A Bayesian Perspective on Commodity Style-Integration

Ana-Maria Fuertes, Ph.D.
Bayes Business School, City, University of London, U.K.

Nan Zhao, Ph.D.
Bayes Business School, City, University of London, U.K.

Available at SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4112383

Commodity style-integration is appealing because by forming a unique long-short portfolio with simultaneous exposure to mildly correlated factors, a larger risk premium can be captured over time than with any of the underlying standalone styles. A practical decision that a commodity style-integration investor faces at each rebalancing time is the relative weight of the predictive- or sorting-signal that underlies each standalone style. The authors of this paper develop a new Bayesian optimized integration (BOI) method that accounts for estimation risk in the style-weighting decision. Focusing on the problem of a commodity investor that seeks exposure to the carry, hedging pressure, momentum, skewness, and basis-momentum factors, they demonstrate that the BOI portfolio outperforms not only a battery of parametric style-integrations motivated by the portfolio optimization literature, but also the highly effective equal-weight integrated portfolio. The findings survive the consideration of transaction costs, alternative commodity scoring schemes, and long estimation windows.

Introduction

Individual factors can undergo time-variation or be arbitraged away; namely, styles that have captured a sizeable premium over a period of time may weaken or completely fade away due to "factor crowding" (see *e.g.*, Bhattacharya *et al*., 2017). One way to mitigate this problem is by constructing a long-short portfolio or style-integrated portfolio according to a combination of predictive signals which is also known as the multi-factor approach. Style-integration is in essence the old adage of *don't put all your eggs in the same basket* applied to factor exposure or style investing. The key idea is to harness the diversification of predictive signals towards capturing a larger and more resilient risk premia over time. A key decision that a style-integration investor faces is the relative weight to give to the styles at each portfolio rebalancing time. With a history of returns on each of the styles, the investor can estimate the style-weights that are defined as "optimal" according to some criteria. However, these optimized style-integrations (OIs) suffer from parameter uncertainty, which is the main reason why the naive equal-weight style integration (EWI) has stood out as very effective. In a structured contest among the EWI method and a battery of sophisticated style-integrations, Fernandez-Perez *et al*. (2019) show that the former is not outperformed by the latter. The authors of this paper thus believe it is worthwhile to pursue the research question of whether embedding the style-integration problem within a Bayesian framework that accounts for estimation risk can be fruitful for investors.

This digest article was contributed by Ana-Maria Fuertes, Ph.D., Professor in Finance and Econometrics at Bayes Business School, City, University of London (U.K.) and Associate Editor of the GCARD.

The authors develop a Bayesian optimized style-integration (BOI) method that expands the parametric mean-variance optimized integration by allowing investors to incorporate their prior beliefs or knowledge about the merit of the different standalone styles. The priors on the style-weights distribution can then be conveniently mapped into priors on the distribution of excess returns for the candidate commodity futures contracts. In an empirical exercise, the authors compare the reward-to-risk and crash risk profiles of the BOI method with those of the challenging EWI benchmark and of several sophisticated parametric optimized integrations (OI).

