise correlation among these factors. We then compute a simple measure of the common component of the correlations as the first principal component extracted from the pairwise correlations. The table shows the percent of total variation captured by this common component measure ('% Expl.'). To shed light on whether the dynamics of the common correlation component are similar to that of common factor volatility, we report the time series correlations between the common correlation measure and a CFV measure based on market-adjusted industry portfolios. The right-hand side of Table A.2 provides similar statistics for correlations constructed using market-adjusted factor returns. This helps convey to what extent common correlation arises due to common market exposure.

The common correlation measures explain much less of the total variation for most portfolio sets relative to the common volatility series described in Table 1. For raw factor correlations, the percent explained is between 10–30% with the exception of the long-only industry portfolios which are subject to common market exposure. Upon examining correlations constructed from market-adjusted returns, the fraction of variation explained by the first principal component of the correlation series is always under 30%. Correlations with the CFV measure are relatively weak. Although factor return correlations exhibit time series variation, there does not appear to be a dominant common factor, and the common component of variation does not consistently closely relate to time-varying market volatility.

A.3. Statistical Factor Methods

This Appendix discusses our implementation of methods to extract statistical factors from stock returns. We assume that returns adhere to an approximate factor model:

$$r_t = Bf_t + \epsilon_t, \tag{A.7}$$

where r_t denotes an $N \times 1$ realization of de-meaned stock returns at time t, f_t denotes a a $K \times 1$ vector of random factors, and B equals a (fixed) $N \times K$ factor loading matrix B. We construct statistical factors using two general approaches in the paper. One approach constructs factors from an underlying set of portfolios formed by sorts on various firm characteristics associated with anomalies. The second approach constructs factors from *individual firm-level stock returns* and closely follows Connor and Korajczyk (1986, 1988); Connor, Korajczyk, and Linton (2006).

Let R equal the $T \times N$ matrix of realized returns over a sample of length T and F denote the $T \times K$ realizations of factors. We assume that returns are generated by an approximate factor model with a fixed $N \times K$ loadings matrix B such that $(1/N)B'B$ approaches a nonsingular limiting matrix as N becomes large. The approximate factor model does not require that the covariance matrix of ϵ is diagonal, as in a so-called strict factor model, but limits the extent of correlation by bounding the largest eigenvalue associated with the error covariance matrix (Chamberlain and Rothschild (1983)). The asymptotic principal components approach performs an eigen-decomposition on the $T \times T$ cross-product matrix:

$$\Sigma = (1/N)r'r. \tag{A.8}$$

Connor and Korajczyk (1986, 1988) show that the first K eigenvectors of the cross-product matrix in Eq. (A.8) are consistent estimates for the space spanned by the factors F, in the sense of converging in probability to LF, where L denotes a $K \times K$ nonsingular matrix.[5] This result assumes a constant limiting average idiosyncratic return variance. In practice, it is likely that idiosyncratic return variance changes over time (consistent with evidence of CIV discussed in the main paper). Jones (2001) provides a modification of the method that is robust to the presence of time-variation in the average limiting idiosyncratic variance.

We apply the asymptotic principal components method each year to individual stocks listed on NYSE and subject to the following screens: 1) the stock has non-missing returns for at least 50% of trading days during the year; and 2) no more than 30% of observed returns equal exactly zero in the year. These screens, as well as our focus on NYSE stocks are intended to address liquidity issues especially given that we perform the decomposition using daily stock returns. We address the unbalanced panel structure by computing the cross-product matrix element-by-element using only the set of securities that have returns for both time dates. As robustness checks, we explore variations on the sample screens applied to firms. We also find that estimated factors using the standard asymptotic principal components method are similar to those we obtain using the Jones (2001) procedure.

In addition to performing statistical factors decompositions for individual stocks, we also extract statistical factors from sets of *portfolios*. In this case, the method applied differs

[5] Subsequent papers, e.g., Bai (2003), extend the approach to settings with N and $T \to \infty$, whilst also permitting heteroskedasticity and limited dependence in the cross-section and time series.

because the sample size T is greater than the number of assets N. In this case we apply the RP-PCA method of Lettau and Pelger (2020). This approach represents a generalized PCA that permits the incorporation of a form of 'penalty term' to account for cross-sectional pricing errors. The approach boils down to applying PCA to the $N \times N$ matrix:

$$\widehat{\text{cov}}(r_t) + \gamma \, \overline{r} \, \overline{r}', \tag{A.9}$$

