Online Appendix for "Hedging macroeconomic and financial uncertainty and volatility"

Ian Dew-Becker, Stefano Giglio, and Bryan Kelly

September 30, 2020

OA.1 Implied volatility and regression forecasts

Implied volatilities are, under certain assumptions, expectations of future realized volatility under the risk-neutral measure. If there is a time-varying volatility risk premium, then implied volatilities will be imperfectly correlated with physical expectations of future realized volatility, which constitutes actual uncertainty. This section compares implied volatilities to regression-based forecasts of future volatility to evaluate the quantitative magnitude of that deviation.

For each market, we estimate the regression

$$RV_{i,t} = a_i + b_i\left(L\right) RV_{i,t-1} + c_i IV_{i,t-1} + \varepsilon_{i,t} \tag{OA.1}$$

where $b_i\left(L\right)$ is a polynomial in the lag operator, L, and a_i and c_i are coefficients. $RV_{i,t}$ is realized volatility in month t for market i – the sum of squared daily futures returns during the month. $IV_{i,t}$ is the (at-the-money) implied volatility at the end of month t in market i.

The table below reports the correlation between the fitted values from that regression – which represent physical uncertainty – and implied volatility. That is, it reports $corr\left(b_i\left(L\right) RV_{i,t-1} + c_i IV_{i,t-1}, IV_{i,t-1}\right)$. Ideally, we would like that correlation to be 1, so that implied volatility is perfectly correlated with physical uncertainty, and hedging implied volatility hedges uncertainty. Note that this does not require that risk premia are constant. If $b_i\left(L\right) = 0$ but $c_i \neq 1$, risk premia are time-varying, but the physical uncertainty is still perfectly correlated with implied volatility. It is only deviations of $b_i\left(L\right)$ from zero that reduce the correlation. To the extent that the implied volatility summarizes all available information, we would expect $b_i = 0$.

Correlations of implied volatility with fitted uncertainty

S&P 500	0.966	Crude oil	0.998	Silver	0.984
Treasuries	0.940	Feeder cattle	0.951	Soybeans	0.970
British Pound	0.987	Gold	0.994	Soybean meal	0.974
Swiss Franc	0.994	Heating oil	0.992	Soybean oil	0.946
Yen	0.976	Lean hogs	0.937	Wheat	0.998
Copper	0.963	Live cattle	0.919		
Corn	0.994	Natural gas	0.949		

The table shows that across the various markets, the correlations are all high, with a minimum of 91.1 percent and a mean of 97.0 percent. So while implied volatility is not literally the same as physical uncertainty, it appears to be fairly close. In the baseline results, we allow for two lags in the polynomial b, but we have experimented with alternative specifications and obtain similar results.

OA.2 Factor models and factor-hedging portfolios

In this section we review a useful result from the algebra of cross-sectional regressions: given a set of K nontradable factors F_t, the cross-sectional estimates of the K risk premia, λ, are the average excess returns of K portfolios, each of which has betas of exactly 1 with respect to one factor, and 0 with respect to the other $K-1$ factors: we refer to these as *factor-hedging portfolios* for the K factors in F_t. The time series of returns for the factor-hedging portfolios are the slopes of period-by-period cross-sectional regressions. These results hold in population.

Consider K nontradable factors F_t, and a vector of N excess returns r_t of test assets. Nontradable factors have a risk premium of λ (a $K\times 1$ vector), so the factor model can be written as:

$$\underbrace{r_t}_{N\times 1} = \underbrace{\beta}_{N\times K}\underbrace{\lambda}_{K\times 1} + \underbrace{\beta}_{N\times K}\underbrace{(F_t - E[F_t])}_{K\times 1} + \underbrace{e_t}_{K\times 1} \quad \text{(OA.2)}$$

Cross-sectional regressions operate in two stages. First, they estimate the $N \times K$ matrix β from time series regressions of the form:

$$r_t = k + \beta F_t + e_t$$

where the constant k also depends on the means of the factors $E[F_t]$, which is not interpretable in general when factors are nontradable, and is irrelevant for computing β. The second step of the cross-sectional regression could either be estimated using *average* returns (in one cross-sectional regression), or as a sequence of period-by-period cross-sectional regressions. The latter approach is often used in practice (as in the Fama-MacBeth version of the two-step regressions) because it makes standard errors calculation easier, but either method yields the same point estimates for risk premia λ. Here, we also follow the second method, but for a different reason: because it generates a time-series of factor-hedging portfolios.

We therefore run, for each period t, cross-sectional regressions of r_t on the estimated β:

$$r_t = a_t + \beta g_t + u_t$$

obtaining a time-series of $K \times 1$ slope vectors g_t. Risk premia λ are then estimated as the time-series average of the slopes: $\lambda = E[g_t]$.

The time-series slopes g_t have a useful interpretation. They are calculated in each period as:

$$g_t = (\beta^{'}\beta)^{-1}\beta' r_t \tag{OA.3}$$

This equation clarifies that g_t are themselves excess returns (they are the returns of portfolios of the underlying N assets, with weights $w = (\beta^{'}\beta)^{-1}\beta'$); the risk premia λ are the (risk premia) average excess returns of these K portfolios g_t. We can now explore the properties of these portfolios. Substituting r_t out from (OA.2) we have:

$$g_t = (\beta^{'}\beta)^{-1}\beta'(\beta\lambda + \beta(F_t - E[F_t]) + e_t) = \lambda + (F_t - E[F_t]) + (\beta^{'}\beta)^{-1}\beta' e_t$$

Under suitable assumptions on the cross-sectional dispersion in the β (see Giglio and Xiu (2019) for a formal analysis) the last term is close to zero for large N (intuitively, the idiosyncratic errors are diversified away, and the g_t are well-diversified portfolios). We therefore can write:

$$g_t \simeq \lambda + (F_t - E[F_t])$$

From this equation, it is clear that, as expected, $E[g_t] = \lambda$. In addition, these K portfolios have the special property of being exposed to exactly one of the underlying factor F_t each: the matrix of exposures of g_t to factor innovations $F_t - E[F_t]$ is simply the identity matrix. So the first portfolio has betas [1, 0, 0 , 0, ...], the second portfolio has betas [0, 1, 0 , 0, ...], and so on. This is why we refer to these portfolios as *factor-hedging portfolios.*

Finally, it is worth pointing out that the latter property also holds in *any* sample: the estimated betas of the factor-hedging portfolios with respect to the nontradable factors will be the vectors [1, 0, 0 , 0, ...], [0, 1, 0 , 0, ...] and so on in every sample.

OA.3 Approximating return sensitivities

This section describes the approximation of option returns used to obtain the rv and iv portfolios. P denotes the price of an at-the-money straddle or strangle. σ is the Black–Scholes volatility, n is the time to maturity, F is the forward price, and K is the strike. N denotes the standard Normal cumulative distribution function.

For the calls and puts, respectively, we set

$$K_{call} = F \exp\left(b\sigma\sqrt{n} + \frac{\sigma^2}{2}n\right) \tag{OA.4}$$

$$K_{put} = F \exp\left(-b\sigma\sqrt{n} + \frac{\sigma^2}{2}n\right) \tag{OA.5}$$

We calculate everything for arbitrary b. A straddle is the special case where $b = 0$, while a strangle has positive b, so that both the put and call are out of the money.

OA.3.1 Prices

We first calculate the price of a strangle. The Black–Scholes formula gives

$$P_{call} = FN(-b) - F\exp\left(b\sigma\sqrt{n} + \frac{\sigma^2}{2}n\right) N\left(-b - \sigma\sqrt{n}\right) \tag{OA.6}$$

$$P_{put} = -FN(-b) + F\exp\left(-b\sigma\sqrt{n} + \frac{\sigma^2}{2}n\right) N\left(-b + \sigma\sqrt{n}\right) \tag{OA.7}$$

So the total price is

$$P = P_{call} + P_{put} = F\left(N(-b) - N(-b)\right) \tag{OA.8}$$

$$-F\left(\exp\left(b\sigma\sqrt{n} + \frac{\sigma^2}{2}n\right) N\left(-b - \sigma\sqrt{n}\right) - \exp\left(-b\sigma\sqrt{n} + \frac{\sigma^2}{2}n\right) N\left(-b + \sigma\sqrt{n}\right)\right) \tag{OA.9}$$

$$\approx FN'(-b)\, 2\sigma\sqrt{n} \tag{OA.10}$$

where the second line uses a first order approximation to $N(x)$ around $-b$ and $\exp\left(b\sigma\sqrt{n} + \frac{\sigma^2}{2}n\right) \approx 1$.

OA.3.2 Return derivatives

The local approximation for returns that we use is

$$\frac{\partial r_{t+1}}{\partial x_{t+1}} = \frac{\partial}{\partial x_{t+1}} \frac{P\left(F_{t+1}, \sigma_{t+1}\right)}{P\left(F_t, \sigma_t\right)} \tag{OA.11}$$

and we evaluate the derivatives at the point $F_{t+1} = F_t$, $\sigma_{t+1} = \sigma_t$.

We have

$$\begin{aligned}
\frac{\partial r_{t+1}}{\partial \sigma_{t+1}} &= \frac{P_{\sigma,t+1}}{P_t} && \text{(OA.12)} \\
&= \frac{N'\left(-b\right) + N'\left(b\right)}{N'\left(-b\right) 2\sigma_t} && \text{(OA.13)} \\
&\approx \frac{1}{\sigma_t} && \text{(OA.14)}
\end{aligned}$$

where $P_{\sigma,t}$ denotes $\partial P\left(F_{t+1}, \sigma_{t+1}\right) / \partial \sigma_{t+1}$ (evaulated at $\sigma_{t+1}^2 = \sigma_t^2$), and using the approximation that $N'\left(b\right) \approx N'\left(-b\right)$. We then have

$$\frac{\partial r_{t+1}}{\partial\left(\Delta \sigma_{t+1} / \sigma_t\right)} \approx 1 \tag{OA.15}$$

Next, for squared returns, we have

$$\begin{aligned}
\frac{\partial r_{t+1}}{\partial F_{t+1}^2} &= \frac{P_{FF,t}}{P_t} && \text{(OA.16)} \\
&= \frac{1}{F_t N'\left(-b\right) 2\sigma\sqrt{n}} \frac{N'\left(-b\right) + N\left(b\right)}{F_t \sigma_t \sqrt{n}} && \text{(OA.17)} \\
&\approx \frac{1}{F_t^2 \sigma_t^2 n} && \text{(OA.18)}
\end{aligned}$$

Again using $N'\left(b\right) \approx N'\left(-b\right)$. Finally, note that $\partial f_{t+1} = \partial F_{t+1} / F_{t+1}$, so that

$$\begin{aligned}
\frac{\partial^2 r_{t+1}}{\partial\left(f_{t+1} / \sigma_t\right)^2} &= \frac{\partial r_{t+1}}{\partial F_{t+1}^2} F_t^2 \sigma_t^2 && \text{(OA.19)} \\
&\approx \frac{1}{n} && \text{(OA.20)}
\end{aligned}$$

OA.3.3 Accuracy

To study how effective the above approximation is, we examine a simple simulation. We assume that options are priced according to the Black–Scholes model. We set the initial futures price to 1 and the initial volatility to 30 percent per year. We then examine instantaneous returns (i.e. through shifts in σ and S) on the iv and rv portfolios for straddles defined exactly as in the main text, allowing the futures return to vary between between $+/-23.53$ percent, which corresponds to variation out to four two-week standard deviations. We allow volatility to move between 15 and 60 percent – falling by half or doubling.

The top two panels of figure OA.3 plot contours of returns on the rv and iv portfolios defined in the main text, while the middle panels plot the contours predicted by the approximations for the partial derivatives. For the iv portfolio, except for very large instantaneous returns – 15–20 percent – the approximation lies very close to the truth. The bottom-right panel plots the error – the middle panel minus the top panel – and except for cases where the underlying has an extreme movement and the implied volatility falls – the exact opposite of typical behavior – the errors are all quantitatively small, especially compared to the overall return.

For the rv portfolio, the errors are somewhat larger. This is due to the fact that we approximate the rv portfolio using a quadratic function, but its payoff has a shape closer to a hyperbola. Again, for underlying futures returns within two standard deviations (where the two-week standard deviation here is 5.88 percent), the errors are relatively small quantitatively, especially when σ does not move far. Towards the corners of the figure, though, the errors grow somewhat large.

These results therefore underscore the discussion in the text. The approximations used to construct the iv and rv portfolios are qualitatively accurate, and except in more extreme cases also hold reasonably well quantitatively. But they are obviously not fully robust to all events, so the factor model estimation, which does not rely on any approximations, should be used in situations where the nonlinearities are a concern.

OA.3.4 Empirical return exposures

To check empirically the accuracy of the expressions for the risk exposures of the straddles, appendix figure OA.2 plots estimated factor loadings for straddles at maturities from one to five months for each market from time series regressions of the form

$$r_{i,n,t} = a_{i,n} + \beta_{i,n}^{f} \frac{f_{i,t}}{IV_{i,t-1}} + \beta_{i,n}^{f^2} \frac{1}{2} \left(\frac{f_{i,t}}{IV_{i,t-1}} \right)^2 + \beta_{i,n}^{\Delta IV} \frac{\Delta IV_{i,t}}{IV_{i,t-1}} + \varepsilon_{i,n,t} \tag{OA.21}$$

The prediction of the analysis above is that $\beta_{i,n}^{f} = 0$, $\beta_{i,n}^{f^2} = 1/n$, and $\beta_{i,n}^{\Delta IV} = 1$.

Across the panels, the predictions hold surprisingly accurately. The loadings on $f_{i,t}$ are all near zero, if also generally slightly positive. The loadings on the change in implied volatility are all close to 1, with little systematic variation across maturities. And the loadings on the squared futures return tend to begin near 1 (though sometimes biased down somewhat) and then decline monotonically, consistent with the predicted n^{-1} scaling.

Table OA.2 reports results of similar regressions for each underlying of the returns on the *rv* and *iv* portfolios on the underlying futures return, the squared futures return, and the change in implied volatility. The table shows that while the Black–Scholes predictions do not hold perfectly, it is true that the *rv* portfolio is much more strongly exposed to realized than implied volatility, and the opposite holds for the *iv* portfolio. The coefficients on $(f_t/\sigma_{t-1})^2$ average 0.78 for the *rv* portfolio and 0.12 for the *iv* portfolio (though that average masks some variation across markets). Conversely, the coefficients on $\Delta\sigma_t/\sigma_{t-1}$ average 0.03 for the *rv* portfolio and 0.81 for the *iv* portfolio Furthermore, the R^2s are large, averaging 70 percent across the various portfolios, implying that their returns are well described by the approximation (4).

OA.3.5 Volume

Figure OA.14 reports the average daily volume of all of the option contracts across maturities 1 to 6 months. For crude oil, which we use here as a reference contract, the figure reports average daily volume in dollars; for all other contracts, it reports the average daily volume relative to crude oil. Empirically, crude oil options have volume numbers of the same order of magnitude as the S&P 500, while there is more heterogeneity across the other markets. Looking across maturities, the general pattern is that dollar volume declines by about a factor of three in almost all the markets between the 1- and 6-month maturities – so the 6-month maturity has less volume, but far from zero.

OA.3.6 Alternative scaling for returns

Because returns have a price in the denominator, if that price is measured with error, returns can be biased upwards. The *iv* portfolio is net long the straddles, while the *rv* portfolio has a total weight of zero, so measurement error in prices would bias *iv* returns up but not *rv* returns. To account for that possibility, this section examines results when all the straddle returns are scaled by the price of the one-month straddle, instead of the price of a straddle with the same maturity.

Specifically, denoting $P_{n,t}$ the price of a straddle of maturity n on date t, the return on an n-month straddle used in the main results is

$$R_{n,t} = \frac{P_{n-1,t+1} - P_{n,t}}{P_{n,t}} \tag{OA.22}$$

We consider returns on a portfolio that puts weight $\frac{P_{n,t}}{P_{1,t}}$ on the n-month straddle and weight $1 - \frac{P_{n,t}}{P_{1,t}}$ on the risk-free asset (which is a tradable portfolio), which is

$$r_{n,t+1}^{rescaled} = \frac{P_{n-1,t+1} - P_{n,t}}{P_{n,t}} \frac{P_{n,t}}{P_{1,t}} + \left(1 - \frac{P_{n,t}}{P_{1,t}}\right) r_{f,t} \tag{OA.23}$$

$$= \frac{P_{n-1,t+1} - P_{n,t}}{P_{1,t}} + \left(1 - \frac{P_{n,t}}{P_{1,t}}\right) r_{f,t} \tag{OA.24}$$

This portfolio is useful for two reasons. First, the one-month maturity has the highest volume in most markets we study, and it is typically considered to be the most accurate. Second, this eliminates differences in bias across maturities since in this specification, the denominator is the same for all n.

