Crises Do Rhyme: Identifying Shocks to Systematic Risk in Times of Crisis*

Jacob Boudoukh Yukun Liu Tobias Moskowitz Matthew Richardson

March 26, 2026

Abstract

Crises such as Covid-19, the GFC, and Brexit reveal a common pattern in global asset returns. Risk escalates through spikes in principal components (PCs) of return variation. Despite different origins, their PC structures converge in strikingly similar ways: PC1 captures broad volatility surges, while PC2 transforms into a "crisis factor" driving cross-asset correlations. These dynamics have asset-pricing and investment implications. Conventional pre-emptive hedges fail under such regime shifts, implying inevitable crisis losses. However, a novel hedge using fixed crisis weights, calibrated ex ante from prior crises, mitigates crisis risk. The price of such risk is estimated at 2–3% per year.

*Jacob Boudoukh (Arison School of Business, RUNI), Yukun Liu (Simon Business School, University of Rochester), Tobias Moskowitz (Yale University, NBER, and AQR Capital Management), Matthew Richardson (Stern School of Business, NYU, NBER and consultant to AQR Capital Management). We would like to thank Brendan Hoffman and Ronen Israel for comments and discussions, as well as seminar participants at Baruch University, Bath University, Carnegie Mellon University, Hong Kong University, New York University, Shanghai Jiaotong University, University of Miami, University of Notre Dame, and Yale University as well as conference participants at CFA New Zealand, FSU Suntrust conference, and GSU CEAR conference for helpful comments. AQR Capital Management is a global investment management firm, which may or may not apply similar investment techniques or methods of analysis as described herein. The views expressed here are those of the authors and not necessarily those of AQR. We would like to thank Niels Gormsen, Brendan Hoffman, Ronen Israel, and Sophia Li for comments and discussions. We would also like to thank Isabella Louise Cortes Malixi for helpful research assistance.

1 Introduction

Over the past 15 years, the global economy has endured a series of major shocks, including the Global Financial Crisis (GFC), the Eurozone sovereign debt crisis, Brexit, and the Covid-19 pandemic. These episodes imposed substantial economic costs, often triggering sharp contractions in GDP. Their substantial impact naturally raises a central question: how do such crises affect global aggregate asset returns? The violent swings in equity, fixed income, and exchange rate prices are immediately visible, yet little is known about the underlying global factor structure of returns that drives these dynamics.

This paper addresses that gap. We provide a detailed study of how the global factor structure of asset returns evolves during crises over the 2007 to 2020 sample period, and how these shifts impact investment portfolios. We find that asset return distributions are shaped by two forces during crises: (i) large increases in portfolio variance associated with a relatively stable first factor, and (ii) sudden changes in the loadings of global financial assets on a second factor. Although crises in our sample arise from varying origins, we find that the structure of this second principal component is surprisingly consistent across crises—crises may not repeat, but they do rhyme. This regularity makes it possible to construct a "crisis portfolio" using ex ante information that offsets losses during these events, that conventional hedging methods fail to capture.

Our analysis uses high-frequency (aggregated to minute-by-minute) data to estimate the correlation matrix of 26 aggregate asset returns: global equity indices, 8 fixed-income futures, and 6 major exchange rates. The high frequency data provides more precision to the factor structure than using daily data, and allows us to detect shocks to the covariance structure with real-time accuracy and precision. Using non-overlapping intervals of the high frequency data, we document the evolution of the factor structure of volatility-adjusted returns, focusing on the first two principal components (PC1 and PC2) drawn from the joint distribution of global asset returns. The data reveal extreme shocks to the factor structure that ex post are identified as the Lehman Brothers bankruptcy in September 2008, the Brexit

vote in June 2016, and the onset of the Covid-19 pandemic in March 2020.

We propose a novel metric to evaluate the magnitude of PC structural changes and what is driving them. PC1's composition remains stable throughout the sample period, capturing only shocks to aggregate volatility. PC2, however, changes its composition dramatically. The structure of PC2 during these extreme times seems to capture a systemic "crisis factor" of global markets, one that re-emerges whenever global stress is acute. The shocks to PC2's structure are remarkably consistent across crises, despite each crisis originating from a different economic source.

Establishing any causal relationship between the factor structure and crisis shocks is difficult since these events are typically not independent of economic or financial conditions at the time. However, the exogeneity of the Covid-19 shock, offers a rare opportunity to uncover the fundamental (causal) relation between factor returns and crises. Exploiting the fact that Covid-19 was an exogenous shock to the global economic and financial system, we use news about Covid across countries over time to link to changes in the factor structure. We extract Covid-related news via natural language processing of media text and relate it to PC1 and PC2 returns. Covid news shocks are associated strongly with factor structure returns. Moreover, changes in the factor structure driving these shocks to returns are driven primarily by PC2 factor composition changes, further identifying it as a crisis factor. Finally, as further identification, we use leading epidemiological forecasts of Covid deaths at the time to relate factor returns by country to that country's short-term forecast error of Covid-19 deaths, exploiting the variation in countries' responses to Covid-19. We find that the component of a country's equity index exposed to PC2, forecasts negatively Covid-19 deaths *after* controlling for leading epidemiological death forecasts, providing compelling evidence of PC2 as a crisis factor.

The factor structures across different crises converge toward a remarkably similar form, despite each crisis being so different in their origins. The sudden alignment of PC structures at the onset of crises underscores a fundamental regularity in how global markets process

systematic shocks. Armed with these insights, we examine the implications for investment portfolios. First, portfolio return correlations based on PC2 loadings are far higher when using crisis-period weights than when using non-crisis weights. As an application, we examine the implications of these findings for global risk parity portfolios (volatility-adjusted equal-weight holdings of the 26 assets). In crisis periods, such portfolios suffer extreme volatility and large drawdowns, with diversification benefits collapsing substantially. Traditional pre-emptive hedging strategies fail, underscoring the inevitability and unpredictability of crisis shocks. However, recognizing the commonality of PC2 across crises allows us to design a novel hedge using fixed "crisis weights" derived from prior crises in history. As an illustration, during Covid-19, the daily volatility and maximum drawdown of the unhedged portfolio is 0.56% and −8.49%, respectively, compared to just 0.12% and −0.74% for the ex ante hedged portfolio. Similarly, during the GFC, volatility and drawdown fell from 0.79% and −8.66% to 0.22% and −1.25% under the hedge.

The factor structure and crisis hedge portfolio have important asset pricing implications as well. The risk parity portfolio unhedged delivers an annualized Sharpe ratio of −0.39 during Covid and −3.0 during the GFC. However, the hedged risk parity portfolio delivers a Sharpe ratio of 0.72 during Covid and 0.90 during the GFC. This hedge comes at a cost, however. The unconditional Sharpe ratio of the unhedged risk parity portfolio over the full sample period is 0.62 per year, whereas the unconditional Sharpe ratio of the crisis-hedged risk parity portfolio is 0.46 per year. Hence, investors give up 0.16 units of Sharpe ratio, or about 2 to 3% per year, to mitigate the crisis drawdowns in the sample. These results highlight the conditional asset pricing dynamics associated with financial crises and what the market will pay to hedge them.

More broadly, our work on identifying high frequency changes to the covariance structure of global asset returns may be a useful area to investigate new asset price dynamics and implications. While our focus is on abrupt shifts in that structure, that happen to coincide with crises, more subtle or gradual changes may prove useful in further understanding asset

price dynamics. We leave these issues for future work.

The rest of the paper is organized as follows. Section 2 describes the global factor structure of asset returns, including our high frequency data, methodology, and estimation. Section 3 focuses on the factor structure during crises, documenting significant changes to the factor structure and identifying the causal link to crises. Section 4 examines the asset pricing and investment implications of these factor structure shifts. Section 5 concludes.

2 The Global Factor Structure of Asset Returns

In standard factor models, an individual asset's variance contribution can shift for several reasons: changes in the variance-covariance structure of the factors, shifts in the asset's factor loadings, or changes in its idiosyncratic risk. A less commonly discussed possibility is that the factor itself evolves over time, with its underlying weights on base assets changing. Little empirical work on how aggregate traded factors change through time exists, however, in part because we do not observe these aggregate factors directly. A widely used approach is to apply statistical techniques such as principal component analysis (PCA) to extract weighted combinations of assets (factors) that capture return co-movement.

A central challenge to this approach is estimation. With many assets and relatively short time series, the statistical properties of PCA can be poor. This challenge intensifies when attempting to estimate PCAs repeatedly over time to track evolving dynamics. Recent advances by Aït-Sahalia and Xiu(2017; 2019) address this issue by developing asymptotic justifications for PCA based on high-frequency data of continuous-time asset price processes.

There is also a growing consensus that a low-dimensional factor structure underpins global asset markets.[1] Several theoretical frameworks—spanning bonds, equities, and exchange rates—seek to explain these common dynamics (*e.g.*, Hau and Rey (2006); Bansal and

[1]For example, within asset classes, see Fama and French (2012) for stock returns, Dahlquist and Hasseltoft (2013) and Jotikasthira et al. (2015) for bond returns, Lustig et al. (2011) and Lustig and Verdelhan (2019) for currency returns, Liu et al. (2022) for cryptocurrency returns, and, across asset classes, see Asness et al., He et al. (2017), Gospodinov and Robotti (2021), and Miranda-Agrippino and Rey (2020, 2022).

Shaliastovich (2013); He and Krishnamurthy (2013); Farhi and Gabaix (2016); Lustig et al. (2019); Koijen and Yogo (2020)). Building on this literature, our paper employs high-frequency global data to measure how asset return variances and covariances evolve through time. Although we analyze the period from 2007 onward, we place particular emphasis on sudden shifts in the factor structure.

2.1 Data

We use minute-by-minute prices of twelve futures on country equity indices (Australia, Brazil, Germany, Canada, Spain, France, India, Italy, Japan, Switzerland, UK and US), eight fixed income futures on 10-year government bonds (Australia, Germany, Canada, France, Italy, Japan, UK and US), and six spot exchange rates (Australia, Euro, Canada, Japan, Switzerland, UK in terms of $US). These 26 assets cover the vast majority (over 70%) of the market values of global equity, fixed income, and exchange rate markets, with the notable exception of China.

An important element of high frequency estimation is the continuous availability of data throughout the day. Unlike Aït-Sahalia and Xiu (2017), who have access only to the 6 ½ US equity trading hours, global futures contracts and spot exchange rates have many more trading hours. Figure 1 documents the coverage of the 26 assets across the 24-hour period over the recent Covid-19 period (January 1, 2020 to June 30, 2020) as an illustration.[2]

We calculate the percentage of 5-minute returns that are available during the hour. (See Section 2.2 for a description of the exact methodology used.) As Figure 1 shows, many assets have nearly 24-hour coverage. For example, all foreign exchange rate markets have coverage for 23 hours with many equity futures markets available for at least 20 hours (Australia, Germany, U.K. and U.S.). Fixed income futures markets have the least amount of coverage with Italy and France at 11 hours and the U.K. at 10 hours. Other fixed income markets have better coverage (*i.e.*, Australia (22 hours), Canada (15 hours), Germany (20 hours),

[2]We use data going back to 2007, which has similar coverage to the Covid-19 sample.

Japan (19 hours) and U.S. (23 hours)).

2.2 Methodology

We perform a PCA using the correlation matrix of the 26 global asset returns. Because equities are much more volatile than either fixed income or exchange rate assets, we analyze the correlation matrix as opposed to the variance-covariance matrix. This is equivalent to weighting each asset return by its volatility, which, in a portfolio sense, matches risk parity-based investment strategies that we examine later. While asset pricing is governed by covariances, as opposed to correlations, the purpose of this paper is to better understand the global cross correlation structure of assets and how it changes over time and in particular during crises. However, our analysis also focuses on measuring changing variances and covariances.[3]

Because news is fast moving, the PCA is applied to biweekly high frequency data from January 1, 2007 through June 30, 2020.[4] In terms of estimating the biweekly correlation matrix, the goal is to use as much of the 24 hours of high frequency data available to obtain a more efficient estimate of asset return correlations than using daily data (*e.g.*, see Barndorff-Nielsen and Shephard (2004) and Andersen et al. (2006; 2013)). As Figure 1 shows, there are many coincident hours for a number of assets, especially exchange rates and equity futures in the larger markets.

We use the following procedure for any two assets, A and B. First, we select the one minute intervals where both A and B have available prices. If no transactions occur for

[3]One drawback is that the PCs (generated from the correlation matrix) applied to variance-covariance space are no longer uncorrelated. We adjust for this issue when estimating asset return betas.

[4]Our data analysis shows that the five-minute high frequency intervals, biweekly horizon and cross-section of assets have sufficient data to back out reasonable estimates of the correlation matrix. Shorter horizons or longer intervals add additional sampling error. In estimating their weekly covariance matrix, Aït-Sahalia and Xiu (2017) estimate the integrated (*i.e.*, average) eigenvalues of the "spot" variance-covariance matrix (using one-minute intervals over 2.5 hours), while we choose to measure the eigenvalues of the integrated variance-covariance matrix (over two weeks). The underlying asymptotic theory implies that in estimating over any fixed period, there are an infinite number of subperiods in continuous-time, allowing the researcher to estimate the "spot" covariance matrix and capture the instantaneous dynamics, and then integrate over these to capture the 2-week covariance matrix. The practical issue is that in finite time each subperiod has a large amount of estimation error due to its short interval and multiple assets.

either A or B within the one minute interval, we move to two minutes; and if no transactions occur for either A or B within two minutes, we move to three minutes, and so on. Second, following standard conventions, we re-index the returns of A and B to five minute intervals in order to minimize the impact of market microstructure. Finally, we estimate the covariance between the return on asset A and B by summing up their five-minute (or longer) cross-terms with a lead-lag adjustment for microstructure nonsyncronicity and bid-ask bias. Similarly, we estimate each asset return's variance by summing over their five-minute squared terms plus a microstructure adjustment for first-order serial autocorrelation. We also impose a liquidity requirement in which we only sample from the hours when more than 50% of the data is available. This requirement necessarily means we ignore periods of market closures, so that if either asset A or B is closed over a given period, we do not count their comovement from close-to-open. However, for each asset's own variance, we incorporate both trading and nontrading returns.