where the first term equals the sample covariance matrix of returns, $\overline{r}$ denotes the $N \times 1$ vector of average portfolio excess returns, and γ reflects the penalty term for pricing errors. Setting $\gamma = -1$ results in standard PCA. Following Lettau and Pelger (2020), we set $\gamma = 10$ and estimate factors using monthly portfolio returns. As shown by Lettau and Pelger (2020), the latter approach produces materially different factors that better fit the cross-section of returns (out-of-sample). The RP-PCA analysis produces a set of K loadings estimates, and factors are then obtained via standard projection methods. Given the loadings estimates from the RP-PCA procedure applied to monthly returns, we then construct daily factor returns that are used in order to compute realized variances and related quantities. For the 'LP factors' based on 74 portfolios corresponding to extreme deciles based on 37 anomaly sorts, we confirm that our implementation produces factors that closely match statistics presented by Lettau and Pelger (2020) for factors based on the same portfolios using their sample period. We then construct an alternative set of RP-PCA factors ('LP-ALT') based on 80 Fama-French industry and characteristics-sorted portfolios that cover the longer sample period 1930–2020. The portfolios consist of 30 industry-sorted portfolios, along with 10 decile portfolios based on each of the following characteristics: size, book-to-market, momentum, short-term reversal, and long-term reversal.

The final set of statistical factors we construct, denoted 'HKS factors' follows Haddad, Kozak, and Santosh (2020) as closely as possible. In particular, we utilize monthly returns for the same underlying set of decile-sorted portfolios for 50 anomaly characteristics. We market- and volatility-adjust these returns following the same procedure as Haddad et al. (2020), and we then estimate factor loadings by applying PCA (equivalent to $\gamma = -1$ in RP-PCA) applied to the first half of the sample. These fixed loadings estimated over the 'in-sample' period are then applied to generate factors over the second 'out-of-sample' half of the total sample period, which is 1974–2020.

A.4. Macroeconomic Mimicking Portfolios

We construct mimicking portfolios for the Chen, Roll, and Ross (1986) factors following the method of Cooper and Priestley (2011). For each of the six risk factors (unexpected inflation, industrial production growth, bond risk premium, term spread, real interest rate and consumption growth), factor realizations are projected onto the space of excess returns. We obtain the risk free rate and 24 portfolios from Kenneth French's website. The 24 portfolios are made up of 6 size and book-to-market sorted portfolios, 6 size and operating profitability sorted portfolios, 6 size and investment sorted portfolios, 6 size and momentum sorted portfolios. Each portfolio's excess return is first regressed on the 6 macro risk factors. That is we run 24 time series regressions to get a 24 by 6 matrix of loadings. We augment this matrix with a column of ones to include an intercept term in the second stage where we define weights. This matrix is denoted B. As in Cooper and Priestley (2011) we define the weight vector for the 24 portfolios by

$$w = (B'V^{-1}B)^{-1}B'V^{-1}$$

where V is the 24 by 24 diagonal matrix with variances of the residuals from the 24 time series regressions along the diagonal. The mimicking portfolios are then defined as the weighted average of excess returns on the 24 portfolios where weights are given by w.

A.5. Volatility Index Models

This section describes the alternative reduced rank regression procedure used to estimate a volatility index model for volatility forecasting.

A.5.1. VARI Model and Reduced Rank Regression Method

As motivation, let Y_t denote a $K \times 1$ vector time series of factor or anomaly portfolio volatility series. These might be realized volatility series, realized variances, or log transformations of realized variance/volatility. A natural starting point to model the joint dynamics of the factor volatility series is the standard vector autoregression (VAR) model:

$$Y_t = \mu + \Phi_1 Y_{t-1} + \cdots + \Phi_p Y_{t-p} + \epsilon_t, \tag{A.10}$$

where Φ_i for $i = 1, \ldots, p$ denote $K \times K$ matrices of coefficients that characterize factor volatility dynamics as a function of the past p volatility realizations, and ϵ_t denotes a (possibly heteroskedastic) vector white noise process.

A drawback of the 'unrestricted' VAR model of Eq. (A.10) is the fact that the number of parameters grows at a rate proportional to K^2. Consequently, a large number of parameters must be estimated even for modest-sized systems. In the context of modeling factor realized volatility series, it is common to ignore dynamic relations among factors and apply a standard autoregressive (AR) specification, or variants such as the heterogeneous autoregressive (HAR) model of Corsi (2009). Many papers apply models of this type, sometimes augmented with additional lagged financial or macroeconomic predictors (see, e.g., Schwert (1989) for market volatility, and Moreira and Muir (2017) for characteristics-based factor volatility).

The class of 'VAR-index' (VARI) models (Reinsel (1983)) offers an approach that efficiently incorporates multivariate volatility information. The simplest version of the VARI model takes the form:

$$Y_t = \mu + \beta I_{t-1} + \epsilon_t, \tag{A.11}$$

where $I_{t-1} = \omega' Y_{t-1}$ is a scalar index constructed as a linear combination of lagged series Y_{t-1} and β reflects the index 'loadings' for each volatility series. Under this model, a one-dimensional index I_t drives all dynamics for the multivariate volatility series Y_t. The model can easily be extended to permit $q < K$ indices, although we focus on single-index models in this paper, and to accommodate additional lags.[6]