For $r_{n,t+1}^{rescaled}$, similar calculations to those above yield the results that

$$\frac{\partial^2 r_{n,t+1}^{rescaled}}{\partial \left(f_{t+1}/\sigma_t\right)^2} \approx \frac{1}{\sqrt{n}} \tag{OA.25}$$

$$\frac{\partial r_{n,t+1}^{rescaled}}{\partial \left(\Delta\sigma_{t+1}/\sigma_t\right)} \approx \sqrt{n} \tag{OA.26}$$

We then calculate alternative rv and iv portfolios as

$$iv_t^{rescaled} = \frac{3}{\sqrt{12}} \left(\sqrt{5/12} r_{5,t}^{rescaled} - \sqrt{1/12} r_{1,t}^{rescaled}\right) \tag{OA.27}$$

$$rv_t^{rescaled} = \frac{5/48}{\sqrt{12}} \left(\sqrt{12} r_{1,t}^{rescaled} - \sqrt{12/5} r_{5,t}^{rescaled}\right) \tag{OA.28}$$

Figure OA.15 replicates figure 3 with the rescaled returns. The results are nearly identical to the baseline for both the Sharpe ratios on the iv and rv portfolios and the estimated factor risk premia. These results show that when we correct for the potential bias induced by low liquidity and measurement error at longer maturities, the estimates are essentially unchanged.

OA.4 Calculating the covariance of the sample mean returns

There are two features of our data that make calculating covariance matrix of sample means difficult: we have an unbalanced panel and the covariance matrix is either singular or nearly so. We deal with those issues through the following steps.

1. For each market, we estimate the two largest principal components, therefore modeling straddle returns for underlying i and maturity n on date t as

$$r_{i,n,t} = \lambda_{1,i,n} f_{1,i,t} + \lambda_{2,i,n} f_{2,i,t} + \theta_{i,n,t} \tag{OA.29}$$

where the λ are factor loadings, the f are estimated factors, and θ is a residual that we take to be uncorrelated across maturities and markets (it is also in general extremely small).

2. We calculate the long-run covariance matrix of all $J \times 2$ estimated factors. The covariance matrix is calculated using the Hansen–Hodrick method to account for the fact that the returns are overlapping (we use daily observations of 2-week returns). The elements of the covariance matrix are estimated based on the available nonmissing data for the associated pair of factors. That means that the covariance matrix need not be positive semidefinite. To account for that fact, we set all negative eigenvalues of the estimated covariance matrix to zero.

Given the estimated long-run covariance matrix of the factors, denoted Σ_f, and given the (diagonal) long-run variance matrix of the residuals θ, denoted Σ_θ, the long-run covariance matrix of the returns is then

$$\Sigma_r \equiv \Lambda \Sigma_f \Lambda' + \Sigma_\theta \tag{OA.30}$$

where Λ is a matrix containing the factor loadings λ.

3. Finally, it is straightforward to show that the covariance matrix of the sample mean returns is

$$\Sigma_{\hat{r}} = M \odot \Sigma_r \tag{OA.31}$$

where $\odot$ denotes the elementwise product and M is a matrix where the element for a given return pair is equal to the ratio of the number of observations in which both returns are available to the product of the number of observations in which each return is available individually (if all returns had the same number of observations T, then we would obtain the usual T^{-1} scaling). We then have the asymptotic approximation that

$$\hat{r} \Rightarrow N\left(\bar{r}, \Sigma_{\hat{r}}\right) \tag{OA.32}$$

where $\hat{r}$ is a vector that stacks the $\hat{r}_i$ and $\bar{r}$ stacks the $\bar{r}_i$ and $\Rightarrow$ denotes convergence in distribution.

To construct $\Sigma_{\widehat{\mathbf{sr}}}$, we simply divide the i, j element of $\Sigma_{\hat{r}}$ by the product of the sample standard deviations of r_i and r_j.

OA.5 Calculating risk prices with unbalanced panels and correlations across markets

In estimating the factor models, we have two complications to deal with: the sample length for each underlying is different, and returns are correlated across underlyings. This section discusses how we deal with those issues.

We have the model

$$E_{T_i}[R_i] = \lambda_i \beta_i + \alpha_i \tag{OA.33}$$

where E_{T_i} denotes the sample mean in the set of dates for which we have data for underlying i, R_i is the vector of returns of the straddles, λ_i is a vector of risk prices, β_i is a vector of risk prices, and α_i is a vector of pricing errors. Note that these objects are all population values, rather than estimates. In order to calculate the sampling distribution for the estimated counterparts, we need to know the covariance of the pricing errors. Note that there is also a population cross-sectional regression with

$$E_{T_i}[R_i] = a_i + \beta_i E_{T_i}[f_i] + E_{T_i}[\varepsilon_i] \tag{OA.34}$$

where ε_i is a vector of residuals and f_i is a vector of pricing factors. That formula can be used to substitute out returns and obtain

$$\alpha_i = a_i + \beta_i E_{T_i}[f_i] + E_{T_i}[\varepsilon_i] - \lambda_i \beta_i \tag{OA.35}$$

Since a_i, λ_i, and β_i are fixed in the true model, the distribution of α_i depends only on the distributions of the sample means $E_{T_i}[f_i]$ and $E_{T_i}[\varepsilon_i]$. Denoting the long-run (i.e. Hansen–Hodrick) covariance matrix of f_i as Σ_{f_i} and that of ε_i as Σ_{ε_i}, we have

$$var(\alpha_i) = \beta_i T_i^{-1} \Sigma_f \beta_i' + T_i^{-1} \Sigma_{\varepsilon_i} \tag{OA.36}$$

Since the λ_i are estimated from a regression, if we denote their estimates as $\hat{\lambda}_i$, we obtain the usual formula for the variance of $\hat{\lambda}_i - \lambda_i$

$$var\left(\hat{\lambda}_i - \lambda_i\right) = \left(\beta_i'\beta_i\right)^{-1}\beta_i' var\left(\alpha_i\right)\beta_i\left(\beta_i'\beta_i\right)^{-1} \tag{OA.37}$$

$$= \Sigma_f + \left(\beta_i'\beta_i\right)^{-1}\beta_i'\Sigma_{\varepsilon_i}\beta_i\left(\beta_i'\beta_i\right)^{-1} \tag{OA.38}$$

Beyond the variance of $\hat{\lambda}_i$, we also need to know the covariance of any pair of estimates, $\hat{\lambda}_i$ and $\hat{\lambda}_j$. Using standard OLS formulas, we have

$$\begin{bmatrix} \hat{\lambda}_i - \lambda_i \\ \hat{\lambda}_j - \lambda_j \end{bmatrix} = \begin{bmatrix} \left(\beta_i'\beta_i\right)^{-1}\beta_i'\alpha_i \\ \left(\beta_j'\beta_j\right)^{-1}\beta_j'\alpha_j \end{bmatrix} \tag{OA.39}$$

$$= \begin{bmatrix} \left(\beta_i'\beta_i\right)^{-1}\beta_i'\left(\beta_i E_{T_i}\left[f_t\right] + E_{T_i}\left[\varepsilon_{j,t}\right]\right) \\ \left(\beta_j'\beta_j\right)^{-1}\beta_j'\left(\beta_2 E_{T_j}\left[f_t\right] + E_{T_j}\left[\varepsilon_{j,t}\right]\right) \end{bmatrix} \tag{OA.40}$$

The covariance between $\hat{\lambda}_i$ and $\hat{\lambda}_j$ is then

$$\frac{T_{12}}{T_1 T_2}\left(\Sigma_{f,i,j} + \left(\beta_1'\beta_1\right)^{-1}\beta_1'\Sigma_{\varepsilon,i,j}\beta_2\left(\beta_2'\beta_2\right)^{-1}\right) \tag{OA.41}$$

where $\Sigma_{f,i,j}$ and $\Sigma_{\varepsilon,i,j}$ are now long-run covariance matrices (again from the Hansen–Hodrick method). Using these formulas, we then have estimates of risk prices in each market individually along with a full covariance matrix of all the estimates.

OA.6 SDF-based analysis

The marginal effects of realized and implied volatility can be estimated using the stochastic discount factor representation of the factor model estimated in the previous section. Specifically, given the set of straddle returns in each market, one can construct a pricing kernel M_t of the form

$$M_t = \bar{M} - m_i^f \frac{f_{i,t}}{IV_{i,t-1}} - m_i^{f^2}\left(\frac{f_{i,t}}{IV_{i,t-1}}\right)^2 - m_i^{\Delta IV}\frac{\Delta IV_{i,t}}{IV_{i,t-1}} \tag{OA.42}$$

where M_t represents state prices (or marginal utility) and $1 = E_{t-1}M_t R_t$ for any return priced by M. The difference between this specification and that in the previous section is that the coefficients $m^{\cdots}$ represent the marginal impact of each term on marginal utility, whereas the $\gamma^{\cdots}$ coefficients represent the premium for total exposure to the factors. Cochrane (2005) discusses the distinction extensively.

Denoting the covariance matrix of the factors in market i by Σ_i, the m coefficients can be recovered as

$$\left[m_i^f, m_i^{f^2}, m_i^{\Delta IV}\right]' = \Sigma_i^{-1}\left[\gamma_i^f, \gamma_i^{f^2}, \gamma_i^{\Delta IV}\right]' \tag{OA.43}$$

The m's now represent Sharpe ratios on portfolios with exposure to each of the individual factors, orthogonalized to the other two. That is, $m_i^{\Delta IV}$ is the Sharpe ratio for a portfolio exposed to the part of $\frac{\Delta IV_{i,t}}{IV_{i,t-1}}$ that is orthogonal to $\frac{f_{i,t}}{IV_{i,t-1}}$ and $\left(\frac{f_{i,t}}{IV_{i,t-1}}\right)^2$.

Figure OA.12 reports the results of this exercise. The findings are qualitatively consistent with the main results in figure 3 and in fact even stronger quantitatively. The marginal effect of an increase in uncertainty on marginal utility, holding realized volatility fixed, is consistently negative, while an increase in realized volatility increases marginal utility. The fact that these results are close to the benchmark case is a consequence of the weak correlation between innovations in realized and implied volatility, so that the rotation by Σ_i^{-1} has small effects.

Figure OA.12 also reports premia on orthogonalized versions of the rv and iv portfolios.[1] Again, the results are similar to the main analysis.

OA.7 Robustness: ETF options

This section provides an alternative check on the results for crude oil options by examining returns on straddles for options on two exchange traded funds. The first is the United States Oil Fund (USO), which invests in short-term oil futures. USO has existed since 2006, and Optionmetrics reports quotes for options beginning in May, 2007. The second fund is the Energy Select Sector SPDR fund (XLE), which tracks the energy sector of the S&P 500. XLE has existed since 1998 and Optionmetrics reports data since December, 1998.

We eliminate observations using the following filters:

1. Volume less than 10 contracts
2. Time to maturity less than 15 days
3. Bid-ask spread greater than 20 percent of bid/ask midpoint
4. Initial log moneyness – log strike divided by the futures price – greater than 0.75 implied volatility units in absolute value (where implied volatility is scaled by he square root of time to maturity).

We then calculate straddle returns as in the main text over two-week periods and average

[1] These are constructed simply through a rotation. The $rv_\perp$ portfolio has a positive correlation with rv and zero correlation with iv, whie the $iv_\perp$ portfolio has zero correlation with rv and a positive correlation with iv.

across the two straddles nearest to the money for each maturity, weighting them by the inverse of their absolute moneyness.

The top section of table OA.7.1 reports the number of (potentially overlapping) two-week straddle return observations across maturities for USO, XLE, and the CME Group futures options used in the main analysis. Since the CME data goes back to 1983, there are far more observations for that series than the other two. More interestingly, though, the number of observations only declines by about 10 percent between the 1- and 6-month maturities, while it falls by more than 2/3 for the XLE and USO samples. The CME data therefore has superior coverage at longer horizons, which justifies its use in our main analysis.

The bottom section of table OA.7.1 reports the correlations of the USO and XLE straddle returns with those for the CME on the days where they overlap. The correlations are approximately 90 percent at all maturities for USO and 50 percent for XLE. The 90-percent correlations for USO and the CME sample provide a general confirmation of the accuracy of the CME straddle returns, since we would expect the USO and CME options to be highly similar as USO literally holds futures. The lower correlation for XLE is not surprising given that it holds energy sector stocks rather than crude oil futures.

Table OA.7.1.

	Maturity:	1	2	3	4	5	6
# obs.	USO	1640	1616	1721	1679	1118	525
	XLE	2612	2545	2454	1928	1134	369
	CME	6762	6645	6817	6801	6606	5998
Corr. w/	USO	0.93	0.96	0.95	0.92	0.89	0.83
CME	XLE	0.43	0.48	0.50	0.49	0.50	0.53

In the main text, the RV and IV portfolio returns are calculated using 5- and 1-month straddles. Since the number of observations drops off substantially between 4 and 5 months for both XLE and USO, though, here we examine returns on RV and IV portfolios using both 5- and 4-month straddles for the long-maturity side.

Figure OA.16 plots estimated annualized Sharpe ratios along with 95-percent confidence bands for the RV and IV portfolios using 4- and 5-month straddles for the three sets of options. In all four cases, the three confidence intervals always overlap substantially. The fact that the sample for the CME options is far larger is evident in its confidence bands being much narrower than those for the other two sources. For the IV portfolios, USO has returns that are close to zero, but its confidence bands range from -1 to greater than 0.5, indicating that it is not particularly informative about the Sharpe ratio.

Table OA.7.2 reports confidence bands for the *difference* between the IV and RV average returns constructed with the CME data and the same portfolios constructed using USO and XLE. The top panel shows that the differences for the IV portfolios are negative for USO and positive for XLE, but only the difference for USO constructed with the 4-month straddle is statistically significant. The bottom panel similarly shows mixed results for the point estimates for the differences for the RV portfolios, with none of the differences being statistically significant.

Table OA.7.2. Differences between CME and USO, XLE mean returns

	USO - CME, 4mo.	USO - CME, 5mo.	XLE - CME, 4mo.	XLE - CME, 5mo.
IV return	-2.2	-2.2	-0.8	-1.4
	[-3.9,-0.2]	[-4.8,0.4]	[-2.5,4.1]	[-4.1,6.3]
RV return	0.43	0.47	-0.27	0.67
	[-0.6,1.4]	[-0.6,1.4]	[-1.8,1.3]	[-1.5,2.6]

Notes: the table reports percentage (two-week) returns on USO and XLE minus returns on CME RV and IV portfolios. 95-percent confidence intervals are reported in brackets.

The fact that the USO and CME straddle returns are highly correlated does not necessarily mean that the CME data is accurate for the mean return on the straddles. To check whether the difference in the means observed in the USO and XLE data would affect out main results, we ask how the Sharpe ratios of the RV and IV portfolios in the CME data would change if we shifted their means by the average differences reported in table OA.7.2. The bars labeled "CME, USO adj." and CME, XLE adj." show how the confidence bands would change if we shifted them by exactly the point estimates from table OA.7.2. Note that this is not the same as shifting the Sharpe ratio for the CME data to match that for the XLE or USO data. The reason is that the difference in table OA.7.2 is calculated only for the returns on matching dates, whereas the Sharpe ratio calculated in figure OA.16 is calculated using the full sample for the CME data. So the two adjusted bands take the full-sample band and then shift it by the mean difference calculated on the dates that overlap between the CME data and XLE or USO.

Figure OA.16 shows that the economic conclusions drawn for the crude oil straddles are not changed if the mean returns are shifted by the differences observed in table OA.7.1. The RV portfolio returns remain statistically significantly negative in all four cases, the changes in the point estimates are well inside the original confidence intervals. The top panel shows that the IV returns using 5-month straddles are similarly unaffected. For the 4-month straddles, the only difference is that with the USO options, the estimated Sharpe ratio falls by about half and is no longer statistically significantly greater than zero. So, again, out of eight cases – IV and RV with 4- and 5-month straddles – in only one is there a nontrivial change in the conclusions, and even there the Sharpe ratio on the IV portfolio does not become negative, it is simply less positive.

Overall, the period in which the USO and XLE options are traded is too short to use them for our main analysis. This section shows that the USO straddle returns are highly correlated with the CME returns. The mean returns on the XLE and CME straddles are highly similar, while they differ somewhat more for CME and USO. However, shifting the means used for the CME options in the main analysis by the observed difference between the CME and USO options does not substantially change any of the conclusions.

