Forecasting Factor Returns

Chief Investment Officer, Two Sigma Advisers
Geoff Duncombe

Head of Client Solutions Research
Mike Nigro

Head of Thematic Research
Bradley Kay

Executive Summary

In our recent paper, Introducing the Two Sigma Factor Lens, we proposed a parsimonious set of actionable factors that collectively explains the majority of risk in institutional portfolios.[1] This paper takes the next step, proposing a methodology to estimate the long-term return premium associated with each of these factors. This paper introduces a handful of innovations intended to improve the accuracy of our long-term return forecasts:

- We use new asset class return proxies to extend our analysis much further back than the daily return histories of most modern indices.
- We separate the most heterogeneous of the prior paper's factors, Commodities, into six sector-based factors for which the long-term premia are individually estimated.
- We apply (what we believe to be) common sense adjustments to long-term histories — slightly overweighting recent returns and applying empirically-based shrinkage across the observed historical Sharpe ratios to generate our forward-looking estimates of each factor's premium.

Ultimately, this paper identifies five unique and orthogonal factors across asset classes that we believe carry a positive historical return premium: Equity, Interest Rates, Credit, Energy, and Industrial Metals. Our estimates for the long-term Sharpe ratios of these factors (and others in the Two Sigma Factor Lens) may be found in Exhibit 1. We believe these five compensated factors could collectively form the basis of an asset allocation strategy with substantial investment capacity.

1 For more detail on the construction of these factors and the principles by which they were derived, please see "Introducing the Two Sigma Factor Lens", by Duncombe and Kay (2018).

Two Sigma Investments | Two Sigma Advisers
www.twosigma.com
NEW YORK HOUSTON LONDON HONG KONG TOKYO

Table of Contents

Exhibit 1 | Factor Descriptions and Expected Premia as of Dec 31, 2018

	FACTOR	ESTIMATED ORTHOGONAL FACTOR SHARPE RATIO	DESCRIPTION
Core Macro	Interest Rates	0.25	Exposure to the time value of money (inflation risk and future interest rate changes)
	Equity	0.35	Exposure to the long-term economic growth and profitability of companies
	Credit	0.21	Exposure to default and illiquidity risks specific to developed market corporate bonds
	Energy	0.23	Exposure to changes in prices for hard assets and any embedded futures premium
	Industrial Metals	0.19	
	Other Commodities	0.00[2]	
Secondary Macro	Emerging Markets	0.00[3]	Exposure to the sovereign and economic risks of emerging markets relative to developed markets
	Foreign Currency	0.00[4]	Exposure to moves in developed market currency values versus the portfolio's local currency
	Local Equity	0.00[4]	Exposure to the "home bias" of holding disproportionate risk in the local equity market
	Local Inflation	---[5]	Exposure to inflation-linked rates relative to fixed nominal rates within the local currency area

2 Other commodity sectors are found to have no actionable and statistically significant orthogonal factor premium, thus their forward-looking estimates are fixed to 0.0. See Section V for more details.

3 Emerging Markets assets are found to have no statistically significant orthogonal factor premium, thus the EM factor forward-looking estimate is fixed to 0.0. See Section VI for more details.

4 We believe the Foreign Currency and Local Equity factors should provide no orthogonal factor premium, thus their forward-looking estimates are fixed to 0.0. See discussion in Introduction for more details.

5 Although part of the Two Sigma Factor Lens where supported, we can observe only very short histories for the Local Inflation factor due to the recent introduction of inflation-linked sovereign bonds as an asset class. Hence we do not estimate a long-term expected premium in this paper.

I. Introduction

Our prior paper proposed the Two Sigma Factor Lens, a parsimonious set of factors derived from the performance of major asset classes that collectively explains the majority of risk in institutional portfolios. While the academic and practitioner literature includes many fine efforts to identify risk factors across asset classes,[6] we believe our approach contained two key differentiators.

First, we built our factors directly from traded assets rather than from economic variables like growth or inflation, which can have a surprisingly tenuous link with the assets themselves.[7] The factors were also each derived from individual asset classes rather than statistically-determined combinations such as principal components, which eases the interpretability of which assets in a portfolio are contributing to individual factor exposures and can make it easier to translate between a factor-based analysis and asset allocations.

Second, we asserted a hierarchy across the factors beginning with the most liquid asset classes, and statistically separated (i.e., orthogonalized) the unique returns for each set of less liquid assets from its exposure to the higher-order factors. This helps aggregate common risk factors across multiple asset classes to better identify concentrations of exposure to, say, the Equity factor from many different assets also sensitive to economic growth and investor risk aversion. For investors seeking to implement a factor-based portfolio, this approach can also preferentially tilt optimal factor exposures toward the lowest-cost factors such as Equity and Interest Rates,[8] with factors from less-liquid assets only selected for an optimized portfolio if they appear to provide significant diversifying return benefits.

A factor lens will suffice to describe a portfolio, providing insight on the attribution of its historical risk and returns. Choosing a desired allocation of factors or assets is a far trickier issue, and requires forecasts of risk and return expectations. In this paper, we propose a methodology using historical data to derive long-term forecasts of the return premia for major asset-class-based factors. Our methodology for identifying premia-bearing factors and estimating their respective long-term returns forecast follows five steps:

1. **Examine the maximum available return histories**, as decades of data provide greater insight on the likely return premium through different market and economic regimes.
2. **Consolidate individual assets and asset classes wherever feasible**, as a single risk factor driving multiple assets should carry the same long-term premium regardless of which individual assets provide a portfolio with exposure to that factor (i.e., there should be no easy arbitrages).
3. **Estimate the premium of each new factor *relative* to more liquid factors**, to assess whether less liquid factors have shown sufficient orthogonal returns to justify their addition (and inform sizing) in an efficient portfolio.

6 Please see the introduction and bibliography of Duncombe and Kay (2018) for examples of alternative approaches to factor identification.

7 Ilmanen (2011, Chapter 16) presents correlations of monthly changes in consensus growth and inflation forecasts with a variety of asset and market factor returns, finding changes in growth expectations have a maximum absolute correlation of 0.45 with listed private equity firms while changes in inflation expectations have a maximum absolute correlation of 0.50 with the S&P GSCI index of commodity futures.

8 We believe these factors may be considered the "lowest cost" in two important but distinct senses: they have the lowest transaction costs for direct investment as public equities and developed market sovereign bonds are among the most liquid marketable assets, and indexed investments in sovereign bonds and public equities have among the lowest fees of all managed investment vehicles.

4. **Select factors that exhibit positive long-term premia, backed by empirical and fundamental evidence**. Factors without orthogonal premia may be useful to identify for risk management purposes, but we believe they should not be included in an efficient portfolio.

5. **Finally, estimate the return premia jointly across all selected factors**, overweighting recent history and shrinking Sharpe ratio estimates toward a reasonable prior to make the forecasts more robust.

This paper bases its long-term forecasts on historical returns, with only light guidance from theory, founding its methodology in a belief that human nature is the fundamental force underlying all risk premia. Whether factors are rational compensation for exposure to "bad times" or rooted in the common behavioral biases of the marginal human investor, we believe that typical investors remain as human today as they were in 1900.[9] Neither our perceptions of bad times nor our behavioral reactions are likely to be very different from the investors in past market cycles, and though the marginal investor may shift over time due to structural changes, these changes should typically be gradual and minor rather than sudden and cataclysmic. This inescapably human element of the "price" of risk means that although it may be nearly impossible to predict next month's return for a factor, the return over the next few decades should resemble the price investors demanded for decades past.

We understand that any methodology for forecasting long-term premia, including that proposed in this paper, will have inevitable flaws. The investor may be tempted to throw up her hands and assume, for example, that all factors provide equal premia in the face of this uncertainty. This would even be a reasonably robust assumption,[10] and we find in Section VII that shrinking historically-observed asset class Sharpe ratios toward their cross-sectional average improves predictions of future premia. However, we find evidence in this paper of significantly differing premia across our identified factors, including a few that appear to provide no significant long-term returns. Given that nobody investing for the long run can avoid the implicit inclusion of some form of forward-looking views in their portfolio, we would rather examine and test each factor in turn, with more than a century of data behind us, than turn a blind eye to theory and history.

Before proceeding, we should briefly note a few things this paper does not do. We do not provide analysis of returns to the Foreign Currency, Local Equity, or Local Inflation factors, as these three factors seem to have neither fundamental nor empirical justification for a premium. Foreign currency exposures across global investors are net-neutral, so static holdings of foreign currency risk should not carry a premium unless they are exposed to systematic style factor risks such as currency carry or momentum.[11] For similar reasons, the Local Equity factor, which captures the orthogonal returns to an investor's local equity market relative to global equities, has a net zero return across all countries in the global portfolio and should not provide a premium beyond exposure to style premia such as cross-country value or momentum. Finally, the Local Inflation factor should theoretically provide a negative return premium on average, as

9 Though we do appreciate the irony of this sentiment coming from an investment firm employing almost entirely systematic strategies.

10 Rappoport and Nottebohm (2012) analyze simulated performance of portfolio optimization with forecast uncertainty under a variety of assumptions, finding that risk parity's implicit assumption of equal Sharpe ratios across asset classes performs well under conditions of high uncertainty. They also test a methodology to improve forecasts by combining the equal Sharpe ratio assumption with a priori estimates of expected asset class returns.

11 Although Siegel's paradox suggests that all investors may have positive excess return expectations from holding static foreign currency exposure due to Jensen's inequality, Campbell et al. (2010) show that the expected premium from this mathematical curiosity is negligible and we will not all get rich by trading currency exposures with one another.

it represents the returns from hedging inflation risk in sovereign bonds by holding securities with inflation-linked coupon payments rather than their maturity-matched nominal counterparts.

This paper also does not estimate returns for style factors, such as carry strategies or selling equity volatility, instead focusing on major asset class returns. Styles do play an important part in understanding the cross-section of returns and risk within asset classes, however, and we plan to address these in a future report.

The rest of the paper is structured as follows: Sections II through VI provide case studies analyzing asset class factors for potential risk premia; Section VII illustrates the utility of long-term historical data and shrinkage toward "average" risk-adjusted returns to prevent overfitting to the observed history; and Section VIII concludes.

.

II. Value of Long Histories: Looking Past the Interest Rates Mountain

Our dive into history begins with the Interest Rates and Equity factors, which are derived from returns to assets (sovereign bonds and common stock, respectively) that have been trading in relatively modern forms for over a century in many countries. Recent research has managed to reconstruct historical returns and pricing data for new markets and much older periods than previously available, providing useful perspective on whether the commonly-studied asset returns of recent decades appear truly representative of long-term expectations.

One of the most striking features of these long histories is the "interest rates mountain" of the 1950s through to the present day (see Exhibit 2). Over the course of six decades, long-term bond yields in most developed markets climbed from the low single digits to peaks around 15%, and then back down to near zero. This ascent and decline are astonishing in their scale, a treacherous alpine ridge with sheer drops to each side of 1980. But perhaps the more important part of the chart is the broad plain of mid single digit rates we see preceding the mountain, extending back to the pre-Industrial Age origins of modern central banks and sovereign bonds in northern Europe.