Specifically, for day T, we use five-minute bars (re-sampled from one-minute bars to avoid microstructure noise) from the previous day $T-14$ (16:00:01 EST) to previous day $T-1$ (16:00:00 EST). Let r_t denote the log return from t minus 5 minutes to t. We compute the high frequency realized variance as of day T as follows:

$$\sigma^2_{T,adjusted} = \left(\sum_{t=T-14}^{T-1} r_t^2 + 2 r_t r_{t-1} \right). \tag{1}$$

The first term is the integrated variance over the lookback period; the second term adjusts for the one-period (5 minute) autocorrelation, so $\sigma_{T,adjusted}$ is an unbiased estimator of the previous biweekly realized volatility. When five-minute data is not available (for example, overnight hours), we use the gap returns. For example, if there is no trading between 3:00am and 3:20am, we use the twenty-minute return for r_t.

We compute the high frequency covariance as of day T as follows:

$$cov_{T,adjusted} = (\sum_{t=T-14}^{T-1} r_{A,t}r_{B,t} + r_{A,t-1}r_{B,t} + r_{A,t}r_{B,t-1}). \quad (2)$$

The first term is the integrated covariance over the lookback period. The second and third terms adjust for one-period (5 minute) lead-lag autocovariances. The five-minute returns are only sampled over the periods where both securities have five-minute returns available. For example, if security A trades from 9am to 6pm while security B trades from 1pm to 9pm, then the covariance will be computed using the data from 1pm to 6pm. In estimating the correlation between returns A and B, the covariance is divided by their integrated standard deviations over the coincident period. However, when working with the full variance-covariance matrix of returns, the correlation matrix is scaled up by each assets' 24-hour biweekly return variance, *i.e.*, including both trading and nontrading hours.

In the next section, we perform a PCA on the biweekly high frequency correlation matrix of the returns on the 26 global assets, constructed from the pairwise correlation estimates described above. The implicit assumption is that the correlations over the coincident periods are representative of those over 24-hour periods. For some assets, such as exchange rates and equity markets in the U.S. and Europe, the coincident periods cover much of the 24 hours.

Figure 2 documents the mean and standard deviation of biweekly correlations of global assets within and across equities, fixed income, and exchange rates using high frequency versus daily data over the period 2007 to 2020. Both within and across asset classes, the stan- dard deviation of the correlation estimates is much smaller using high frequency than daily data. For example, the standard deviation of the biweekly correlation using daily versus high frequency data is 0.280 versus 0.145 within equities, 0.307 versus 0.172 for fixed income, and 0.339 versus 0.194 for exchange rates. Similar results hold for cross-correlations across asset classes. Second, for the most part, the mean correlation estimates are similar across and within asset classes using either high frequency or daily data, providing support

for the coincident period assumption for assets in the high frequency setting. For example, comparing daily versus high frequency correlation estimates, the average correlation is 0.550 versus 0.568 for equity, 0.461 versus 0.500 for fixed income and 0.406 versus 0.385 for exchange rates. Hence, we get similar point estimates (no bias) using high frequency data, but much more precision. The added precision should help us detect significant changes to the correlation structure, which we investigate in the next section.

2.3 Estimating the Factor Structure of Returns

We estimate the high-frequency correlation matrix of the 26 asset returns and corresponding standard deviations using two-week periods from 2007 to 2020. For each of these biweekly periods, we extract the principal components from the correlation matrix of the 26 asset returns, resulting in a time-series of 364 data points of the principal components. Figure 3 graphs the average of the corresponding scree plots, *i.e.*, the variance explained by the factor against the principal component number.[5] The first two principal components capture 58% of the variance, with PC1 capturing 43% and PC2 capturing 15%. Since no other component explains more than 5%, we choose to focus on the first two principal components.

Figure 4 plots factor loadings for each of the three asset classes for PC1 and PC2, defined respectively as $w_{i1,t}$ and $w_{i2,t}$ for each asset in their respective asset class over the sample period. We graph quarterly estimated loadings and highlight three crisis periods by a dashed red line. Equities, fixed income, and exchange rates are graphed separately for each PC. As is standard, the factor loadings are normalized so that $\sum_i w_{i1,t}^2 = 1$ and $\sum_i w_{i2,t}^2 = 1$. In order to make the sign of the factor loadings comparable across equities, fixed income, and exchange rates through time, the cross-sectional average of equity is assumed to have a positive loading for PC1 and that of fixed income is assumed to have a positive loading for PC2.[6]

[5]The scree plot also graphs the amount of variance explained during crisis and pre-crisis periods.

[6]For PC2, there is an instance we keep the average fixed income loading to be negative because PC2 clearly flips sign. This can be seen by the fact that requiring average equity loadings to be positive results in large average absolute differences in factor loadings from the adjunct periods, while not requiring positive equity loadings leads to markedly smaller differences.

The first principal component of global asset returns represents a long position in global equities financed by a short position in global bonds and is long carry in currencies—a levered exposure to global equity risk plus carry. Although PC1's loadings trend modestly toward smaller magnitudes over time, they remain remarkably persistent despite each period being estimated from an independent sample of asset returns. This persistence suggests the presence of a fundamental and stable underlying factor structure, rather than mere sampling noise or measurement error.

The exchange rate loadings in PC1 are also persistent but display greater variability across time. The equal-weighted average of these currencies represents a "dollar basket"—a portfolio long foreign currencies and short the U.S. dollar. Figure 4 illustrates that variations in the dollar basket's PC1 loading are strongly linked to interest rate differentials between the U.S. and the average of the six foreign economies in the sample.[7] The correlation between the PC1 loadings and the interest rate spread is –0.32 (statistically significant at the 5% level), indicating that when U.S. interest rates are relatively low, the dollar basket is positively exposed to the global carry trade, and vice versa. Thus, fluctuations in the exchange rate loadings primarily reflect shifts in the profitability of the carry strategy. Consistent with this interpretation, the yen—long considered a funding currency—stands apart from other currencies, exhibiting persistently negative loadings throughout the sample period.

To explore this relationship further, we reorganize the dollar-basket currencies into a "carry basket," long in the three currencies with the highest interest rates and short in the three with the lowest. Figure 5 Panel A shows that the average PC1 loadings for this carry portfolio are substantially more stable than those for the dollar basket. Moreover, the PC1 loadings on the carry factor remain positive across the entire sample, with no discernible trend. In untabulated results, the correlation between PC1's carry loadings and the underlying interest rate spread is 0.42 (statistically significant at conventional levels).

Taken together, these findings reveal that the dollar basket plays a pivotal and evolving

[7]We use each countries' 3-month interbank lending rate from the FRED dataset available from the St. Louis Fed.

role in the global correlation structure of asset returns—and that this role depends crucially on the dollar's carry status. This result complements the work of Lustig et al. (2014), who document a positive premium for the "dollar carry trade"—being long the dollar when it is a carry currency, and short it when it is a funding currency—that earns returns above and beyond those of a conventional carry strategy. Figure 5 Panel B reinforces this insight by showing that, conditional on the dollar's role as a carry or funding currency, the loadings on the dollar basket differ markedly. Comparing Panels A and B in Figure 5 suggests that carry dynamics, rather than pure exchange rate effects, are the dominant determinant of PC1's FX loadings. Overall, PC1 loads positively on equities and carry exposure, and negatively on government bonds and the yen—clearly characterizing PC1 as a global "risk-on" factor.

As shown in Figure 4, the swings in PC2 loadings are considerably larger than those of PC1. While equities exhibit small positive loadings, the dominant stylized fact is that fixed income carries large positive loadings—the mirror image of PC1. In other words, PC2 represents a long exposure to the global bond market. Both equity and bond weights in PC2 display more pronounced dynamics than in PC1, and this pattern is especially striking for exchange rates. The dollar basket begins the sample with large positive loadings during the onset of the Global Financial Crisis (GFC), then trends steadily downward, reaching negative territory by mid-2016. These changes occur gradually rather than abruptly, suggesting a slow-moving evolution in the global risk structure.

Similar to PC1, PC2's variation is strongly correlated with the relative level of U.S. interest rates—that is, the dollar's status as a carry or funding currency. The correlation between the PC2 loadings on the dollar basket and the U.S.–foreign interest rate differential is –0.57, statistically significant at the 5% level. Thus, whether the U.S. dollar occupies a high- or low-yield position plays a key role in shaping PC2's FX loadings.

When the PC2 dollar-basket loadings are reorganized into a formal carry strategy, Figure 5 Panel A shows that the carry loadings are close to zero, exhibiting no clear trend and only modest variation over time. The correlation between these carry loadings and the underlying

interest rate spread is just 0.17 and statistically insignificant. At first glance, this suggests that PC2's FX exposures are not tied directly to conventional carry dynamics—they appear to respond to the dollar's role itself rather than to interest differentials per se. Figure 5 Panel B supports this interpretation: conditional on the interest rate spread, the dollar-basket loadings are negative but turn sharply positive around crises. This pattern implies that the dollar, independent of its carry status, exerts a distinctive influence on PC2.

Across all three asset classes—equities, bonds, and exchange rates—PC2's loadings also display sharp spikes in their time-series patterns. These surges coincide closely with three major global crises of the past 14 years: the Global Financial Crisis, Brexit, and the Covid-19 pandemic.[8]

Figure 4 marks these episodes with vertical dashed lines corresponding to September 15, 2008 (Lehman Brothers' bankruptcy), June 23, 2016 (the Brexit referendum), and March 11, 2020 (the World Health Organization's pandemic declaration). Several key observations emerge. Although both PC1 and PC2 exhibit significant movement during these episodes, PC2's loadings change much more dramatically. The onset of a crisis appears to exert a stronger structural effect on the second PC factor. For example, during crises, PC2's fixed-income loadings experience large temporary spikes, declining sharply from their typically positive levels. This pattern suggests that the nature of government bonds temporarily shifts in crises: while demand for safe assets rises, sovereign debt itself may become less "safe" as investors reassess the fiscal health of governments. Likewise, PC2's dollar-basket loadings spike upward during crises, often reversing from negative to positive. A positive loading indicates that assets co-moving with the PC are effectively short the dollar basket—consistent with the dollar acting as a global hedge when systemic risk intensifies. We explore these interesting correlation dynamics more in depth during crises.

[8]There are several large spikes outside of these three periods. For example, one of the larger spikes occurs for fixed income PC1 loadings over several quarters starting in May 2013. This period includes the so-called "taper tantrum" in which the Fed announced it would be ending its quantitative easing program. Of some note, Figure 3 also includes the Euro debt crisis which arose from the GFC but contained two periods of severe disruption, the spring of 2011 and summer of 2012.

3 Global Factor Structure in Times of Crisis

Section 2.3 documents how the factor structure shifts though time, including times of crisis. The largest changes in the factor structure appear to be related to three events: the global pandemic, the global financial crisis and, to a lesser extent, Brexit. In this section, we document three main findings. First, we confirm this finding and then show that the first two PCs change in different ways. Changes to the first PC are mostly driven by changes in the underlying asset return volatility, while the second PC is driven by changes in factor loadings, *i.e.*, shifts in the correlations of asset returns. Second, we investigate how the PC structure changes around the onset of each crisis. Most interesting, the changes seem to be only temporary in nature. Third, we provide a detailed analysis of one of these crises, namely the Covid-19 global pandemic. Because the Covid-19 shocks were exogenous, we are able to identify causation and gain a deeper understanding of how the timing and characteristics of the changes to the first two PCs are related to crisis shocks. In particular, we relate shocks to the two PCs to news about Covid-19 during the crisis period.

3.1 Changes in the Factor Structure through Time

To provide evidence for how significant the factor structure shift is over the sample period, Figure 6 graphs the absolute value of differences in PC1 and PC2 through time. Each day, we take the absolute value of that quarter's PC weights minus last quarter's weights times that day's asset return, *i.e.*,

$$|PC1_{qt-(q-1)t,k}| \equiv |\sum_i \left(w_{i1qt} - w_{i1(q-1)t}\right) R_{iqt+k}| \tag{3}$$

and

$$|PC2_{qt-(q-1)t,k}| \equiv |\sum_i \left(w_{i2qt} - w_{i2(q-1)t}\right) R_{iqt+k}| \tag{4}$$

where R_{iqt+k} is the return on the k-th day of the i-th asset in quarter q_t, and w_{i2qt}, for example, is the current quarter loading of PC2 on asset i. The returns of both $PC1_t$ and $PC2_t$ are standardized to have zero mean and unit variance over the full sample period.

Two observations stand out from Figure 6. First, and most important, for both PC1 and PC2, the three largest changes occur during the GFC and the Covid-19 pandemic (marked by the dashed lines in the figure). The Covid-19 period represents the largest change in both PC1 and PC2 over the entire sample period. Moreover, the absolute difference is a magnitude higher than other periods. Second, the absolute value of the PC differences is a much higher scale for PC2 than PC1, in particular, approximately four times higher. To the extent the PC1 and PC2 returns are standardized, this suggests that the change in the factor structure is much greater for PC2.

If the weights do not change from quarter to quarter, *i.e.*, $w_{i1qt} \approx w_{i1(q-1)t}$ then $|PCJ_{qt-(q-1)t,k}| \approx 0$. However, even if $w_{i1qt} \neq w_{i1(q-1)t}$, then $|PCJ_{qt-(q-1)t,k}| \gg 0$ only if asset returns R_{iqt+k} are highly variable. In other words, both conditions must be true. While the PC weights are constructed from the correlation matrix of global asset returns, the variance of PC1 and PC2 also depend on the variance levels of these asset returns. To better understand the driver of the changing PC factor return structure, Figure 7 graphs the high frequency realized biweekly volatility of PC1 and PC2. While Figures 6 and 7 support the premise that both PC1 and PC2 loadings change the most and their return volatilities reach the highest level during crisis periods, there is an important difference. Figure 6 shows that the scale of the changes in PC2 loadings are four times those of PC1, while Figure 7 shows that return volatility shock is four times higher for PC1 than for PC2.

Across the 364 biweeks over the sample period, the top ten biweeks for PC1 and PC2 in terms of largest absolute value of PC differences are dominated by the crisis events. For example, three of the four largest changes for PC1 and two of the largest changes for PC2 are during the Covid-19 period. In terms of magnitude, the largest change in PC1 and PC2 relative to the third largest change (GFC for PC1 and Brexit for PC2) is, respectively, 2.55

and 1.77 times. Compared to the average change, these changes are an order of magnitude more extreme. For PC1, five of the largest 10 changes are the GFC, while, for PC2, four of the largest changes are the GFC and three of the others are Brexit. Figures 6 and 7 also confirm the peak crisis events occur during the biweek of the Lehman bankruptcy, the Brexit vote, and the WHO declaration of the global Covid-19 pandemic.