OA.8 Model

To help provide some context for the empirical results and fit them into a standard framework, this section describes results from a simple extension of the standard long-run risk model of Bansal and Yaron (2004). The technical analysis is in Section OA.9; here we report the specification and key results.

Agents have Epstein–Zin preferences over consumption, C_t, with a unit elasticity of substitution, where the lifetime utility function, v_t, satisfies

$$v_t = (1-\beta)\log C_t + \frac{\beta}{1-\alpha}\log E_t \exp\left((1-\alpha)\, v_{t+1}\right) \tag{OA.44}$$

where α is the coefficient of relative risk aversion. Consumption growth follows the process

$$\Delta c_t = x_{t-1} + \sqrt{\sigma^2_{B,t-1} + \sigma^2_{G,t-1}}\varepsilon_t + Jb_t \tag{OA.45}$$

$$x_t = \phi_x x_{t-1} + \omega_x \eta_{x,t} + \omega_{x,G}\eta_{\sigma,G,t} - \omega_{x,B}\eta_{\sigma,B,t} \tag{OA.46}$$

$$\sigma^2_{j,t} = (1-\phi_\sigma)\,\bar{\sigma}^2_j + \phi_\sigma \sigma^2_{j,t-1} + \omega_j \eta_{\sigma,j,t}, \text{ for } j \in \{B, G\} \tag{OA.47}$$

where ε_t and the $\eta_{\cdot,t}$ are independent standard normal random variables. x_t represents the consumption trend. We have two deviations from the usual setup. First, we include jump shocks, Jb_t, where b_t is a Poisson distributed random variable with intensity λ and J is the magnitude of the jump. This addition allows for random variation in realized volatility and is drawn from Drechsler and Yaron (2011). Second, there are two components to volatility, which we refer to as bad and good. Bad volatility, σ^2_B, is associated with low future consumption growth, while good volatility, σ^2_G, is associated with high future growth (where all of the $\omega_\cdot$ coefficients are nonnegative).

Define realized volatility to be the realized quadratic variation in consumption growth, while implied volatility is the conditional variance of consumption growth (these are formalized in the appendix).

Proposition 1 *The average excess returns on forward claims to realized and implied volatility for consumption growth in this model are,*

$$E\left[RV_{t+1} - P_{RV,t}\right] = J^2 \lambda \left(1 - \exp\left(-\alpha J\right)\right) \tag{OA.48}$$
$$E\left[IV_{t+1} - P_{IV,t}\right] = \left(\alpha - 1\right) \left(\upsilon_{Y,x}\left(\omega_{x,G}\omega_G - \omega_{x,B}\omega_B\right) + \upsilon_{Y,\sigma}\left(\omega_G^2 + \omega_B^2\right)\right) \tag{OA.49}$$

where $P_{x,t}$ is the forward price for x. $E\left[IV_{t+1} - P_{IV,t}\right] > 0$ for $\omega_{x,G}$ sufficiently larger than $\omega_{x,B}$. Furthermore, the sign of $E\left[RV_{t+1} - P_{RV,t}\right]$ is the same as the sign of J and of the conditional skewness of consumption growth (i.e. the skewness of Δc_{t+1} conditional on date-t information).

Proposition 1 contains our key analytic results. We analyze premia for realized and implied volatility on consumption – real activity – consistent with the focus in the empirical analysis on macro volatility and uncertainty. The negative premium on realized volatility is driven by downward jumps, similar to the literature on the volatility risk premium in equities (Drechsler and Yaron (2011), Wachter (2013)). The sign of the premium on implied volatility depends on the contribution of good versus bad volatility. When good volatility shocks, where high volatility is associated with high future growth (e.g. due to learning about new technologies), are relatively larger than bad volatility shocks ($\omega_{x,G}\omega_G > \omega_{x,B}\omega_B$) the premium on implied volatility can be positive.

Section OA.9 provides a numerical calibration of the model using values close to those in Bansal and Yaron's (2004) original choices. It shows that the model generates quantitatively realistic Sharpe ratios for implied and realized volatility in addition to a reasonable equity premium.

The key economic mechanism for the positive pricing of uncertainty shocks is that high volatility is sometimes associated with higher long-term growth. Intuitively, that mechanism contributes positive skewness to consumption growth, while the jumps contribute negative skewness. The appendix provides novel evidence on the skewness of consumption growth consistent with the model. In particular, conditional skewness in the model, which depends only on the jumps, is more negative than the skewness of expected consumption growth, which depends on the relationship of volatility and long-run growth (x). We show that consumption growth displays exactly the same pattern in US data.

So a simple version of the long-run risk model with good and bad volatility shocks and jumps in consumption can match our key empirical facts. Furthermore, the empirical results are sharp, in the sense that the sign of the premium on implied volatility identifies the relative importance of the bad and good volatility shocks, while the sign of the premium on realized volatility identifies the sign of consumption jumps.

OA.9 Model details

OA.9.1 Dynamics

Consumption growth follows

$$\Delta c_t = x_{t-1} + \sqrt{\sigma_{B,t-1}^2 + \sigma_{G,t-1}^2}\varepsilon_t + Jb_t \tag{OA.50}$$

$$x_t = \phi_x x_{t-1} + \omega_x \eta_{x,t} + \omega_{x,G}\eta_{\sigma,G,t} - \omega_{x,B}\eta_{\sigma,B,t} \tag{OA.51}$$

$$\sigma_{j,t}^2 = (1-\phi_\sigma)\,\bar{\sigma}_j^2 + \phi_\sigma \sigma_{j,t-1}^2 + \omega_j \eta_{\sigma,j,t} \tag{OA.52}$$

for $j \in \{G, B\}$. The shocks ε, η_x, η_G, η_B are independent and Gaussian with unit variances. The ω coefficients are all assumed to be positive. b_t is a Poisson random variable with intensity λ.

The dynamics can also be written as

$$\begin{bmatrix} x_t \\ \sigma_t^2 - \bar{\sigma}^2 \end{bmatrix} = \begin{bmatrix} \phi_x & 0 \\ 0 & \phi_\sigma \end{bmatrix} \begin{bmatrix} x_{t-1} \\ \sigma_{t-1}^2 - \bar{\sigma}^2 \end{bmatrix} + \begin{bmatrix} \omega_x & \omega_{x,G} & 0 \\ 0 & \omega_G & \omega_B \end{bmatrix} \begin{bmatrix} \eta_{x,t} \\ \eta_{G,t} \\ \eta_{B,t} \end{bmatrix} \tag{OA.53}$$

$$\Delta c_t = x_{t-1} + \sigma_{t-1}^2 \varepsilon_t + Jb_t \tag{OA.54}$$

$$Y_t = FY_{t-1} + G\eta_t \tag{OA.55}$$

where $Y_t = [x_t, \sigma_t^2 - \bar{\sigma}^2]'$, etc. The fact that the model can be rewritten with only a single variance process follows from the linearity of the two processes, the fact that they have the same rate of mean reversion, and the fact that they appear additively. We can then write consumption and dividend growth as

$$\Delta c_t = c_Y' Y_{t-1} + \sqrt{\bar{\sigma}^2 + g_Y' Y_{t-1}}\varepsilon_t + Jb_t \tag{OA.56}$$

$$\Delta d_t = \gamma\left(c_Y' Y_{t-1} + \sqrt{\bar{\sigma}^2 + g_Y' Y_{t-1}}\varepsilon_t + Jb_t\right) + \omega_d \varepsilon_{d,t} \tag{OA.57}$$

for vectors c_Y and g_Y. Δd_t is log dividend growth, which we will use for modeling equities. It satisfies $\Delta d_t = \gamma \Delta c_t + \omega_d \varepsilon_{d,t}$ ($\varepsilon_{d,t} \sim N(0,1)$), where γ determines the leverage of equities.

OA.9.2 Preferences

We assume agents have Epstein–Zin preferences with a unit IES,

$$v_t = (1-\beta)\,c_t + \frac{\beta}{1-\alpha}\log E_t \exp\left((1-\alpha)\,v_{t+1}\right) \tag{OA.58}$$

$$vc_t = \frac{\beta}{1-\alpha}\log E_t \exp\left((1-\alpha)\left(vc_{t+1} + \Delta c_{t+1}\right)\right) \tag{OA.59}$$

where vc_t is the log utility/consumption ratio, $vc_t = v_t - c_t$. We look for a solution to the model of the form

$$vc_t = \bar{v} + v_Y' Y_t \tag{OA.60}$$

Inserting into the recursion for vc,

$$vc_t = \frac{\beta}{1-\alpha}\log E_t \exp\left((1-\alpha)\left(\bar{v} + v_Y' Y_{t+1} + c_Y' Y_t + \sqrt{g_Y' Y_t}\,\varepsilon_{t+1} + J b_{t+1}\right)\right) \tag{OA.61}$$

$$= \frac{\beta}{1-\alpha}\log E_t \exp\left((1-\alpha)\left(\bar{v} + v_Y'\left(F Y_t + G\eta_{t+1}\right) + c_Y' Y_t + \sqrt{\bar{\sigma}^2 + g_Y' Y_t}\,\varepsilon_{t+1} + J b_{t+1}\right)\right) \tag{OA.62}$$

$$= \beta\left(\bar{v} + \left(v_Y' F + c_Y'\right) Y_t\right) + \beta\frac{1-\alpha}{2}\left(v_Y' G G' v_Y + \bar{\sigma}^2 + g_Y' Y_t\right) + \frac{\beta}{1-\alpha}\lambda\left(\exp\left((1-\alpha)\,J\right) - 1\right) \tag{OA.63}$$

Matching coefficients,

$$v_Y' = \beta\left(v_Y' F + c_Y'\right) + \beta\frac{1-\alpha}{2} g_Y' \tag{OA.64}$$

$$v_Y' = \beta\left(c_Y' + \frac{1-\alpha}{2} g_Y'\right)\left(I - \beta F\right)^{-1} \tag{OA.65}$$

$$\bar{v} = \frac{\beta}{1-\beta}\left(\frac{1-\alpha}{2}\left(v_Y' G G' v_Y + \bar{\sigma}^2\right) + \frac{1}{1-\alpha}\lambda\left(\exp\left((1-\alpha)\,J\right) - 1\right)\right) \tag{OA.66}$$

The pricing kernel is then

$$M_{t+1} = \beta\frac{\exp\left((1-\alpha)\left(vc_{t+1}\right)\right)}{E_t \exp\left((1-\alpha)\left(vc_{t+1} + \Delta c_{t+1}\right)\right)}\exp\left(-\alpha\Delta c_{t+1}\right) \tag{OA.67}$$

$$m_{t+1} = -\log\beta + (1-\alpha)\,vc_{t+1} - \alpha\Delta c_{t+1} - \log E_t \exp\left((1-\alpha)\left(vc_{t+1} + \Delta c_{t+1}\right)\right) \tag{OA.68}$$

Or, equivalently,

$$m_{t+1} = m_0 + m_Y' Y_t + m_\eta \eta_{t+1} - \alpha\sqrt{\bar{\sigma}^2 + g_Y' Y_t}\,\varepsilon_{t+1} - \alpha J b_{t+1} \tag{OA.69}$$

$$m_0 = -\log\beta - \frac{(1-\alpha)^2}{2}\left(v_Y' G G' v_Y + \bar{\sigma}^2\right) - \lambda\left(\exp\left((1-\alpha)J\right) - 1\right) \quad \text{(OA.70)}$$

$$m_Y' = -c_Y - \frac{(1-\alpha)^2}{2} g_Y \quad \text{(OA.71)}$$

$$m_\eta = (1-\alpha) v_Y' G \quad \text{(OA.72)}$$

OA.9.3 Pricing equities

We have the usual Campbell–Shiller approximation for the return on equities, r_{t+1}, with

$$r_{t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + \Delta d_{t+1} \quad \text{(OA.73)}$$

where z_t is the log price/dividend ratio of equities. We look for a solution of the form $z_t = z_0 + z_Y' Y_t$, which leads to the pricing equation

$$0 = \log E_t \exp\left(\begin{array}{c} m_0 + m_Y' Y_t + m_\eta \eta_{t+1} - \alpha\sqrt{\bar{\sigma}^2 + g_Y' Y_t}\varepsilon_{t+1} - \alpha J b_{t+1} \\ +\kappa_0 + (\kappa_1 - 1) z_0 + \kappa_1 z_Y' (F Y_t + G\eta_{t+1}) - z_Y' Y_t \\ +\gamma\left(c_Y' Y_t + \sqrt{\bar{\sigma}^2 + g_Y' Y_t}\varepsilon_{t+1} + J b_{t+1}\right) + \omega_d \varepsilon_{d,t+1} \end{array}\right) \quad \text{(OA.74)}$$

The solution satisfies

$$z_0 = (1-\kappa_1)^{-1}\left(\begin{array}{c} m_0 + \kappa_0 + \lambda\left(\exp\left((\gamma-\alpha)J\right) - 1\right) \\ +\frac{1}{2}\left((m_\eta + \kappa_1 z_Y' G)(m_\eta + \kappa_1 z_Y' G)' + (\gamma-\alpha)^2 \bar{\sigma}^2 + \omega_d^2\right) \end{array}\right) \quad \text{(OA.75)}$$

$$z_Y' = \left(m_Y' + \gamma c_Y' + \frac{1}{2}(\gamma-\alpha)^2 g_Y'\right)(I - \kappa_1 F)^{-1} \quad \text{(OA.76)}$$

OA.9.3.1 Average excess returns

To get average returns, on equities, first note that

$$\log E_t\left[\exp\left(r_{t+1} - r_{f,t}\right)\right] = \log E_t\left[\exp\left(\begin{array}{c} \kappa_0 + (\kappa_1 - 1) z_0 + \kappa_1 z_Y'(F Y_t + G\eta_{t+1}) - z_Y' Y_t \\ +\gamma\left(c_Y' Y_t + \sqrt{\bar{\sigma}^2 + g_Y' Y_t}\varepsilon_{t+1} + J b_{t+1}\right) + \omega_d \varepsilon_{d,t+1} \\ -r_{f,0} - r_{f,1}' Y_t \end{array}\right)\right] \quad \text{(OA.77)}$$

$$= \kappa_0 + (\kappa_1 - 1) z_0 - r_{f,0} + \left(\kappa_1 z_Y' F - z_Y' + \gamma c_Y' - r_{f,1}'\right) Y_t \quad \text{(OA.78)}$$

$$+\frac{1}{2}\left(\kappa_1^2 z_Y' G G' z_Y + \gamma^2\left(\bar{\sigma}^2 + g_Y' Y_t\right)\right) + \frac{1}{2}\omega_d^2 + \lambda\left(\exp(\gamma J) - 1\right) \quad \text{(OA.79)}$$

The risk-free rate is of the form $r_{f,t} = r_{f,0} + r'_{f,1} Y_t$, with

$$r_{f,0} = \log\beta + \frac{(1-2\alpha)}{2}\bar{\sigma}^2 + \lambda\left(\exp\left((1-\alpha)J\right) - \exp\left(-\alpha J\right)\right) \quad \text{(OA.80)}$$

$$r'_{f,1} = c'_Y - \frac{1}{2}\alpha^2 g'_Y \quad \text{(OA.81)}$$

which allows for the calculation of the average excess return on equities. The conditional standard deviation of equity returns is

$$\sqrt{\kappa_1^2 z'_Y G G' z_Y + \gamma^2\bar{\sigma}^2 + \gamma^2 J^2 \lambda} \quad \text{(OA.82)}$$

OA.9.4 Pricing realized volatility

Since our empirical work estimates premia for realized and implied volatility for macro variables, we examine here the pricing of realized and implied volatility for Δc_{t+1}. The cumulative innovation in consumption between dates t and $t+1$ is

$$\Delta c_{t+1} - E_t \Delta c_{t+1} = \sigma_t^2 \varepsilon_{t+1} + J\left(b_{t+1} - \lambda\right)$$

The first part is typically thought of as a diffusive component. That is, we can think of $\varepsilon_{t+1} = B_{t+1} - B_t$, for a standard (continuous-time) Brownian motion B_t. Similarly, b_{t+1} is the innovation in a pure jump process, $b_{t+1} = N_{t+1} - N_t$, where N_t is a (continuous-time) Poisson counting process. Now consider measuring the total quadratic variation in those two processes (i.e. as though we were measuring realized volatility from daily futures returns, as in our empirical analysis). The quadratic variation in B between dates t and $t+1$ is exactly 1, while the quadratic variation in N is exactly $N_{t+1} - N_t = b_{t+1}$. We then say that the realized volatility in consumption growth between period t and $t+1$ is

$$RV_{t+1} = \sigma_t^2 + J^2 b_{t+1} \quad \text{(OA.83)}$$

In this case, the diffusive part of the realized volatility is entirely predetermined. This is a typical result. It is only the jumps that contribute an unexpected component to realized volatility. The pricing of realized volatility will therefore depend on the pricing of jumps.