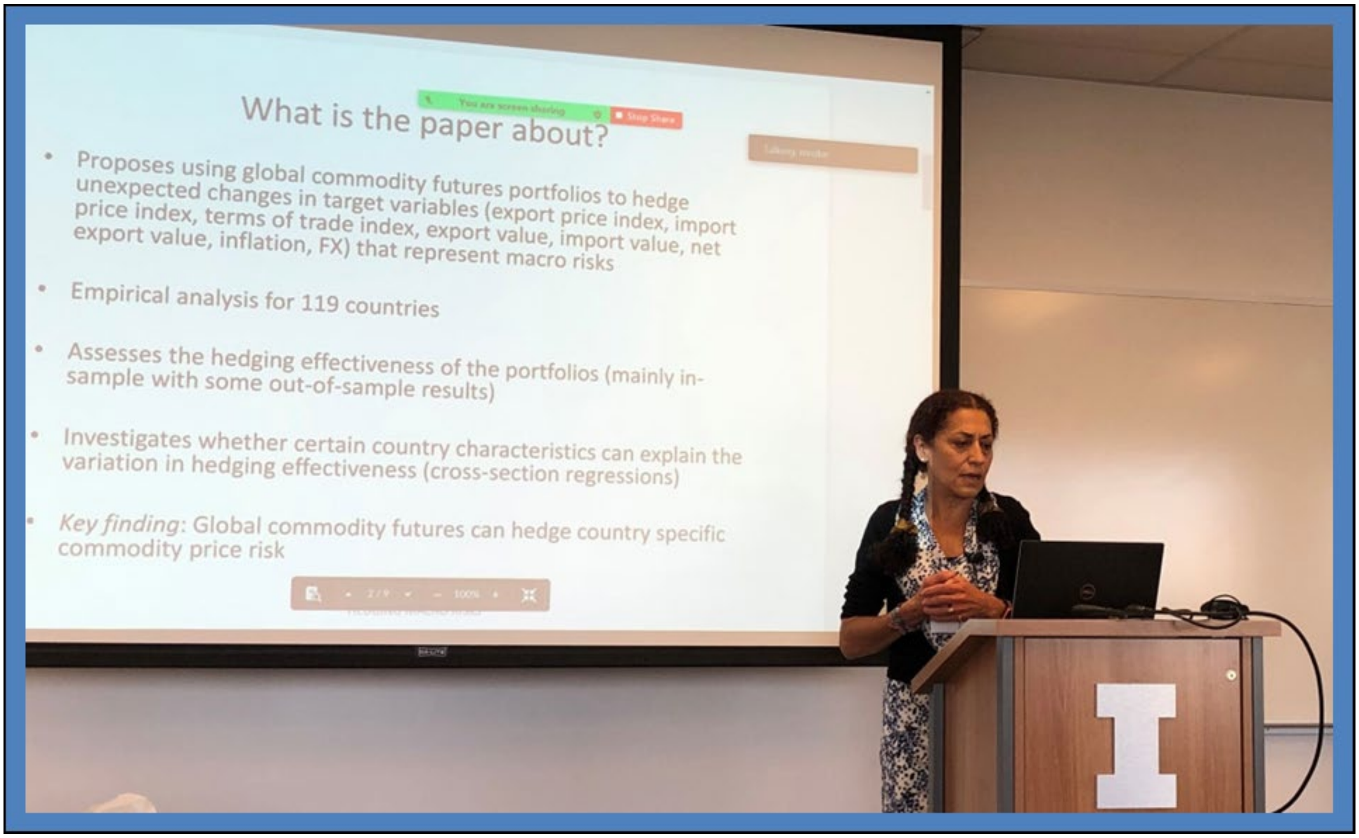

Professor Ana-Maria Fuertes of Bayes Business School, City, University of London, U.K., lecturing during the Commodities & Energy Markets Association (CEMA) conference at the University of Illinois' Illini Center in Chicago. This conference took place on June 21st and 22nd, 2022.

Why the Paper's Research Question is Important

Research over the last few years has established that a number of factors can explain return performance in commodity futures but the corresponding style premia are not constant over time. Rewarding factors over specific periods can temporarily weaken. Improving the return profile through mixing styles is, in fact, currently the critical issue for many commodity investors. This paper seeks to assist investors by developing a BOI strategy that seeks efficiently (that is, with a low noise-to-signal ratio) to construct a unique long-short portfolio with exposure to multiple commodity risk. The BOI approach is flexible enough to facilitate integration of any number of styles using an investor-chosen criteria for the optimal estimation of the style-exposures. The research question is also relevant for academics because it allows the authors to advance the Bayesian statistics literature towards commodity style-integration.

Style-Integration Methodology

The investor's decision at portfolio formation time t about the relative wealth to allocate to each commodity futures and the nature of the position, long versus short, can be represented by the $N \times 1$ commodity allocation vector $\boldsymbol{\phi}_t$ obtained as

$$\boldsymbol{\phi}_t \equiv \boldsymbol{\Theta}_t \times \boldsymbol{\omega}_t = \begin{pmatrix} \theta_{1,1,t} & \cdots & \theta_{1,K,t} \\ \vdots & \ddots & \vdots \\ \theta_{N,1,t} & \cdots & \theta_{N,K,t} \end{pmatrix} \begin{pmatrix} \omega_{1,t} \\ \vdots \\ \omega_{K,t} \end{pmatrix} = \begin{pmatrix} \phi_{1,t} \\ \vdots \\ \phi_{N,t} \end{pmatrix} \tag{1}$$