Table 1 provides formal tests of factor structure changes. The first test compares $|\boldsymbol{PCJ}_{qt-(q-1)t,k}|$ during the GFC, Brexit, and Covid-19, versus all other periods for both PC1 and PC2. The differences for PC1 are 0.145, 0.020, and 0.099 for the three crises versus 0.640, 0.567, and 0.516 for PC2 with corresponding t-statistics of 15.46, 2.04, and 9.40 for PC1 and 17.34, 15.16, and 12.45 for PC2. While these results imply changes in the factor structure for both PCs, they include the variability of returns. The second test simply calculates the average difference in weights from quarter to quarter, $|\sum_i \left(w_{iJqt} - w_{iJ(q-1)t}\right)|$, for both PC1 and PC2 in the three crises versus the rest of the sample. The differences are also significantly positive, being 0.017 (t-stat 1.85) , 0.016 (t-stat 2.29) and 0.018 (t-stat 3.45) for PC1 and 0.054 (t-stat 3.72), 0.150 (t-stat 7.77), and 0.033 (t-stat 2.29) for PC2. While the factor structure changes during crises are significant for both PC1 and PC2, the magnitudes are considerably higher for PC2. For example, during the Covid-19 crisis, the estimated differences of 0.018 and 0.033 for PC1 and PC2 represent 9.0% and 16.5% changes, respectively, in factor weights.

3.2 Factor Loadings During Crises

Figure 8 plots the principal-component loadings of all 26 global assets—equities, fixed income, and exchange rates—before and during three major crises: Covid-19 (Panel A), GFC (Panel B), and Brexit (Panel C) at a granular level. The figure graphs loadings during the average of the biweeks in the quarter pre-crisis and then through the 13 biweeks throughout the crisis periods. The crisis periods are defined as January 1, 2020 to June 30, 2020 for Covid, July 1, 2008 to December 31, 2008 for the GFC, and March 25, 2016 to September 22, 2016 for Brexit. These crisis periods are chosen based on the data described in the

previous Subsection 3.1. Loadings for equities, fixed income, and exchange rates are graphed respectively in blue, red, and green, with the x-axis and y-axis representing PC1 and PC2, respectively.

The first plot in each panel, where each panel represents a unique crisis, are the factor loadings for PC1 and PC2 in the quarter immediately prior to each crisis. PC1 loads consistently positively on equities, negatively on fixed income, and slightly negative to neutral on exchange rates. These weights are the normal PC1 weights observed over the full sample. While the stability of PC1's loadings prior to each crisis is evident, PC2 exhibits mixed signs across asset classes that vary prior to each crisis. However, as each crisis unfolds, the factor structure, particularly for PC2, changes significantly. PC1's factor loadings largely remain stable throughout each crisis, consistent with our previous results. However, for PC2, the factor loadings shift and converge toward a common configuration in each criss—equities become largely neutral, while fixed income and exchange rates acquire strongly positive loadings. Hence, PC1 looks essentially the same before, during, and after crises. The only thing that changes for PC1 is that its volatility (as shown in the last subsection) spikes during crises. PC2, however, changes throughout the sample, where its configuration differs prior to each crisis, but converges to a singular configuration during every crisis, despite their different origins and causes. The common convergence in each crisis occurs despite PC2 having different factor structure prior to each crisis. This convergence implies that global crises generate a similar reordering of cross-asset relationships. Looking at the loadings at the height of each crisis (highlighted in bold), the dollar's behavior is central to this process—PC2 loadings on the dollar basket rise sharply, signifying a flight to dollar safety, while the yen continues to bear negative loadings, consistent with its role as a funding and safe-haven currency. These findings are economically intuitive.

During each crisis, the shifts in factor loadings are abrupt but short-lived. Within several biweeks, the crisis structure begins to dissipate, with PC1 and PC2 gradually reverting toward their pre-crisis configurations. This pattern suggests that PC2 becomes a "crisis

factor" during these episodes that is a transient but recurrent phenomenon. When systemic stress erupts, the evidence points to a common, temporary transformation in global risk dynamics that defines the onset and resolution of financial turmoil.

Table 2 documents the average PC1 and PC2 loadings for equity, fixed income, and currencies during the period prior to each of the crisis periods and for the six biweek periods starting from the defining events of each crisis. Across the three crisis periods, the average loadings of PC1 for equities are 0.257, 0.246, and 0.270 for Covid-19, Brexit and the GFC, respectively, during the crisis itself, compared to 0.243, 0.249, and 0.270 in the quarter prior to each crisis. Hence, the equity loadings on PC1 do not change pre- versus during crises and do not change across crises, either. For fixed income, the story is similar. Crisis loadings are -0.154, -0.126, and -0.177, and pre-crisis loadings are -0.191, -0.193, and -0.187. For FX loadings, however, the weighting on PC1 changes, albeit in a small way. FX loadings on PC1 during each crisis are 0.064, 0.029 and 0.075, but before the crisis are -0.020, -0.011, and -0.0112.

Turning to PC2, the loadings differ substantively between pre-crisis and crisis periods. For example, pre-crisis PC2 loadings for equities are 0.157, 0.106, and 0.116 for Covid, Brexit, and the GFC, respectively. During the crisis, the equity PC2 loadings are lower at 0.082, 0.073, and 0.009. For fixed income, pre-crisis PC2 loadings are 0.277, 0.271, and 0.067, but during the crisis are 0.259, 0.244, and 0.226. For currencies, the PC2 loadings pre-crisis are -0.055, -0.050, and 0.337, but during crises are 0.168, 0.091, and 0.262. During the crisis periods, there is greater convergence to similar loadings than there is pre-crisis for PC2. During each crisis, equities converge to zero/neutral loadings, while fixed income and exchange rate loadings increase significantly and are positive. It is an interesting finding that, even with different factor structures pre-crisis, there is common movement towards a similar factor structure at the height of each crisis, despite these crises originating from different sources and being quite different in nature. We explore and exploit this finding further in Section 4.

As a final comment, factor loading changes during the crisis periods could result from changing PC weights (constructed from the correlation matrix of global asset returns) or changing asset return variances. Figure 7 showed that the volatility of PC1 and PC2 reached their highest level during the GFC, Brexit, and Covid-19 periods. How much of this variance increase is due to changing PC weights versus changing asset return variances? To address this question, we decompose the changes of the realized biweekly variances during crisis periods into two components as follows:

$$\begin{aligned} var\left(PCi_{t+1}\right) - var\left(PCi_t\right) &= w_{t+1}^{\prime i}\Sigma_{t+1}w_{t+1}^{i} - w_t^{\prime i}\Sigma_t w_t^{i} \\ &= \left(w_{t+1}^{\prime i}\Sigma_{t+1}w_{t+1}^{i} - w_t^{\prime i}\Sigma_t w_t^{i}\right) + \left(w_t^{\prime i}\Sigma_{t+1}w_t^{i} - w_t^{\prime i}\Sigma_{t+1}w_t^{i}\right) \\ &= \underbrace{w_t^{\prime i}\left(\Sigma_{t+1} - \Sigma_t\right)w_t^{i}}_{\text{Portion due to changing var-cov}} + \underbrace{\left(w_{t+1}^{\prime i}\Sigma_{t+1}w_{t+1}^{i} - w_t^{\prime i}\Sigma_{t+1}w_t^{i}\right)}_{\text{Portion due to changing PC weights}} \end{aligned} \tag{5}$$

The first component captures the changes due to changing variances and covariances, and the second component captures the changes due to changing PC weights. Figure 9 plots the fraction of total variance changes for PC1 and PC2 due to each component for the two biweeks prior to and four biweeks during the defining event of each crisis. Panel A1 reports the percentage contribution to the change in variance of PC1 from just the variance-covariance changes and Panel A2 reports the contribution to PC1's changes in variance coming from changing weights. For all three crises, the increase in PC1 variance comes almost entirely from a change in the underlying asset returns' variance-covariance matrix, Σ_t, and not changes in the loadings of PC1. Panels B1 and B2 repeat the same exercise for PC2. In contrast, the changing variance of PC2 over the three crisis periods is much more a mix of the two components. Changing weights play a much larger role, and especially for the GFC and Covid-19 periods, changing weights (Δw_t) is just as important as changing variance-covariances ($\Delta\Sigma_t$).

3.3 Identifying Systematic Risk Shocks Using Covid-19

In the previous Sections 3.1.and 3.2, we show dramatic changes in the global financial structure, especially for PC2, that coincide with the bankruptcy of Lehman Brothers in September 2008, the Brexit vote in June 2016, and the declaration of the global pandemic in March 2020. While it is tempting to interpret causality of these events affecting abrupt risk structure changes, since these events were so impactful, the challenge is that crises are identified ex post and are often not independent of prevailing economic and financial conditions at the time. To better identify the causal link between systematic risk changes and crises, we use the unique features of the Covid-19 pandemic. Covid-19 presents a unique opportunity for empirically identifying the causal link between a crisis and systematic risk since the Covid-19 pandemic represents a true exogenous (and unanticipated) shock to the global and economic financial system. We use Covid-19 news events over time and across countries to identify shocks to that crisis and link them to changes in the factor structure of returns to identify a causal relationship to systematic risk changes.

3.3.1 Using Covid-19 News to Identify Systematic Risk Changes

We measure Covid-19 news using data from Ravenpack, who employ natural language processing (NLP) algorithms to measure specific events and sentiment generated from unstructured text contained in a variety of media sources and company reports. Specifically, starting in January 2020, in response to the coronavirus outbreak, Ravenpack produced a set of related Covid-19 news indices for every country, as well as a worldwide index. Specifically, they report on a daily basis (i) Coronavirus Panic Index, (ii) Coronavirus Media Hype Index, and (iii) Coronavirus Media Coverage Index for each country. These measures capture the amount of news related to Covid-19 by country over time. Appendix A2 provides a detailed description of the data.

We first test whether the factor structure of returns is related to these news shocks by examining whether changes in Covid news are contemporaneously related to the first

two PCs of returns. Table 3 Panel A reports results from regressions of weekly changes in each of the three Ravenpack Coronavirus worldwide indices on contemporaneous PC1 and PC2 returns. The first four columns report the regression results for PC1. The coefficients are negative and statistically significant for each Covid news index (t-statistics are -2.76, -3.03, -2.96, and -2.98) with R^2s of 48.0%, 45.7%, 31.1%, and 44.1%, respectively for the Panic, Media Hype, Media Coverage and the composite index. These results indicate that increases in Covid news are contemporaneously correlated with decreases in PC1's returns, suggesting a negative economic shock to the first aggregate risk factor (which looks like an equity market factor in excess of government bonds plus carry exposure to currencies). The next four columns report results for PC2 returns, which also shows a decrease in returns with Covid news that is statistically significant. (t-statistics are -2.97, -2.91, -3.08, and -3.02 with R^2s of 31.1%, 41.8%, 50.4%, and 44.7%). Hence, both factor returns (PC1 and PC2) are highly correlated with changes in Covid-19 news during the Covid-19 period. Moreover, the relationship between Covid-19 news and PC1 and PC2 are negative because the sign of the loadings on PC1 and PC2 are restricted to have positive average equity loadings (for consistency through time). The last four columns report regression results including both PC1 and PC2 in the regression. The coefficients remain negative on both, but PC2 seems to be the stronger effect, having a more significant relation to Covid news.[9]

Panel B of Table 3 investigates how the two PCs are related to Covid news by examining whether the PC's changing variance during Covid is related to either the changing variance-covariance matrix of returns or changing PC weights as we did in Section 3.2. To examine this question, we calculate the difference between the PC returns using the PC weights of the Covid-19 period versus the weights from the "normal" period (*i.e.*, prior quarter), and denote the difference $PC - PC^{norm}$. The returns of the underlying assets are measured daily and applied to each of the PC weights calculated during Covid versus normal times. There is clear evidence that the change in PC weights from the "normal" to Covid-19 period are

[9]Recall that PC1 and PC2 are correlated because the PC factor structure is generated from the correlation matrix as opposed to the variance-covariance matrix.

correlated with Covid-19 news shocks, for both PC1 and PC2. These results on the difference in weights appear to be even stronger statistically than those in Panel A. Turning to the last four columns that report the multivariate regressions of changes in the Covid-19 indices on $PC1 - PC1^{norm}$ and $PC2 - PC2^{norm}$, $PC1 - PC1^{norm}$ is largely drowned out, suggesting that a significant part of the changes in the structure of the second PC factor is explained by Covid-19 news. These results indicate that PC2 is more related to Covid-19 news and more representative of a "Covid-19 factor," which further supports that PC2 is characterized as a crisis factor, consistent with our earlier results.

3.3.2 Systematic Risk Changes and Death Forecast Errors

In the wake of the Covid-19 pandemic, leading epidemiological groups within research institutions began to produce forecast models of infection rates and deaths both within and across countries. One of the better-known research groups, the MRC Centre of Global Infectious Disease Analysis and Abdul Latif Jameel Institute for Disease and Emergency Analytics at the Imperial College of London provided one-week ahead forecasts of cumulated deaths for up to 69 countries as reported by the European Centre for Disease Prevention and Control (ECDC). A detailed description of the data is provided in Appendix A3. As a further test of the crisis causing systematic risk structure changes, we relate each country's specific factor return component to short-term forecast errors of each country's Covid-19 deaths. We define a country's systematic factor return component for equity from time t to $t+1$, as $F_{ij;t,t+1}$. Decomposing the country's equity return as follows:

$$R_{ij;t,t+1} = \alpha_{ij} + \underbrace{\beta_{ij1;t-1,t+1} PC1_{t,t+1} + \beta_{ij2;t-1,t+1} PC2_{t,t+1}}_{F_{ij;t,t+1}\text{, systematic factor return component}} + \epsilon_{ij;t,t+1} \tag{6}$$

where the corresponding returns for country i's equity are measured weekly, and the betas, $\beta_{ij1;t-1,t+1}$ and $\beta_{ij2;t-1,t+1}$, are biweekly estimates using the high frequency covariance matrix.

We run a panel regression of the log growth rate in weekly cumulative Covid-19 deaths for

each country against the Imperial College of London's ensemble forecast model along with the country's lagged systematic factor return for equity, as well as its idiosyncratic equity return. Specifically,

$$\ln\left(\frac{D_{it+1}}{D_{it}}\right) = \mu + \gamma_1 \ln\left(\frac{E[D_{it}^{t+1}]}{D_{it}}\right) + \gamma_2 F_{ij;t-k,t-k+1} + \eta_{it,t+1} \tag{7}$$

where D_{it} is the cumulative Covid-19 deaths at time t for country i, $E[D_{it}^{t+1}]$ is the forecast at time t for the cumulative Covid-19 death total at time $t+1$, $F_{ij;t-k,t-k+1}$ is the country's systematic factor return component for country j from time $t-k$ to $t-k+1$, defined in equation (6) and time t is measured in weeks.