This presents a tricky issue for anyone seeking to forecast the Interest Rates premium: 60 years of data might sound like plenty of historical perspective for any factor, but interest rates appear to have been in "anomalous" territory for nearly that long. Trickier still, maybe the mountain isn't the anomaly after all. The shift by most developed market central banks from the gold standard to a fiat standard in the mid-20th century (culminating with the breakup of Bretton Woods in 1971) may render any data prior to the mountain obsolete, leaving the mountain as the only data we have that describes the current fiat regime.

The challenge of identifying the relevant period for estimating the Interest Rates premium puts in stark relief some of the issues faced in relying on historical data as a lens for the future, and especially highlights the potential benefit of using longer windows and applying sensible prior assumptions where possible.

Exhibit 2 | Historical Long-Term Bond Yields by Country

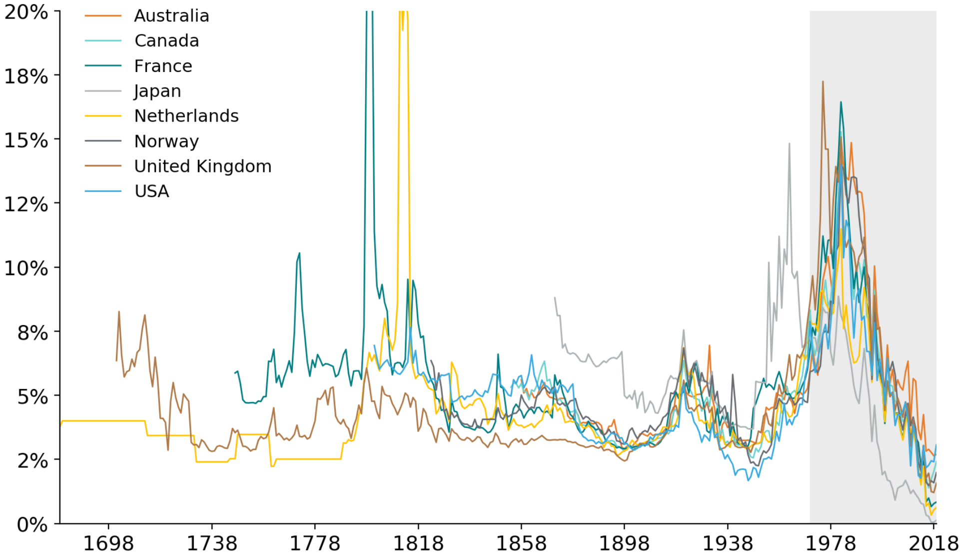

Exhibit 2: Historical long-term bond yields for major developed markets show the pronounced "Interest Rates Mountain" from the 1950s to today. Previous to this period, yields tended to hover in the mid-single-digits outside of extreme periods of sovereign instability. The shaded region highlights yields after 1970, representing the "modern era" for global financial markets (see footnote 16 for details). See Appendix B for data sources and other details.

We ultimately opt for an estimate of the very long-term historical average of the Interest Rates premium as more representative of our future expectations than the experience of recent decades. While declining yields have provided a strong tailwind to bond returns since 1980, forward-looking estimates of bond premia have been declining across the world based on both econometric and survey-based estimation approaches.[12] This decline appears tied to two notable, and likely interlinked, phenomena: inflation uncertainty and volatility appear to have significantly fallen from their 1980 highs, while correlations between stocks and bonds have shifted from positive to negative.[13] The former suggests a smaller risk premium should be priced into long-term bonds due to lower perceived risk of inflation-driven losses, while the latter suggests that bonds should carry a lower (or even negative) premium from providing a partial hedge against equity-led asset price shocks. Until we see evidence that expected bond premia have begun to rise from current levels, we choose to base our forecasts upon the lower forward-looking premium estimate derived from very long-run returns.[14]

III. Equities: The Long-Term Returns Champ

The value of taking a longer perspective becomes clearer when we add equity markets to our historical analysis. To wit, since 1900, risk-adjusted returns on equities are significantly higher than those on bonds. Yet if we only looked at the five decades since 1970 (i.e., as we rode down the rates mountain and bond prices rose), the historical record would suggest the opposite. Even a question as simple as "which factor has a higher premium" can take over a century of data to answer with a reasonable degree of certainty.

Exhibit 3 shows histograms of the historical risk-adjusted returns by country to equities and 10-year sovereign bonds over the two different periods. The top panel shows the results for 10 major developed markets since 1970.[15,16] Note that this corresponds to the grey area in Exhibit

12 The decline of implied long-term bond premia has been documented by Wright (2011) for G10 sovereign bonds since 1990 (using both affine curve model estimates and survey-based estimates), and by Adrian et al. (2014) in US Treasuries since 1961 (using affine curve model estimates).

13 Campbell, Pflueger, and Viceira (2018) find that shocks to US inflation have significantly decreased in size and changed from positive to negative correlation with changes in the output gap since 2000, suggesting that a more stable macroeconomic environment with better-anchored inflation expectations can explain both the lower expected bond premium and changing correlation between bonds and equities in a consumption-based macroeconomic model with habit formation. D.E. Shaw (2019) shows that surveyed inflation expectations and correlations between bonds and equities have shifted lower not only in the US, but also contemporaneously in Japan, Germany, and the United Kingdom. The authors also find that there still appears to be some long-term inflation-related premium priced into US Treasuries via a positive roll-down yield around 15- to 20-year maturities.

14 See Appendix B: Exhibits 1 and 14 for more details of the authors' calculations of the long-term average returns to the Interest Rates factor.

15 Countries included in analysis since 1970: Australia, Canada, France, Germany, Italy, Japan, Netherlands, Spain, United Kingdom, and United States.

16 Several developments point toward financial markets modernizing around 1970, including but not limited to the end of the Bretton Woods system in 1971-1973 as many developed markets shifted to floating exchange rates; the first listing of financial futures (contracts on foreign currencies) on the International Monetary Market (now part of CME Group) in 1972; and some of the earliest global stock market indices being first licensed by Capital International (now part of MSCI Inc.) in 1969.

2, when sovereign yields around the developed world were mostly declining from their peaks. Over this period, the average 10-year bond outperforms the average equity market on a risk-adjusted basis.

However, the dataset of historical returns provided by Dimson, Marsh, and Staunton (2016) allows us to extend this historical analysis to cover the risk-adjusted returns since 1900, for 21 countries.[17] This is the bottom panel of Exhibit 3. In contrast to the top panel, here we see that Equity risk has generated a substantially higher premium than bonds on average over this period, once the tailwind of falling rates from the past 48 years is sufficiently diluted. This surprising result shows that taking a longer view can not only enhance results — it can flip them.

Exhibit 3 | Comparisons of Risk-Adjusted Returns for Equities and Bonds

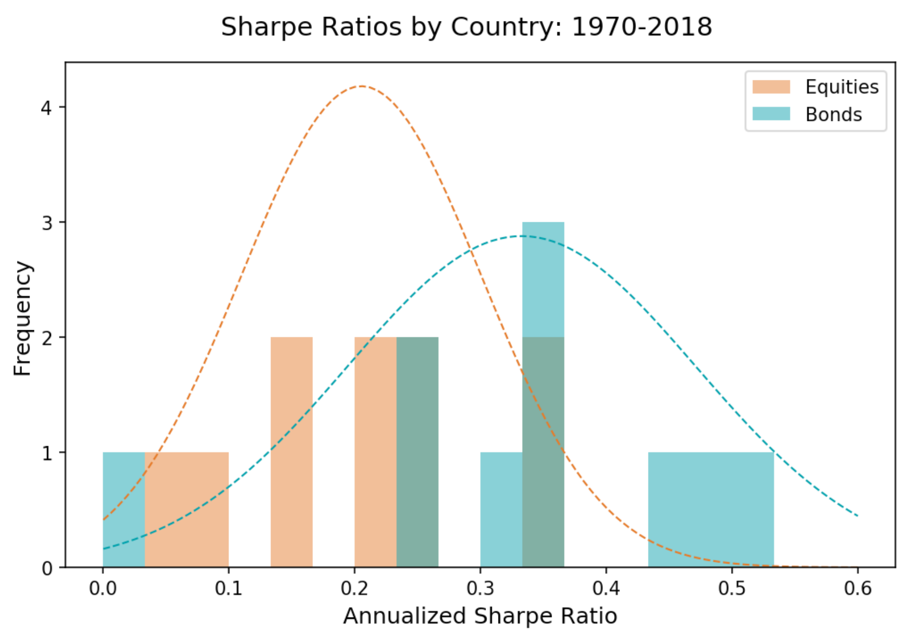

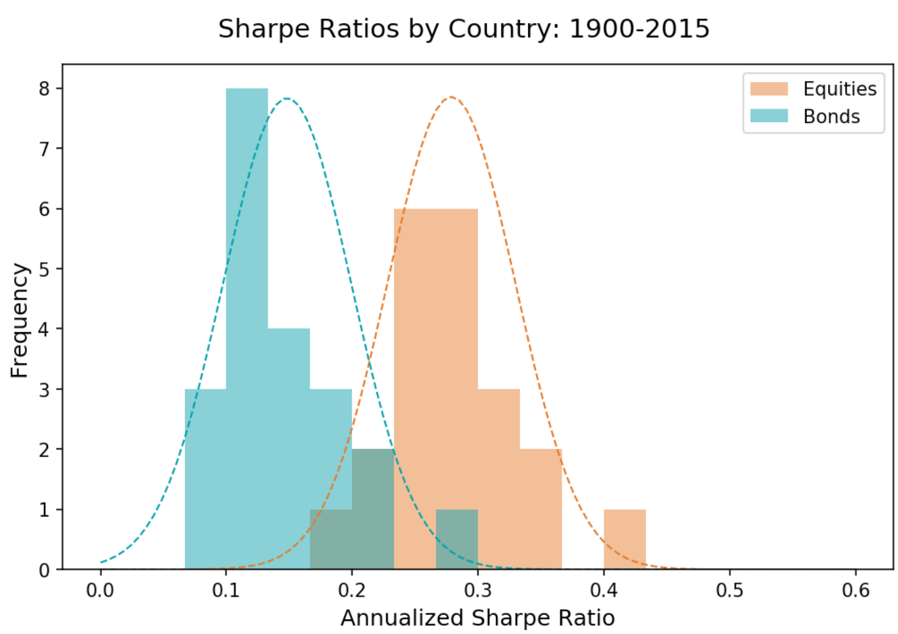

Exhibit 3: Histograms of realized Sharpe ratios for local equity markets and 10-year sovereign bonds by country, for a panel of 10 developed markets since 1970 in the top panel and 21 developed and emerging markets in the bottom panel. Best-fit normal distributions are drawn for the cross-section of realized bond and equity Sharpe ratios in each panel. Note that the best-fit curves overlap heavily in the commonly-examined period since 1970, with bonds showing a higher average Sharpe ratio, while there is clearer separation of the two distributions and a much higher average Sharpe ratio for equities in the longer sample. See Appendix B for data sources and other details.

17 Countries included in analysis since 1900: Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Netherlands, New Zealand, Norway, Portugal, South Africa, Spain, Sweden, Switzerland, United Kingdom, and United States.