The price of a forward claim on RV_{t+1} is

$$\begin{aligned}
P_{RV,t} &= E_t\left[\frac{\exp(m_{t+1})}{E_t\exp(m_{t+1})}RV_{t+1}\right] \\
&= E_t\left[\exp\left(\begin{array}{c}(1-\alpha)\,v_Y' G\eta_{t+1}-\alpha\left(\sqrt{\bar{\sigma}^2+g_Y' Y_t}\varepsilon_{t+1}+Jb_{t+1}\right)\\ -\left[\frac{1}{2}\left((1-\alpha)^2 v_Y' GG' v_Y'+\alpha^2\left(\bar{\sigma}^2+g_Y' Y_t\right)\right)+\lambda\left(\exp(-\alpha J)-1\right)\right]\end{array}\right)\left(\sigma_t^2+J^2 b_{t+1}\right)\right] \\
&= \sigma_t^2+J^2\lambda\exp(-\alpha J)
\end{aligned}$$

The average excess return on that forward is then

$$E_t\left[RV_{t+1}-P_{RV,t}\right] = \sigma_t^2+J^2\lambda-\sigma_t^2-J^2\lambda\exp(-\alpha J) \quad \text{(OA.84)}$$

$$= J^2\lambda\left(1-\exp(-\alpha J)\right) \quad \text{(OA.85)}$$

The sign of this object is equal to the sign of J. Note also that this is the sign of the conditional skewness of consumption growth.

OA.9.5 Pricing uncertainty

We define uncertainty on date t as expected realized volatility on date $t+1$. That is, it is the conditional variance for Δc_{t+1}. So we say

$$IV_t \equiv \sigma_t^2+J^2\lambda \quad \text{(OA.86)}$$

We now consider the price and excess return for a forward claim to IV_{t+1}.

$$\begin{aligned}
P_{IV,t} &= E_t\left[\frac{\exp(m_{t+1})}{E_t\exp(m_{t+1})}IV_{t+1}\right] \\
&= J^2\lambda+\bar{\sigma}^2+\phi_\sigma\hat{\sigma}_t^2+E_t\left[\exp\left(\begin{array}{c}(1-\alpha)\,v_Y' G\eta_{t+1}\\ -\frac{1}{2}\left((1-\alpha)^2 v_Y' GG' v_Y'\right)\end{array}\right)g_Y'\eta_{t+1}\right] \\
&= J^2\lambda+\bar{\sigma}^2+\phi_\sigma\hat{\sigma}_t^2+\frac{E_t\left[\exp\left((1-\alpha)\,v_Y' G\eta_{t+1}\right)g_Y' G\eta_{t+1}\right]}{\exp\left(\frac{1}{2}\left((1-\alpha)^2 v_Y' GG' v_Y'\right)\right)} \\
&= J^2\lambda+\bar{\sigma}^2+\phi_\sigma\hat{\sigma}_t^2+(1-\alpha)\left(\begin{array}{c}\omega_G\left(v_{Y,x}\omega_{x,G}+v_{Y,\sigma}\omega_G\right)\\ +\omega_B\left(v_{Y,\sigma}\omega_B-v_{Y,x}\omega_{x,B}\right)\end{array}\right)
\end{aligned}$$

where the last line follows from straightforward but tedious algebra. The average return on the claim on uncertainty is then

$$E\left[IV_{t+1}\right]-P_{IV,t} = J^2\lambda+\bar{\sigma}^2+\phi_\sigma\hat{\sigma}_t^2-\left(J^2\lambda+\bar{\sigma}^2+\phi_\sigma\hat{\sigma}_t^2+(1-\alpha)\left(\begin{array}{c}\omega_G\left(v_{Y,x}\omega_{x,G}+v_{Y,\sigma}\omega_G\right)\\+\omega_B\left(v_{Y,\sigma}\omega_B-v_{Y,x}\omega_{x,B}\right)\end{array}\right)\right) \quad \text{(OA.87)}$$

$$= -(1-\alpha)\left(\begin{array}{c}\omega_G\left(v_{Y,x}\omega_{x,G}+v_{Y,\sigma}\omega_G\right)\\+\omega_B\left(v_{Y,\sigma}\omega_B-v_{Y,x}\omega_{x,B}\right)\end{array}\right) \quad \text{(OA.88)}$$

In the standard case from Bansal and Yaron (2004), we would have $\omega_{x,G}=\omega_{x,B}=0$, so this would be

$$E\left[IV_{t+1}\right]-P_{IV,t}=(\alpha-1)\,v_{Y,\sigma}\left(\omega_G^2+\omega_B^2\right) \quad \text{(OA.89)}$$

Since $v_{Y,\sigma}<0$, the premium for IV will be negative in that case. Now when $\omega_{x,G}$ can be positive, we have

$$E\left[IV_{t+1}\right]-P_{IV,t}=(\alpha-1)\left(v_{Y,x}\left(\omega_{x,G}\omega_G-\omega_{x,B}\omega_B\right)+v_{Y,\sigma}\left(\omega_G^2+\omega_B^2\right)\right) \quad \text{(OA.90)}$$

Since $v_{Y,x}>0$, if $\omega_{x,G}$ is sufficiently large relatively to $\omega_{x,B}$, the premium can be positive.

The Sharpe ratio on this object depends on the standard deviation of $IV_{t+1}-P_{IV,t}$, which is exactly $\sqrt{\omega_G^2+\omega_B^2}$.

OA.9.6 Calibration

The calibration is relatively close to Bansal and Yaron's (BY; 2004) choices, with a few changes. For the preferences, we set $\beta=0.998$ and $\alpha=15$. β is as in BY, while α is set somewhat higher to help match the equity premium. We study post-war data here, in which the volatility of consumption growth is lower, thus necessitating higher risk aversion to match the equity premium. Leverage, γ, is set to 3.5, on the upper end of the range of values studied by BY.

The jump intensity is 1/18, implying jumps occur on average once every 18 months, while the jump size $J=-0.015$.

The persistence of x and σ^2 are 0.979 and 0.987, as in BY.

$\bar{\sigma}=0.0039$, which is half the value used in BY in order to match the lower consumption volatility noted above. The standard deviation of innovations to x is set to $0.06\times\sigma$, which is somewhat higher than the value of 0.044 in BY. Of that, $\omega_x=\omega_{x,G}=0.0129$ and $\omega_{x,B}=0$. Similarly, $\omega_G=\omega_B=1.62\times10^{-6}$, so that the standard deviation of innovations to σ^2 is 0.23×10^{-5}, as in BY. Finally, $\omega_d=0.01$.

OA.9.7 Results

The table below lists key moments from the model along with analogs from the data. The model moments are based on a monthly simulation of the model that is aggregated to the quarterly frequency to match quarterly data observed empirically (see also BY).

The first three rows on the left show that the model is able to generate realistic values for mean, standard deviation, and Sharpe ratio for equity returns. The top row on the right shows that the volatility of consumption growth is somewhat higher than in the data. However, this value is still smaller than that used by Bansal and Yaron (2004) by 40 percent. Our calibration of 0.87 percent is the midpoint between Bansal and Yaron's (2004) original value and the value in the post-war data. Using a smaller volatility would require either increasing some other form of risk (e.g. long-run risk or stochastic volatility) or risk aversion in order to generate a realistic equity premium.

Next, the table shows that the Sharpe ratios for claims on RV and IV are approximately -0.21 and 0.19, respectively, which agree well with the empirical values (which are calculated as the overall means across all 19 markets we study; see figure 3). These are the key moments that the model was designed to match. They show that it is able to generate quantitatively realistic premia for uncertainty and realized volatility shocks.

As discussed in the main text, the economic mechanism behind the negative premium on RV is negative conditional skewness in consumption growth, while the mechanism behind the positive premium for IV – the good volatility shocks that raise future consumption growth – pushes in the direction of positive skewness. That implies that the skewness of the conditional expectation of consumption growth should be less negative than conditional skewness. To test that idea, we examine skewness in the model and data. The information set used for conditioning here is lagged consumption growth. That is, we look at results involving regressions of consumption growth on three of its own lags in both the model and the data.

The table shows that the data and model both share the feature that the conditional expectation of consumption growth is much less negatively skewed than the surprise in consumption growth, consistent with the main mechanism in the model. This is not a moment that the model was explicitly designed to match. The model was meant to match the premia on RV and IV, so this represents an additional test of the proposed mechanism.

To be clear, the main contribution of the paper is not meant to be this model, but nevertheless this section shows that the empirical results can be rationalized in a standard structural asset pricing model.

Summary statistics from the model and empirical data, 1947–2018

	Model	Data		Model	Data
$E\left[r_m - r_f\right]$	0.077	0.056	$std\left(\Delta c\right)$	0.0087	0.0052
$std\left(r_m - r_f\right)$	0.14	0.11	$skew_t\left(\Delta c_{t+1}\right)$	-0.32	-0.15
$\frac{E\left[r_m - r_f\right]}{std\left(r_m - r_f\right)}$	0.53	0.52	$skew\left(E_t \Delta c_{t+1}\right)$	-0.10	-0.07
$\frac{E\left[RV_{t+1} - P_{RV,t}\right]}{std\left[RV_{t+1} - P_{RV,t}\right]}$	-0.21	-0.32			
$\frac{E\left[RV_{t+1} - P_{RV,t}\right]}{std\left[RV_{t+1} - P_{RV,t}\right]}$	0.19	0.26			

References

Bansal, Ravi and Amir Yaron, "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 2004, *59* (4), 1481–1509.

Cochrane, John H., *Asset Pricing*, Princeton University Press, 2005.

Drechsler, Itamar and Amir Yaron, "What's Vol Got to Do with it?," *The Review of Financial Studies*, 2011, *24(1)*, 1–45.

Giglio, Stefano and Dacheng Xiu, "Asset pricing with omitted factors," *Chicago Booth Research Paper*, 2019, (16-21).

Wachter, Jessica A., "Can time-varying risk of rare disasters explain aggregate stock market volatility?," *Journal of Finance*, 2013, *68(3)*, 987–1035.

Figure OA.1: Fit to realized volatility indexes

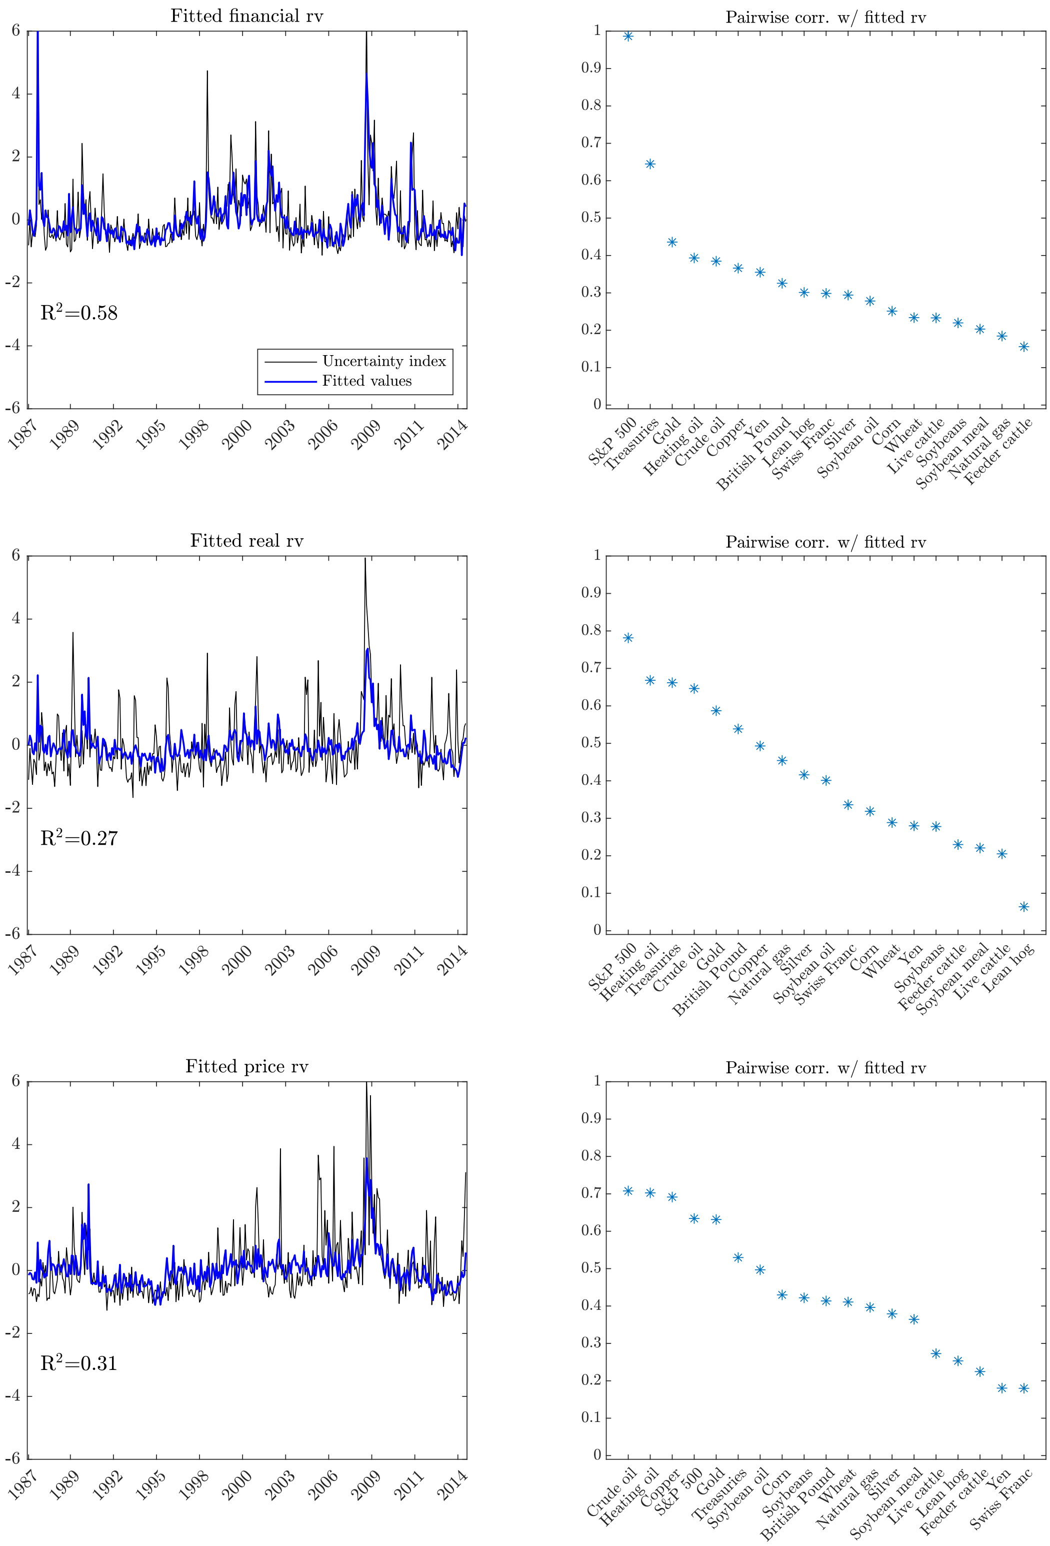

Note: See figure 2. This figure uses the JLN realized volatility series instead of uncertainty.