Table 4 Panel A reports the results from this regression. First, and not surprisingly because much of the relevant information for next week's Covid-19 death totals is known, (*i.e.*, from reported cases from weeks back and fatality rates), the first column reports that the coefficient on $\ln\left(\frac{E[D_{it}^{t+1}]}{D_{it}}\right)$ is close to one (0.86) with a high R^2 (83%). Second, and more interesting, is that when the model forecast and the country factor return are included together in the regression (second column), the coefficient on the one-week lagged factor return is -2.32 and statistically significant (t-stat $= -2.35$) and increases the R^2 to 84%. The factor return component on a country's equity index helps forecast next week's Covid-19 deaths above and beyond leading epidemiological forecasts. The fact that this coefficient is negative is consistent with a country's factor return component being inversely related to bad Covid-19 news—low return portends higher than expected Covid-19 death totals. Moreover, the coefficient on the epidemiological model forecast declines from 0.86 to 0.84, suggesting the factor returns capture some of the forecast. The third column reports regression results replacing the factor component of the country's equity return with its idiosyncratic return component. The idiosyncratic equity return is not significant, providing further support for the PCs capturing the crisis news.

In addition, factor returns weeks ago may contain information about Covid-19 infection rates weeks earlier which in turn determine the following week's Covid-19 death rates. Con-

sistent with this idea, the coefficients on the two- and three-week lagged equity factor return (in columns 4 and 5) are still negative, albeit smaller with weaker significance levels (-2.01 and -1.57 with t-statistics of -2.79 and -2.12).

These results provide further evidence of a definitive link between crisis shocks and factor returns. The fact that each country's own Covid-19 news impacts its specific equity factor return component is compelling and helps further identify the link between this crisis and systematic risk by exploiting cross-sectional variation across countries in addition to the time-series variation we showed previously.

4 Asset Pricing and Investment Implications

We examine the asset pricing and investment implications of the risk structure changes we find, particularly during crises. We begin by looking at the PC return series we construct and compare them during, versus out of, crises. We then examine hedging implications for a risk parity portfolio applied across asset classes. Risk parity portfolios became a popular investment vehicle more than a decade ago to provide more resilient portfolios that can withstand macroeconomic shocks, including economic crises of the sort studied in this paper.[10] The portfolio is constructed as an ex ante equal risk (standard deviation) weighted portfolio to balance risk contribution across asset classes. We implement this portfolio using the high frequency variance estimates of the assets from our data as described earlier. Such a portfolio is consistent with our modeling of the high frequency correlation matrix, which would represent the PCs of the covariance matrix under the assumption that all the global asset returns have been rescaled to have equal variance. We investigate here whether global risk parity portfolios can be hedged during crisis periods and the price of that hedge, which provides asset pricing implications for dynamic risk and the role of crises.

[10]See, for example, Maillard et al. (2010), Qian (2011), and Asness et al. (2012) for a theoretical discussion of risk parity portfolios and their statistical properties.

4.1 Summary statistics on PC factors and risk parity portfolios

Table 5 Panel A documents the mean and volatility of the annualized returns of PC1 and PC2 and the risk parity portfolio during crisis and non-crisis periods. The crisis period is defined as the six biweeks starting with the defining event of each crisis. We calculate PC returns using three different weighting schemes. The first one is the time-varying weights of PCs throughout the sample period, estimated ex ante from the high frequency correlation matrix as described previously. This is the standard time-varying PC weights. We also look at two other PC weighting schemes: (i) fixed weights calculated as the average PC weights during crisis periods, and (ii) fixed weights calculated as the average PC weights during the pre-crisis or "normal" periods. Since the PC weights change significantly during crises, especially for PC2, we want to see how different the returns are for the changing PCs over the sample period. In order to make the PCs comparable, we normalize all PC returns to have an annualized standard deviation of 15% over the entire sample.

The first row of Panel A reports the mean, standard deviation, and Sharpe ratio of PC1 during crises and outside of crises. Averaged across the three crisis periods, PC1 has an annualized mean return of −44.1% with an annualized volatility of 36.7% (Sharpe ratio of −1.23). As shown previously, PC1 is a highly leveraged position long in equities, short fixed income, and, to a lesser extent, long carry. All of these asset positions lost significant value during crisis episodes. Outside of crises, PC1 has a positive mean of 1.0% with 12.8% volatility.[11]

The second row of Panel A reports the returns for PC2, whose mean is −9.2% with a standard deviation of 27.3% (Sharpe ratio of −0.37) during crises. Recall that PC2 on average offsetts long positions in equity and fixed income with holdings in $-based exchange rates flipping signs throughout the period. Outside of crises, PC2 has a 12.8% mean return with 14.1% volatility (Sharpe ratio of 0.84).

[11]The relatively poor performance of PC1 is likely less to do with an ex ante risk-return tradeoff and instead by the well-documented ex post higher Sharpe ratios of fixed income governments over this particular sample period, driven by low inflation and low economic growth rates.

Recall from Section 3 that PC2's factor loadings are quite similar across all the crises and effectively becomes a "crisis" factor during these times. We therefore use fixed asset weights to construct a time series of PC2 returns using the average crisis weights, $\boldsymbol{PC2}^{crisis}$. We do the same for PC1, and report their returns across crisis and non-crisis periods in the third and fourth rows of Panel A of Table 5. For PC1, using crisis weights, the return is −44.9% during crises, which is very similar to the returns of PC1 using the dynamic weights, and the −1.23 Sharpe ratio is identical. Outside of crises, the PC1 returns are 1.8% with a standard deviation of 12.8%, which is the same as the PC1 returns with dynamic weights.

The fourth row reports the returns of PC2 using fixed crisis weights. Here, the mean return is −3.2% during crises (Sharpe ratio = −0.15) and 14.1% outside of crises (Sharpe ratio = 0.94), both of which are significantly different from the PC2 returns we get from dynamic weights.

We repeat this exercise using fixed weights from "normal" times, which we define as the average PC weights in the quarter prior to each crisis, where we take the average weight across all three pre-crisis periods and apply them to the 26 asset returns over the entire sample period. Again, PC1 using normal weights looks the same as PC1 using dynamic weights, which looks the same as PC1 using crisis-only weights, producing similar returns and risk to each other during crisis periods and non-crisis periods. The composition of PC1 doesn't change much. PC2, however, exhibits different returns using normal weights than using crisis weights. These results indicate that PC1 largely stays the same in terms of its composition and weights, while PC2 looks quite different during versus outside of crises.

Panel B of Table 5 reports the correlation matrix of the PC return series over the entire sample period using the fixed crisis and normal weights for each crisis separately. For PC1, return series across the six different fixed weights have 0.99 correlation, whether using crisis or normal weights. In other words, consistent with the results in Section 3, there is little difference in the factor structure in- and out-of-crises for PC1. In contrast, the correlation of the six PC2 return series ranges from 0.22 (*i.e.*, pre-crisis Brexit and pre-crisis Covid-19)

to 0.98 (Brexit and Covid-19). Moreover, all the crisis PC2 returns are highly correlated (ranging from 0.77 to 0.98), whereas the pre-crisis PC2 returns are much less so and the cross-correlation of pre-crisis and crisis PC2 returns is also much lower. Hence, the loadings of PC2 during crises are similar, but outside of crises they are quite different, consistent with the results in Section 3.

4.2 Hedging a Risk Parity Portfolio During Crises

What do these findings imply for the risk mitigation properties of investment portfolios? Given the sharp rise in systematic risk during crises, can such risk be effectively hedged using the identified principal components (PCs), and what is the price of such a hedge? To address these questions, we evaluate whether investors can hedge the systematic risk of a global risk parity portfolio during the Covid-19, Global Financial Crisis (GFC), and Brexit episodes, and the cost of this hedge.

Specifically, we compare three hedging approaches that use either PC1 alone or both PC1 and PC2 as hedging instruments. The first approach applies a traditional regression-based hedge. Here, we estimate betas using daily return data of the risk parity portfolio and the daily returns to the PC factor returns using rolling regressions over the prior quarter at each point in time. The estimated coefficients from this regression are then used as hedge ratios. We call this approach the "traditional" hedging approach. The second approach employs a high-frequency hedge, where we extract the implied betas directly from the most recent high-frequency variance–covariance matrix of global asset returns. This approach may provide even more timely estimates of the betas and hedge ratios, but may also be less stable. As shown in previous results (Figures 6 and 8 and Table 1), factor loadings shift abruptly at the onset of crises, so using the high frequency estimates may be able to capture those swift changes better than rolling regressions. However, even the most recently estimated high frequency betas may fail to capture the rapidly changing dynamics. Hence, as a final approach, we also construct a "crisis hedge" by using the average weights derived from other

crisis periods to form the PCs and then hold them fixed for a particular crisis. In our case, for example, to calculate the hedge for Covid ex ante, we use the average crisis weights from the GFC and Brexit crisis periods to compute PC1 and PC2, and then back out the implied betas of the risk parity portfolio on those crisis PCs from the most recent high frequency covariance matrix. We repeat this procedure for each crisis by using the average weights from the other two crises not being hedged to form the PCs so that the analysis is fully out of sample. The high-frequency beta between the risk parity portfolio and the crisis PC is then used to implement the hedge. This approach uses information from the covariance structure of previous crises to estimate the hedge for a current crisis. Doing so will only be effective if crises contain similar risk dynamics. Moreover, this approach will only do better than the rolling regression or high frequency approach if in addition to crises behaving similarly, their risk dynamics are also difficult to forecast statistically with recent data—that is, the risk structure of crises changes too abruptly. By comparing these three approaches, we assess how predictable risk dynamics are during a crisis.

Table 6 reports the results from the tests. Panel A reports results that just try to hedge PC1. We examine each crisis (Covid, GFC, and Brexit) separately. The first row of each subpanel reports the unhedged risk parity portfolio. The table reports the daily mean, standard deviation, and minimum return of the risk parity portfolio, along with its annualized Sharpe ratio and maximum drawdown (DD) and length of DD of the portfolio during each crisis. The unhedged risk parity portfolio has a −1 bp daily return during Covid with a 56 bp daily standard deviation, a minimum daily return of −2.62% and a Sharpe ratio (annualized) of −0.38. The max DD during Covid was −8.5% lasting 14 days. During the GFC, the numbers are even more dramatic, with a −14 bp daily average return, 79 bps of volatility, and a −3.0 Sharpe ratio (annualized) with a max DD of −8.7% lasting 38 days. The returns during Covid and the GFC are much more negative than during Brexit, where the risk parity portfolio only has a max DD of −1.12% lasting four days.

The remaining rows of each subpanel of Panel A report results for each of the three

hedging approaches: traditional regression-based, high frequency, and fixed crisis weights. Here, in Panel A we are only hedging PC1. As the table shows, hedging PC1 provides only modest improvement. For Covid-19, the volatility (and maximum drawdown) declines to 0.50% (−7.71%) with the traditional hedge, 0.39% (−7.11%) with the high-frequency hedge, and 0.35% (−5.97%) using fixed crisis weights. The traditional regression hedge performs worst, likely due to stale betas and measurement error. The high-frequency hedge is more responsive but cannot offset the sudden structural shifts in factor relationships that accompany crises. The fixed crisis-weight hedge performs better. However, the results are modest because, as shown in Section 3, PC1's loadings remain largely stable through time, even during crises. Dramatic shifts in factor structure are driven by PC2.

Panel B of Table 6 extends the analysis to hedging the global risk parity portfolio using *both* PC1 and PC2 simultaneously. Incorporating PC2 should help if PC2 acts like a "crisis factor," as suggested by our previou results. During Covid 19, portfolio volatility (and maximum drawdown) declines from 0.56% (−8.65%) to 0.41% (−6.25%) using the traditional hedge, to 0.22% (−3.20%) using the high frequency hedge, and to 0.12% (−0.74%) using fixed crisis weights. Clearly, hedging PC2 helps. Moreover, the high frequency hedge improves significantly on the traditional rolling regression approach, but the best results come from using the fixed crisis weights. In fact, the fixed crisis weight hedged portfolio is the only one that delivers positive returns (9 bp) over the Covid period with a negligible drawdown. It seems to almost completely immunize the risk parity portfolio from the Covid crisis and this is an ex ante hedge using only the crisis weights from the other two crises (GFC and Brexit) to hedge Covid. Repeating this analysis for the GFC and Brexit, we get similar results. In the GFC, the risk parity portfolio volatility (and maximum drawdown) decline from 0.79% (−8.7%) to 061% (−6.3%) using the traditional hedge, to 0.35% (−3.3%) using the high frequency hedge, and to 0.22% (−1.25%) using the fixed crisis weights, which are the average PC weights of from the other two crisis (Brexit and Covid in this case). Again, the crisis weight hedged portfolio is the only one that has postive returns (1.6%) over the GFC

and seems to be completely hedged during this crisis. For Brexit, the risk parity portfolio did not fare as badly, but nevertheless, we get similar results that hedging improves from traditional to high frequency to using the fixed crisis weights, steadily reducing both the volatility and maximum drawdown of the risk parity portfolio. The fixed crisis weights give by far the best results.

These results indicate that using fixed crisis loadings—constructed from other crises out of sample—effectively eliminates the majority of systematic risk for each crisis in our sample. Across the crises, volatility is reduced by 80% and drawdowns virtually disappear. For investors seeking protection against unforeseen systematic shocks, this crisis-weight approach offers a promising and practical risk mitigation tool. Taken together with the preceding evidence, these findings further reinforce the interpretation of PC2 as a distinct "crisis factor" that dominates global asset return dynamics during periods of financial stress, and using other crises to capture this factor can be a powerful hedging tool for future crises.

Given the efficacy of hedging PC1 and PC2 during crises, an interesting question is what do the fixed crisis weights look like? Table A1 in the appendix reports the weights of each asset during each crisis in PC1 and PC2. Consistent with previous results, PC1 is largely long equity and short fixed income and long carry, while PC2 is long fixed income and somewhat short the dollar. Combining these two, PC1 + PC2 during a crisis looks a lot like a risk parity portfolio. Hence, hedging against this reduces a risk parity portfolio's exposure to the risk during a crisis. Panel B shows that outside of a crisis, PC1 remains the same, but PC2 looks very different and varies considerably over different pre-crisis periods. Hence, PC1 + PC2 does not look like a risk parity portfolio outside of crises. These results come purely from the correlation structure of asset returns and its dynamics.