Of course, any analysis of long-term equity returns must account for survivorship bias, and longer histories are more prone to this than shorter ones. For example, neither Russia nor China are included in the lower panel of Exhibit 3, and both countries experienced complete wipeouts for equity holders during their respective revolutions in 1917 and 1949. Fortunately, Dimson, Marsh, and Staunton (2016) provides summary returns for an index of global equity returns since 1900 that includes the total losses in both Russia and China and also weights all countries according to their market capitalizations at the start of the period. This index, revolutions and all, has an arithmetic mean excess return of 5.5% over bills, with a volatility of 17.5%, corresponding to a Sharpe ratio of 0.31 for equities, while a similarly constructed index for bonds realized a Sharpe ratio of only 0.12.[18] Assuming a near zero correlation, this equity/bond difference corresponds to a t-stat right around 2.0 given the 115 year history.[19]

Historical data provide the numbers, but what do we make of them? Why do equities outperform bonds on a volatility-adjusted basis? As one might expect, markets are hardly providing a free lunch for equity holders. We believe equities *should* provide a higher volatility-adjusted return than bonds because their return stream is more correlated to human capital and consumption, and because they compose the majority of risk in portfolios of wealthy stockholders and institutions.[20] In simple terms, this means equity downturns are extra painful because they have a disproportionate impact on portfolio values and happen at times when people find themselves simultaneously unemployed and short on cash. Self-aware investors should consider whether this "when it rains it pours" correlation applies to them at the margin (it probably does), and adjust accordingly.

IV. Credit: From Many Assets, One Factor

Corporate bonds represent a tremendously heterogeneous asset class in terms of geography, issuers, and expected risk of default. Yet, underneath these differences lie two key commonalities: lower average liquidity per issue than equities or sovereign bonds, and the negatively-skewed payoff risk from potential default. Both of these commonalities imply that credit might provide a premium over and above its embedded exposure to the factor risk of more liquid asset classes like equities.

First, we need to survey the credit landscape and condense it into one factor. Corporate bonds across the United States and Europe represent the most liquid of cash credit bonds and derivatives (based on the CDX and iTraxx indices of corporate credit default swaps for their respective geographies), and they form the basis of our combined Credit factor. Fortunately, the returns to corporate bond indices, after extracting the Equity and Interest Rates factors, appear to explain the residual returns to many other classes of credit instruments as well, as shown in Exhibit 4. We believe this justifies our parsimonious use of a single orthogonal Credit factor based on the most liquid indices, yet with explanatory power across a wide range of less liquid sub-classes of credit bonds.

18 The global asset class returns estimated by Dimson, Marsh, and Staunton (2016) are calculated in a common currency (USD) at spot rates for currency conversion. This inclusion of currency risk in the asset class returns will generate a relatively unbiased estimate of average return (since currency exposure tends to carry zero long-term return, see Boudoukh et al. (2015)), but will tend to bias the estimated volatility upward, and thus the long-term Sharpe ratio downward, relative to the currency-hedged Equity and Interest Rates factors.

19 Using the country-level results can provide greater statistical significance, as the lack of perfect correlation across countries implies that we are picking up some semi-independent information in observing that each country in Dimson, Marsh, and Staunton (2016) exhibited a higher risk-adjusted return for local equities than for local government bonds.

20 For models tying the Equity Risk Premium (ERP) to consumption risk of concentrated shareholders, see Guvenen (2009) and Malloy, Moskowitz, and Vissing-Jørgensen (2009). For empirical analysis of equity risk in institutional portfolios, see Duncombe and Kay (2018) and Leibowitz, Bova, and Hammond (2010). For more on the connection of the ERP to macroeconomic variables and theoretical motivations, excellent overviews may be found in Ang (2014), Damodaran (2018), and Norges Bank Investment Management (2016).

Exhibit 4 | Correlation Heatmap of Residualized Credit Indices and Combined Credit Factor

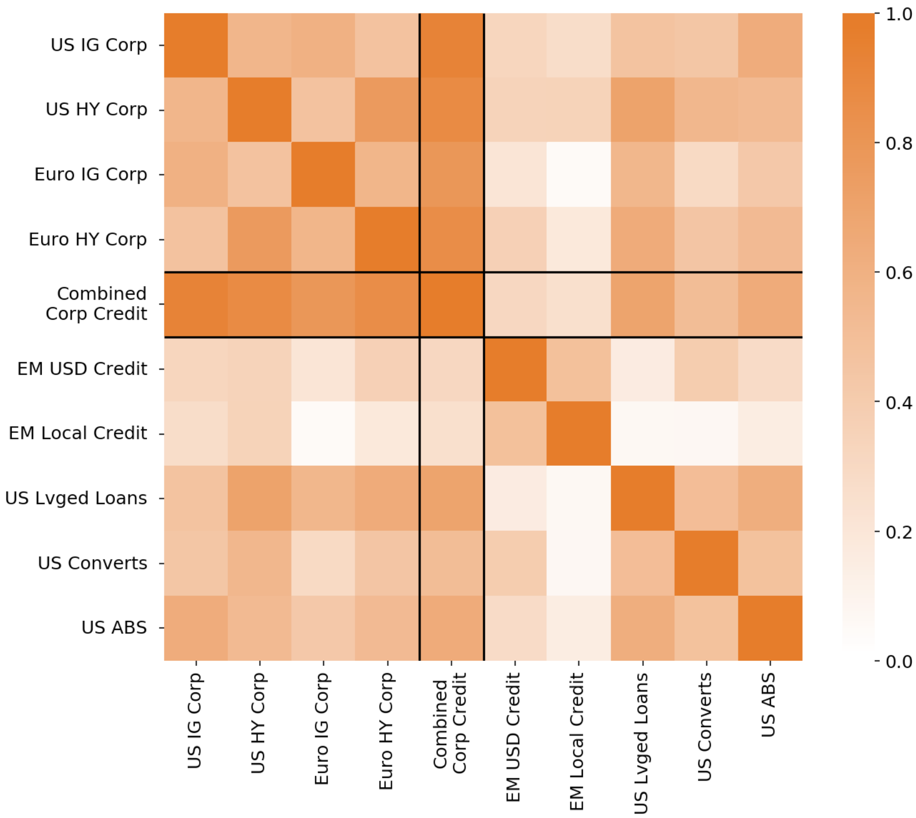

Exhibit 4: After stripping out exposures to the Equity and Interest Rates factors, residual returns for corporate bonds in the US and Europe show a high degree of correlation suggestive of a common risk factor (represented by "Combined Corp Credit"). Other US credit sectors show moderate correlations to the residual returns of corporate credit, while emerging market bonds have lower correlations due to their large EM-specific risk (see Section VI, especially Exhibit 11). See Appendix B for data sources and other details.

Credit bonds are among the oldest classes of securities,[21] yet the modern heterogeneity of the asset class makes it difficult to assume that past returns say much about today's Credit factor premium. For centuries, only the most developed nations and creditworthy corporations were capable of issuing bonds, and most classes of debt traded today have shifted from residing on bank balance sheets to securitized bond issues only in the past few decades. High yield bonds only became acceptable new issues in the mid-1980s in the United States, and they remained a very niche market in Europe and elsewhere until the late 1990s. Emerging-market sovereign bonds were also a miniscule market until the 1989 creation of Brady Bonds to help restructure and offload illiquid sovereign loans from bank balance sheets. Although the common risk factor underlying all these forms of credit allows us to extend the factor returns back further than many individual sub-classes of credit bonds, we would still like to ensure that our factor encompasses the broad range of default risk in today's credit markets.

For this reason, we have opted to calculate our extended Credit factor only for periods when both investment grade and high yield indices are available, even if we need to sacrifice some geographic diversity to achieve a longer-term view. Exhibits 5 and 6 show the results of extending our Credit factor back to 1990 through the use of US-only investment grade and high yield indices.

21 The Dutch East India Company issued corporate bonds as early as the 17th century, while Giesecke et al. (2011) documents a robust corporate bond market in the United States from 1866 onward.

It should be acknowledged that our Credit factor does not have a long enough history for its mean return to appear positive with statistical significance. At an expected Sharpe ratio of 0.15 to 0.35, typical of most factors derived from long-only asset classes, it would take up to a century of returns to meet the commonly accepted p < 0.05 cut-off for statistical significance. High yield bonds as a mature asset class have simply not been around long enough for a broad Credit factor to support this length of analysis.

However, we take some comfort from the longer perspective of researchers who have focused on the investment grade credit markets. Asvanunt and Richardson (2017) analyzed returns to US corporate credit bonds going back to 1936, and found their spliced proxy for credit returns to outperform a combination of duration-matched US Treasuries and the S&P 500 index with a t-statistic of 2.17, meeting the conventional bar for statistical significance. As they were analyzing a 79-year history (1936-2014), this is equivalent to a residualized annual Sharpe ratio of approximately 0.24 for their corporate credit proxy[22] — in line with our longer-term historical and forward-looking estimates.

Exhibit 5 | Cumulative Summed Returns to Orthogonal Credit Factor

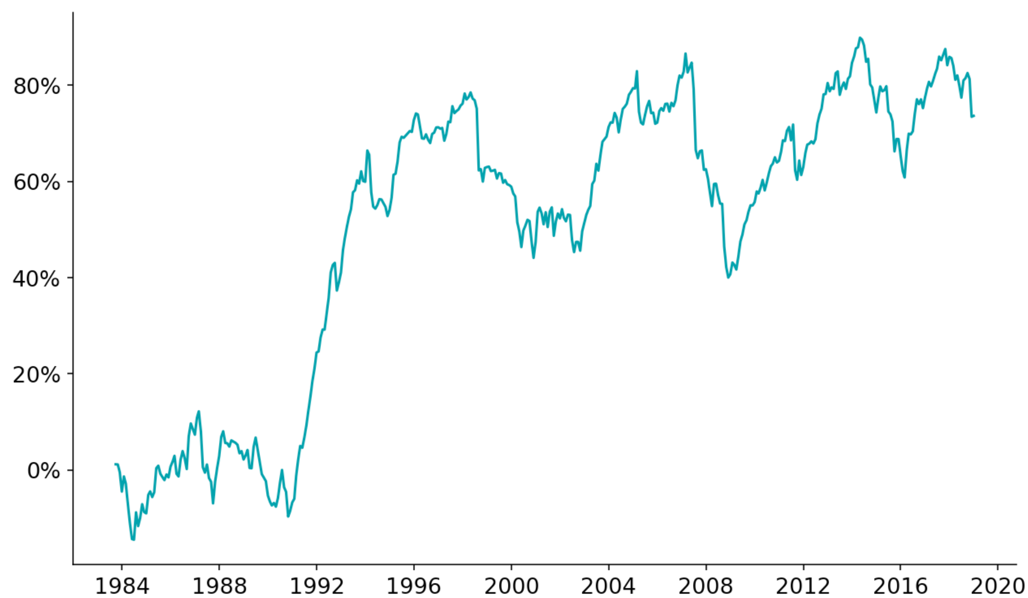

Exhibit 6 | Performance of Orthogonal Credit Factor

	SHARPE RATIO (EQUALLY WEIGHTED)	SHARPE RATIO (EXPONENTIALLY WEIGHTED)
Broad Credit Factor (European and US Corporate Credit)	0.13	0.11
Long History Credit Factor (US only prior to 1999)	0.23	0.17

Exhibit 6: The exponentially-weighted estimates of the historical Sharpe ratio use a 20-year half-life, in line with the findings of Section VII. The historical performance shown in Exhibit 6 is roughly in line with the forward-looking estimates in Exhibit 1 for the longest credit time series, though the shorter-term exponentially-weighted residualization process used for Exhibits 5 and 6 may slightly decrease the Sharpe ratio of the orthogonal Credit factor due to sharp shifts in the estimated Equity and Interest Rates exposures of credit indices around market crisis periods.