The assumption that a single scalar volatility index fully captures volatility dynamics across a potentially large set of factor or anomaly portfolios is strong. We therefore consider a more general specification that permits dynamics to depend both on the lagged index as well as the lagged own-factor volatility. The expanded model is:

$$Y_t = \mu + \phi^D Y_{t-1} + \beta I_{t-1} + \epsilon_t, \tag{A.12}$$

where ϕ^D is a $K \times K$ *diagonal* matrix. The null hypothesis that $\beta = 0$ for all volatility series is of particular interest. This reflects the case in which there exists no added value

[6]It is straightforward to show that the indices themselves follow a corresponding VAR process under the VARI model. Cubadda, Guardabascio, and Hecq (2017) show how the VARI model can be extended to incorporate heterogeneous autoregressive (HAR) structures in the spirit of Corsi (2009).

from incorporating the index in the volatility forecasting model in addition to the lag of own-factor volatility.

Given the weights ω that define the volatility index $I_t = \omega' Y_t$ in Eq. (A.12), other key parameters (μ, ϕ^D, and β) can be estimated via standard linear regression. But how to obtain the weights? We take a simple approach in the main paper that directly associates the volatility index I_t with CFV. Implicitly, this amounts to taking ω to equal the weighting vector for individual volatility series in CFV. Below we discuss an alternative reduced rank regression procedure that directly estimates the weights along with other model parameters. An advantage of this approach is that the volatility index weights are chosen in order to maximize predictive power for volatility within the model. Drawbacks are that the estimation approach is more involved and weights must be estimated recursively in order to conduct out-of-sample forecast analysis.

The reduced rank estimation procedure essentially amounts to canonical correlation analysis (Hotelling (1936)) and is analogous to the procedure used to estimate error correction models related to cointegrated systems (Engle and Granger (1987)). To describe the approach formally, we focus on the simpler pure VARI model of Eq. (A.11). The reduced rank regression model can be written in the general form:

$$Y = \mu + A\,B\,X + \epsilon, \tag{A.13}$$

where Y, μ, and ϵ are $K \times 1$ vectors, X is an $R \times 1$ vector, A is $K \times q$, B is $q \times r$ (both full rank). In our application, X represents the lagged dependent variable, so that $R = K$ and the matrix $C \equiv A\,B$ is $K \times K$ but of reduced rank q.[7] It is assumed that the vector

[7]The decomposition $C = AB$ is not unique, because for any nonsingular conformable matrix G it is the case that $C = (AG^{-1})(GB)$, which is an alternative decomposition of the same form.

$Z = (Y', X')'$ is (covariance) stationary. Let μ_Y and μ_X denote the means of Y and X, respectively. Σ_{XX} denotes the covariance matrix of X and Σ_{YX} and Σ_{XY} denote the cross-covariance matrices. Dealing initially with the population model, it can be shown that (see, e.g,. Izenman (2013)), under the assumption that the true rank of C is q, the mean sum of squared error criterion is minimized by setting

$$C^{(q)} = \sum_{j=1}^{q} v_j v_j' \Sigma_{YX} \Sigma_{XX}^{1/2}, \tag{A.14}$$

where v_j denotes the j-th eigenvector associated with the j-th largest eigenvalue of the $K \times K$ symmetric matrix

$$\Sigma_{YX} \Sigma_{XX}^{-1} \Sigma_{XY}. \tag{A.15}$$

Let $V_q = [v_1 \dots v_q]$ denote the $K \times q$ matrix consisting of the first j eigenvectors associated with the matrix of Eq. (A.15) This gives rise to the following (population) RRR estimation formulas:

$$A^{(q)} = V_q \tag{A.16}$$

$$B^{(q)} = V_q' \Sigma_{YX} \Sigma_{XX}^{-1} \tag{A.17}$$

$$\mu^{(q)} = \mu_Y - A^{(q)} B^{(q)} \mu_X, \tag{A.18}$$

where μ_Y and μ_X denote the means of Y and X, respectively. Given a sample of size T from $Z = (X, Y)$, feasible versions of the estimators may be obtained by replacing population moments by the corresponding sample moments. This involves replacing μ_Y and μ_X with sample means, and population covariances and cross-covariances with corresponding sample moments, e.g., $\hat{\Sigma}_{XX} = (1/T) \sum_{t=1}^{T} X_{c,t} X_{c,t}'$, where $X_{c,t}$ denotes the centered version of X_t.

Because in our application the explanatory variable X consists of the lagged dependent variable, our implementation of RRR relates closely to the so-called canonical correlation analysis of Hotelling (1936). See Izenman (2013), for example, for a detailed discussion of the connection.