Since PC2 is the crisis factor that changes its composition during a crisis, we investigate what asset returns correlate with it by using the fixed crisis weights for PC2. Figure 10 plots the correlation of an equal-weighted portfolio of equities, fixed income, a dollar basket of currencies (equal-weighted average of all currencies relative to the $US) and returns to

gold futures and the Japanese Yen, which are often touted as "safe" assets during a crisis.[12] The correlations are estimated using the high frequency correlation matrices. As the figure shows, Crisis PC2 has a positive correlation with fixed income, gold and Yen, is negatively correlated to the US dollar, and has a fluctuating albeit hovering around zero correlation with equity. These correlations are consistent with Crisis PC2 capturing the behavior of some of the "safe assets" described in the literature. This finding is also consistent with our results that PC2 becomes a crisis factor during a crisis.

4.3 Asset Pricing Implications

Hedging crises appears to be ex ante feasible using high frequency data (with limited success) and especially using the insights on the changing covariance structure from other crises. But, what are the asset pricing implications of these hedges? Specifically, what is the price of crisis risk and how much would hedging it cost an investor? There is a substantive literature modeling crises as rare events, and then documenting that a significant fraction of risk premia can be explained by these low probability disasters, e.g., Barro (2009), Barro and Jin (2011), Gabaix (2012) and Farhi and Gabaix (2016) to name a few. These papers are theoretically relevant for this paper because they generally look at the risk premia across global assets.

To address the above questions, Table 7 reports the returns of the hedged risk parity portfolios *outside* of crises. The average return of the hedged portfolio outside of crises relative to the average return unhedged outside of crises is a proxy for the price of the hedge—the returns given up during normal times from the hedge.

Panel A reports the results from just using PC1. The returns to the risk parity portfolio hedged to PC1 are no different than the returns without hedging—0.012% versus 0.013%—only 0.1 bp difference. The hedged returns are the same across all hedging methods as

[12]For example, for the role of (i) fixed income, and U.S. government bonds in particular, see Krishnamurthy and Vissing-Jorgensen (2012), (ii) the dollar, see Lustig et al. (2011) and Jiang et al. (2021), (iii) gold, see Baur and Lucey (2010) and Baur and McDermott (2010), and (iv) yen and other funding currencies, see Brunnermeier et al. (2008), Ranaldo and Söderlind (2010) and Hossfeld and MacDonald (2015)

well (traditional, high frequency, and fixed crisis weights). These results are consistent with previous results showing that hedging PC1 did little to avert the adverse effects of each crisis. The lack of difference between the hedged and unhedged returns outside of crises is consistent with markets indicating that these hedges have negligible value.

Panel B reports results for risk parity portfolios hedged to PC1 and PC2. Here, the hedged risk parity portfolios perform significantly worse outside of crises, consistent with them performing much better during crises. This is the cost of the hedge, which on average is about 0.009% per day or about 2.25% per year given up for the hedge. Since the risk parity portfolio unhedged averages 3.25% per year over our sample period, the hedge costs about 69% of the risk parity premium to avoid the crisis drawdowns.

Figure 11 plots the cumulative returns and maximum drawdowns of the unhedged and hedged risk parity portfolios with respect to PC1 and PC2. The first two panels report the cumulative return and drawdowns using the high frequency hedge, and the third and fourth panels report results using the fixed crisis-weight hedge. As the first two figures show, the high frequency hedged risk parity portfolio grows at a much lower rate than the unhedged risk parity portfolio. However, as shown in the second panel, it significantly reduces the drawdowns, especially during crises. The third and fourth panels show even more dramatic results. The hedged risk parity portfolio is almost flat over the sample period, indicating that it gives up almost all of the premium associated with the risk parity portfolio. However, it essentially eliminates the drawdowns of the portfolio, especially during crises. These results are intuitive. The hedge is successful at substantially ameliorating the risk parity portfolio's drawdowns during crises, but the cost of this hedge is significant, as it eliminates most of the unconditional return premium as well. These findings are consistent with the aforementioned asset pricing theory and suggest that markets price the crisis hedge dearly. From an asset pricing perspective, one interpretation of these results is that the risk premium for a crisis factor, that can hedge crises is about 2.5 to 3% per year, with an annual Sharpe ratio of 0.35.

5 Conclusion

Over the past two decades, the global economy has faced a sequence of profound disruptions—the Global Financial Crisis, Brexit, and the Covid-19 pandemic. Each event brought severe economic dislocation and dramatic swings in global asset prices. We study how the underlying factor structure of global asset returns evolves through such crises, using high-frequency data on twenty-six major equity, bond, and currency markets from 2007 to 2020. Examining the first two principal components (PC1 and PC2) of volatility-adjusted returns, we uncover how systemic shocks reshape the relationships between global assets and challenge the stability of conventional diversification frameworks.

Our results reveal that crises share a strikingly consistent pattern despite their differing origins. Portfolio risk surges mainly through two forces: an increase in the variance of a largely stable first factor and abrupt changes in asset loadings on a second, crisis-sensitive factor. This second principal component acts as a "crisis factor" in global markets, emerging whenever systemic stress intensifies, regardless of its origin. We show that the structure of PC2 is surprisingly uniform across crises. Traditional risk parity portfolios, designed to balance risk across asset classes, suffer severe drawdowns during crises, but leveraging the regularity in the factor structure allows us to construct an ex ante (based on past crises) crisis hedge portfolio that effectively eliminates these losses.

Our results highlight a fundamental rhythm in global financial behavior, where crises do not repeat, but their risk structure echoes, revealing a recurring architecture of risk that investors can learn to anticipate and manage. Measuring this risk and forming a crisis factor allows us to estimate the risk premia for crisis exposure, which we find is about 2.5% per year with a Sharpe ratio of 0.35. This premium is similar to that of the risk parity portfolio itself. Hence, one interpretation of these findings is that the risk parity premium represents compensation for risk during these extreme times.

More generally, exploring risk dynamics in global asset returns may reveal more useful insights into asset pricing and the global structure of asset returns that has importance for

investment, risk sharing, and compensation for risk in the economy.

References

Aït-Sahalia Y, Xiu D. 2017. Using principal component analysis to estimate a high dimensional factor model with high-frequency data. *Journal of Econometrics* **201**: 384–399.

Aït-Sahalia Y, Xiu D. 2019. Principal component analysis of high-frequency data. *Journal of the American Statistical Association* **114**: 287–303.

Andersen TG, Bollerslev T, Christoffersen PF, Diebold FX. 2006. Volatility and correlation forecasting. In *Handbook of Economic Forecasting*, volume 1. Elsevier, 777–878.

Andersen TG, Bollerslev T, Christoffersen PF, Diebold FX. 2013. Financial risk measurement for financial risk management. In *Handbook of the Economics of Finance*, volume 2. Elsevier, 1127–1220.

Asness CS, Frazzini A, Pedersen LH. 2012. Leverage aversion and risk parity. *Financial Analysts Journal* **68**: 47–59.

Asness CS, Moskowitz TJ, Pedersen LH. 2013. Value and momentum everywhere. *The Journal of Finance* **68**: 929–985.

Atkeson A. 2020. What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios. Technical report, National Bureau of Economic Research.

Bansal R, Shaliastovich I. 2013. A long-run risks explanation of predictability puzzles in bond and currency markets. *The Review of Financial Studies* **26**: 1–33.

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

Barro RJ. 2009. Rare disasters, asset prices, and welfare costs. *American Economic Review* **99**: 243–264.

Barro RJ, Jin T. 2011. On the size distribution of macroeconomic disasters. *Econometrica* **79**: 1567–1589.

Baur DG, Lucey BM. 2010. Is gold a hedge or a safe haven? an analysis of stocks, bonds and gold. *Financial Review* **45**: 217–229.

Baur DG, McDermott TK. 2010. Is gold a safe haven? international evidence. *Journal of Banking & Finance* **34**: 1886–1898.

Brunnermeier MK, Nagel S, Pedersen LH. 2008. Carry trades and currency crashes. *NBER Macroeconomics Annual* **23**: 313–348.

Dahlquist M, Hasseltoft H. 2013. International bond risk premia. *Journal of International Economics* **90**: 17–32.

Eichenbaum MS, Rebelo S, Trabandt M. 2021. The macroeconomics of epidemics. *The Review of Financial Studies* **34**: 5149–5187.

Fama EF, French KR. 2012. Size, value, and momentum in international stock returns. *Journal of Financial Economics* **105**: 457–472.

Farhi E, Gabaix X. 2016. Rare disasters and exchange rates. *The Quarterly Journal of Economics* **131**: 1–52.

Gabaix X. 2012. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics* **127**: 645–700.

Gospodinov N, Robotti C. 2021. Common pricing across asset classes: Empirical evidence revisited. *Journal of Financial Economics* **140**: 292–324.

Hau H, Rey H. 2006. Exchange rates, equity prices, and capital flows. *The Review of Financial Studies* **19**: 273–317.

He Z, Kelly B, Manela A. 2017. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics* **126**: 1–35.

He Z, Krishnamurthy A. 2013. Intermediary asset pricing. *American Economic Review* **103**: 732–770.

Hossfeld O, MacDonald R. 2015. Carry funding and safe haven currencies: A threshold regression approach. *Journal of International Money and Finance* **59**: 185–202.

Jiang Z, Krishnamurthy A, Lustig H. 2021. Foreign safe asset demand and the dollar exchange rate. *The Journal of Finance* **76**: 1049–1089.

Jones CJ, Philippon T, Venkateswaran V. 2021. Optimal mitigation policies in a pandemic: Social distancing and working from home. *The Review of Financial Studies* **34**: 5188–5223.

Jotikasthira C, Le A, Lundblad C. 2015. Why do term structures in different currencies co-move? *Journal of Financial Economics* **115**: 58–83.

Koijen RS, Yogo M. 2020. Exchange rates and asset prices in a global demand system. Technical report, National Bureau of Economic Research.

Krishnamurthy A, Vissing-Jorgensen A. 2012. The aggregate demand for treasury debt. *Journal of Political Economy* **120**: 233–267.

Liu Y, Tsyvinski A, Wu X. 2022. Common risk factors in cryptocurrency. *The Journal of Finance* **77**: 1133–1177.

Lustig H, Roussanov N, Verdelhan A. 2011. Common risk factors in currency markets. *The Review of Financial Studies* **24**: 3731–3777.

Lustig H, Roussanov N, Verdelhan A. 2014. Countercyclical currency risk premia. *Journal of Financial Economics* **111**: 527–553.

Lustig H, Stathopoulos A, Verdelhan A. 2019. The term structure of currency carry trade risk premia. *American Economic Review* **109**: 4142–4177.

Lustig H, Verdelhan A. 2019. Does incomplete spanning in international financial markets help to explain exchange rates? *American Economic Review* **109**: 2208–2244.

Maillard S, Roncalli T, Teïletche J. 2010. The properties of equally weighted risk contribution portfolios. *Journal of Portfolio Management* **36**: 60.

Miranda-Agrippino S, Rey H. 2020. Us monetary policy and the global financial cycle. *The Review of Economic Studies* **87**: 2754–2776.

Miranda-Agrippino S, Rey H. 2022. The global financial cycle. In *Handbook of International Economics*, volume 6. Elsevier, 1–43.

Qian E. 2011. Risk parity and diversification. *Journal of Investing* **20**: 119.

Ranaldo A, Söderlind P. 2010. Safe haven currencies. *Review of Finance* **14**: 385–407.

Stock JH. 2020. Data gaps and the policy response to the novel coronavirus. Technical report, National Bureau of Economic Research.

Tables and Figures

Table 1: Differences in PC Structure During Crisis Periods

Reported are results for differences in the PC structure for crisis periods (GFC, Brexit and Covid-19) versus normal periods from 2007 to 2020. Panel A reports the absolute value of differences between each PC return during the crisis quarter (q) and during the preceeding quarter immediately before the crisis ($q-1$), $|PCJ_{qt-(q-1)t,k}|$ from equations (3) and (4). Panel B reports the average difference in weights from quarter to quarter, $|\sum_i \left(w_{iJqt} - w_{iJ(q-1)t}\right)|$, for both PC1 and PC2 for each crisis. Statistical tests on the differences are reported with t-statistics (in parentheses) and *, **, *** indicating significance levels at the 10, 5, and 1% levels.

Differences Between Crisis vs. Pre-Crisis Period		
Panel A: Absolute Differences in PC Returns		
	PC1	PC2
Covid	0.099***	0.516***
	(9.40)	(12.45)
	PC1	PC2
GFC	0.145***	0.640***
	(15.46)	(17.34)
	PC1	PC2
Brexit	0.020**	0.567***
	(2.04)	(15.16)
Panel B: Absolute Differences in PC Loadings		
	PC1	PC2
Covid	0.017***	0.034**
	(3.49)	(2.29)
	PC1	PC2
GFC	0.018*	0.053***
	(2.05)	(3.60)
	PC1	PC2
Brexit	0.015**	0.151***
	(2.18)	(7.69)

Table 2: PC1 and PC2 Loadings Prior and During Crises

Reported are the average PC1 and PC2 loadings for EQ, FI, and FX during each crisis and during the "normal" period prior to each crisis (quarter before the defining quarter of each crisis). We define the crises periods for six biweeks starting from the defining events of each crisis, which are biweek 5 for Covid, 19 for the GFC, and 13 for Brexit. The differences between the crisis and pre-crisis periods are reported at the bottom of the table.