22 See Exhibit 8 of Asvanunt and Richardson (2017). Approximate annualized Sharpe ratio is calculated by taking the t-statistic of 2.17 for the intercept from regression III (monthly excess returns to their spliced corporate credit series regressed against excess returns of duration-matched US Treasuries and the S&P 500), dividing it by the square root of the approximate degrees of freedom in the regression (79×12 - 2), and multiplying by the square root of 12 to annualize.

V. Commodities: Rorschach's Factor Test

Commodities present an asset class with a much less monolithic risk structure than credit, requiring us to trade off parsimony (pushing us toward fewer total factors) for holism (pushing us to incorporate every unique observable risk in some factor). Even if this asset class appears too heterogeneous to fold into a single cohesive risk factor, we still find it is possible to consolidate a mere six or so representative factors from dozens of underlying commodities. Unfortunately, the evidence supporting a risk premium for each of the factors is not so clear cut, especially due to the relatively short time frame of available data. This leaves the commodities factors something like the Rorschach inkblots, with suggestions of premia yet still requiring much interpretation by the researcher.

Exhibit 7 | Correlations of Commodity Futures Returns

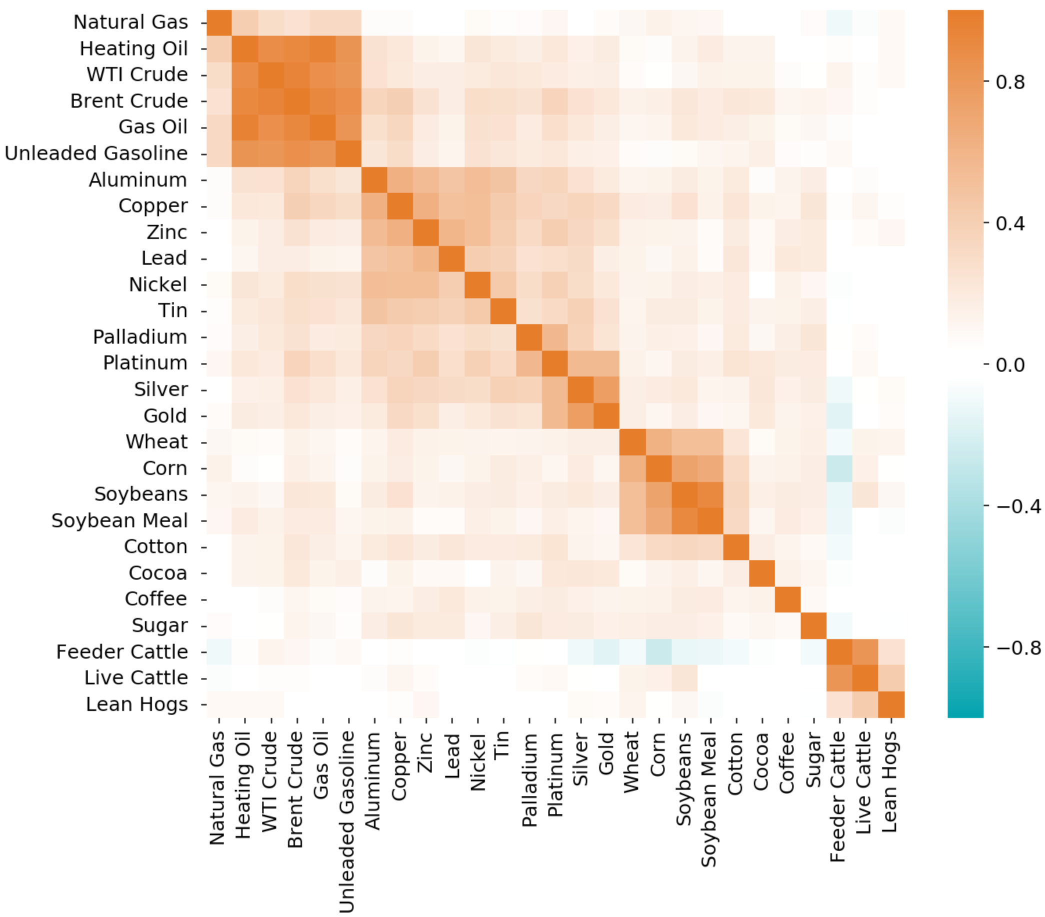

Exhibit 7: Correlations of monthly excess returns for single-commodity S&P GSCI indices from January 1970 (or index inception) to December 2018. Note the clusters of higher average cross-correlations within sector groups and relatively low correlations across commodities in different sectors, suggestive of a sector-related factor structure. See Appendix B for data sources and other details.

To start identifying potential risk factors that cut across multiple commodities, we show the long-term correlations of major commodity futures and forwards in Exhibit 7, where discernable clusters of highly-correlated assets suggest common underpinning risk factors. These statistical clusters align well with the classic commodity sectors, suggesting a fundamental explanation that matches the empirical evidence, and motivates our creation of the following risk factors:

- **Energy**, consisting of oil, natural gas, and distillate products
- **Industrial Metals**, including copper, aluminum, nickel, and other metals predominately used as industrial inputs
- **Precious Metals**, including gold, silver, and the less liquid platinum and palladium
- **Grains**, including wheat, corn, soybeans, and related products such as soybean meal
- **Livestock**, including lean hogs and cattle

The remaining "soft" commodities of cocoa, coffee, cotton, and sugar present a harder issue, as they are each highly idiosyncratic agricultural goods with their own unique supply and demand characteristics as shown by their very low correlations to other futures. However, we do not wish to test each of those individual commodities as a separate risk factor.[23] We instead opt to test the four soft commodities collectively for a return premium despite the relative lack of empirical evidence for a common risk process.

Exhibits 8 and 9 show the historical performance of our orthogonal commodity factors after extracting their time-varying loadings on the Interest Rates and Equity factors. Of the six consolidated commodity factors, only Energy and Industrial Metals showed substantial historical evidence of long-term actionable return premia.[24] As with the Credit factor in Section IV, none of the orthogonal commodity series meet the cut-off for a statistically significant return premium due to the insufficient history of many futures series, though their historical returns would be economically significant in a diversified portfolio.

Exhibit 8 | Cumulative Summed Returns to Orthogonal Commodity Factors

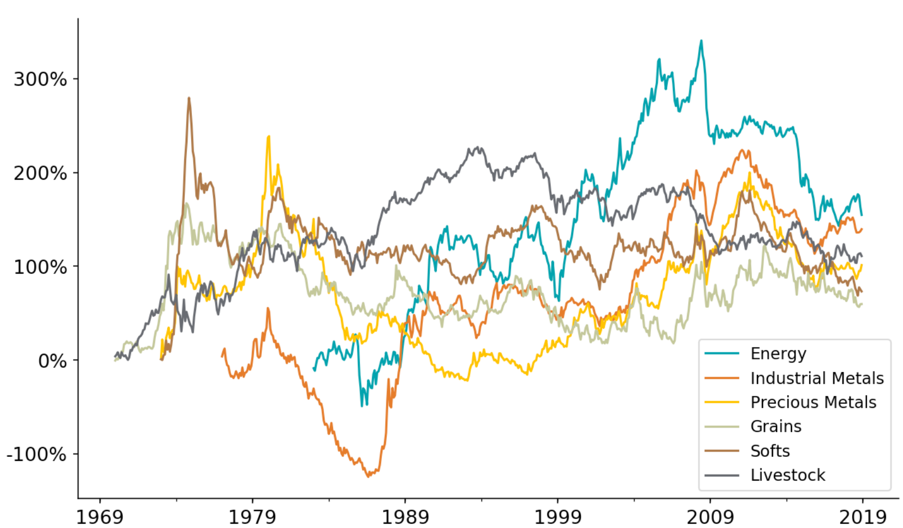

23 The decision to consolidate the soft agricultural commodities as a single factor to test for a return premium was based on two criteria. First, each individual risk factor tested for long-term profitability increases the chance of false positives. Thus, unless we have a strong fundamental reason for why assets would have unique risk premiums, we can mitigate this risk of false positives by collectively testing the assets' returns for a premium. Second, these contracts have relatively low liquidity, so the practical limits on long-term investment capacity of these futures means they would each individually have a negligible risk weight in a multi-asset portfolio.

24 Livestock futures show positive historical performance but are too illiquid to be seriously considered as an actionable premium.

Exhibit 9 | Performance of Orthogonal Commodity Factors

	SHARPE RATIO (EQUALLY WEIGHTED)	SHARPE RATIO (EXPONENTIALLY WEIGHTED)
Energy	0.15	0.05
Industrial Metals	0.15	0.15
Precious Metals	0.09	0.08
Grains	0.05	0.01
Softs (Other Agricultural)	0.07	-0.04
Livestock	0.12	-0.02

Exhibit 9: The exponentially-weighted estimates of the historical Sharpe ratio use a 20-year half-life, in line with the findings of Section VII. The historical performance shown in Exhibit 9 tends to be lower than the forward-looking estimates in Exhibit 1 due to the cross-sectional shrinkage increasing our forward-looking estimates for the individual commodity portfolios underpinning each factor.

Since the empirical evidence in Exhibit 9 does not meet the typical bar of statistical significance, we need to find further theoretical and empirical support for the strong assertion that only specific commodity futures sectors carry a risk premium. Fortunately, recent academic literature based upon the Theory of Storage points us toward a fundamental explanation and stronger statistical evidence for the premia found in Energy, Industrial Metals, and Livestock futures.

The Theory of Storage and Commodity Futures Premia

The Theory of Storage posits that the expected premium for long positions in commodity futures is not constant over time, but varies inversely with the level of inventories for the underlying commodity.[25] When inventories are low, the risk of "stock-out" or scarcity for consumers of the commodity rises, thus elevating spot prices and inducing higher expected price volatility in response to any further supply or demand shocks. Risk-averse producers will then have greater demand for hedges on their future production, being willing to pay some premium to futures holders as compensation for reducing exposure to the higher price volatility. In equilibrium models of the Theory of Storage, this time-varying expected premium to futures positions is known as the "convenience yield", as it equates to the premium paid by commodity holders willing to pay (or forgo) elevated spot prices to have inventory on hand immediately.

While recent development of the Theory of Storage leads to predictions of how commodity futures' premia change through time in response to inventories and suppliers' characteristics, it still says little about how average premia should differ across commodities. The inventory shock channel posited by the theory suggests that relatively "hard to store" commodities should be more prone to low inventory levels and more likely to have a premium, but the high cost of storage for those commodities could wipe out any average level of convenience yield.

More usefully, the Theory of Storage also suggests that any return premium to commodity futures should come from the difference in returns to a futures position plus short-term bonds (as the futures price is discounted in equilibrium by the interest rate to expiry) versus changes in the underlying commodity's spot price. These return differences, in expectation, are equal to the unobservable convenience yield minus the cost of storage.