A.5.2. Results and Diagnostics

We estimate the reduced rank regression model for alternative assumptions regarding the reduced rank number q. Intuitively, when $q = 1$ there is a single relevant volatility index, whereas when $q = 2$ there are two relevant volatility indices, and so on. In order to assess the quality of fit and determine the appropriate choice of q, we rely on the 'rank trace' criterion (see, e.g., Izenman (2013))). Let $\hat{C}^*$ denote the unrestricted (full rank) estimate of C and $\hat{\Sigma}^*_{\epsilon\epsilon}$ the corresponding estimated residual covariance matrix. In contrast, $\hat{C}^{(q)}$ and $\hat{\Sigma}^{(q)}_{\epsilon\epsilon}$ denotes the analogs for the reduced rank regression estimate assuming a rank of q. The rank trace diagnostic focuses on the following metrics for coefficient and model fit, respectively:

$$\Delta\hat{C}^{(q)} = \frac{||\hat{C}^* - \hat{C}^{(q)}||}{||\hat{C}^*||} \tag{A.19}$$

$$\Delta\hat{\Sigma}^{(q)}_{\epsilon\epsilon} = \frac{||\hat{\Sigma}^*_{\epsilon\epsilon} - \hat{\Sigma}^{(q)}_{\epsilon\epsilon}||}{||\hat{\Sigma}^*_{\epsilon\epsilon} - \hat{\Sigma}_{YY}||}, \tag{A.20}$$

where $||A||$ denotes the classical Euclidean norm.

Figure A.2 presents plots characterizing the fit of the VARI(1) model for various choices of rank (i.e., different numbers of indices q) as applied to quarterly volatility series for 10 factor portfolios. Panel A applies the 'rank trace' approach to diagnose the effective rank of the coefficient matrix (see, e.g., Izenman (2013))). Reduced rank coefficient estimates for models with the correct rank q (or greater) should be 'close' to full-rank estimated coefficients. In addition, the fit for models with appropriate rank or greater should be 'close' to the full-rank fit. The particular metrics take values in the interval [0,1] with 0 representing ideal fit. The plot area of Panel A of Figure A.2 is the unit square. The upper right (1,1) point corresponds to a zero rank model (mean model). The (0,0) point corresponds to the full rank unrestricted VAR(1) model. Points between correspond to models with different ranks, several of which are labeled. A rank one or 'single index' model achieves roughly 85% accuracy relative to an unrestricted VAR(1) model. Increasing the rank improves fit, but incremental improvements are modest. The coefficient criterion indicates that low rank models are imperfect, as they do not closely approximate the full rank coefficient estimates. Intuitively, the rank trace analysis suggests that a single index model captures much concerning volatility dynamics, but there exist 'idiosyncratic' dynamics among the factor volatilities that are imperfectly captured. It is for this reason that we ultimately use the expanded forecasting model of Eq. (A.12) that includes own-lagged volatility as well as the lagged volatility index.

The VARI(1) model yields a single index that drives the volatility system dynamics. This raises an important question: how different are measures of CFV from the volatility index that emerges from the reduced rank regression forecasting approach described here? Panels B and C of Figure A.2 address this question. Panel B contrasts the (normalized) weights corresponding to the index identified by RRR in a single-index model with the weights corresponding to the static PC measure of CFV. The PC weights are uniformly positive and relatively close to equal weights. In contrast, the RRR single volatility index applies relatively different weights to the factor volatility series, including several negative weights. Despite the notable difference in weights, however, Panel C shows that the dynamics of the resulting volatility index under RRR approach are in fact very similar to those of a measure based on the first PC from the volatility series. This is reassuring, in these sense that alternative static and dynamic approaches to obtaining a linear summary measure of commonality in factor volatility produce highly correlated 'CFV' measures.

A.6. Impulse Responses in VARI Models for Volatility

Figure A.3 depicts impulse response functions (IRFs) associated with various types of volatility shocks for the AR(1)-Index model of Eq. (A.12) applied to 10 factor portfolio volatility series. To enhance the visibility of the plots, we report results only for the first six factors (MKT, SMB, HML, CMA, RMW, UMD); however results for the omitted factors (STR, LTR, BAB, QMJ) are qualitatively similar and are available upon request. We consider general types of shocks. The first is a unit shock to a particular factor volatility series, with all other volatility series shocks set to zero. The orange dashed line in each panel of Figure A.3 illustrates the response for that factor to a one unit 'own-volatility' shock. Responses are consistent with the effects of shocks in standard AR(1) models: the response at lag j equals ρ_i^j, where ρ_i is the AR(1) slope coefficient, and the impact therefore decays geometrically. The green dash-dot line shows for each factor the responses to a one unit shock to *market* volatility (for all factors). Because the upper left panel pertains to the MKT factor, the shock is in this case equivalent to an own-volatility shock. For the other five factor volatility series, however, the impact of the market shock occurs via its relation to the volatility index. As a benchmark, note that in a 'pure AR' model such that all VAR coefficient matrices are diagonal, the market volatility shock would have no impact on other factor volatility series. In the AR(1)-Index model, the market volatility shock increases the index. This leads to persistent positive impact responses for other volatility series. The magnitude of the response depends upon the extent to which the volatility series loads on the index. For example, the CMA factor exhibits the smallest loading, and consequently CMA volatility responses to the market volatility shock are minimal. UMD and HML load relatively heavily on the index factor and consequently the 'cross factor' impact of the market volatility shock is larger.