		COVID		Brexit		GFC		Average	
		PC1	PC2	PC1	PC2	PC1	PC2	PC1	PC2
Crisis	EQ	0.257	0.082	0.246	0.073	0.270	0.009	0.258	0.054
	FI	-0.154	0.259	-0.126	0.244	-0.177	0.226	-0.153	0.243
	FX	0.064	0.168	0.029	0.091	0.075	0.262	0.056	0.174
Pre-Crisis	EQ	0.243	0.157	0.249	0.106	0.270	0.116	0.254	0.126
	FI	-0.191	0.277	-0.193	0.271	-0.187	0.067	-0.191	0.205
	FX	-0.020	-0.055	-0.011	-0.050	-0.112	0.337	-0.048	0.077
Difference	EQ	0.014	-0.075	-0.003	-0.033	0.000	-0.107	0.004	-0.072
	FI	0.037	-0.018	0.067	-0.027	0.010	0.159	0.038	0.038
	FX	0.084	0.223	0.040	0.141	0.187	0.075	0.104	0.097

Table 3: Factor Returns and Covid-19 News Shocks

Reported are results from regressions of weekly changes in each of the three Ravenpack Coronavirus worldwide indices (*Panic*, *Media Hype*, and *Media Coverage*) plus the composite index on contemporaneous PC1 and PC2 returns (generated using the methodology described in Section 2). Employing a natural language processing (NLP) algorithm to measure specific events and sentiment generated from unstructured text contained in a variety of media sources and company reports, Ravenpack constructs on a daily basis (i) *Coronavirus Panic Index*, (ii) *Coronavirus Media Hype Index*, and (iii) *Coronavirus Media Coverage Index* for each country. All of these measures are meant to capture the amount of news related to Covid-19 (such as references to panic/hysteria, percent of news coverage, *etc...*). Results are reported based on 7-day cumulative overlapping. The standard errors are Newey-West adjusted with 5 lags. *, **, and *** denote significance levels at the 10%, 5%, and 1% levels.

	Panel A: Regressions on PC1 and PC2											
	ΔPanic	ΔHype	ΔCoverage	ΔComposite	ΔPanic	ΔHype	ΔCoverage	ΔComposite	ΔPanic	ΔHype	ΔCoverage	ΔComposite
$PC1$	-0.021***	-0.021***	-0.005***	-0.052***					-0.017*	-0.014**	-0.002	-0.033*
	(-2.762)	(-3.025)	(-2.958)	(-2.980)					(-1.954)	(-2.100)	(-1.327)	(-1.902)
$PC2$					-0.014***	-0.016***	-0.005***	-0.042***	-0.006*	-0.010***	-0.004***	-0.027***
					(-2.969)	(-2.906)	(-3.082)	(-3.024)	(-1.869)	(-2.869)	(-3.401)	(-3.093)
R^2	0.4800	0.4568	0.3108	0.4405	0.3105	0.4182	0.5036	0.4469	0.5203	0.5597	0.5395	0.5669
	Panel B: Regressions on Changes to PC1 and PC2 Pre- Versus During Crises											
	ΔPanic	ΔHype	ΔCoverage	ΔComposite	ΔPanic	ΔHype	ΔCoverage	ΔComposite	ΔPanic	ΔHype	ΔCoverage	ΔComposite
$PC1 - PC1^{norm}$	-0.064***	-0.067***	-0.021***	-0.181***					-0.053*	-0.034	-0.006	-0.006
	(-3.489)	(-2.924)	(-2.762)	(-3.004)					(-1.707)	(-1.152)	(-0.765)	(-0.765)
$PC2 - PC2^{norm}$					-0.014***	-0.017***	-0.006***	-0.045***	-0.003	-0.010*	-0.005***	-0.005***
					(-3.080)	(-3.030)	(-3.103)	(-3.117)	(-0.526)	(-1.810)	(-2.903)	(-2.903)
R^2	0.2709	0.3013	0.3418	0.3292	0.2076	0.3149	0.4269	0.3432	0.2748	0.3419	0.4368	0.4368

Table 4: Epidemiological Forecast Errors and Country Factor Returns

The table reports results from a panel regression of the log growth rate in weekly cumulative Covid-19 deaths for each country against the Imperial College of London's ensemble epidemiological death forecast model along with the country's lagged systematic factor return and idiosyncratic return for equity. Specifically, $\ln\left(\frac{D_{it+1}}{D_{it}}\right) = \mu + \gamma_1 \ln\left(\frac{FD_{it}^{t+1}}{D_{it}}\right) + \gamma_2 F_{ij;t-k,t} + \eta_{it,t+1}$, where D_{it} is the cumulative Covid-19 deaths at time t for country i, FD_{it}^{t+1} is the forecast at time t for the cumulative Covid-19 death total at time $t+1$, $F_{ij;t-k,t}$ is the aforementioned country's systematic factor return component for asset class j from time $t-k$ to t where $k \in (1,2,3)$, and time t is measured in weeks. The systematic and idiosyncratic component of equity returns is calculated from $R_{ij;t,t+1} = \alpha_{ij} + \underbrace{\beta_{ij1;t-1,t+1} PC1_{t,t+1} + \beta_{ij2;t-1,t+1} PC2_{t,t+1}}_{F_{ij;t,t+1},\ \text{systematic factor return component}} + \epsilon_{ij;t,t+1}$. *, **, and *** denote significance at the 10%, 5%, and 1% levels.

	Dependent variable = actual deaths				
	(1)	(2)	(3)	(4)	(5)
Forecasted Deaths	0.859***	0.847***	0.854***	0.803***	0.756***
	(12.15)	(11.69)	(11.71)	(12.13)	(8.51)
Systematic EQ ret_{t-1}		-2.324**			
		(-2.35)			
Idiosyncratic EQ ret_{t-1}			1.388		
			(0.71)		
Systematic EQ ret_{t-2}				-2.008**	
				(-2.79)	
Systematic EQ ret_{t-3}					-1.567*
					(-2.12)
Index FE	Y	Y	Y	Y	Y
R^2	0.830	0.841	0.831	0.839	0.831

Table 5: Correlations of PCs Using Normal Loadings versus Crisis Loadings

Reported are summary statistics of different $PC1$ and $PC2$ portfolios using different weighting schemes. $PC1$ and $PC2$ are the realized $PC1$ and $PC2$ from the dynamic correlation matrix estimated from high frequency data. $PC1^{norm}$ and $PC2^{norm}$ use the average of the loadings of the pre-crises periods (quarter before each crisis) applied to the entire sample period. $PC1^{crisis}$ and $PC2^{crisis}$ use the average of the loadings of each crisis period applied to the entire sample period. We define the crises periods for six biweeks starting from the defining events of each crisis. All PCs are normalized to have an annualized standard deviation of 15% over the full sample. We also report a risk parityy portfolio that is an equal volatility weighted portfolio of all assets, using a rolling volatility estimate for each asset from the high frequency data. Panel A reports the mean, standard deviation, and Sharpe ratio of each return series in and outside of crises. Panel B reports the correlations of each PC with different weights from each crisis across the weighting schemes, for the PC1s and PC2s separately.

Panel A: Summary Statistics (annualized)

	Mean	Stdev	SR	Mean	Stdev	SR
		Crisis			Outside of crisis	
$PC1$	-0.441	0.367	-1.226	0.010	0.128	0.008
$PC2$	-0.092	0.273	-0.370	0.128	0.141	0.844
$PC1^{crisis}$	-0.449	0.368	-1.245	0.018	0.128	0.070
$PC2^{crisis}$	-0.032	0.266	-0.154	0.141	0.141	0.936
$PC1^{norm}$	-0.435	0.361	-1.230	0.008	0.130	-0.008
$PC2^{norm}$	-0.108	0.350	-0.334	0.170	0.131	1.229
Risk parity	-0.359	0.384	-0.958	0.124	0.126	0.917

Panel B: Correlations of PC Using Different Weights from Each Crisis

	COVID	GFC	Brexit	Pre-COVID	Pre-GFC	Pre-Brexit
	Correlation of PC1s with each other					
COVID	1.000					
GFC	0.999	1.000				
Brexit	0.992	0.991	1.000			
Pre-COVID	0.995	0.993	0.990	1.000		
Pre-GFC	0.986	0.984	0.983	0.995	1.000	
Pre-Brexit	0.990	0.986	0.995	0.992	0.983	1.000
	Correlation of PC2s with each other					
COVID	1.000					
GFC	0.863	1.000				
Brexit	0.975	0.770	1.000			
Pre-COVID	0.754	0.343	0.842	1.000		
Pre-GFC	0.644	0.688	0.508	0.310	1.000	
Pre-Brexit	0.766	0.382	0.863	0.978	0.222	1.000

Table 6: Hedging a Global Risk Parity Portfolio During Crises

The table compares three hedging approaches for a risk parity portfolio that use either PC1 alone (Panel A) or both PC1 and PC2 (Panel B) as hedging instruments. The first approach applies a traditional regression-based hedge using daily return data of the risk parity portfolio and the daily returns ("traditional" hedging approach). The second approach employs a high-frequency hedge, where we extract the implied betas directly from the most recent high-frequency variance–covariance matrix of global asset returns. The third approach uses a "crisis hedge" by using the average weights derived from other crisis periods to form the PCs and then hold them fixed for a particular crisis. The high-frequency beta between the risk parity portfolio and the crisis PC is then used to implement the hedge. The mean, standard deviation, and minimum returns as well as the annual Sharpe ratio, maximum drawdown, and duration of the maximum drawdown are reported.

Panel A: Hedging Risk Parity Portfolio with PC1 only

					Drawdown	
	Mean	Stdev.	Min	SR	Max	Duration
COVID						
No Hedging	-0.010%	0.559%	-2.616%	-0.389	-8.487%	14
Trad Hedge	-0.006%	0.495%	-2.450%	-0.301	-7.707%	12
HF Hedge	-0.015%	0.391%	-2.017%	-0.747	-7.111%	12
Crisis Hedge	-0.003%	0.345%	-1.651%	-0.294	-5.968%	12
GFC						
No Hedging	-0.146%	0.790%	-1.938%	-3.002	-8.663%	38
Trad Hedge	-0.130%	0.701%	-1.863%	-3.021	-7.836%	38
HF Hedge	-0.061%	0.448%	-1.555%	-2.282	-5.240%	25
Crisis Hedge	-0.049%	0.417%	-1.445%	-1.995	-4.420%	25
Brexit						
No Hedging	0.045%	0.259%	-1.069%	2.550	-1.123%	4
Trad Hedge	0.031%	0.205%	-0.527%	2.137	-0.751%	8
HF Hedge	0.032%	0.203%	-0.462%	2.237	-0.723%	8
Crisis Hedge	0.025%	0.188%	-0.461%	1.824	-0.640%	5

Panel B: Hedging Risk Parity Portfolio with PC1 and PC2

	Mean	Stdev.	Min	SR	Drawdown Max	Drawdown Duration
COVID						
No Hedging	-0.010%	0.559%	-2.616%	-0.389	-8.487%	14
Trad Hedge	-0.001%	0.413%	-2.052%	-0.169	-6.501%	12
HF Hedge	-0.009%	0.215%	-0.790%	-0.916	-3.244%	12
Crisis Hedge	0.009%	0.123%	-0.521%	0.723	-0.737%	22
GFC						
No Hedging	-0.146%	0.790%	-1.938%	-3.002	-8.663%	38
Trad Hedge	-0.106%	0.606%	-1.38%	-2.866	-6.257%	46
HF Hedge	-0.015%	0.351%	-0.924%	-0.832	-3.329%	36
Crisis Hedge	0.016%	0.222%	-0.644%	0.901	-1.252%	3
Brexit						
No Hedging	0.045%	0.259%	-1.069%	2.550	-1.123%	4
Trad Hedge	0.020%	0.110%	-0.174%	2.396	-0.407%	21
HF Hedge	0.015%	0.096%	-0.249%	1.918	-0.489%	6
Crisis Hedge	0.005%	0.049%	-0.116%	0.518	-0.250%	27

Table 7: Performance of Hedged Risk Parity Portfolios Outside of Crises

The performance of the hedged risk parity portfolios from Table 6 are reported outside of all crises.

	Mean	Stdev	Min	SR
Panel A: Hedging PC1 only				
No Hedging	0.013%	0.267%	-1.662%	0.569
Trad Hedge	0.012%	0.200%	-1.145%	0.671
HF Hedge	0.012%	0.195%	-1.162%	0.682
Crisis Hedge	0.012%	0.195%	-0.971%	0.726
Panel B: Hedging both PC1 and PC2				
No Hedging	0.013%	0.267%	-1.662%	0.569
Trad Hedge	0.004%	0.128%	-0.804%	0.119
HF Hedge	0.003%	0.109%	-0.696%	-0.015
Crisis Hedge	0.005%	0.081%	-0.559%	0.216

Figure 1: Asset Coverage

Figure 1 documents the coverage of the 26 global equity, fixed income and exchange rate assets, as well as four commodities, across the 24-hour period around the Covid-19 pandemic, January-June 2020. The first row of the table starts from 12am EST to 12:59am to the last row from 11:00pm to 11:59pm EST. We calculate the percentage of 5-minute returns that are available during the hour. (See Section 2.2 for a description of the exact methodology used.) Along with documenting the actual % covered, the table provides a color scale to visually reflect the coverage, i.e., from full coverage in dark green to lesser and lesser shades of green as coverage declines, with no coverage (0%) in no color.

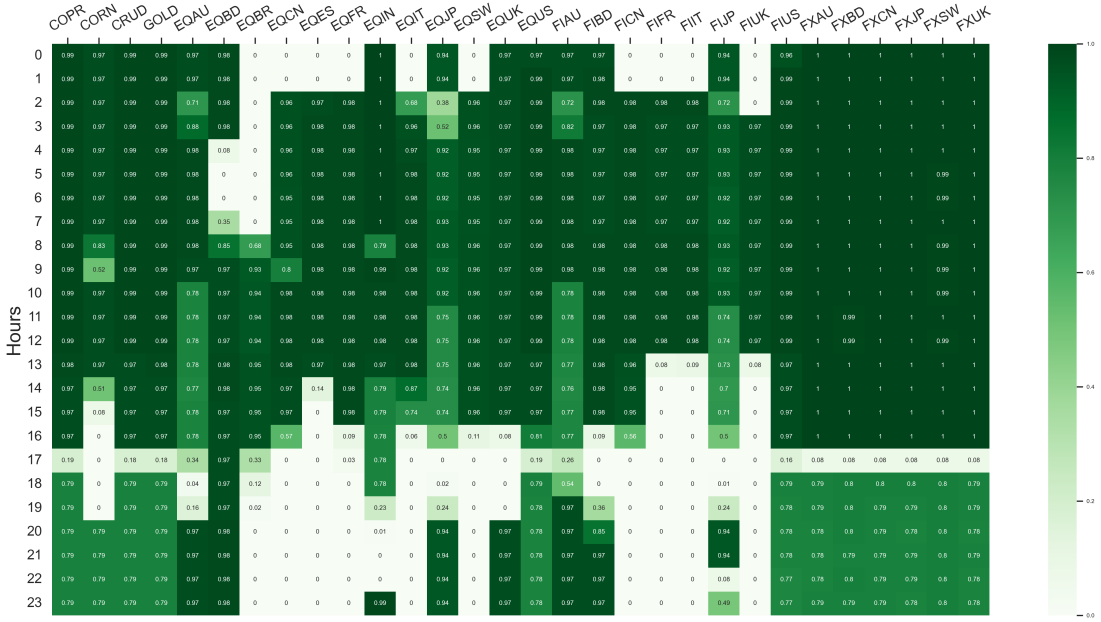


Figure 2: High Frequency Versus Daily Data for Correlation Estimates

Plotted are the mean and standard deviation of biweekly correlations of global assets within and across equities, fixed income, and exchange rates using high frequency (HF) and daily data over the relevant Covid-19 quarter (January 29 to May 5, 2020) and prior quarter (October to December 2019). The error bars show 95% confidence intervals around the mean correlation estimates.