25 Although the Theory of Storage was first outlined in Kaldor (1939), Working (1949), and Brennan (1958), it is only in recent years that theoretical models have extended to endogenously derive spot prices, futures prices, and the risk premium for long-only futures speculators accounting for the presence of stock-out risk. The discussion in this section relies heavily on the theoretical model and empirical results in Gorton, Hayashi, and Rouwenhorst (2013). Acharya, Lochstoer, and Ramadorai (2013) derive and test a similar model, though they find an additional important relationship between producers' time-varying risk-aversion and firm-level measures of financial distress.

We build off this insight to generate a more statistically powerful test for potential return premia in commodity futures, by looking at the monthly differences between S&P GSCI single commodity index total returns (which tracks the returns to a fully-cash-collateralized rolling futures position) and spot price changes for the underlying commodity. This analysis eliminates the main source of variance in commodity futures returns, namely spot price changes. When we pool these monthly realized return differences over time and across the several commodities in each sector, we can test the estimated sector-level premia (if any) embedded in the historical futures returns with much greater statistical power. The results of this analysis are shown in Exhibit 10, and align well with the findings of the residualized time-series returns in Exhibits 8 and 9. Only the Energy, Industrial Metals, and Livestock sectors showed statistically significant average premia through time.

Exhibit 10 | Results of Pooled Regression on Commodity Futures Returns minus Spot Price Changes

	ANNLZD FUTURES RETURN OVER SPOT	T-STATISTIC	# OF OBSERVATIONS
Energy	2.0%	3.54	1962
Industrial Metals	4.8%	5.87	1771
Precious Metals	0.8%	0.95	1042
Grains	0.0%	-0.02	2051
Softs	0.1%	0.08	1928
Livestock	3.1%	3.00	1306

Exhibit 10: Estimated sector-level return premia and t-statistics after controlling for monthly time effects in a panel of 24 GSCI single-commodity futures indices over the period January 1970 to December 2018. Individual contracts' observations are weighted based on an equal blend of their target 2019 weights in the Bloomberg Commodity Index and the S&P GSCI Index to avoid overweighting the estimated premia of relatively illiquid commodity futures. See Appendix B for details of return series used and time periods available by commodity.

This analysis, with its greater statistical power, gives us more comfort in singling out the Energy and Industrial Metals commodity sectors as having actionable long-term premia, as livestock futures tend to be too illiquid for large scale investment. But it would be better still to also have theoretical justification for why Industrial Metals appear to have a risk premium while harder-to-store commodities such as many agricultural products do not. Although we have been unable to find pre-existing academic literature addressing this precise question, we do believe that several recent papers and empirical findings collectively point in a promising direction.

Several papers have studied the fluctuation of rolling futures returns and the futures basis over the course of the business cycle, finding that the apparent risk premium from holding commodity futures is correlated with economic conditions and highest around business cycle peaks.[26] This accords with the Theory of Storage's prediction that risk premia should be driven by the risk of stock-out in commodities, which would presumably be highest at times of high economic growth. Assuming that risk-averse speculators are more capacity constrained and demand higher premia when more commodities are close to stock-out, this evidence suggests a procyclical common factor driving risk premia across commodities that would be most present in futures returns for the individual commodities most closely tied to economic growth: Energies and Industrial Metals.

26 In particular, Hong and Yogo (2012) find that a diversified basket of commodity futures has procyclical expected returns after controlling for the lagged average futures basis; Fama and French (1988) find that demand shocks for metals around business-cycle peaks result in negative futures basis and high convenience yields, supporting predictions of the theory of storage; and Kucher and Kurov (2014) find the futures basis for most energy commodities grows more negative at business cycle peaks while expected spot returns rise, implying a higher expected return for futures investors.

As the numerous assumptions and caveats highlighted above suggest, this remains a fertile area for further research. We will defer more detailed tests of the cross-section of commodity futures premia and a potential procyclical common factor to future publications.

We conclude with one final caveat regarding commodities as an asset class: the exact method one uses to get exposure matters greatly, as the Theory of Storage predicts that premia may be available in futures contracts but not by holding commodity inventories. Furthermore, even the rules one uses for rolling futures can have a material impact on returns.[27] Yet the empirical and theoretical findings above give us some comfort in recommending diversified Energy and Industrial Metals commodity futures as part of an efficient asset allocation.

VI. Emerging Markets: Risk without Reward?

We continue our process of testing newer asset classes for their marginal returns with emerging markets (EM). EM assets have grown considerably in size and liquidity over the past couple of decades, though expected transaction costs for EM bonds, currencies, or stocks typically remain higher than for their developed market counterparts. As a newer geographic class of investments exposed to unique political risks and greater illiquidity, EM assets appear to have a reasonable case for providing a positive factor premium to compensate investors who take on these risks.

Exhibit 11 shows the risk decomposition of individual EM asset class returns since inception, stripping out the more liquid developed market factors embedded in the classes to isolate the (sizeable) orthogonal risk associated with EM bonds, currencies, or equities above and beyond leveraged exposure to global Equity or Credit risk. The orthogonal returns to each EM asset class do show statistically significant cross-correlations supporting their combination into a single factor, even if their cross-correlations are not as high as those observed for corporate credit residual returns in Section IV.[28]

Exhibit 11 | Risk Decomposition of Emerging Market Asset Class Returns

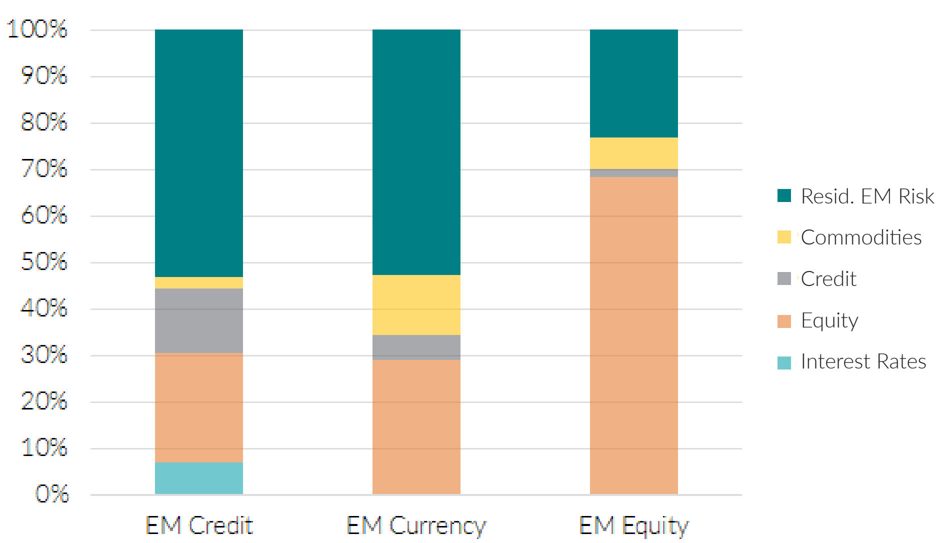

Exhibit 11: The global Equity factor explains a significant percentage of time-series risk across multiple EM asset classes, with smaller contributions from factors related to global corporate credit and commodities. The residual, idiosyncratic returns of all three EM asset classes explain a sizeable portion of their respective variances, and show significant cross-correlations supportive of a common EM risk factor. See Appendix B for data sources and other details.

27 Mouakhar and Roberge (2010); Rallis, Miffre, and Fuertes (2012)

28 Correlations across equally-weighted residual returns to the EM Credit, Currency, and Equity proxies used in this paper varied from 0.31 to 0.36, with t-statistics of 3.4 to 7.0. All correlation estimates were higher when using exponentially-weighted returns with a 20-year half-life, suggesting that the common factor across emerging market assets may have increased in importance over the analysis period

The Emerging Markets factor, however, has thus far fallen short in the most important test for long-term investment: profitability. Exhibits 12 and 13 show that, once the leveraged exposures to global Equity, Interest Rates, and Credit risk factors have been stripped out of the EM asset class returns, the orthogonal factor candidates show negative or near-zero long-term returns, at least over this relatively short sample.

Despite a reasonable case for a premium, the empirical evidence suggests that the Emerging Markets factor may provide no extra risk-adjusted return to a globally diversified asset class portfolio, even without accounting for higher expected transaction costs in these less liquid markets. This accords with the findings of Dimson, Marsh, and Staunton (2018) that EM equities have underperformed their developed market counterparts on a cumulative basis since 1900. Although this may seem surprising given the higher recent economic growth rates of emerging markets, the link between economic growth and local asset returns actually appears quite tenuous and has shown little relationship to the returns of EM assets either through time or across countries.[29]

Exhibit 12 | Cumulative Summed Returns to Orthogonal Emerging Markets Factors

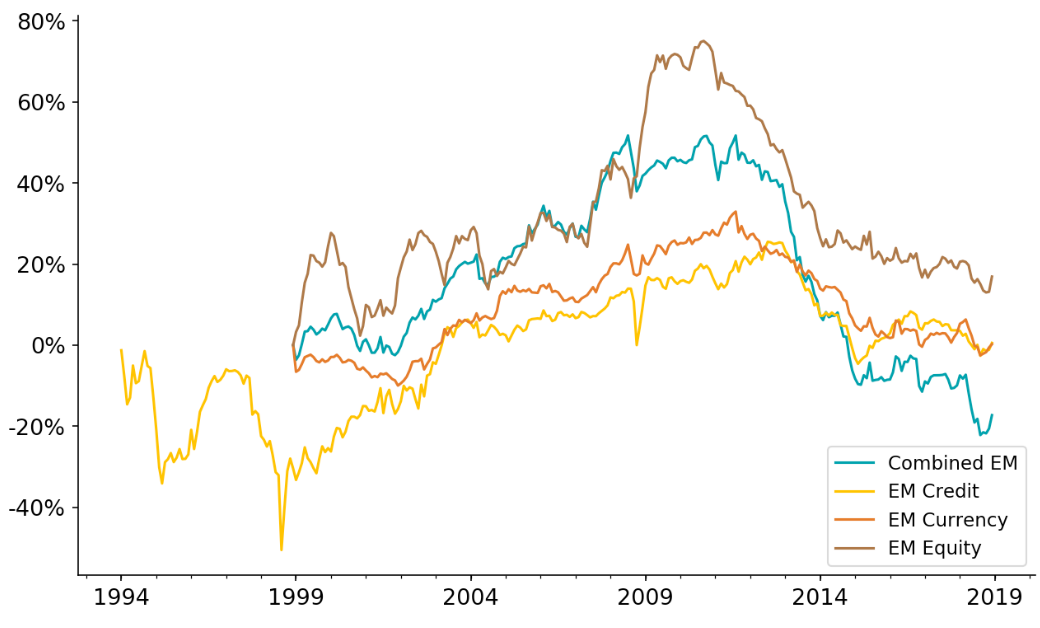

Exhibit 13 | Performance of Orthogonal Emerging Market Factors

	SHARPE RATIO (EQUALLY WEIGHTED)	SHARPE RATIO (EXPONENTIALLY WEIGHTED)
Combined EM	-0.12	-0.24
EM Credit	0.00	0.00
EM Currency	0.00	-0.07
EM Equity	0.10	0.05

Exhibit 13: The exponentially-weighted estimates of the historical Sharpe ratio use a 20-year half-life, in line with the findings of Section VII. Despite periods of outperformance, such as the early 2000s, the orthogonal EM factors extracted from all three asset classes show nearly zero long-term premium after accounting for exposure to global Interest Rates, Equity, Credit, and commodities-related factors.