where $\boldsymbol{\Theta}_t$ is the $N \times K$ *score* matrix (*N* is the number of assets and *K* the number of standalone styles) and $\boldsymbol{\omega}_t$ is the $K \times 1$ signal- (or style-) weighting vector. The sign of the *i*th commodity allocation weight $\phi_{i,t}$ dictates the type of position (long or short). The element $\theta_{i,k,t}$ is the score assigned to the *i*th commodity futures contract according to the *k*th sorting signal (or style) at portfolio rebalancing time *t*. Alternative scoring schemes are plausible such as defining $\theta_{i,k,t}$ as the signals (appropriately standardized) or standardized rankings or binary long-versus-short signals {+1, -1}.

A key element in the integration is the style-weights vector $\boldsymbol{\omega_t} = (\omega_{1,t}, \ldots, \omega_{K,t})$ where the weight $\omega_{k,t}$ reflects the relative importance given to the *k*th individual investment style (or factor) in the integrated portfolio. The naïve EWI strategy assigns equal importance to the *K* styles, *i.e.,* $\boldsymbol{\omega_t} = \left(\frac{1}{K}, \ldots, \frac{1}{K}\right)'$, at each rebalancing time and thus it is parameter-free. Besides the EWI, various OIs have been deployed in the literature.

In an OI strategy the style-weight decision hinges on solving an optimization problem; namely, at each portfolio rebalancing time *t* the investor ought to find the weights that minimize or maximize a property of the style-integrated portfolio return distribution. For instance, quadratic utility or mean-variance maximization (MV), MV maximization with shrinkage (MVshrinkage), variance minimization (MinVar), diversification-ratio maximization (MaxDiv), power utility maximization (PowerU), PowerU with disappointment aversion (PowerDU) or on style-volatility timing (StyleVol); see, *e.g.,* Ledoit and Wolf (2003), Choueifaty and Coignard (2008), Brandt *et al*. (2009), Kirby and Ostdiek (2012) and Fernandez-Perez *et al*. (2019). A common denominator to these OIs is that albeit they can potentially discriminate better among the *K* styles because they allow for time-varying, heterogeneous exposures to the different styles, such an advantage can be largely contaminated by parameter estimation uncertainty.

The key idea behind the BOI method proposed by the authors is to mitigate uncertainty about the parameters describing the distribution of commodity returns by forming priors that are subsequently updated. Investors do not need to directly form a prior on $\boldsymbol{\mu}_t$, the $N \times 1$ commodity mean excess returns. They can instead harness their beliefs (or information) on the past relative performance of the styles to form a prior on $\boldsymbol{\omega}_t$ which can be mapped onto a prior for $\boldsymbol{\mu_t}$. Given the success of the equal-weight rule in portfolio allocation (DeMiguel *et al*., 2009) and in style-integration (Fernandez-Perez *et al*., 2019), the authors adopt $1/K$ as the informative prior for the mean of the distribution of $\boldsymbol{\omega}_t$ which is assumed Gaussian. A history of commodity excess returns over a window of L months is used to update the priors in order to obtain the posterior density of $\boldsymbol{\mu_t}$ using the Gibbs sampling approach that belongs to the family of Markov Chain Monte Carlo (MCMC) methods. With the posterior density of $\boldsymbol{\mu_t}$ at hand, the MV optimization problem is solved at each portfolio rebalancing time *t* to obtain the BOI style-weights $\boldsymbol{\omega}_t$.

Results

The authors carry out an empirical analysis of style-integration methods in the context of data for a cross-section of 28 commodity futures contracts from January 1992 to December 2021. Without loss of generality, the focus is on five fairly well-known commodity investment styles that exploit as predictive- or sorting-signals, respectively, the basis, hedgers' net short positions, momentum, skewness, and basis-momentum.

The naïve EWI strategy outperforms each of the standalone styles in terms of risk-reward (Sharpe ratio, Omega ratio, and Sortino ratio) and crash risk (semi-deviation, 99% Value-at-Risk, and maximum drawdown). This finding confirms the diversification benefits of style-integration. Another important confirmation result is that the naïve EWI portfolio is not challenged by any of the sophisticated OI portfolios.