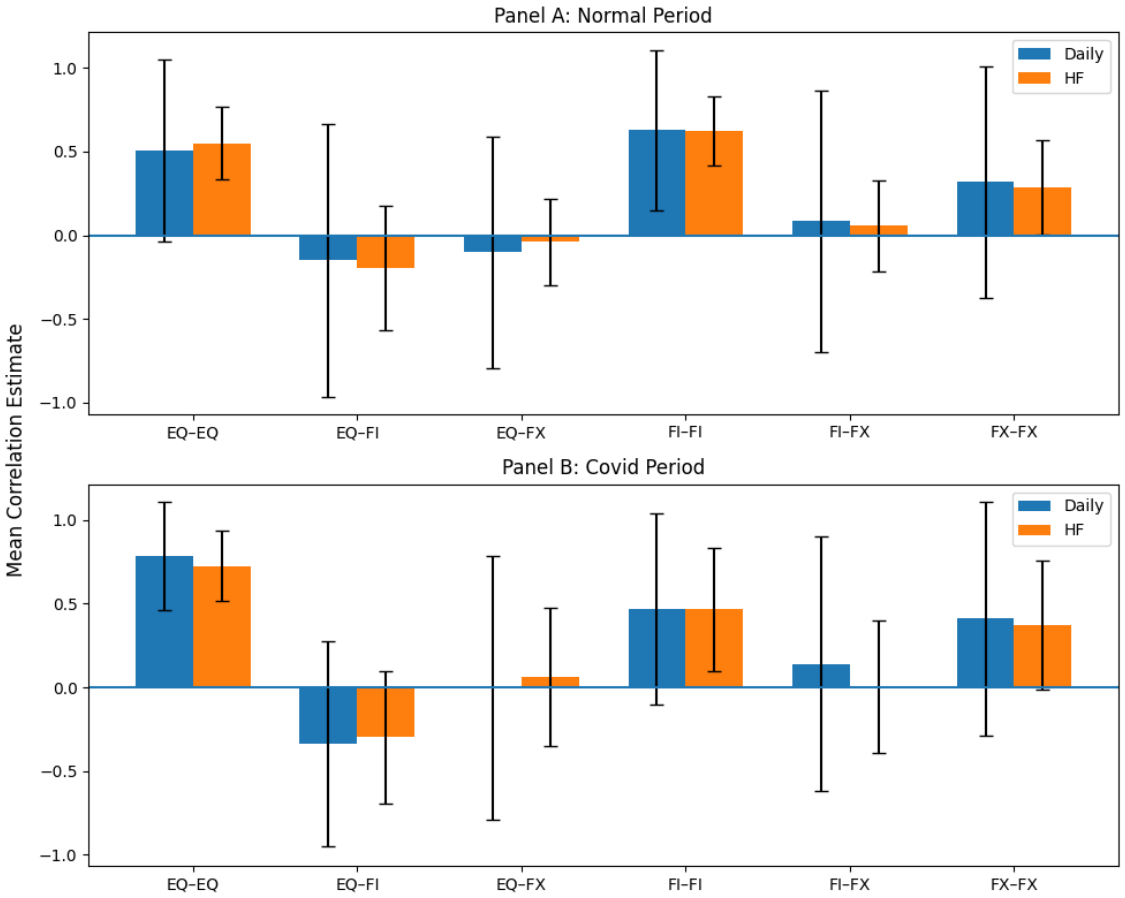


Figure 3: Eigenvalues

This figure shows the average ratios of the eigenvalues to the total eigenvalue for the first ten principal components. The average ratios for all periods, pre-crisis periods, and crisis periods are separately plotted in the figure.

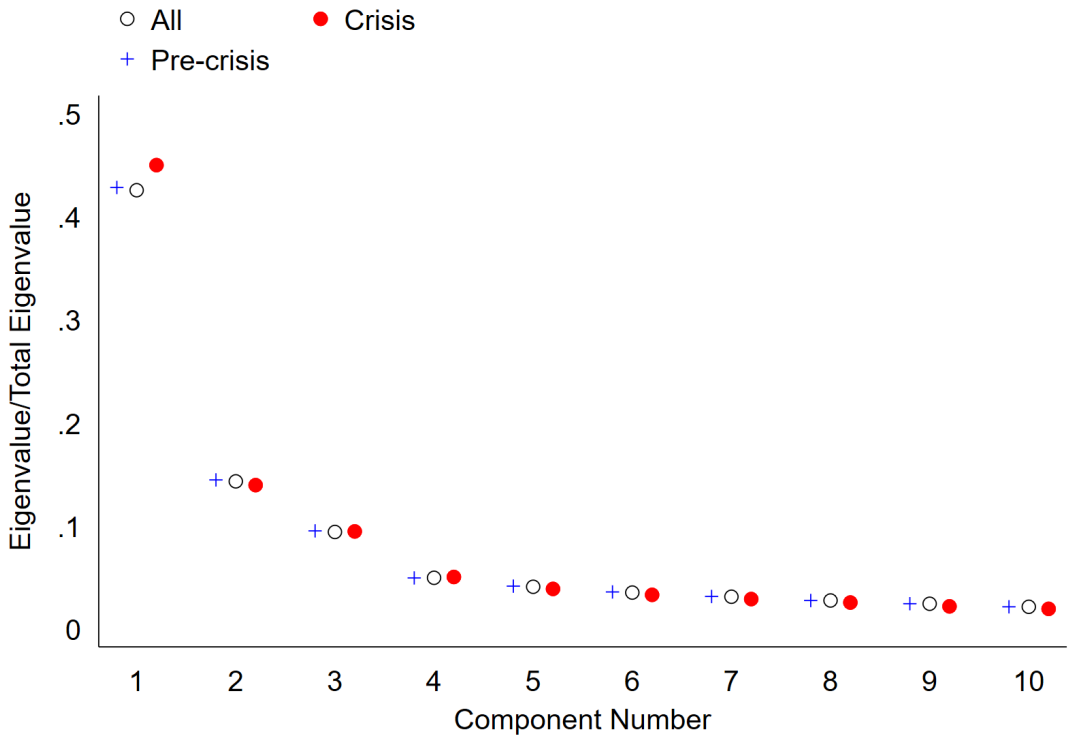
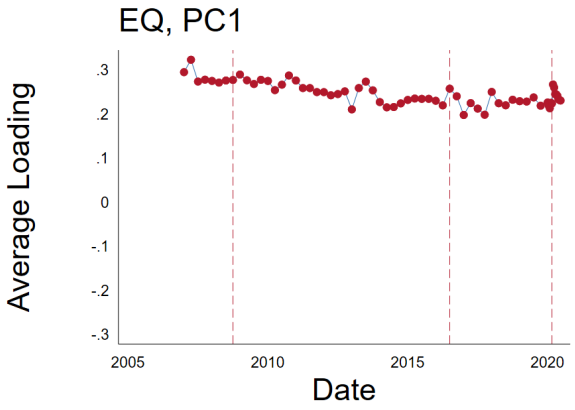


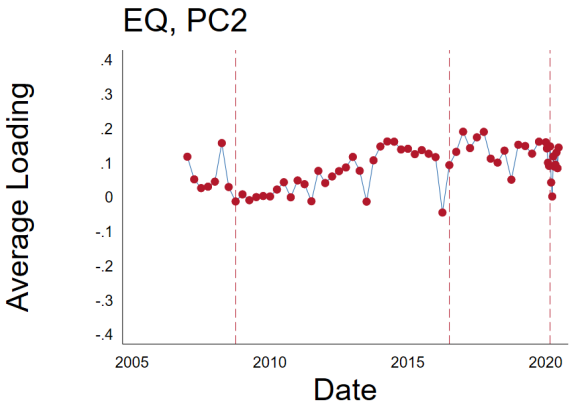
Figure 4: PC Loadings

This graph shows the average PC1 and PC2 loadings of an equal-weighted index of 12 global equity futures (EQ), 8 global fixed income futures (FI), and 6 exchange rates against the dollar (FX) over the sample period 2007-2020. The PCs are estimated with the correlation matrix of the 26 global asset returns using high frequency data. In terms of the average loadings, the averages are reported over quarters pre 2020 and over biweeks for the 2020 Covid-19 crisis period. The dashed lines represent three crisis events, the 5th biweek of 2020 (Covid-19), the 13th biweek of 2016 (Brexit), and the 19th week of 2008 (GFC).

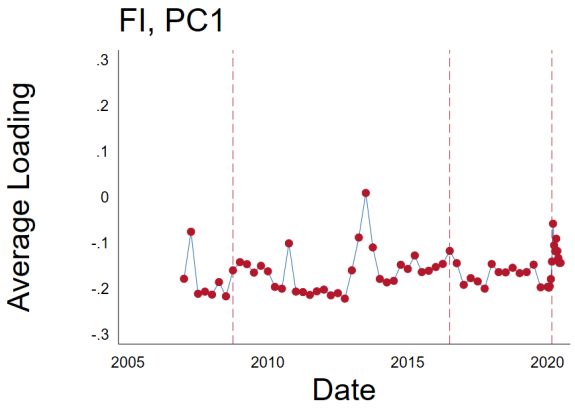
(1) Panel A



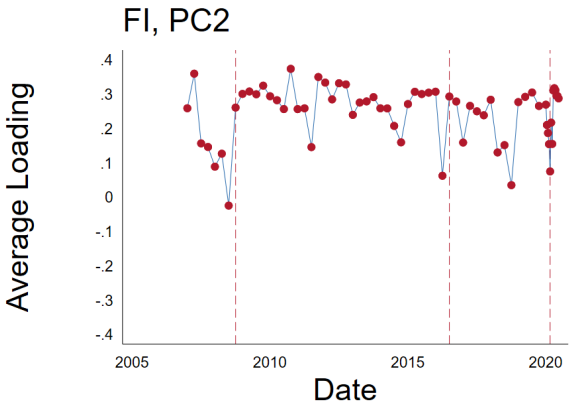
(2) Panel B



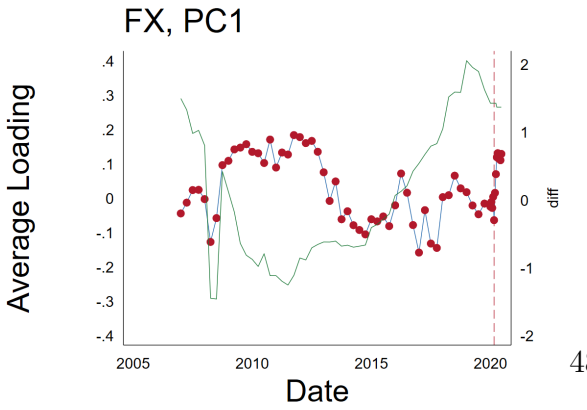
(3) Panel C



(4) Panel D



(5) Panel E



(6) Panel F

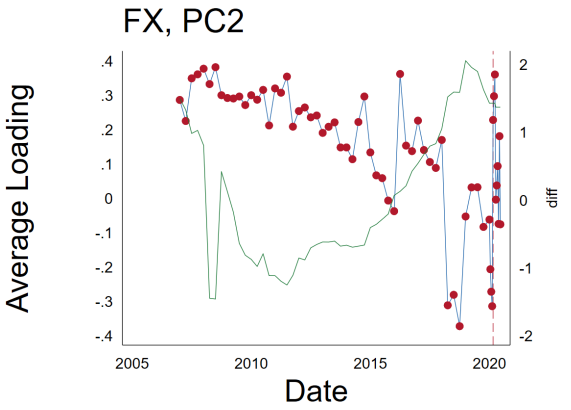
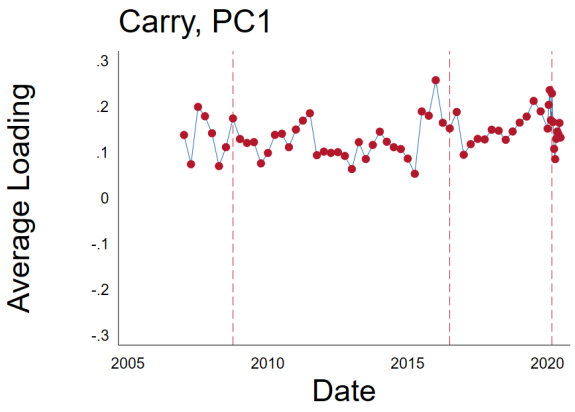


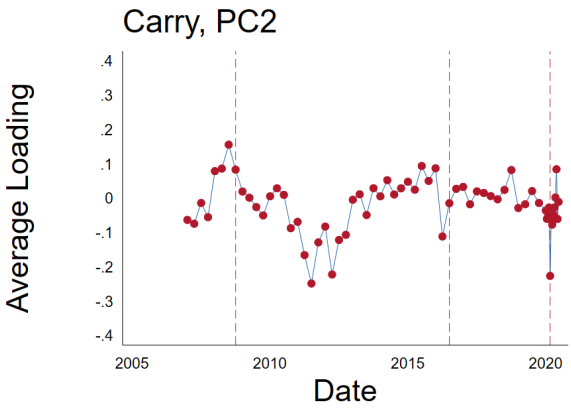
Figure 5: Carry Strategy

This graph plots PC loadings over time related to the FX carry strategy. Panels A and B shows the average PC1 and PC2 loadings of an equal-weighted index of 6 exchange rates against the dollar (FX) over the sample period 2007-2020, graphed alongside the spread between US short-term rates and an average of these rates for the six countries. Panels C and D plot the average PC1 and PC2 loadings of FX, restructured as a carry strategy. The carry strategy longs the exchange rates of the three countries with relatively high interest rates and shorts the other three countries with relatively low interest rates. In terms of the average loadings, the averages are reported over quarters pre 2020 and over biweeks for the 2020 Covid-19 crisis period. The dashed lines represent three crisis events, the 5th biweek of 2020 (Covid-19), the 13th biweek of 2016 (Brexit), and the 19th week of 2008 (GFC).