The final two IRF curves in each panel of Figure A.3 characterize the effects of a *common* shock to volatility. The solid blue line depicts responses to a unit length shock exactly in the direction of the single index weights estimated by reduced rank regression (RRR). The magnitude of the immediate impact of this shock depends largely on the magnitude of the weighting that each volatility receives in the index. Panel B of Figure 7 shows that UMD receives the largest weight, whereas SMB receives a relatively small (absolute) weight. Consequently, the initial response of UMD volatility is much larger than that for SMB. Aside from variation in the initial response to an 'index shock,' it is notable that the responses across all factor volatility series are quite persistent and more persistent than own-volatility shocks. This is attributable to the fact that the index process implied by the model is more persistent than the own-volatility component not captured by the index. The red dotted line depicts responses to alternative 'common' volatility shock that is defined as an equal-weighted positive, unit length shock across volatility series. The effects of this shock are again quite persistent across all factor volatility series. There are; however, some differences between the responses when the common shock is equal-weighted versus in the direction of the RRR index weights. In particular, there is more variation in the response patterns in the latter case, attributable to heterogeneity among factors in the relative weight allocated in the RRR index.

These extended results provide additional insight and intuition regarding the forecasting value of incorporating information concerning the strong common feature in factor volatility.

Common or 'index' volatility shocks have economically significant effects on dynamics across the factor volatility series. Including the index leads to a more nuanced notion of the dynamics of factor volatility. In particular, the impact of factor volatility shocks depends on whether the corresponding shock occurs in isolation, i.e., in 'idiosyncratic' fashion, or whether shocks coincide in the form of an 'index shock'. Shocks to the index reflecting common volatility dynamics turns out to be more persistent than pure own-volatility shocks.

A.7. CFV and Financial and Operating Leverage

This section provides additional details regarding tests briefly discussed in Section 5 of the paper concerning a potential relation between CFV and financial and operating leverage.

We test whether CFV is correlated with measures of financial and operating leverage as well as growth options. Book financial leverage is short-term debt plus long-term debt divided by total assets. Book operating leverage is operating profits divided by total assets, as in Novy-Marx (2010). We also consider market versions of these measures, in which the denominator is market equity plus book total liabilities. Our final measure is the market-to-book ratio as a proxy for growth options. We compute the average of each of these variables every quarter. Accounting variables are lined up with market variables on the fiscal quarter end date. Table A.3 presents results for regressions of innovations in CFV on innovations in each of these variables. (Innovations are residuals from an ARMA (1,1) model.) Innovations in book operating and financial leverage have small correlations with innovations in CFV, with R^2 values of 1%. The next two models include market leverage and contemporaneous average stock returns. Returns are included because innovations in market leverage are positively correlated with returns, and returns are negatively correlated with volatility. (Not including returns results in a negative sign for the leverage measures.) These specifications show that the incremental explanatory power of market leverage measures is small. Finally, the market-to-book ratio only weakly relates to CFV.

References

Aït-Sahalia, Y., Jacod, J., 2014. High-frequency financial econometrics. In: *High-Frequency Financial Econometrics*, Princeton University Press.

Bai, J., 2003. Inferential theory for factor models of large dimensions. Econometrica 71, 135–171.

Barndorff-Nielsen, O. E., Shephard, N., 2004. Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics. Econometrica 72, 885–925.

Chamberlain, G., Rothschild, M., 1983. Arbitrage and mean-variance analysis on large asset markets. Econometrica 51, 1281–1304.

Chen, N., Roll, R., Ross, S., 1986. Ross, 1986, economic forces and the stock market. Journal of Business 59, 383–403.

Connor, G., Korajczyk, R. A., 1986. Performance measurement with the arbitrage pricing theory: A new framework for analysis. Journal of financial economics 15, 373–394.

Connor, G., Korajczyk, R. A., 1988. Risk and return in an equilibrium apt: Application of a new test methodology. Journal of financial economics 21, 255–289.

Connor, G., Korajczyk, R. A., Linton, O., 2006. The common and specific components of dynamic volatility. Journal of Econometrics 132, 231–255.

Cooper, I., Priestley, R., 2011. Real investment and risk dynamics. Journal of Financial Economics 101, 182–205.

Corsi, F., 2009. A simple approximate long-memory model of realized volatility. Journal of Financial Econometrics 7, 174–196.

Cubadda, G., Guardabascio, B., Hecq, A., 2017. A vector heterogeneous autoregressive index model for realized volatility measures. International Journal of Forecasting 33, 337–344.

Engle, R. F., Granger, C. W., 1987. Co-integration and error correction: representation, estimation, and testing. Econometrica: journal of the Econometric Society pp. 251–276.

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

Hotelling, H., 1936. Relations between two sets of variates. Biometrika 28, 321–377.

Izenman, A. J., 2013. Linear dimensionality reduction. In: *Modern Multivariate Statistical Techniques*, Springer, pp. 195–236.