The key novel evidence in this paper is that the BOI approach is able to significantly improve upon the challenging EWI benchmark. With a Sharpe ratio of 1.060, maximum drawdown of -0.174, and 99% of VaR of -0.051, the BOI portfolio is a more attractive proposition than any of the alternative OI portfolios, and also the challenging EWI portfolio as regards both reward-to-risk and crash risk profiles; see Table 1 on the next page.

Table 1
Performance of Commodity Style-Integrated Portfolios

		Optimized Style-Integrations (OI)							
	EWI	MV	MVshrinkage	MinVar	StyleVol	MaxDiv	PowerU	PowerDA	BOI
Panel A: *Static portfolio evaluation*									
Mean	0.080	0.054	0.051	0.075	0.082	0.083	0.052	0.052	0.092
StDev	0.101	0.094	0.094	0.084	0.102	0.096	0.093	0.094	0.087
Semi-deviation	0.272	0.258	0.258	0.209	0.275	0.248	0.258	0.262	0.212
Max Drawdown	-0.243	-0.297	-0.287	-0.158	-0.255	-0.219	-0.296	-0.296	-0.174
99% VaR	-0.061	-0.058	-0.058	-0.050	-0.062	-0.057	-0.058	-0.059	-0.051
Sharpe Ratio (SR)	0.815	0.606	0.577	0.904	0.823	0.886	0.588	0.587	1.060
Sortino ratio	1.393	1.012	0.960	1.677	1.400	1.566	0.976	0.970	1.987
Omega ratio	1.900	1.599	1.563	2.041	1.918	2.023	1.576	1.574	2.309
ΔSR (gain versus EWI)		-0.209	-0.239	0.089	0.008	0.071	-0.227	-0.229	0.245
Ledoit-Wolf test p-value		0.883	0.931	0.222	0.383	0.128	0.902	0.901	0.005
Opdyke test p-value		0.199	0.128	0.888	0.915	0.774	0.173	0.101	0.042
Panel B: *Dynamic Sharpe ratio (style ranking)*									
Jan 1992 - Dec 1997	1.108(7)	1.296(3)	1.107(8)	1.293(4)	1.103(9)	1.265(6)	1.278(5)	1.300(2)	1.373(1)
Jan 1998 - Dec 2003	0.999(4)	1.000(3)	0.860(8)	0.671(9)	1.002(2)	0.902(7)	0.923(6)	0.931(5)	1.005(1)
Jan 2004 - Dec 2009	1.115(2)	0.378(9)	0.464(6)	1.058(5)	1.113(3)	1.076(4)	0.411(7)	0.398(8)	1.314(1)
Jan 2010 - Dec 2015	0.979(4)	0.513(9)	0.604(6)	1.042(3)	0.977(5)	1.055(2)	0.547(8)	0.558(7)	1.180(1)
Jan 2016 - Dec 2021	0.193(5)	0.116(6)	0.089(7)	0.496(2)	0.194(4)	0.381(3)	0.081(8)	0.072(9)	0.583(1)

Notes: The table reports summary statistics for the excess returns of the equal-weight style integrated (EWI) portfolio and optimized style-integrated (OI) portfolios with the style-weight vector estimated at each portfolio rebalancing time by quadratic utility maximization (mean variance; MV), mean-variance with shrinkage maximization (MVshrinkage), variance minimization (MinVar), style-volatility timing (StyleVol), diversification-ratio maximization (MaxDiv), power utility maximization (PowerU), maximization of power utility with disappointment aversion (PowerDA), and Bayesian optimized integration (BOI). The length of the rolling estimation window is 60 months. The style-integrations are based on standardized signals as commodity scores. The reported mean and standard deviation are annualized. The hypotheses of the Ledoit and Wolf (2008) and Opdyke (2007) tests are $H_0: SR_i - SR_{EWI} \leq 0$ vs $H_A: SR_i - SR_{EWI} > 0$ where *i* is an OI strategy.

Panel A reports statistics over the full sample period January 1992 to December 2021. Panel B reports Sharpe ratios over 6-year non-overlapping subperiods and corresponding style-integrated portfolio ranking in parentheses.