29 Saret (2014)

VII. From History to Expectations

Our analysis so far has focused on historical returns to assets and factors, with the implicit assumption that the future will look like the past. As a conservative start, we do not expect a risk factor with decades of flat returns to begin providing a premium tomorrow. But we would like to move beyond simple "zero or non-zero" classifications and provide forward-looking estimates of the expected premium for each factor. We believe that empirical data from long performance histories remains our best guide, but consider two additional steps that might make our forecasts more robust than simply taking the average return over the longest possible history.

First, it seems that it would be a bit silly to claim that returns in 1906 contain just as much information about today's expected premium as returns in 2016. Markets change over time — evolution in the investor base, issuer base, and market structure could all potentially affect the long-term expectations for factor returns. This provides our motivation to overweight more recent returns when estimating the long-term premium of any given factor. However, we believe this rate of change is likely to be relatively modest for major risk factors that have long been present in investors' portfolios. A very long-term exponentially-weighted average return, say with a half-life of decades, may strike a suitable balance between emphasizing more recent data and capturing a broad sample of performance through longer- and shorter-term market cycles.

Second, we also believe that it may help to "shrink" individual asset returns toward a common average. This shrinkage toward a common mean could help prevent overfitting our estimates of long-term factor premia on the returns of the past couple of decades, making them more robust even after periods of extreme individual asset class returns.

Our empirical analysis of historical returns weighting and shrinkage is in Appendix A. Overall, the results suggest that the best balance of overweighting recent periods while keeping relevant historical data for predicting future performance occurs when we use an exponential weighting with a half-life around twenty years. They also suggest that observed historical Sharpe ratios for individual asset classes should be shrunk cross-sectionally to better estimate forward-looking premia, with approximately 50% weight on the asset-specific return premium and 50% weight on the cross-sectional average. We believe these two adjustments allow for the sensible and reasonably robust estimate of forward-looking premia from historical returns.

However...[30]

While we always try to trust in data and follow our models, there are still some occasions when the historical data appears so clearly atypical that common-sense adjustments could improve out-of-sample accuracy. We believe the current state of sovereign yields represents one of these occasions. Although it is impossible to predict the path of global sovereign yields in the next couple decades, we are quite confident they will not decline another 10 percentage points from the current levels of zero to three percent. So for the critical Interest Rates factor, we believe that the equal-weighted long-term average Sharpe ratio going back to the early 20th century would be a more sensible starting point in estimating future performance than an estimate overweighting the past few decades.[31]

30 With our deep apologies, as there always seems to be a "however." Financial and economic data is so noisy that we often commiserate with President Truman's apocryphal desire for a one-handed economist.

31 For details on the derivation of our estimate for the long-term average Sharpe ratio of the Interest Rates factor, please see notes on Exhibit 1 and 14 in Appendix B.

VIII. Conclusion: Forward-Looking Returns by Factor

In this paper, we propose a methodology for generating long-term, forward-looking views on asset-class-based factors. By examining long histories of performance and statistically extracting the diversifying return of each new asset class, we seek to identify key common factors and determine which appear to carry a return premium additive to a diversified institutional portfolio. After a study of major asset classes, we find five factors that appear to provide unique return premia with high potential investment capacity: Interest Rates, Equity, Credit, Energy futures, and Industrial Metals futures.

Slightly overweighting recent returns and shrinking historical performance of asset classes toward the average also appears to improve the explanatory power of past returns in the future. Our one discretionary exception to the overweighting of recent returns is the Interest Rates factor, where an unsustainable decline of global yields in the past four decades leads us to suggest using the longest available returns history as the best guide for the future.

We conclude by recapping our procedure for estimating the long-term risk premia of the factors underpinning asset class returns:

1. Calculate historical Sharpe ratios for broad asset class indices that reasonably represent currency-hedged returns to global equities, bonds, credit, and commodities, applying exponential weights with a half-life of 20 years to overweight more recent returns.

2. Apply haircuts to the global bond and credit indices' Sharpe ratios, reflecting a go-forward expected Sharpe ratio for the Interest Rates factor equal to its historical average since 1900.

3. Shrink all post-adjustment asset class Sharpe ratios by 50% toward their cross-sectional mean.

4. Residualize the shrunken asset classes' Sharpe ratios to extract the implicit forward-looking Sharpe ratios for the orthogonal factors.

The historical Sharpe ratio estimates and intermediate steps of our procedure are shown in Exhibit 14 for illustration purposes, using returns data through December 2018. The second column provides the long-term observed historical Sharpe ratio for the asset class proxies used to construct each factor. The third column applies our suggested cross-sectional shrinkage detailed in Section VII. The fourth column extracts the expected Sharpe ratios we estimated for the unique, orthogonal premia of each less liquid factor relative to their embedded exposure to the two most liquid premia: Interest Rates and Equity.

The findings of modest explanatory power for past returns on a three- to five-year forecast horizon should encourage humility in the application of these forecasts. However, we believe the "inescapably human element" behind factor premia allows us to draw some careful conclusions from the historical record.

This paper identifies five compensated factors that may collectively form the basis of an asset allocation strategy with substantial investment capacity, and provides a method for extracting long-term factor return forecasts from asset class histories. With future research, we will turn to the next question of forecasting risk and translating those forecasts into suggested long-term portfolio allocations.

Exhibit 14 | Calculation Steps for Expected Premia as of Dec 31, 2018

	WEIGHTED HISTORICAL ASSET CLASS SHARPE RATIO	POST-SHRINKAGE ASSET CLASS SHARPE RATIO	POST-SHRINKAGE ORTHOGONAL FACTOR SHARPE RATIO (FINAL ESTIMATE)
Interest Rates	0.20[32]	0.25	0.25
Equity	0.40	0.35	0.35
Credit	0.56[33]	0.46	0.21
Energy	0.17[33]	0.24	0.23
Industrial Metals	0.28[33]	0.29	0.19
Other Commodities	---[34]	---[34]	0.00[35]
Emerging Markets	---[34]	---[34]	0.00[36]
Foreign Currency	---[34]	---[34]	0.00[37]
Local Equity	---[34]	---[34]	0.00[37]
Local Inflation	---	---	---[38]

This paper provides only an overview of the subject matter discussed herein. It does not discuss many important assumptions, methodologies and other aspects of these subjects. All information herein is subject to change without notice.

32 Based on long-term, equally-weighted average returns and risk rather than exponential weights. See Section VII and Appendix B for details.
33 Historical Sharpe ratio estimates for all assets with shorter return series are adjusted to reflect the long-term expected Sharpe ratios of their estimated exposure to Interest Rates and Equity factors.
34 We believe that other commodity sectors, emerging markets assets, foreign currency exposure, and local equity exposure carry no statistically significant orthogonal factor premium, and thus are excluded from the cross-sectional shrinkage calculations applied to asset classes believed to carry a unique premium.
35 Other commodity sectors are found to have no statistically significant orthogonal factor premium, thus their forward-looking estimate is fixed to 0.0. See Section V for more details.
36 Emerging Markets are found to have no statistically significant orthogonal factor premium, thus their forward-looking estimate is fixed to 0.0. See Section VI for more details.
37 We believe the Foreign Currency and Local Equity factors should provide no orthogonal factor premium, thus their forward-looking estimates are fixed to 0.0. See discussion in Introduction for more details.
38 Although part of the Two Sigma Factor Lens where supported, it is only possible to observe very short histories for the Local Inflation factor due to the recent introduction of inflation-linked sovereign bonds as an asset class. Hence we do not estimate a long-term expected premium in this paper.

Appendix A: Regression Analysis to Estimate Cross-Sectional Shrinkage and Historical Returns Estimation

Though we lack the centuries of data across all of our selected factors necessary to rigorously test these hypotheses, we still have nearly a century of data across three key asset classes that can be applied to testing long-term returns predictability. By looking at the returns to global equities, global sovereign bonds, and US credit going back to 1925, we find that both a multi-decade lookback period and cross-sectional shrinkage appear to provide a modest, but statistically significant, forecast of asset class returns.

To check whether differences in historical asset class premia persist in the future, we used an exponentially-weighted moving average of historical returns and volatility for global equities, global sovereign bonds, and US corporate credit to estimate their historical Sharpe ratio on a rolling basis for each month starting in December 1940 (so we could begin with at least 15 years of data in the lookback period). We then ran panel regressions of the historical differences in Sharpe ratio against the three- and five-year forward-looking Sharpe ratios. Two main specifications were considered:

1. In the top row of each set of results, we tested a model of shrinking historical asset class Sharpe ratios toward a constant value, persisting through time and across asset classes:
   $$fsr_{i,t} = \beta_0 + \beta_1 \cdot psr_{i,t} + \varepsilon$$

2. In the second row, we tested a model of shrinking historical asset class Sharpe ratios toward their cross-sectional average at each point in time, allowing the overall return premium across all assets to fluctuate over the years:
   $$fsr_{i,t} = \beta_0 + \beta_1 \cdot \left(psr_{i,t} - \overline{psr}_t\right) + \varepsilon$$

In the equations above, $psr_{i,t}$ represents the exponentially-weighted Sharpe ratio of past excess returns for asset i at time t, $\overline{psr}_t$ represents the cross-sectional average of the exponentially-weighted past return Sharpe ratios across all assets at time t, and $fsr_{i,t}$ represents the equally-weighted Sharpe ratio of future three- or five-year returns for asset i at time t.

Exhibit 15 shows the results of our regression analyses, with standard errors adjusted to reflect the large degree of autocorrelation in these slow-moving historical returns. As can be seen from the tables of r-squared values and t-statistics for the asset-specific betas (β_1 in the model equations above), an assumption of shrinking historical asset class performance toward a constant Sharpe ratio provided virtually zero explanatory or statistical power regardless of the weighting used for historical data or the length of forecast period. Shrinking historical asset class Sharpe ratios toward their cross-sectional average (in the second row of each set of results) showed statistically significant predictive power over both forecast horizons, especially when historical Sharpe ratios were calculated using exponential weights with a 20-year half-life.

Exhibit 15 | Results of Long-Term Regressions over Multiple Lookback Periods and Forecast Periods

Exhibit 15: Test results of regressions based on the assumption of shrinking historical asset class Sharpe ratios toward either a constant Sharpe ratio or the cross-sectional average. Results across each row reflect regressions where the independent variable was based on historical asset class Sharpe ratios calculated using exponential weights with the specified half-life. T-statistics were calculated using Driscoll-Kraay (1998) robust standard errors accounting for up to 36 months of lagged effects for the three-year forecast regressions and up to 60 months of lagged effects for the five-year forecast regressions. See Appendix B for data sources and other details.