Jones, C. S., 2001. Extracting factors from heteroskedastic asset returns. Journal of Financial economics 62, 293–325.

Kalnina, I., Tewou, K., 2019. Cross-sectional dependence in idiosyncratic volatility. Working paper .

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

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

Novy-Marx, R., 2010. Operating leverage. Review of Finance 15, 103–134.

Pelger, M., 2020. Understanding systematic risk: a high-frequency approach. The Journal of Finance 75, 2179–2220.

Reinsel, G., 1983. Some results on multivariate autoregressive index models. Biometrika 70, 145–156.

Schwert, G. W., 1989. Why does stock market volatility change over time? The journal of Finance 44, 1115–1153.

Vetter, M., 2015. Estimation of integrated volatility of volatility with applications to goodness-of-fit testing. Bernoulli 21, 2393–2418.

Zhang, L., Mykland, P. A., Aït-Sahalia, Y., 2005. A tale of two time scales: Determining integrated volatility with noisy high-frequency data. Journal of the American Statistical Association 100, 1394–1411.

Table A.1: **Correlations Among Alternative CFV Proxies**

This table presents the (sample) correlation matrix for the set of alternative CFV proxies depicted in Figure 3 of the paper. Five quarterly CFV proxies are compared. The first four correspond to CFV proxies from characteristics-based factors, anomaly long-short portfolios, industry-sorted portfolios, and macroeconomic tracking factors, respectively. The 'Char.-Based' proxy equals the average realized volatility across 9 characteristics-based factors. The 'Anomalies' proxy equals the average quarterly realized volatility across the the GHZ anomaly portfolio series. 'Industries' equals the average quarterly realized volatility across the Fama-French 17 value-weighted industry portfolio volatility series. The 'Macro' proxy equals the average quarterly realized volatility for the Chen, Roll, Ross factors. All underlying realized volatility series are computed using market-adjusted daily returns. The thick blue line depicts a 'Pooled CFV' measure that equals the equal-weighted average of quarterly realized volatility series across the pooled set of portfolio returns from all of the aforementioned sets. The sample period is 1964.3–2020.1 (quarterly).

	Pooled CFV	Char.-Based	Anomalies	Industries	Macro
Pooled CFV	1.00	0.95	0.99	0.97	0.97
Char.-Based	0.95	1.00	0.94	0.91	0.95
Anomalies	0.99	0.94	1.00	0.95	0.96
Industries	0.97	0.91	0.95	1.00	0.94
Macro	0.97	0.95	0.96	0.94	1.00

Fig. A.1: Similarity among Alternative CFV Measures

This figure compares two quarterly proxies for common factor volatility (CFV) based on a set of 43 anomaly portfolios (KNS portfolios). The blue line depicts a quarterly CFV measure based on 'realized variances' computed as the sum of squared daily factor returns over the corresponding quarter. We apply an ARMA (1,1) model to the raw CFV measure based on realized volatility in order to produce a time series of expected (common) volatility that is comparable to the GARCH-based CFV measure. The orange line depicts an alternative CFV measure based on conditional factor variances estimated using a GARCH (1,1) model with constant mean applied to quarterly returns for each factor. Given the underlying quarterly factor volatility series, the CFV proxy is computed as the cross-sectional average of individual volatility series across the 43 different anomaly series. Both CFV measures are then standardized to facilitate comparison.