Adding statistical significance to these results, the Ledoit and Wolf (2008) and Opdyke (2007) tests suggest at the 5% significance level or better that the Sharpe ratio of the BOI portfolio is notably larger than that of the naive EWI portfolio. These key findings are obtained both with fixed-length rolling windows of $L = 60$ months to determine the style-weights, and also with long estimation such fixed $L = 120$ months (rolling) or expanding windows starting from 60 months. Likewise, the superiority of the BOI portfolio survives the consideration of transaction costs and the use of alternative scoring schemes.

Conclusions

A large number of factor models have been suggested to explain returns in commodity markets. Forming a unique long-short portfolio with simultaneous exposure to mildly correlated risk factors is an intuitive "style diversification" idea but it requires a choice of style-weights at each portfolio rebalancing time. To date, the different sophisticated style-integrations attempted have not been as effective as the naïve equal-weights style integration. The reason is that, by contrast with parametric methods, the EWI is not contaminated by estimation risk. This paper develops a novel Bayesian optimized style-integration that alleviates estimation risk. Focusing on well-known commodity styles – basis, hedging pressure, momentum, skewness, and basis momentum – the authors provide evidence to suggest that the BOI portfolio significantly outperforms a battery of sophisticated OIs and the challenging EWI. The main take away of this research is that embedding extant OI methods into a Bayesian framework to account for estimation risk allows investors to harness multiple commodity factor exposures more efficiently towards capturing a larger and more resilient risk premium over time.

References

Bhattacharya, D., Li, W.-H. and G. Sonaer, 2017, "Has Momentum Lost Its Momentum?", *Review of Quantitative Finance and Accounting*, Vol. 48, No. 1, January, pp. 191-218.

Brandt, M., Santa-Clara, P. and R. Valkanov, 2009, "Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns," *Review of Financial Studies*, Vol. 22, No. 9, September, pp. 3411-3447.

Choueifaty, Y. and Y. Coignard, 2008, "Toward Maximum Diversification," *Journal of Portfolio Management*, Vol. 35, No. 1, Fall, pp. 40-51.

DeMiguel, V., Garlappi, L. and R. Uppal, 2009, "Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?", *Review of Financial Studies*, Vol. 22, No. 5, May, pp. 1915-1953.

Fernandez-Perez, A., Fuertes, A.-M. and J. Miffre, 2019, "A Comprehensive Appraisal of Style-Integration Methods," *Journal of Banking and Finance*, Vol. 105, August, pp. 134-150.

Kirby, C. and B. Ostdiek, 2012, "It's All in the Timing: Simple Active Portfolio Strategies that Outperform Naive Diversification," *Journal of Financial and Quantitative Analysis*, Vol. 47, No. 2, April, pp. 437-467.

Ledoit, O. and M. Wolf, 2003, "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection," *Journal of Empirical Finance*, Vol. 10, No. 5, December, pp. 603-621.

Ledoit, O. and M. Wolf, 2008, "Robust Performance Hypothesis Testing with the Sharpe Ratio," *Journal of Empirical Finance*, Vol. 15, No. 5, December, pp. 850-859.

Opdyke, J. D., 2007, "Comparing Sharpe Ratios: So Where are the p-Values?", *Journal of Asset Management,* Vol. 8, No. 5, December, pp. 308-336.

Author Biographies

ANA-MARIA FUERTES, Ph.D.
Bayes Business School, City, University of London, U.K.

Dr. Ana-Maria Fuertes is a Professor in Finance and Econometrics at Bayes Business School, City, University of London. Her research interests are in commodity markets, financial forecasting, empirical asset pricing and credit risk modelling. Dr. Fuertes' articles have been published in the *Review of Finance*, *Journal of Banking and Finance, International Journal of Forecasting* and the *Journal of International Money and Finance*. She has been awarded research grants by the Economic and Social Research Council (ESRC), British Academy/Leverhulme Trust, and INQUIRE U.K. Her research has been featured in practitioner and policy-oriented publications such as *Hedge Funds Review*, *Investment & Pensions Europe*, CEPR *VoX*, CEPS.eu and on the Bank of England's website.

Dr. Fuertes' previous articles for the *GCARD* are available here.

NAN ZHAO, Ph.D.
Bayes Business School, City, University of London, U.K.

Dr. Nan Zhao is a Research Fellow at Bayes Business School, City, University of London. He is also an algorithm trading quant at Barclays corporate and investment bank. His research focuses on empirical asset pricing and financial econometrics.