Figure OA.2: Factor loadings

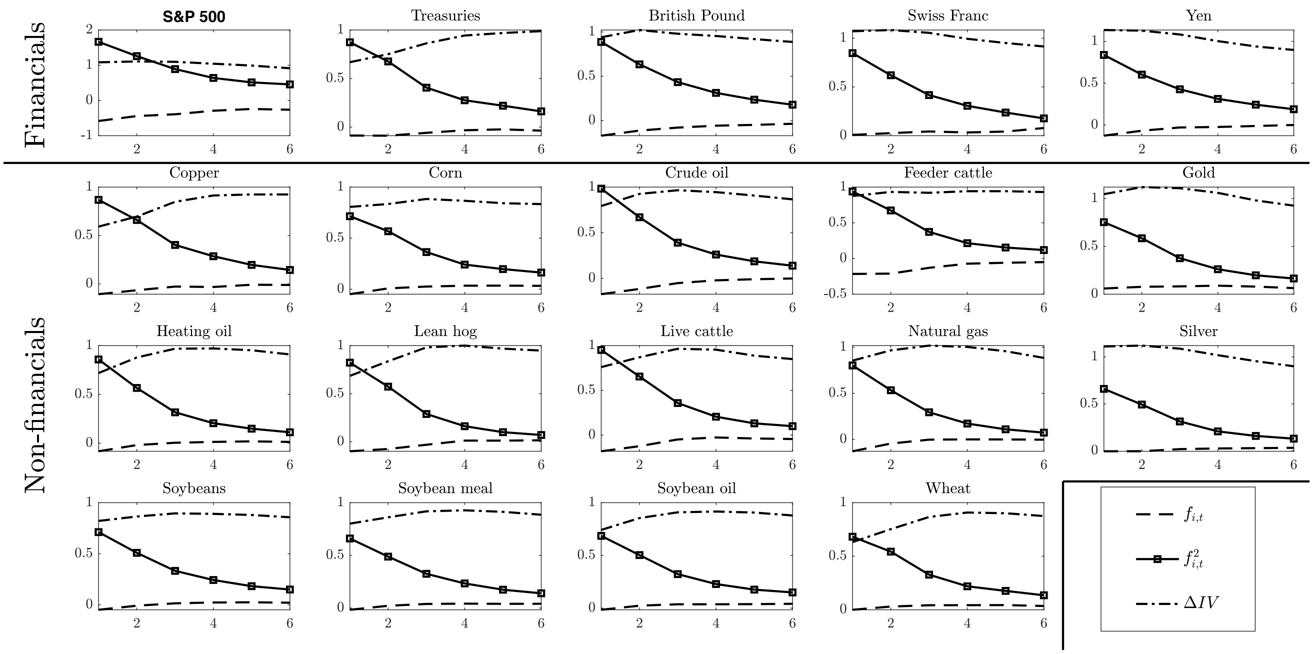

Note: Loadings of two-week straddle returns on the three risk factors. The factors are all scaled by current IV, as in equation 2. The loadings are scaled so that if the Black-Scholes approximation was exact, the loading on ΔIV would be 1 at all maturities, the loading on $f_{i,t}$ would be 0 at all maturities, and the loading on $f_{i,t}^2$ would be $1/n$ where n is the maturity in months.

Figure OA.3: rv and iv portfolio approximation errors

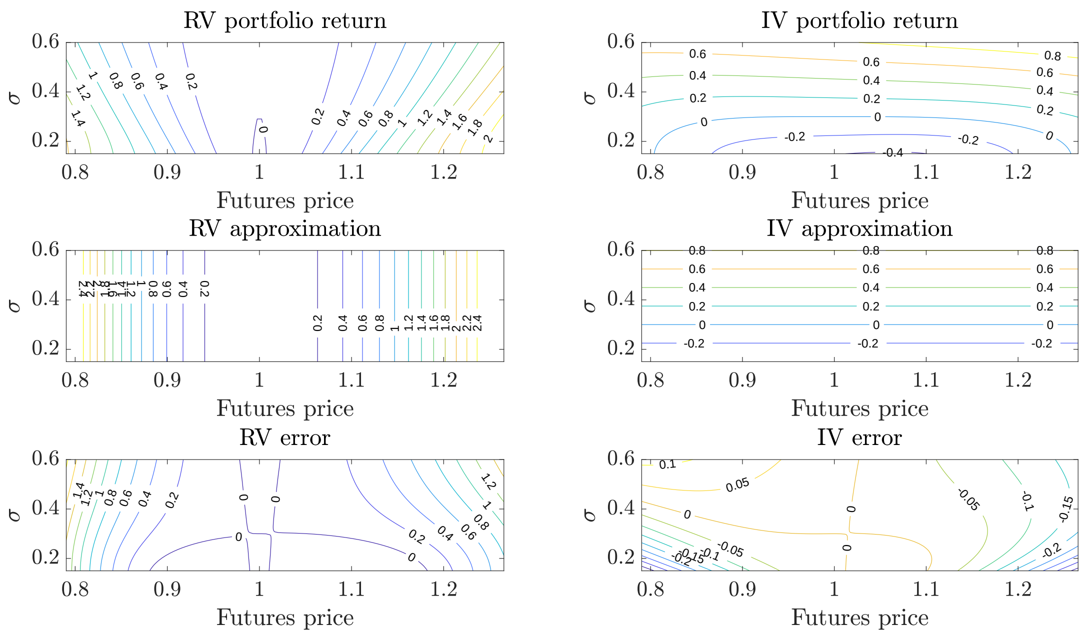

Note: The initial futures price is 1 and the initial volatility, σ, is 0.3. The top panels calculate the return on the rv and iv portfolios given an instantaneous shift in the futures price and volatility to the values reported on the axes under the assumption that the Black–Scholes formula holds. The middle panels plot returns under the approximations used in the text. The bottom panels are equal to the middle minus the top panels. All returns and errors are reported as decimals.

Figure OA.4: Straddle and strangle returns

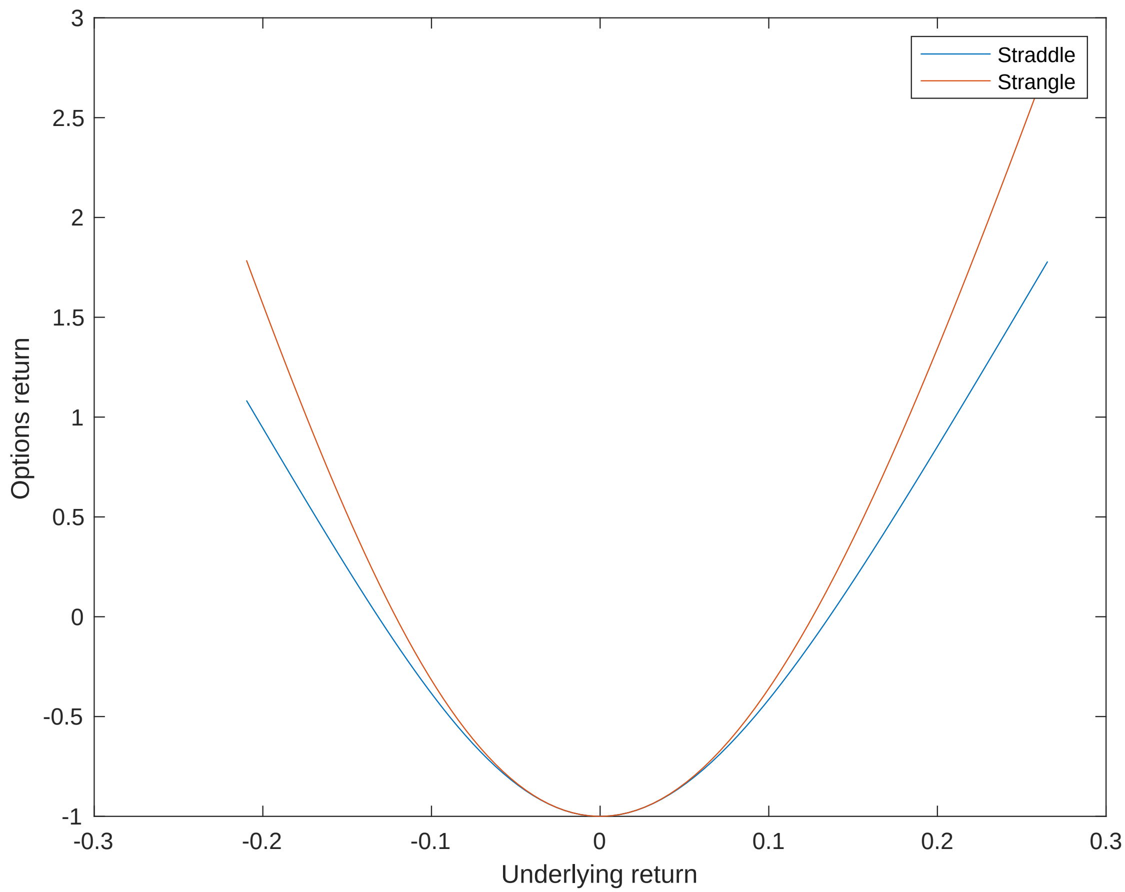

Note: Returns of 1-standard deviation strangles and straddles as function of the underlying's return.

Figure OA.5: Imposing a filter on volume

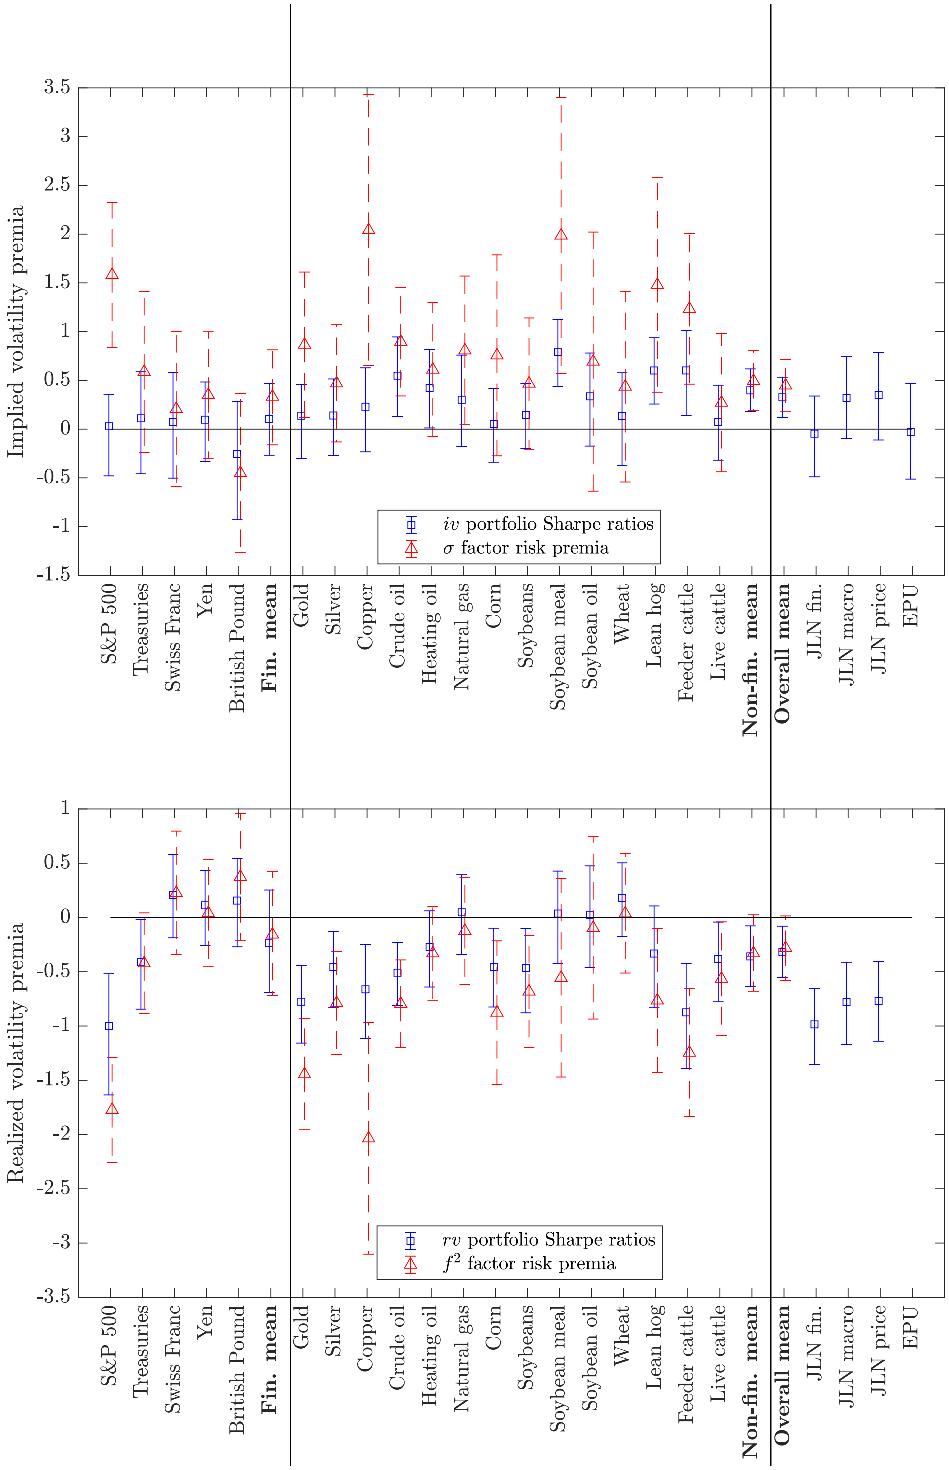

Note: Same as figure 3, but using only options for which volume is neither zero nor missing.

Figure OA.6: RV and IV portfolio Sharpe ratios and factor risk premia, one-week holding period

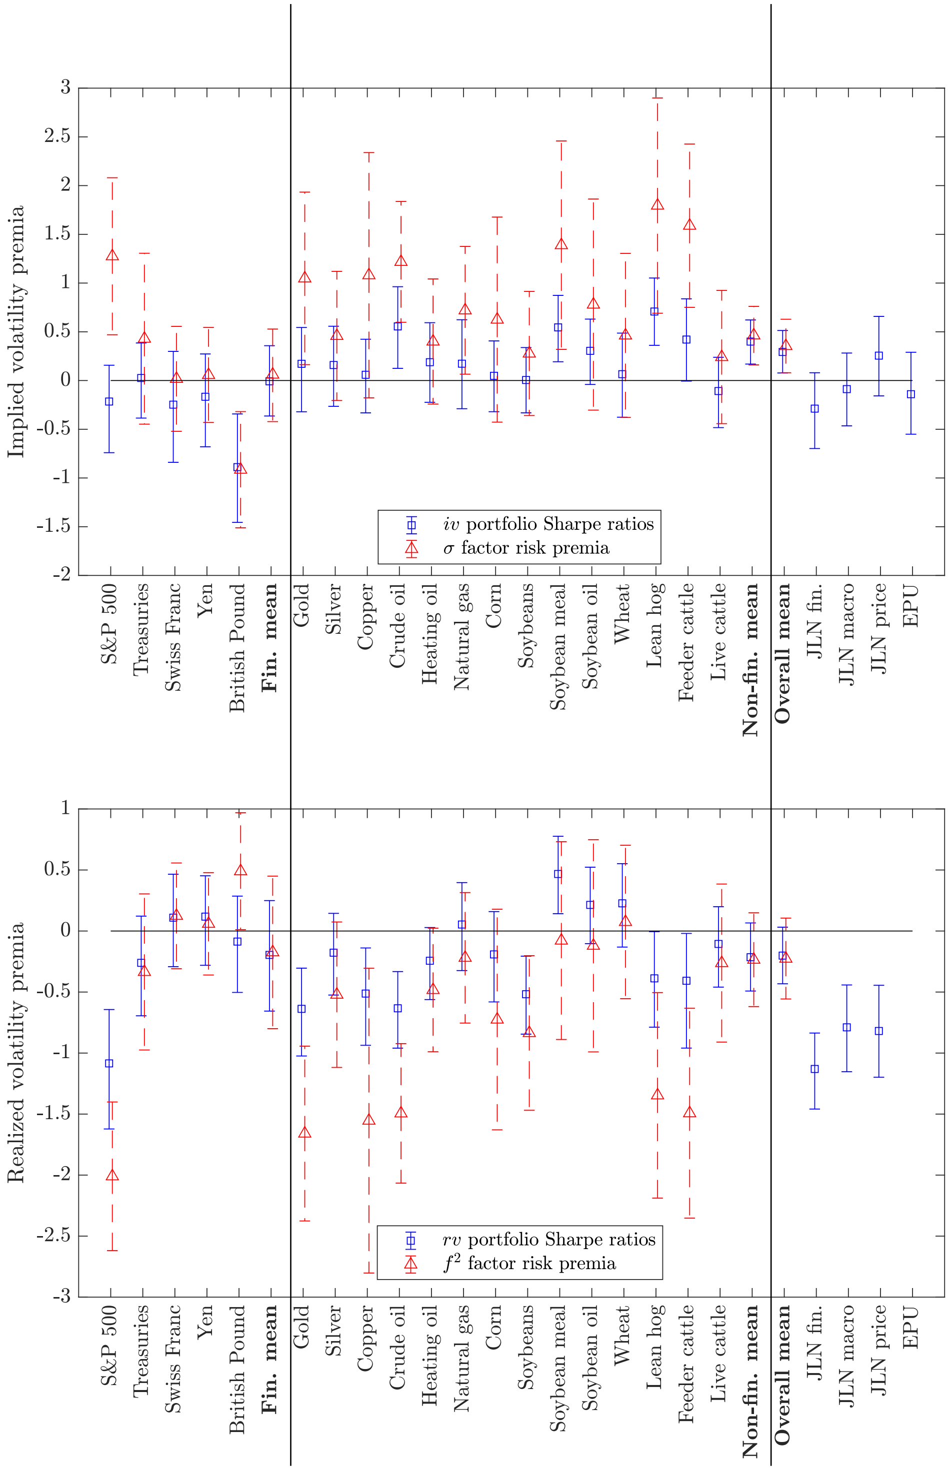

Note: Same as figure 3, but using one-week holding periods.