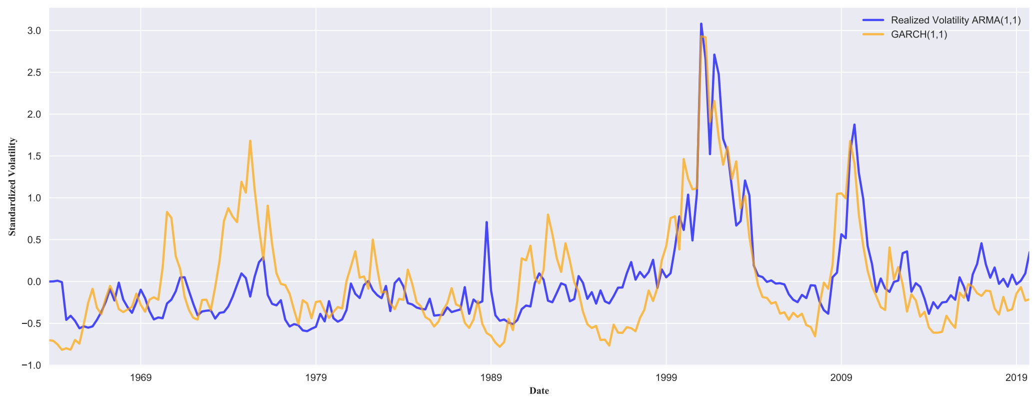

Table A.2: **Commonality in Correlations**

The table shows statistics associated with quarterly correlation measures for factor and anomaly portfolios. The first column specifies the set of factors. The second column indicates the sample period. N equals the number of correlation pairs for the corresponding factor set. Factor sets indicated with an asterisk (*) are such that total correlation pairs exceed the number of time series observations T. In these cases we randomly selected T pairs to analyze. The column (% Expl.) shows the percent of total variance explained by the first principal component (PC) extracted from the quarterly correlation series for the set of portfolios. $\hat{\rho}$(CFV) denotes the time-series correlation between the PC measure of common correlation for the factor set and a measure of common factor volatility (CFV). Results are shown both for raw correlations (columns 4 and 5) and for correlations computed using residuals from a market model regressions (columns 6 and 7) estimated each calendar year using daily returns, except for the high-frequency factor set (denoted with a †) for which regressions are monthly using 5-minute intraday returns. See Section 2 for definitions of other portfolio sets.

			Raw		Market-Adjusted	
Portfolios	Sample	N	% Expl.	$\hat{\rho}(CFV)$	% Expl.	$\hat{\rho}(CFV)$
Panel A: Characteristics-Based Factors						
SMB,HML,UMD,BAB,(S/L)TR	1930.4-2020.1	15	15.95	0.06	15.92	0.03
FF5 + UMD	1963.3-2020.1	10	19.34	0.26	24.50	0.37
FF5 Augmented	1963.3-2020.1	45	13.74	0.19	17.26	0.25
All Factors	1963.3-2016.4	55	16.91	0.36	13.53	0.31
HF FF5 + UMD†	1996.1-2017.4	10	25.32	0.31	31.05	0.40
Panel B: Anomaly Long-Short Portfolios						
GHZ Anomalies*	1964.3-2018.4	217	10.58	0.54	8.81	0.45
KNS Anomalies*	1963.4-2017.4	216	15.80	0.32	14.80	0.35
Panel C: Industry Portfolios						
FF 12 Industries	1926.3-2020.1	66	58.46	0.04	13.22	0.28
FF 30 Industries*	1926.3-2020.1	374	50.94	0.02	10.42	0.24
Panel D: Macroeconomic Factors						
CRR Factors	1963.3-2018.4	21	12.51	0.15	14.71	0.03
CRR - (OIL + CONS.)	1963.3-2018.4	10	18.52	0.21	20.67	0.14
Panel E: Statistical Factors (Excluding 'Market')						
FF Sorted Ports.	1931.1-2020.1	45	18.85	0.02	22.47	0.06
KNS Decile Ports.	1973.1-2018.2	45	13.38	0.37	14.17	0.16

Table A.3: **Potential explanations for the commonality in volatility**

The table presents results of regressions of innovations in CFV on innovations in measures of financial and real options. Book (Mkt) leverage is total long and short term debt divided by book (market) value of total assets. Book (Mkt) operating leverage is operating expense divided by book (market) total assets, Market total assets are market equity + total assets - book equity. M/B is the ratio of market to book equity. All financial variables are averages across all firms reporting earnings in the quarter for which volatility is measured. Returns are the average returns of all firms in the quarter over which volatility is measured. Standard errors are Newey-West with 3 lags.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Intercept	-0.00	0.00	0.02	0.01	0.02	-0.04	-0.04
	(-0.01)	(0.00)	(1.71)	(1.23)	(1.60)	(-0.89)	(-0.93)
Financial Leverage (Book)	2.37						1.52
	(1.62)						(1.42)
Operating Leverage (Book)		0.42					0.03
		(2.45)					(0.18)
Returns			-0.65	-0.44	-0.61	-0.67	-0.65
			(-3.84)	(-2.41)	(-3.32)	(-4.04)	(-3.93)
Financial Leverage (Mkt)				2.90			
				(2.03)			
Operating Leverage (Mkt)					0.17		
					(1.04)		
M/B						0.02	0.03
						(1.53)	(1.90)
R^2	0.01	0.01	0.15	0.17	0.15	0.16	0.17
N	188.00	188.00	188	188	188	188.00	188.00

Fig. A.2: VAR Model with Reduced Rank Restrictions: Factor Portfolios

This figure shows plots characterizing the fit of first order vector autoregression (VAR(1)) models that impose reduced rank restrictions on the coefficient matrix. The models are estimated via reduced rank regression for rank choices ranging from one to full rank (10) for quarterly volatility series for the following factor portfolios: MKT, SMB, HML, RMW, CMA, UMD, STR, LTR, BAB, and QMJ. The sample period is 1964–2017. Panel A shows the 'rank trace' analysis to help determine the appropriate rank. See the discussion in section 3 and the Appendix for details and interpretation. Panel B contrasts the weights of the 10 factor volatility series in a single 'index' or rank 1 model with those for the first (static) principal component (PC) extracted from the same volatility series. Portfolios numbered 1 through 10 correspond to the 10 factors as listed above. Panel C plots standardized time series for the single index based on a reduced rank model with rank 1 with the CFV measure based on the first PC.