R-Squared of Regressions

Regressions vs 3-Year Forward SRs

Lookback EWMA Half-Life in Years	5	10	20	30	50
Shrink to Constant SR	0.003	0.004	0.003	0.003	0.003
Shrink to XS Avg SR	0.015	0.016	0.014	0.013	0.012

Regressions vs 5-Year Forward SRs

Lookback EWMA Half-Life in Years	5	10	20	30	50
Shrink to Constant SR	0.005	0.004	0.002	0.001	0.001
Shrink to XS Avg SR	0.025	0.024	0.019	0.017	0.015

Asset-Specific Robust T-Statistics

Regressions vs 3-Year Forward SRs

Lookback EWMA Half-Life in Years	5	10	20	30	50
Shrink to Constant SR	0.73	0.95	1.02	1.03	1.04
Shrink to XS Avg SR	1.89	2.57	2.92	2.93	2.86

Regressions vs 5-Year Forward SRs

Lookback EWMA Half-Life in Years	5	10	20	30	50
Shrink to Constant SR	0.81	0.78	0.56	0.45	0.36
Shrink to XS Avg SR	2.50	3.15	3.09	2.92	2.74

Asset-Specific Beta Estimates

Regressions vs 3-Year Forward SRs

Lookback EWMA Half-Life in Years	5	10	20	30	50
Shrink to Constant SR	0.12	0.19	0.22	0.22	0.22
Shrink to XS Avg SR	0.35	0.45	0.50	0.51	0.51

Regressions vs 5-Year Forward SRs

Lookback EWMA Half-Life in Years	5	10	20	30	50
Shrink to Constant SR	0.12	0.15	0.12	0.10	0.08
Shrink to XS Avg SR	0.33	0.41	0.43	0.42	0.41

Appendix B: Data Sources and Exhibit Details

Exhibits 1 and 14

Historical Estimated Sharpe Ratio for Interest Rates factor is based upon equally-weighted average returns and risk since 1870/1900, while Historical Estimated Sharpe Ratios for all other factors with an expected positive risk premium are based upon exponentially-weighted average returns and volatility with a 20-year half-life following the methodology outlined in Section VII.

Long-term average returns and risk figures for the Interest Rates factor are based on data from Dimson, Marsh, and Staunton (2014); Jorda et al. (2018); and Piketty and Zucman (2014). The estimate of 0.20 for the long-term historical Sharpe ratio of the Interest Rates factor is the authors' own, based upon two separate estimates derived from the aforementioned papers:

1. Jorda et al. (2018) in Table 3 provide pooled estimates for excess returns (1.53%) and standard deviation (8.38%) of local-currency bonds across 16 countries for the period 1870-2015. This methodology should eliminate the volatility impact of currency movements, similar to our currency-hedged Interest Rates factor, and provides a long-term Sharpe ratio estimate of **0.18**.

2. Dimson, Marsh, and Staunton (2014) report 21 country-level real returns and standard deviations for bills and long-term bonds over the period 1900-2015 in Appendices 1.4 and 1.5, respectively. The authors used these to derive country-level estimates of the excess return and standard deviation for long-term bonds based on an assumed average correlation of 0.5 between real returns for bills and bonds within each country. By assuming a correlation of 0.5 for bond excess returns across countries and applying historical GDP weights derived from Piketty and Zucman (2014), we come to an estimate for the Sharpe ratio of a global currency-hedged bond portfolio of **0.21**.

Exhibit 2

Annual 10-year government bond yields from Global Financial Data. Developed markets represented were selected for geographic diversity and length of historical yields available. Yields for Dutch and French 10-year bonds peaked at 66.67% and 25.64%, respectively, around the turn of the 18th century.

LABEL	GFD SYMBOL
Australia	IGAUS10D
Canada	IGCAN10D
France	IGFRA10D
Japan	IGJPN10D
Netherlands	IGNLD10D
Norway	IGNOR10D
United Kingdom	IGGBR10D
USA	_TNXD

Historical Sharpe ratio estimates for all assets with shorter return series are adjusted to reflect the long-term expected Sharpe ratios of their estimated exposure to Interest Rates and Equity factors. This adjustment led to lower estimates of the historical Sharpe ratios for credit assets and a slightly higher estimate of the historical Sharpe ratio for energy commodity futures.

Exhibit 3

Returns data for January 1970 to December 2018 are total returns series from Global Financial Data (bonds and bills) and MSCI local currency net return indices from Bloomberg (equities). The five European countries with the largest 1970 GDP in current USD were selected to avoid overrepresentation of Europe in the dataset due to greater data availability (GDP data from World Bank Open Data). Symbols/tickers for the total return indices used for country-level equity, government bond, and government bills are provided in the table below.

TICKER	BILLS GFD SYMBOL	BONDS GFD SYMBOL	EQUITIES BLOOMBERG
Australia	TRAUSBIM	TRAUSGVM	NDDLAS
Canada	TRCANBIM	TRCANGVM	NDDLCA
France	TRFRABIM	TRFRAGVM	NDDLFR
Germany	TRDEUBIM	TRDEUGVM	NDDLGR
Italy	TRITABIM	TRITAGVM	NDDLIT
Japan	TRJPNBIM	TRJPNGVM	NDDLJN
Netherlands	TRNLDBIM	TRNLDGVM	NDDLNE
Spain	TRESPBIM	TRESPGVD	NDDLSP
United Kingdom	TRGBRBIM	TRGBRGVM	NDDLUK
United States	TRUSABIM	TRUSG10M	NDDUUS

Sharpe ratio calculations for 1900 to 2015 are the authors' own, based upon the country-level real returns and standard deviations for equities, bonds, and bills provided in Appendices 1.2, 1.4, and 1.5 of Dimson, Marsh, and Staunton (2016). For estimating the standard deviation of excess returns over bills, real returns for stocks and bills were assumed to be uncorrelated and the real returns for bonds and bills were assumed to have 0.5 correlation, the latter estimate based on historical correlation data provided in figure 4 of Jorda et al. (2018).

Exhibit 4

Correlations from index inception through December 2018 of monthly residual returns to credit sector total return indices after extracting estimated loadings on the global Equity and Interest Rates factors. "Combined Corp Credit" is an equal-risk-weighted combination of the residual returns to the four US and European corporate credit indices, mirroring the construction of the Credit factor in this paper. All returns data from Bloomberg.

	INDEX NAME	BLOOMBERG TICKER	FIRST MONTHLY RETURN
US IG Corp	Bloomberg Barclays US Corporate	LUACTRUU	Feb 1973
US HY Corp	Bloomberg Barclays US Corporate High Yield	LF98TRUU	Sep 1983
Euro IG Corp	Bloomberg Barclays Corporate Pan-European Aggregate Hedged to USD	LP05TRUH	Feb 1999
Euro HY Corp	Bloomberg Barclays Pan-European High Yield Hedged to USD	LP01TRUH	Feb 1999
Combined Corp Credit	Authors' calculations (see above)	---	Sep 1983
EM USD Credit	Bloomberg Barclays Emerging Markets USD Aggregate	EMUSTRUU	Feb 1993
EM LC Credit	Bloomberg Barclays Emerging Market Local Currency Government	EMLCTRUU	Jul 2008
US Lvged Loans	Credit Suisse Leveraged Loan	CSLLLTOT	Jan 1992
US Converts	Bloomberg Barclays US Convertibles Composite	LUCCTRUU	Feb 2003
US ABS	Bloomberg Barclays US Aggregate ABS	LUABTRUU	Jan 1992

Exhibits 5 and 6

Credit factor returns are an equal-risk-weighted combination of the rolling residual returns relative to the Equity and Interest Rate factors for US investment grade and high yield corporate credit and European investment grade and high yield corporate credit, hedged to USD. All returns data from Bloomberg, using the corporate credit indices listed in the Exhibit 4 data table.

Credit factor returns are calculated daily since January 1990 and monthly to September 1983. European credit indices are only included in the factor from August 2000, when daily data becomes available for both. ICE BofAML High Yield Master II returns are used as the US high yield credit proxy for daily data from January 1990 to August 1998, with Bloomberg Barclays US Corporate High Yield used for monthly returns before and daily returns after.

Exhibit 7

All returns data from Bloomberg, using the commodity futures indices listed in the Exhibit 10 data table.

Exhibits 8 and 9

Commodity sector returns are an equally-weighted combination of futures index returns for the commodities in each sector, residualized on a rolling basis to the Equity and Interest Rate factors. Qualitatively similar results were found when using liquidity-related weights within sectors. All returns data from Bloomberg, using the commodity futures indices listed in the Exhibit 10 data table.

Exhibit 10

All index data from Bloomberg, with tickers supplied in the table below. Regression was run controlling for both monthly time-period and sector-level effects, constraining the monthly time effects to sum to 0.0 so any average futures premium would be present in the sector-level roll return estimates (rather than the time effect estimates). Regression weights do not sum to 100% due to the lack of Soy Oil futures in our dataset, which are a current constituent of the Bloomberg Commodity Index but not the S&P GSCI.

	COMMODITY	FUTURES TICKER	SPOT TICKER	REGRESSION WEIGHT	START MONTH
Energy	WTI Crude	SPGSCLTR	SPGSCL	17.0%	Feb 1987
	Brent Crude	SPGSBRTR	SPGSBR	13.0%	Feb 1999
	Natural Gas	SPGSNGTR	SPGSNG	5.7%	Feb 1994
	Gas Oil	SPGSGOTR	SPGSGO	4.1%	Feb 1999
	Unleaded Gasoline	SPGSHUTR	SPGSHU	3.4%	Feb 1988
	Heating Oil	SPGSHOTR	SPGSHO	3.3%	Feb 1983
Industrial Metals	Copper	SPGSICTR	SPGSIC	5.9%	Feb 1977
	Aluminum	SPGSIATR	LMAHDY	4.1%	Feb 1991
	Zinc	SPGSIZTR	SPGSIZ	2.2%	Feb 1991
	Nickel	SPGSIKTR	SPGSIK	1.7%	Feb 1993
	Lead	SPGSILTR	SPGSIL	0.4%	Feb 1995
	Tin	SPGSISTR	SPGSIS	0.0%	Feb 1995
Precious Metals	Gold	SPGSGCTR	SPGSGC	8.0%	Feb 1978
	Silver	SPGSSITR	SPGSSI	2.2%	Feb 1973
	Palladium	SPGSPATR	SPGSPA	0.0%	Feb 1995
	Platinum	SPGSPLTR	SPGSPL	0.0%	Feb 1984
Grains	Corn	SPGSCNTR	SPGSCN	5.1%	Jan 1970
	Soybeans	SPGSSOTR	SPGSSO	4.6%	Jan 1970
	Wheat	SPGSWHTR	SPGSWH	4.2%	Jan 1970
	Soybean Meal	SPGSSMT	SPGSSM	1.7%	Feb 1995
Softs	Sugar	SPGSSBTR	SPGSSB	2.3%	Feb 1973
	Coffee	SPGSKCTR	SPGSKC	1.6%	Feb 1981
	Cotton	SPGSCTTR	SPGSCT	1.4%	Feb 1977
	Cocoa	SPGSCCTR	SPGSCC	0.2%	Feb 1984
Livestock	Live Cattle	SPGSLCTR	SPGSLC	3.8%	Jan 1970
	Lean Hogs	SPGSLHTR	SPGSLH	1.9%	Feb 1976
	Feeder Cattle	SPGSFCTR	SPGSFC	0.6%	Feb 2002

Exhibits 11, 12, and 13

All returns from Bloomberg, with proxy indices used for EM assets in the table below. EM Equity returns were calculated net of underlying currency movements (i.e. on a currency-hedged basis) by subtracting daily returns to the MSCI Emerging Markets Currency Index.