Figure OA.7: RV and IV portfolio Sharpe ratios and factor risk premia (first half of the sample)

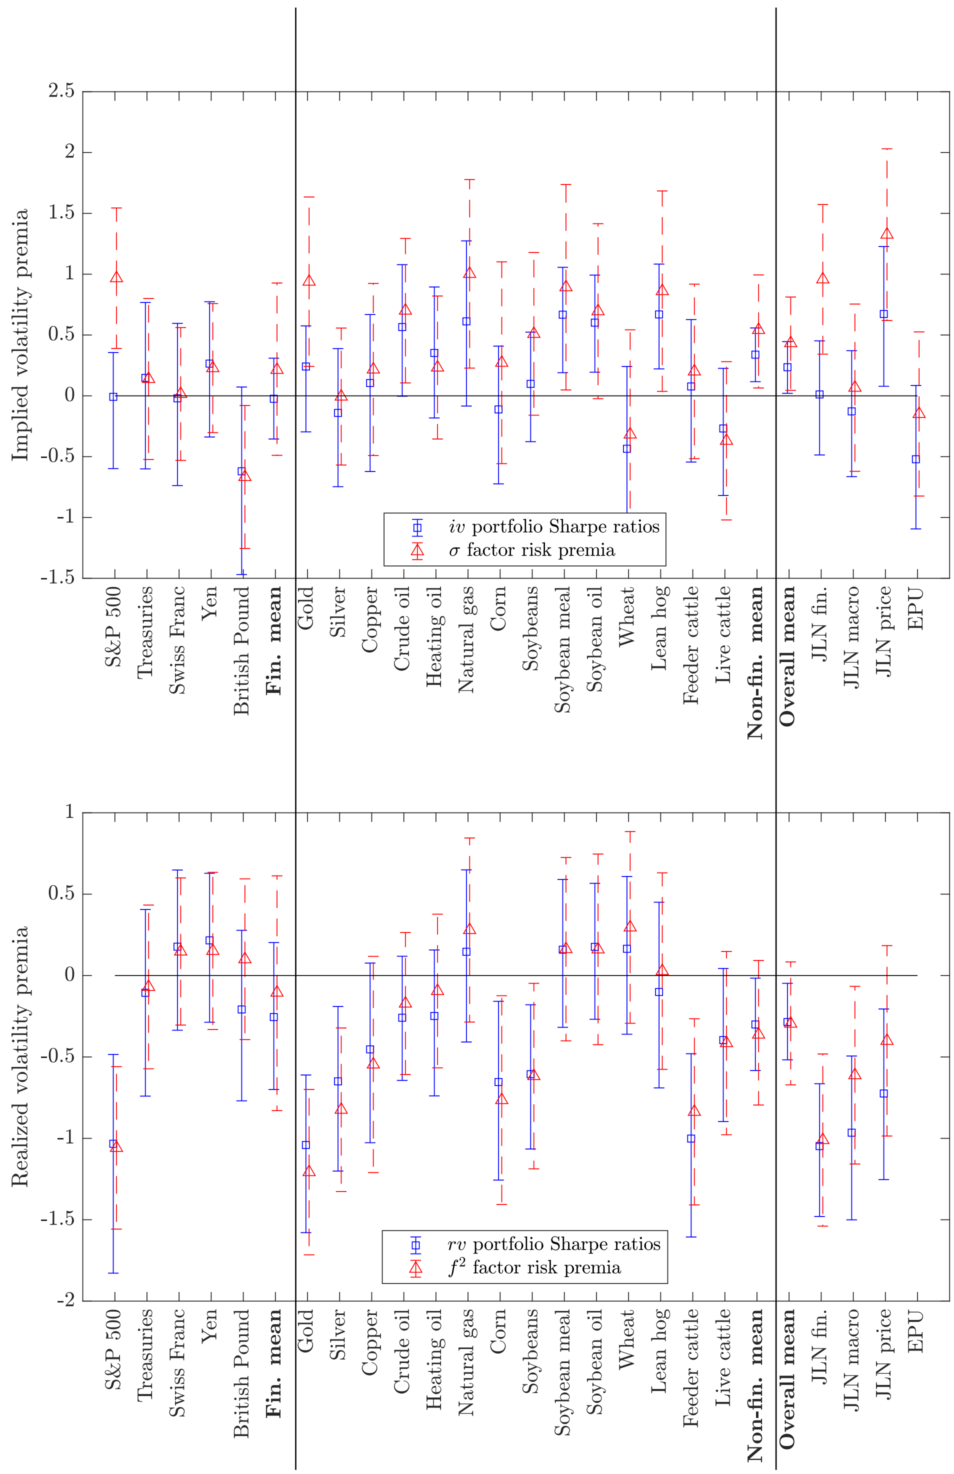

Note: Same as Figure 3, but only using the first half of the sample (up to June 2000).

Figure OA.8: RV and IV portfolio Sharpe ratios and factor risk premia (second half of the sample)

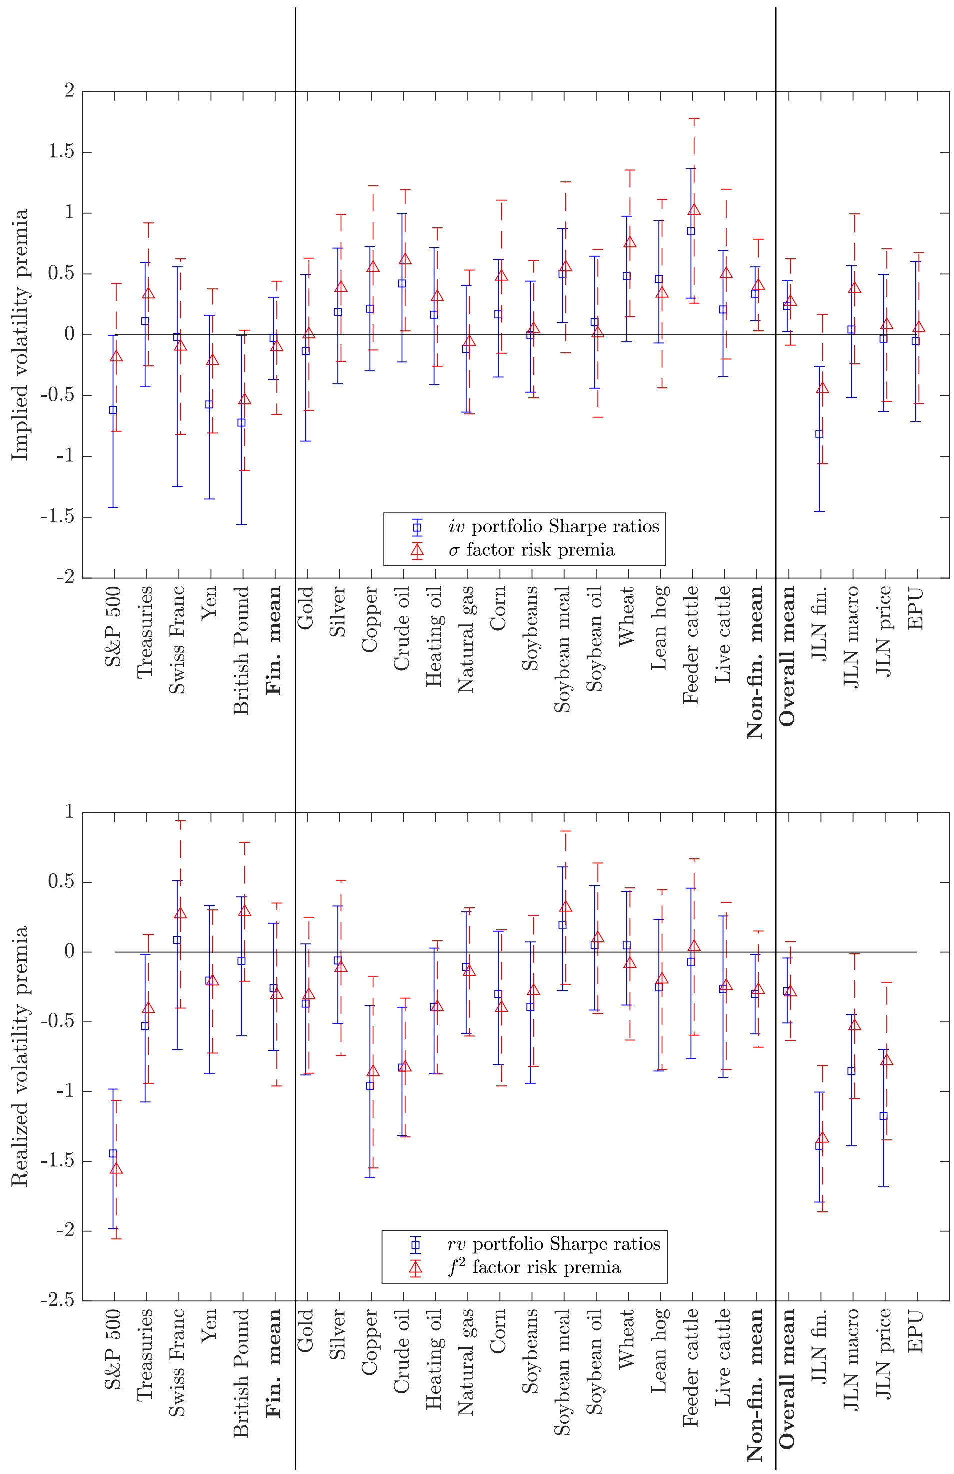

Note: Same as Figure 3, but only using the second half of the sample (after June 2000).

Figure OA.9: Rolling Sharpe ratios of RV and IV portfolios

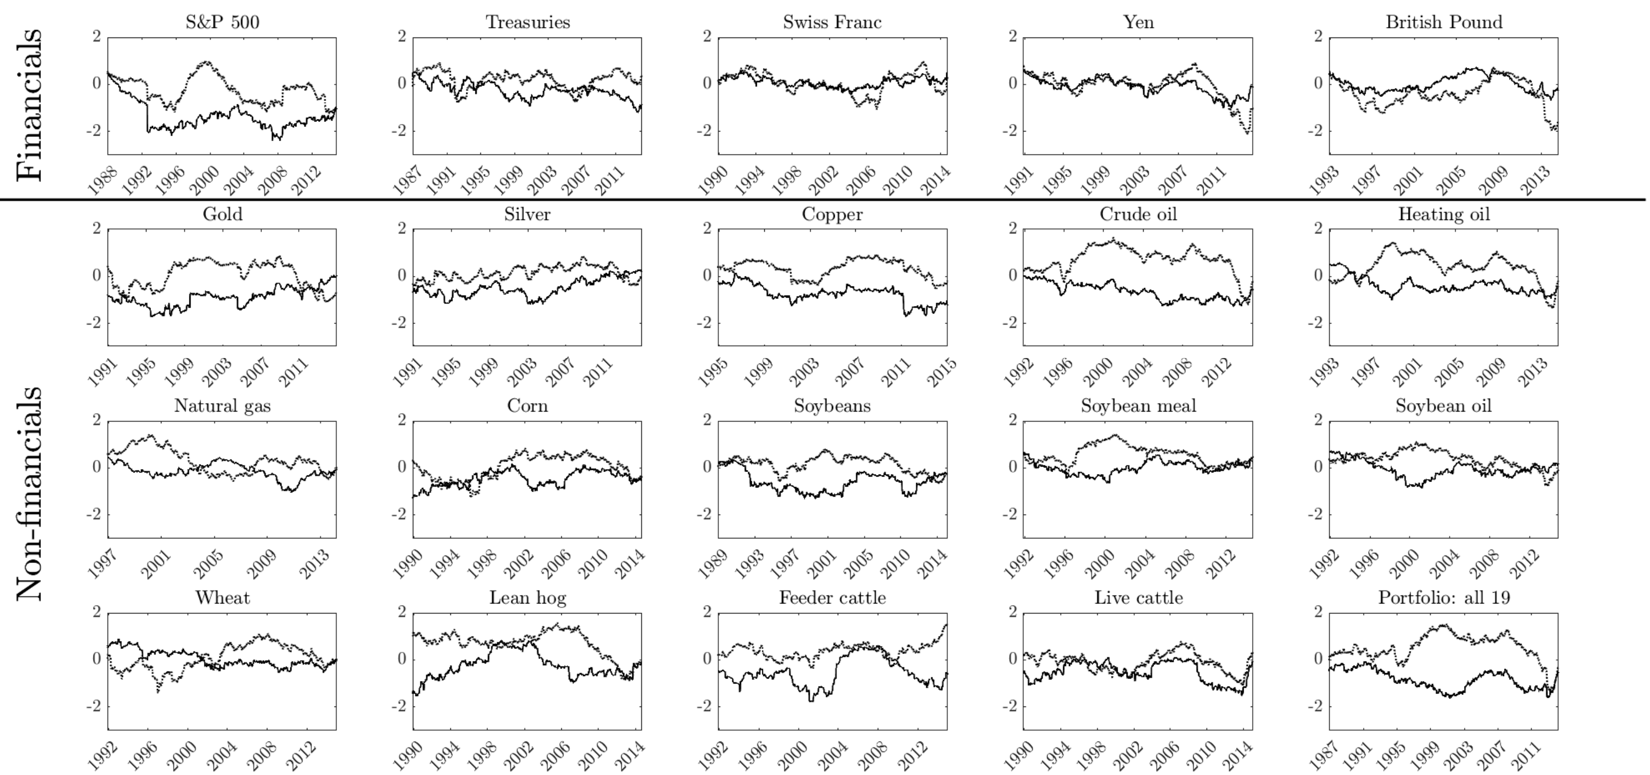

Note: 5-year rolling Sharpe ratios for RV portfolios (solid line) and IV portfolios (dotted lines). The bottom-right panel reports the rolling Sharpe ratio for RV and IV portfolio of all available markets.