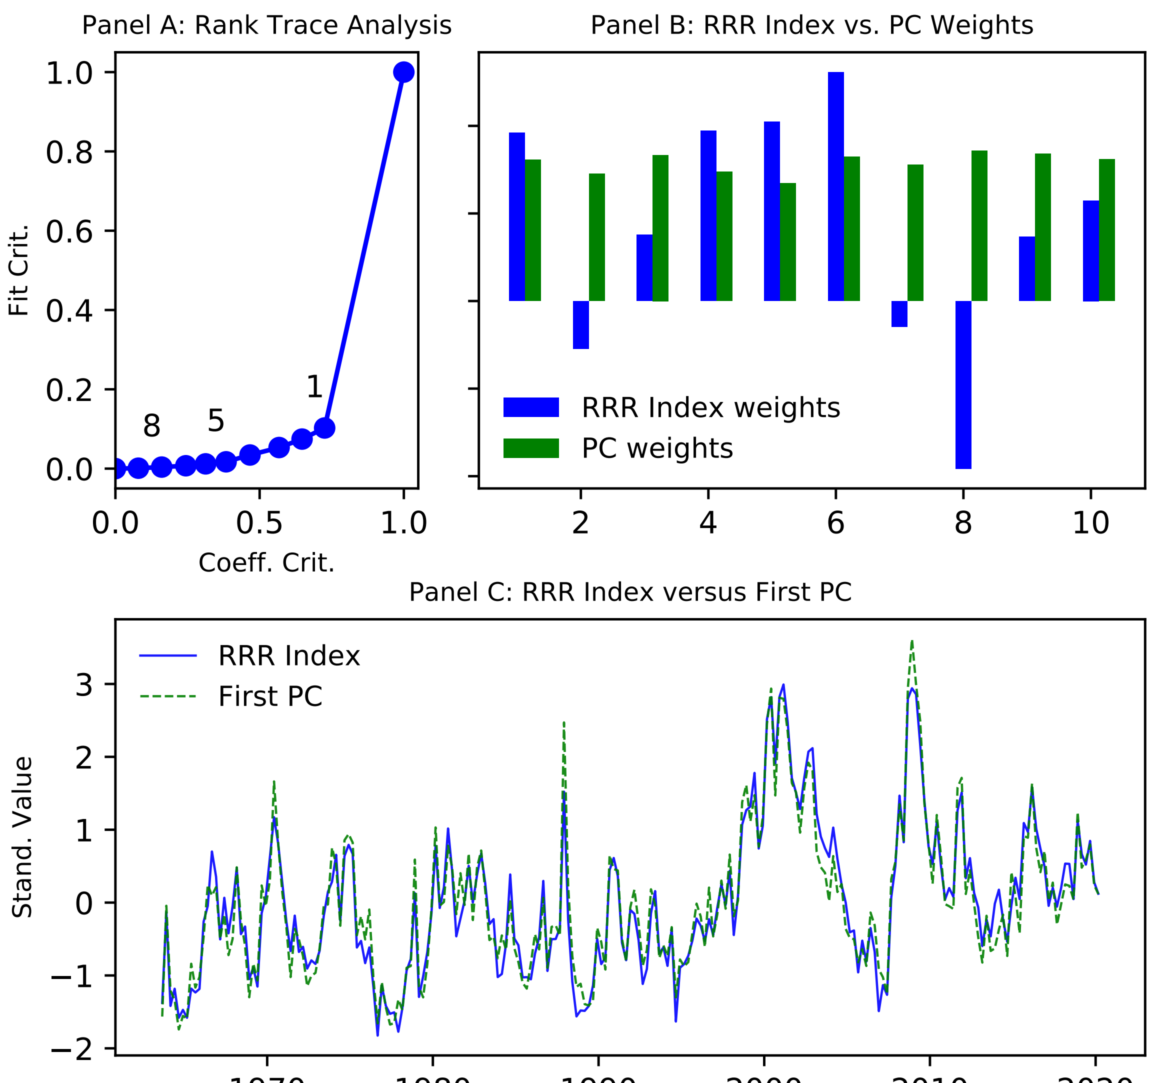

Fig. A.3: Impulse Response Functions: VAR-Index Model for Factors

This figure plots impulse response functions (IRFs) corresponding to various shocks in the VAR(1)-Index model applied to 10 factor portfolio volatility series. To enhance visibility, we show IRFs for only the first 6 among the 10 factors. These are MKT, HML, SMB, CMA, RMW, UMD. Results for the four omitted factor volatilities (STR, LTR, BAB, and QMJ) are qualitatively similar. The solid blue line depicts responses to a unit length shock in the direction of the single index estimated by reduced rank regression (RRR). The red dotted line depicts responses to alternative 'common' volatility shock that is equal-weighted across volatility series (and unit length). The orange dashed line shows for each factor the responses to a one unit 'idiosyncratic' own-volatility shock. Finally, the green dash-dot line shows for each factor the responses to a one unit shock to *market* volatility.