Risk decomposition in Exhibit 11 uses an exponentially-weighted regression with 20-year half-life to determine the variance contribution of the Equity, Interest Rates, Credit, Energy, and Industrial Metals factors to the monthly returns of each of the EM asset class proxies. The contributions to variance of the Energy and Industrial Metals factors are combined as contribution from "Commodities" in the chart.

Exhibits 12 and 13 are based on daily returns to the EM asset proxies residualized on a rolling basis to the Equity, Interest Rates, Credit, Energy, and Industrial Metals factors.

	INDEX NAME	BLOOMBERG TICKER
EM Credit	JP Morgan EMBI Global	JPEIGLBL
EM Currency	MSCI Emerging Markets Currency Index	MXEF0CX0
EM Equity (after Jan 2001)	MSCI Emerging Markets Net Return	M1EF
EM Equity (before Jan 2001)	MSCI Emerging Markets Gross Return	M2EF

Exhibit 15

Monthly returns data from Global Financial Data. Symbols for indices used in this analysis: _DJCBTD for US corporate credit returns, GFWLDM for global equity price returns, SYWLDYM for global equity dividend yields, and TRWLDGVM for global government bonds.

Bibliography

Acharya, V.V., L.A. Lochstoer, and T. Ramadorai (2013). "Limits to Arbitrage and Hedging: Evidence from Commodity Markets", Journal of Financial Economics, 109(2), 441-465.

Adrian, T., R. Crump, B. Mills, and E. Moench (2014). "Treasury Term Premia: 1961-Present", Liberty Street Economics blog post. Accessible at https://libertystreeteconomics.newyorkfed.org/2014/05/treasury-term-premia-1961-present.html

Ang, A. (2014). Asset Management. Oxford University Press, New York.

Asvanunt, A. and S. Richardson (2017). "The Credit Risk Premium", Journal of Fixed Income, 26(3), 6-24.

Brennan, M.J. (1958). "The Supply of Storage", American Economic Review, 48(1), 50-72.

Campbell, J.Y., K.S. Medeiros, and L.M. Viceira (2010). "Global Currency Hedging", Journal of Finance, 65(1), 87-121.

Campbell, J.Y., C.E. Pflueger, and L.M. Viceira (2018). "Macroeconomic Drivers of Bond and Equity Risks", Harvard Business School Finance Working Paper.

Damodaran, A. (2018). "Equity Risk Premiums (ERP): Determinants, Estimation and Implications – The 2018 Edition", Working paper, New York University.

D.E. Shaw Group Market Insights (2019). "Positively Negative: Stock-Bond Correlation and Its Implicaitons for Investors", DE Shaw Group White Paper. Accessible at https://www.deshaw.com/assets/articles/DESCO_Market_Insights_20190208_1.pdf

Dimson, E., P. Marsh, and M. Staunton (2016). "Long-Term Asset Returns" in D. Chambers & E. Dimson (Eds.) Financial Market History: Reflections on the Past for Investors Today (pp. 2-26). CFA Institute Research Foundation.

Dimson, E., P. Marsh, and M. Staunton (2018). Credit Suisse Global Investment Returns Yearbook 2018. Credit Suisse Group AG, Zurich.

Driscoll, J.C. and A.C. Kraay (1998). "Consistent Covariance Estimation with Spatially Dependent Panel Data", Review of Economics and Statistics, 80(4), 549-560.

Duncombe, G. and B. Kay (2018). "Introducing the Two Sigma Factor Lens", Two Sigma White Paper. Accessible at https://www.twosigma.com/insights/article/introducing-the-two-sigma-factor-lens/

Fama, E.F. and K.R. French (1988). "Business Cycles and the Behavior of Metals Prices", Journal of Finance, 43(5), 1075-1093.

Giesecke, K., F.A. Longstaff, S. Schaefer, and I. Strebulaev (2011). "Corporate Bond Default Risk: A 150-Year Perspective", Journal of Financial Economics, 102(2), 233-250.

Gorton, G.B., F. Hayashi, and K.G. Rouwenhorst (2012). "The Fundamentals of Commodity Futures Returns", Review of Finance, 17(1), 35-105.

Guvanen, F. (2009). "A Parsimonious Macroeconomic Model for Asset Pricing", Econometrica, 77(6), 1711-1750.

Hong, H. and M. Yogo (2012). "What Does Futures Market Interest Tell Us about the Macroeconomy and Asset Prices?", Journal of Financial Economics, 105(3), 473-490.

Ilmanen, A. (2011). Expected Returns: An Investor's Guide to Harvesting Market Rewards, Wiley, Chichester, UK.

Jorda, O., K. Knoll, D. Kuvshinov, M. Schularick, and A.M. Taylor (2018). "The Rate of Return on Everything, 1870-2015", CESifo Working Paper.

Kaldor, N. (1939). "Speculation and Economic Stability", Review of Economic Studies, 7(1), 1-27.

Leibowitz, M.L., A. Bova, and P.B. Hammond (2010). The Endowment Model of Investing: Risk, Return, and Diversification. Wiley Finance, Hoboken, NJ.

Malloy, C.J., T.J. Moskowitz, and A. Vissing-Jørgensen (2009). "Long-Run Stockholder Consumption Risk and Asset Returns", Journal of Finance, 64(6), 2427-2479.

Mouakhar, T. and M. Roberge (2010). "The Optimal Approach to Futures Contract Roll in Commodity Portfolios", Journal of Alternative Investments, 12(3), 51-60.

Norges Bank Investment Management (2016). "The Equity Risk Premium", NBIM Discussion Note. Accessible at https://www.nbim.no/en/publications/discussion-notes/2016/the-equity-risk-premium/

Piketty, T. and G. Zucman (2014). “Capital Is Back: Wealth-Income Ratios in Rich Countries 1700-2010”, Quarterly Journal of Economics, 129(3), 1255-1310. Online Appendix at http://piketty.pse.ens.fr/fr/capitalisback accessed on Nov 27, 2018.

Rallis, G., J. Miffre, and A-M. Fuertes (2012). “Strategic and Tactical Roles of Enhanced Commodity Indices”, Journal of Futures Markets, 33(10), 965-992.

Rappoport, P. and N. Nottebohm (2012). “Improving on Risk Parity: Hedging Forecast Uncertainty”, J.P. Morgan Asset Management research report.

Ross, S.A. (1976a). “The Arbitrage Theory of Capital Asset Pricing”, Journal of Economic Theory, 13(3), 341-360.

Ross, S.A. (1976b). “Risk, Return and Arbitrage” in Risk Return in Finance. I. Friend and J.L. Bicksler (Eds.). Ballinger, Cambridge, Mass.

Saret, J.N. (2014). “Emerging and Frontier Market Equities as Asset Class(es)”, Two Sigma White Paper. Accessible at https://www.twosigma.com/insights/article/emerging-frontier-markets-asset-classes/

Working, H. (1949). “The Theory of Price of Storage”, American Economic Review, 39(6), 1254-1262.

Wright, J.H. (2011). “Term Premia and Inflation Uncertainty: Empirical Data from an International Panel Dataset”, American Economic Review, 101(4), 1514-1534.

IMPORTANT DISCLAIMER AND DISCLOSURE INFORMATION

This document is distributed for informational and educational purposes only and is not an offer to sell or the solicitation of an offer to buy any securities or other instruments. The information contained herein is not intended to provide, and should not be relied upon for, investment, accounting, legal, or tax advice. Exposure to risk factors is not a guarantee of increased performance or of decreased risk. Diversification and asset allocation may not protect against market risk or loss of principal. Past performance does not guarantee future results. A factor lens is merely a resource supplementing investment, financial, tax, accounting, legal and other advisors and due diligence. This document does not purport to advise you personally concerning the nature, potential, value or suitability of any particular sector, geographic region, security, portfolio of securities, transaction, investment strategy or other matter. No consideration has been given to the specific investment needs or risk-tolerances of any recipient. The recipient is reminded that an investment in any security is subject to a number of risks including the risk of a total loss of capital, and that discussion herein does not contain a list or description of relevant risks. The recipient hereof should make an independent investigation of the information described herein, including consulting its own tax, legal, accounting and other advisors about the matters discussed herein. This document does not constitute any form of invitation or inducement by Two Sigma to engage in investment activity.

The views expressed herein are not necessarily the views of Two Sigma Investments, LP or any of its affiliates (collectively, "Two Sigma"). Such views (i) may be historic or forward-looking in nature, (ii) may reflect significant assumptions and subjective judgments of the authors of this document, and (iii) are subject to change without notice. Two Sigma may use, construct, and maintain factors and methodologies differently than as discussed in this document, including when making investment decisions on behalf of or when trading for purposes of investable products.

No representation is made as to the accuracy of any information contained herein and the use of information from third-party sources in no way implies an endorsement of the source of such information or its validity. Any forward-looking statements are inherently subject to significant business, economic and competitive uncertainties and contingencies, many of which are beyond our control. In addition, these forward-looking statements are subject to assumptions with respect to future business strategies and decisions that are subject to change. Factors which could cause actual results to differ materially from those anticipated include, but are not limited to: competitive and general business, economic, market and political conditions in the United States and abroad from those expected; changes in the legal, regulatory and legislative environments in the markets in which Two Sigma operates; and the ability of management to effectively implement certain strategies. Words like "believe," "expect," "anticipate," "promise," "plan," and other expressions or words of similar meanings, as well as future or conditional verbs such as "will," "would," "should," "could," or "may" are generally intended to identify forward-looking statements. Two Sigma makes no representations, express or implied, regarding the accuracy or completeness of this information, and the recipient accepts all risks in relying on this document for any purpose whatsoever.

Important Notice to Venn Subscribers: This document is not an offer to, or solicitation of, any subscribers of Venn or otherwise for the provision of investment management, advisory or any other services. Nothing in this document should be considered as a representation of how any research or methodology is used on Venn. Importantly, Venn's use of the Two Sigma Factor Lens and other research or methodologies will differ from the subject matter discussed herein. This document is subject to change without notice.

This document is being furnished to the recipient on a confidential basis and is not intended for public use or distribution. By accepting this document, the recipient agrees to keep confidential the existence of this document and the information contained herein. The recipient should not disclose, reproduce, distribute or otherwise make available the existence of and/or all or any portion of the information contained herein to any other person (other than its employees, officers and advisors on a need-to-know basis, whom the recipient will cause to keep the information confidential) without Two Sigma's prior written consent. This document shall remain the property of Two Sigma and Two Sigma reserves the right to require the return of this document at any time. Some of the images, logos or other material used herein may be protected by copyright and/or trademark. If so, such copyrights and/or trademarks are most likely owned by the entity that created the material and are used purely for identification and comment as fair use under international copyright and/or trademark laws. Use of such image, copyright or trademark does not imply any association with such organization (or endorsement of such organization) by Two Sigma, nor vice versa.