Figure OA.10: RV and IV portfolio Sharpe ratios and factor risk premia (using 2-month IV)

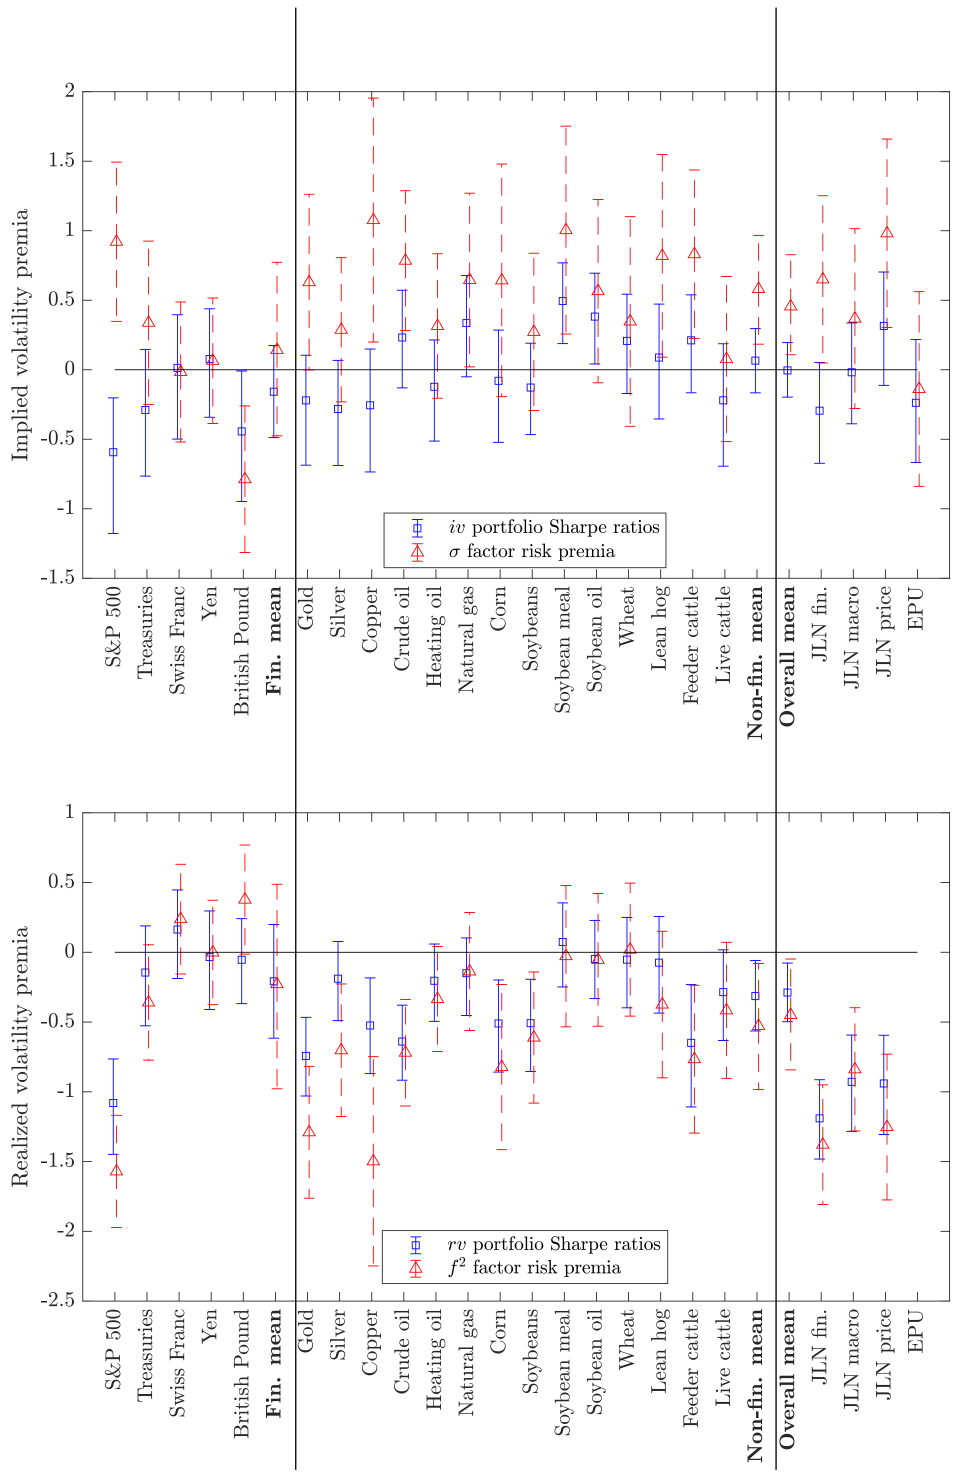

Note: Same as Figure 3, but using 2-month instead of 5-month IV.

Figure OA.11: RV and IV risk premia estimates with and without weighting

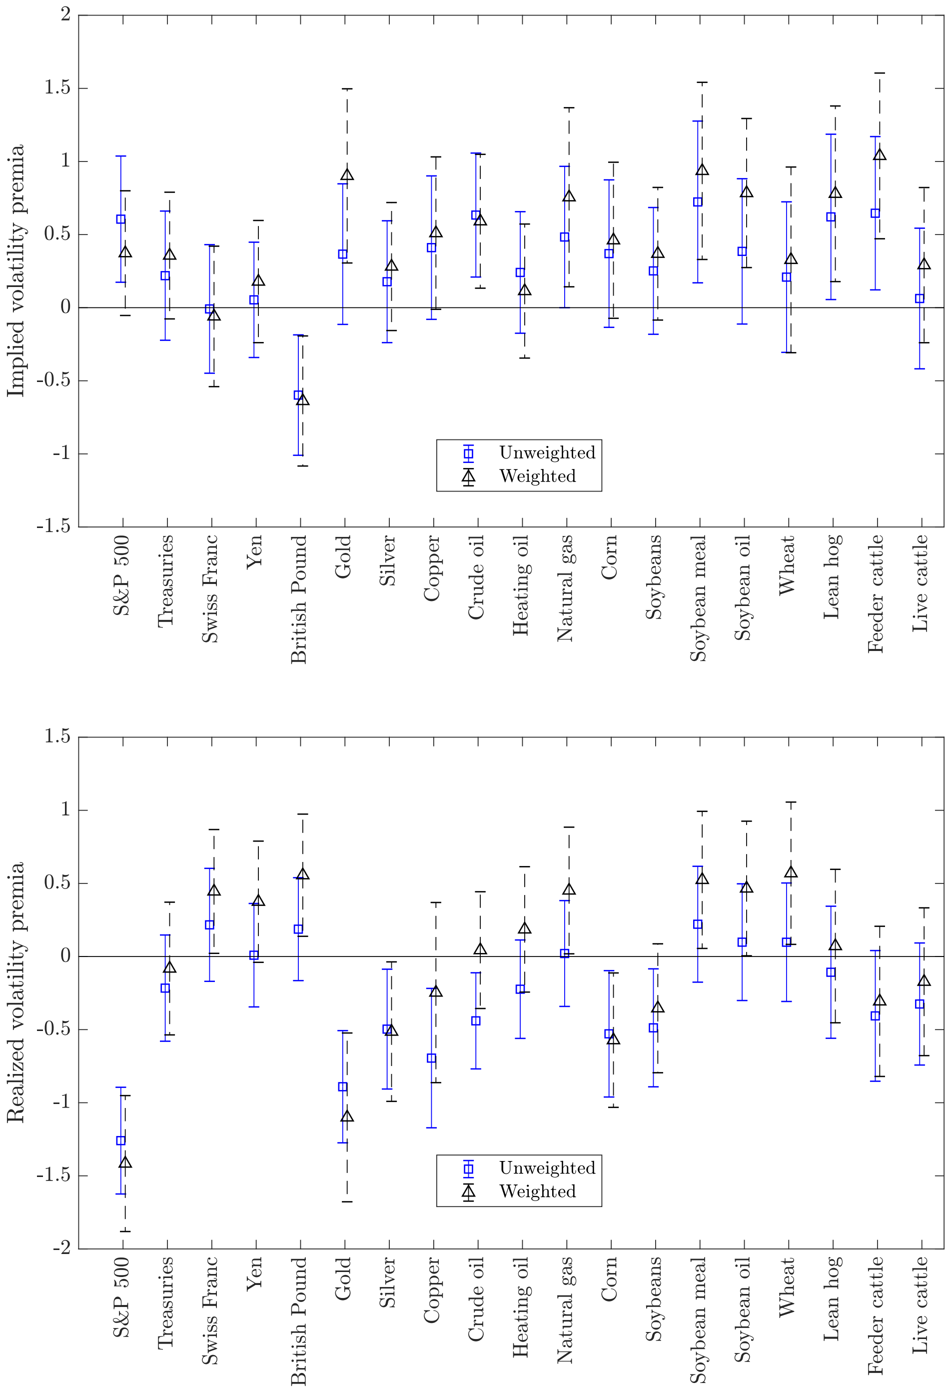

Note: The figure reports risk premia for the factor model, unweighted (as in figure 3) or weighting each observation by the implied volatility.

Figure OA.12: SDF loadings on RV and IV (Sharpe ratios)

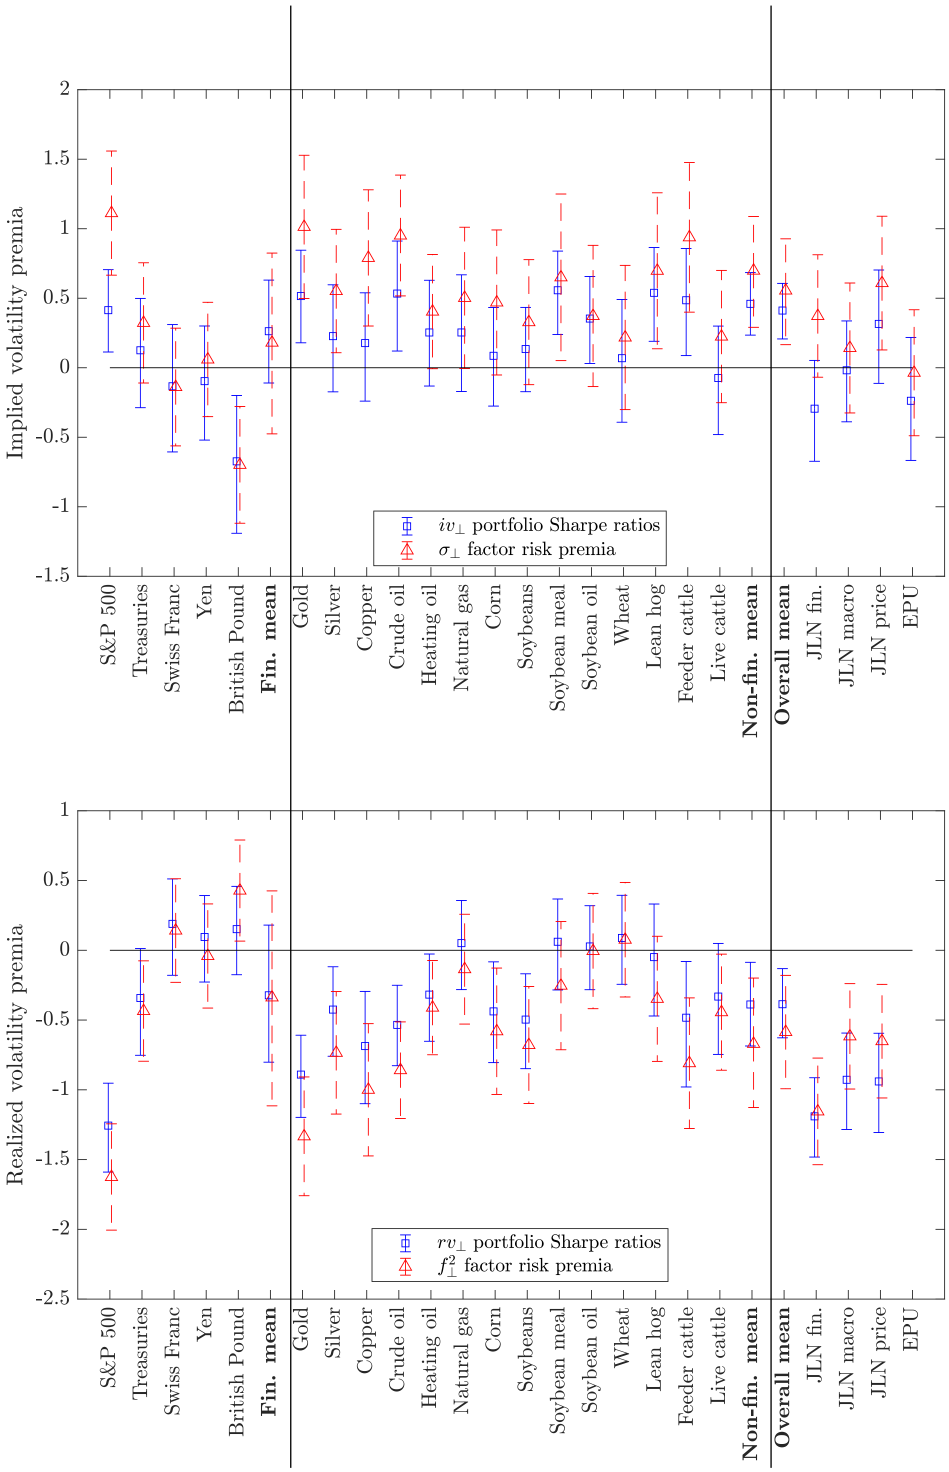

Note: The figure reports the stochastic discount factor (SDF) loadings on IV and RV. The loadings are scaled to correspond to Sharpe ratios of orthogonalized RV and IV portfolios, whose risk premia is equal to the corresponding SDF loading.

Figure OA.13: Bid-ask spreads on 8/4/2017

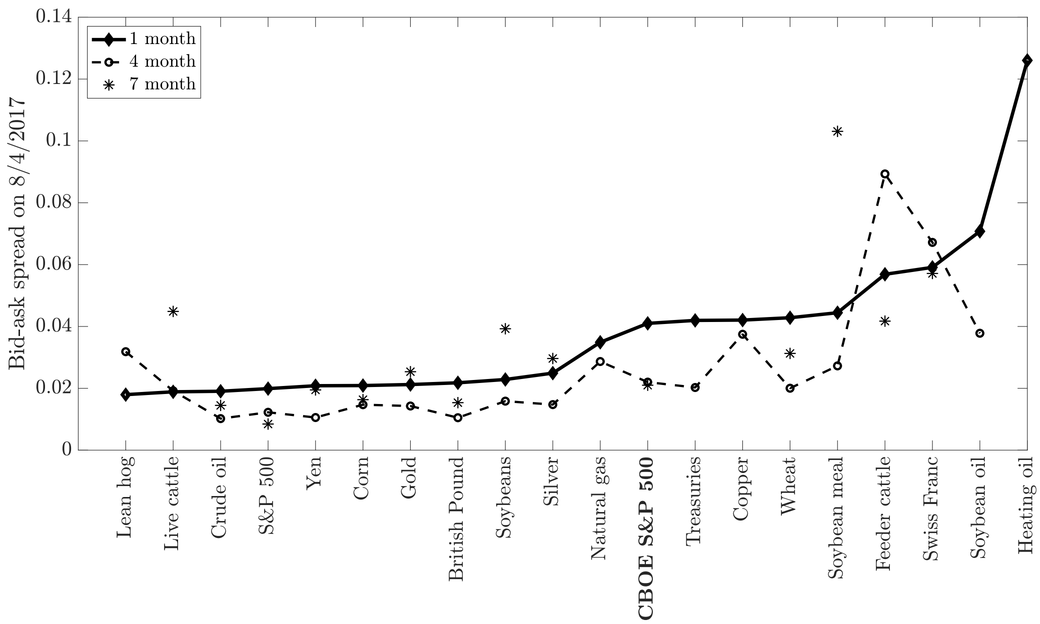

Note: The figure reports posted bid-ask spreads for at-the-money straddles obtained from Bloomberg on of August 4, 2017 (the CBOE S&P 500 spreads on that date are also obtained from Optionmetrics).

Figure OA.14: Volume across markets and maturities

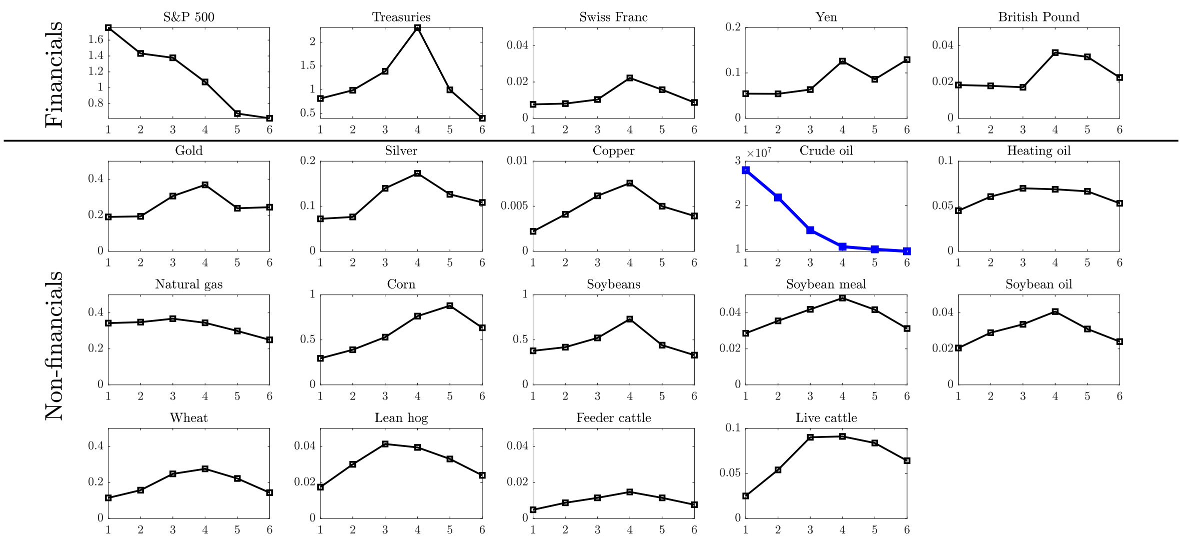

Note: Average daily volume of options in different markets. The panel corresponding to crude oil reports values in dollars. All other panels show values relative to the volume in the crude oil market, matched by maturity.

Figure OA.15: RV and IV portfolio Sharpe ratios and factor risk premia (robust to measurement error)

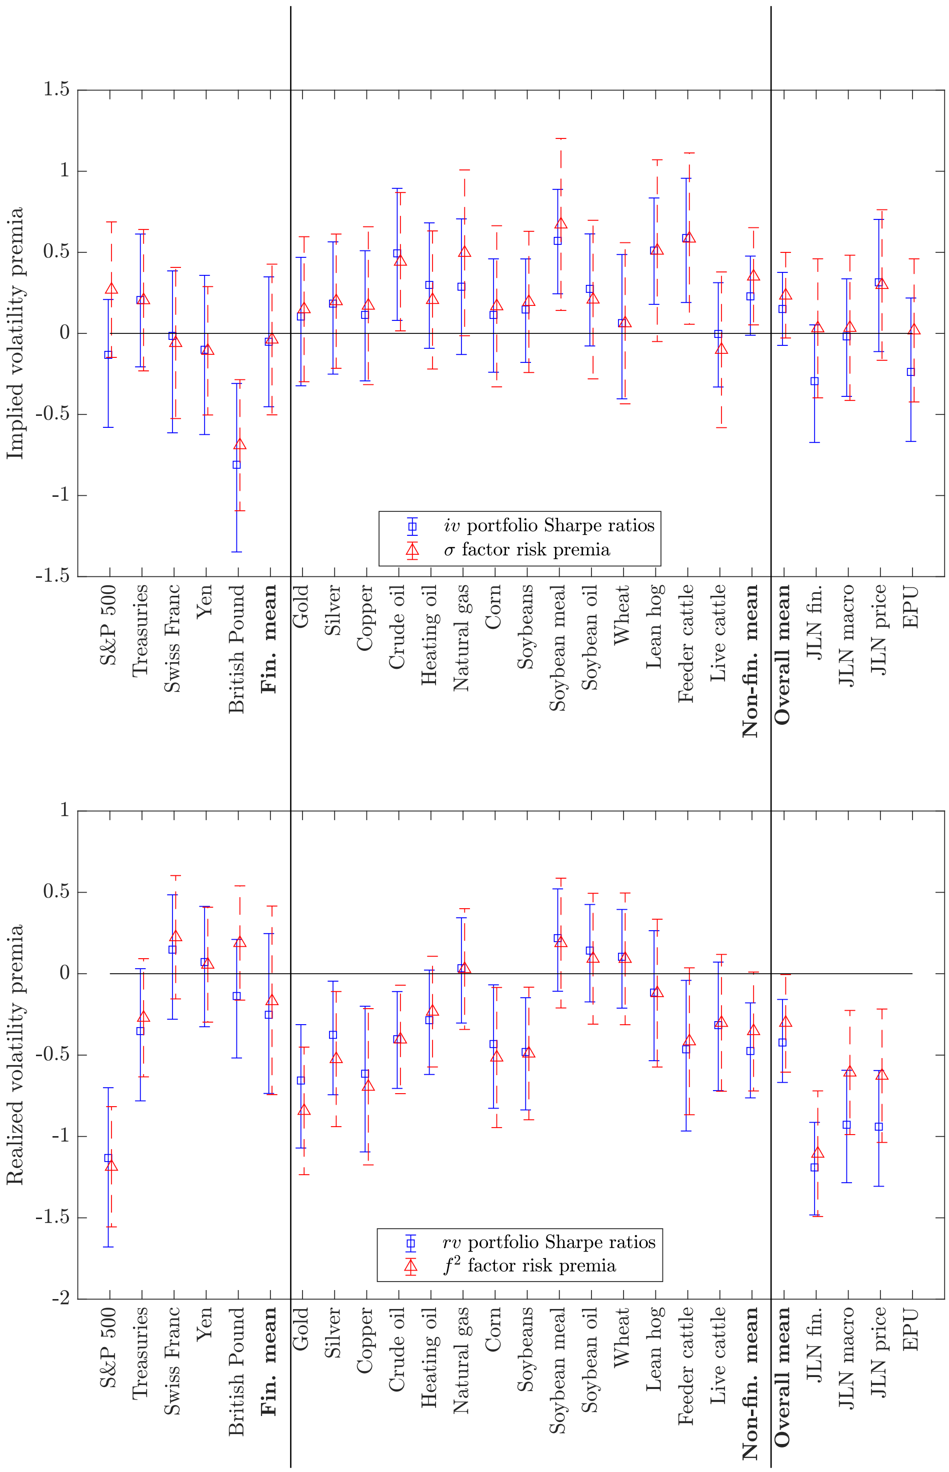

Note: Same as Figure 3, but returns are computed using the same denominator at all maturities, to provide robustness with respect to measurement error in the prices (see section OA.3.6).

Figure OA.16: Options on crude futures vs ETFs

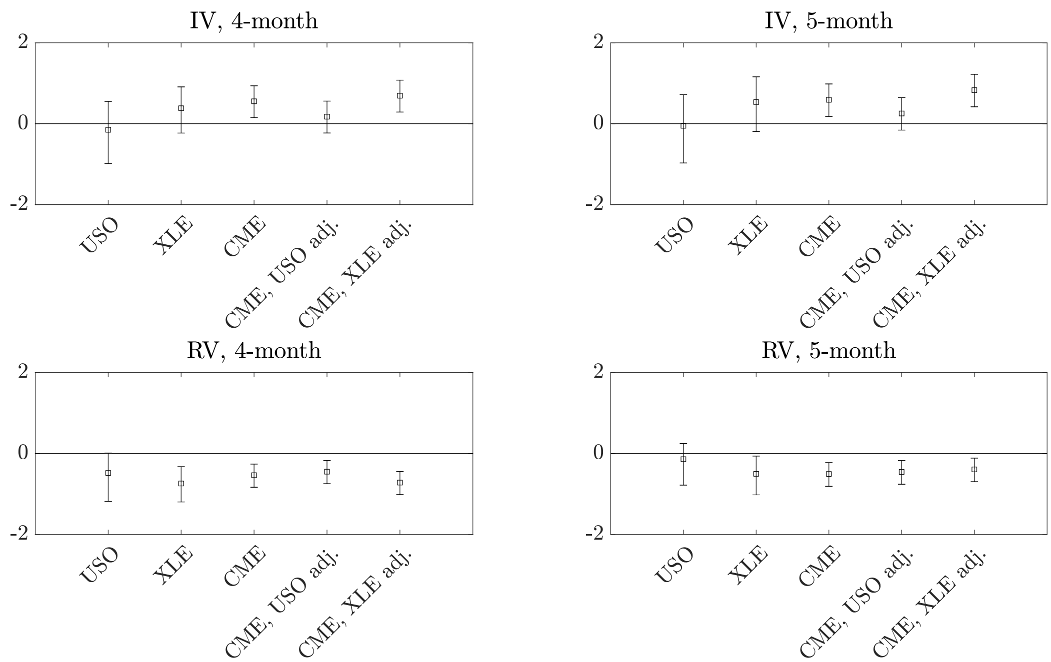

Note: Sharpe ratios on *rv* and *iv* portfolios using straddles for CME crude oil futures and the XLE and USO exchange traded funds. "4-month" and "5-month" refers to the longer of the two maturities used to construct each portfolio (the short maturity is always one month). The squares are point estimates based on the full sample available for each series. The lines are 95-percent confidence bands constructed with a 50-day block bootstrap. "CME, USO adj." and "CME, XLE adj." are identical to the "CME" numbers but with the mean return in the denominator of the Sharpe ratio shifted by the point estimate for the mean difference from table A.6.2.

Table OA.1: χ^2 test of the factor model

	p-value
S&P 500	0.22
T-bonds	0.02
GBP	0.01
CHF	0.38
JPY	0.75
Copper	0.75
Corn	0.00
Crude oil	0.08
Feeder cattle	0.25
Gold	0.44
Heating oil	0.14
Lean hog	0.19
Live cattle	0.80
Natural gas	0.30
Silver	0.68
Soybeans	0.21
Soybean meal	0.41
Soybean oil	0.11
Wheat	0.29

Note: For each market, the table reports bootstrapped p-values for the χ^2 of on the squared fitting errors of the factor model (bootstrapped following Constantinides, Jackwerth, and Savov (2013).

Table OA.2: Risk exposures of rv and iv portfolios

rv portfolio	f	f^2	ΔIV	R^2	**iv portfolio**	f	f^2	ΔIV	R^2	Corr(rv,iv)
S&P 500	-0.07	1.44	0.02	0.68	S&P 500	-0.16	1.37	0.96	0.75	0.48
T-bonds	-0.01	0.81	-0.06	0.75	T-bonds	-0.01	0.35	1.05	0.78	0.13
GBP	-0.03	0.81	0.00	0.82	GBP	-0.02	0.44	0.91	0.86	0.47
CHF	0.00	0.75	0.03	0.73	CHF	0.05	0.52	0.91	0.72	0.64
JPY	-0.02	0.74	0.04	0.80	JPY	0.02	0.57	0.89	0.87	0.63
Copper	-0.01	0.79	-0.06	0.62	Copper	0.01	0.23	1.00	0.85	0.07
Corn	-0.02	0.65	-0.01	0.69	Corn	0.06	0.41	0.85	0.75	0.08
Crude oil	-0.03	1.00	-0.02	0.75	Crude oil	0.03	-0.07	0.93	0.77	0.06
Feeder cattle	-0.03	0.98	-0.01	0.66	Feeder cattle	-0.02	-0.25	0.96	0.78	0.02
Gold	0.00	0.70	0.01	0.68	Gold	0.08	0.35	0.97	0.68	0.48
Heating oil	-0.02	0.88	-0.04	0.76	Heating oil	0.04	-0.17	1.00	0.77	-0.02
Lean hog	-0.02	0.90	-0.06	0.75	Lean hog	0.04	-0.49	1.03	0.64	-0.24
Live cattle	-0.03	1.03	-0.03	0.72	Live cattle	0.00	-0.44	0.92	0.78	-0.12
Natural gas	-0.03	0.87	-0.02	0.80	Natural gas	0.03	-0.38	0.98	0.64	-0.17
Silver	-0.01	0.63	0.03	0.71	Silver	0.04	0.20	0.92	0.85	0.45
Soybeans	-0.02	0.66	-0.01	0.71	Soybeans	0.04	0.30	0.89	0.80	0.18
Soybean meal	-0.01	0.61	-0.02	0.74	Soybean meal	0.05	0.31	0.93	0.69	0.19
Soybean oil	-0.01	0.64	-0.02	0.73	Soybean oil	0.05	0.29	0.94	0.77	0.20
Wheat	-0.01	0.63	-0.05	0.78	Wheat	0.05	0.30	0.97	0.78	0.16
Average	-0.02	0.82	-0.01	0.73	Average	0.02	0.20	0.95	0.76	

Note: The table reports regression coefficients of the rv and iv portfolios for each market onto three market-specific factors: the futures return, the squared futures return, and the change in IV. The column on the right reports the correlation between the rv and iv portfolio returns.

Table OA.3: Risk exposures of rv portfolio to IV innovations at different maturities

	Maturity of IV shock				
rv portfolio	1	2	3	4	5
S&P 500	0.08	0.08	0.07	0.05	0.02
T-bonds	0.07	0.06	0.03	0.00	-0.06
GBP	0.07	0.07	0.06	0.04	0.00
CHF	0.07	0.07	0.07	0.06	0.03
JPY	0.07	0.07	0.07	0.06	0.04
Copper	0.08	0.08	0.05	0.00	-0.06
Corn	0.08	0.08	0.07	0.05	-0.01
Crude oil	0.06	0.06	0.04	0.01	-0.02
Feeder cattle	0.09	0.08	0.06	0.03	-0.01
Gold	0.07	0.07	0.07	0.05	0.01
Heating oil	0.07	0.07	0.05	0.02	-0.04
Lean hog	0.09	0.08	0.06	0.01	-0.06
Live cattle	0.08	0.08	0.06	0.02	-0.03
Natural gas	0.08	0.08	0.07	0.03	-0.02
Silver	0.09	0.09	0.09	0.07	0.03
Soybeans	0.07	0.07	0.06	0.03	-0.01
Soybean meal	0.07	0.07	0.05	0.03	-0.02
Soybean oil	0.07	0.07	0.05	0.02	-0.02
Wheat	0.05	0.05	0.03	-0.01	-0.05

	Maturity of IV shock				
RV-hedging	1	2	3	4	5
S&P 500	0.05	0.04	0.04	0.02	0.00
T-bonds	0.11	0.11	0.10	0.07	0.00
GBP	0.08	0.08	0.07	0.05	0.00
CHF	0.07	0.07	0.06	0.04	0.00
JPY	0.07	0.07	0.06	0.04	0.00
Copper	0.13	0.13	0.12	0.07	0.00
Corn	0.12	0.12	0.12	0.08	0.00
Crude oil	0.07	0.07	0.06	0.03	0.00
Feeder cattle	0.09	0.09	0.07	0.04	0.00
Gold	0.10	0.10	0.08	0.05	0.00
Heating oil	0.09	0.09	0.08	0.06	0.00
Lean hog	0.11	0.11	0.10	0.06	0.00
Live cattle	0.09	0.09	0.07	0.04	0.00
Natural gas	0.10	0.10	0.09	0.06	0.00
Silver	0.12	0.12	0.11	0.07	0.00
Soybeans	0.11	0.11	0.09	0.06	0.00
Soybean meal	0.12	0.12	0.10	0.07	0.00
Soybean oil	0.12	0.12	0.11	0.07	0.00
Wheat	0.11	0.11	0.10	0.06	0.00

Note: The table reports the loading of the rv portfolio (top panel) and of the RV-hedging portfolio built using the factor model (bottom panel) on shocks to IV of different maturity, from 1 to 5 months.

Table OA.4: Risk exposures of rv and iv portfolios, 2-month IV

rv portfolio	f	f^2	ΔIV	R^2	**iv portfolio**	f	f^2	ΔIV	R^2	Corr(rv,iv)
S&P 500	-0.04	0.74	0.04	0.38	S&P 500	-0.33	4.79	0.78	0.66	0.27
T-bonds	0.00	0.37	0.00	0.39	T-bonds	-0.08	2.44	0.84	0.72	0.14
GBP	-0.02	0.45	0.03	0.50	GBP	-0.05	2.04	0.70	0.73	0.27
CHF	0.00	0.40	0.03	0.44	CHF	0.07	2.16	0.74	0.69	0.40
JPY	-0.02	0.42	0.04	0.54	JPY	-0.01	2.04	0.72	0.82	0.46
Copper	-0.01	0.33	0.02	0.25	Copper	0.00	2.18	0.78	0.68	-0.04
Corn	-0.02	0.27	0.03	0.32	Corn	0.10	2.17	0.64	0.72	0.07
Crude oil	-0.01	0.58	0.01	0.50	Crude oil	-0.06	1.72	0.77	0.71	0.19
Feeder cattle	0.00	0.45	0.04	0.36	Feeder cattle	-0.21	2.07	0.77	0.58	0.05
Gold	-0.01	0.28	0.02	0.35	Gold	0.09	2.21	0.84	0.66	0.23
Heating oil	-0.02	0.54	0.01	0.49	Heating oil	0.04	1.31	0.80	0.62	0.09
Lean hog	-0.01	0.45	0.03	0.45	Lean hog	-0.01	1.59	0.74	0.59	0.08
Live cattle	-0.02	0.53	0.02	0.47	Live cattle	-0.06	1.75	0.75	0.68	0.18
Natural gas	-0.03	0.50	0.02	0.55	Natural gas	0.03	1.28	0.77	0.67	0.18
Silver	0.00	0.28	0.04	0.41	Silver	-0.01	1.70	0.79	0.76	0.30
Soybeans	-0.02	0.38	0.03	0.50	Soybeans	0.05	1.52	0.69	0.77	0.29
Soybean meal	-0.01	0.31	0.03	0.47	Soybean meal	0.07	1.60	0.66	0.75	0.26
Soybean oil	-0.01	0.33	0.02	0.43	Soybean oil	0.07	1.63	0.73	0.72	0.20
Wheat	-0.01	0.26	0.00	0.33	Wheat	0.07	2.16	0.70	0.79	0.14
Average	-0.01	0.41	0.02	0.43	Average	-0.01	2.02	0.75	0.70	

Note: Same as table OA.2, but 2-month IV is used as one of the factors (as opposed to 5-month IV) and in the construction of the rv and iv portfolios.