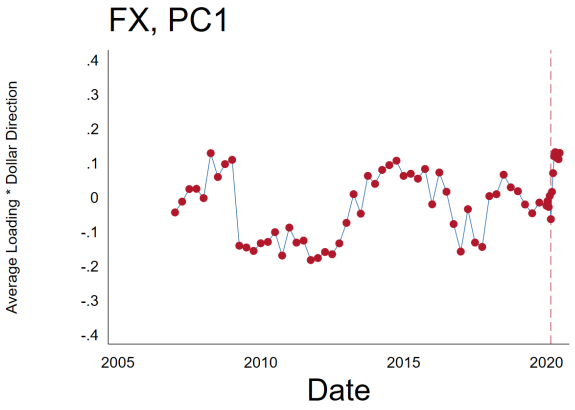
(1) Panel A



(2) Panel B



(3) Panel C



(4) Panel D

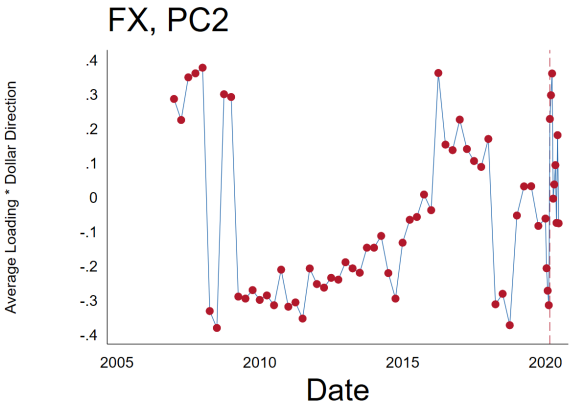
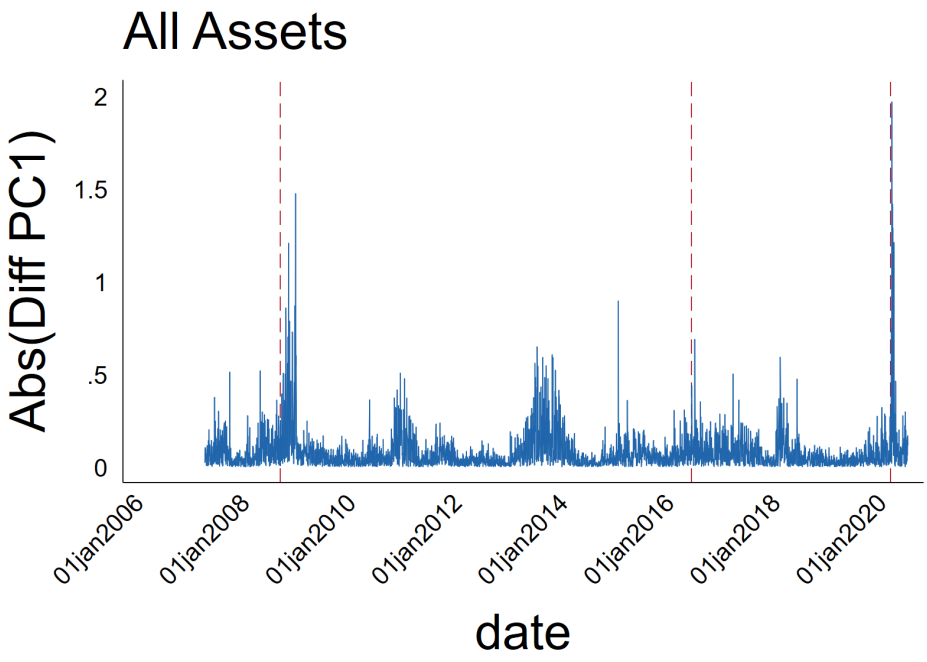


Figure 6: Changes in Factor Structure Over Time

The figure graphs the absolute value of differences in PC1 and PC2 through time. Each day, we take the absolute value of that quarter's PC weights minus last quarter's weights times that day's asset return, *i.e.*, $|PC1_{qt-(q-1)t,k}| \equiv |\sum_i (w_{i1qt} - w_{i1(q-1)t}) R_{iqt+k}|$ and $|PC2_{qt-(q-1)t,k}| \equiv |\sum_i (w_{i2qt} - w_{i2(q-1)t}) R_{iqt+k}|$ where R_{iqt+k} is the return on the k-th day of the i-th asset in quarter q_t, and w_{i2qt}, for example, is the current quarter loaindg of PC2 on asset i. The returns of both $PC1_t$ and $PC2_t$ are standardized to have zero mean and variance 1 using the full time-series of $PC1_t$ and $PC2_t$ over the 2007-2020 period.

(1) Panel A



(2) Panel B

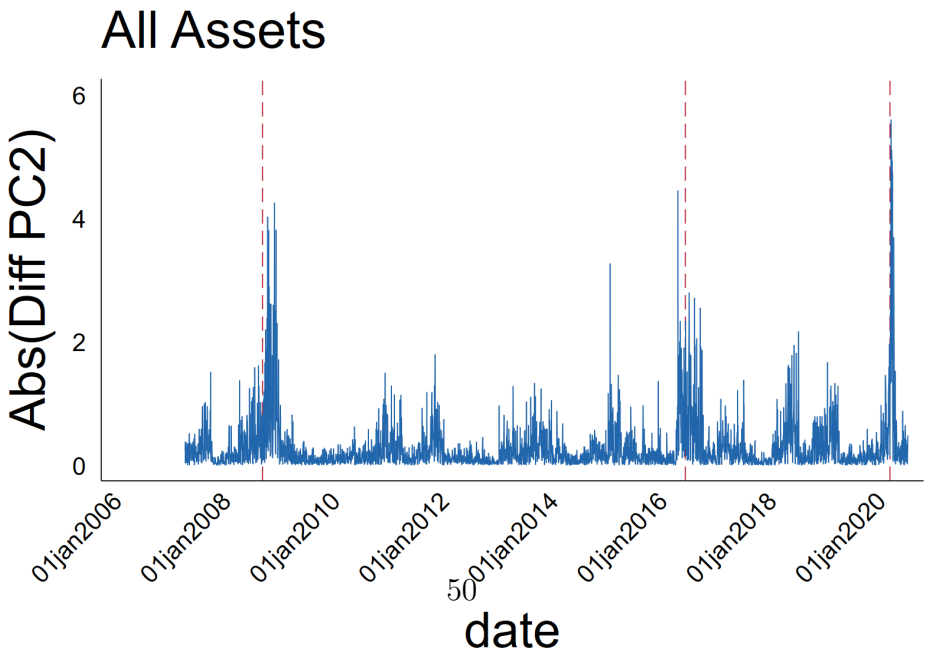
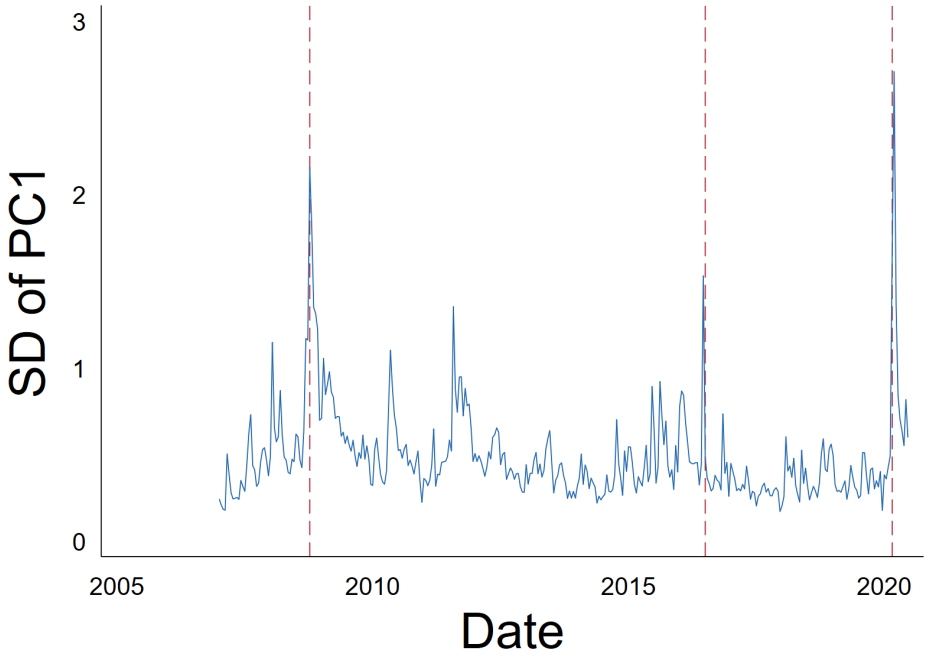


Figure 7: PC1 and PC2 Volatility

The figure graphs the high frequency realized biweekly volatility of PC1 and PC2, where the PCs are constructed from the correlation matrix of 26 global asset returns using high frequency data over biweekly periods in the period 2007-2020. The dashed lines represent three crisis events, the 5th biweek of 2020 (Covid-19), the 13th biweek of 2016 (Brexit), and the 19th week of 2008 (GFC).

(1) Panel A



(2) Panel B

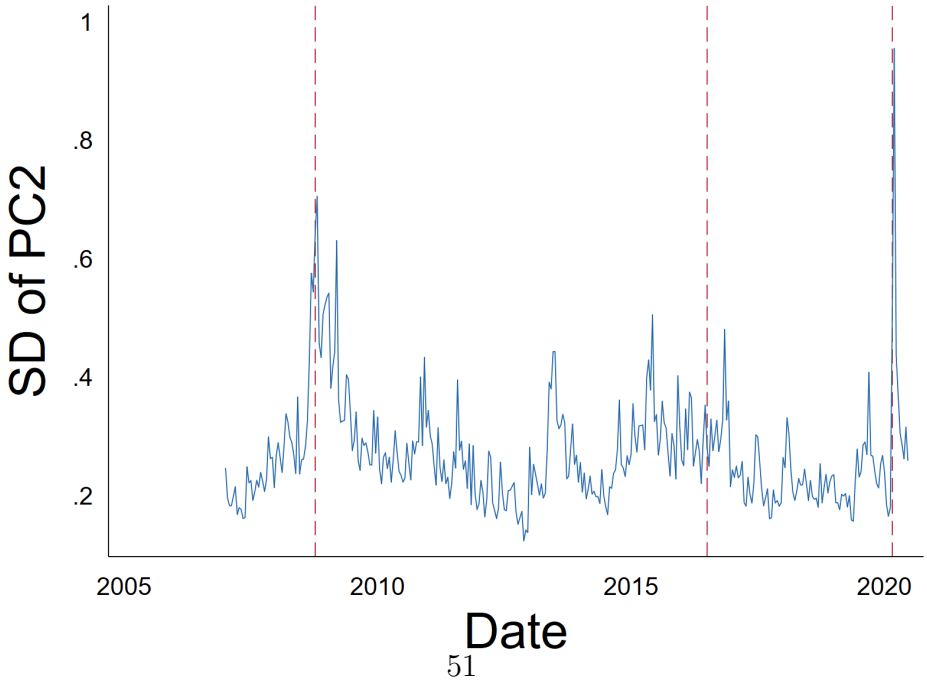
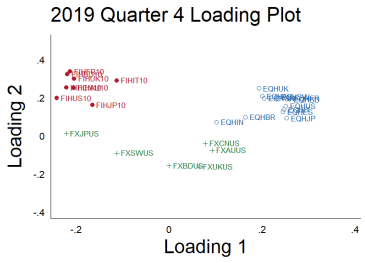


Figure 8: Factor Loadings of Asset Classes During Crisis

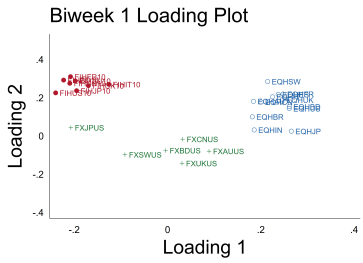
Panel A shows biweekly PC1 and PC2 loadings of EQ, FI, and FX future indexes during the period around the Covid-19 period as measured over January to June, 2020. The indices include 12 equity futures (Australia, Brazil, Germany, Canada, Spain, France, India, Italy, Japan, Switzerland, UK and US) in blue, 8 fixed income futures on 10-year government bonds (Australia, Germany, Canada, France, Italy, Japan, UK and US) in red, and six spot exchange rates (Australia, Euro, Canada, Japan, Switzerland, UK in terms of $US) in green. Panels B and C shows biweekly PC1 and PC2 loadings of EQ, FI, and FX future indexes around the peak of the Global Financial Crisis (GFC) (biweek 14 through biweek 26 of 2008, with biweek 19 being the Lehman bankruptcy) and the Brexit vote (biweek 7 through biweek 19 of 2016, with biweek 13 being the vote). The indices include 12 equity futures (Australia, Brazil, Germany, Canada, Spain, France, India, Italy, Japan, Switzerland, UK and US) in blue, 8 fixed income futures on 10-year government bonds (Australia, Germany, Canada, France, Italy, Japan, UK and US) in red, and six spot exchange rates (Australia, Euro, Canada, Japan, Switzerland, UK in terms of $US) in green. The PCs are constructed from the correlation matrix of the 26 global asset returns using high frequency data over the relevant biweekly periods.

Panel A: Covid-19 Period

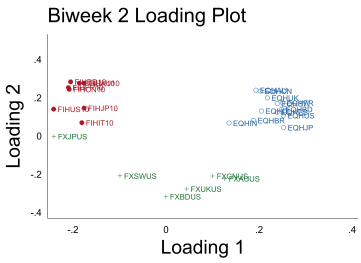
(1) 2019 10/1 – 12/31



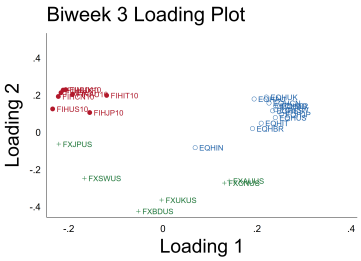
(2) 2020 1/1 – 1/14



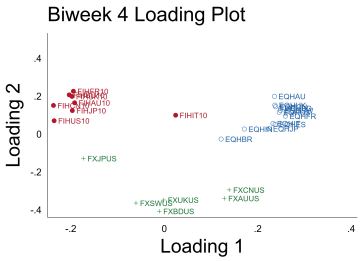
(3) 2020 1/15 – 1/28



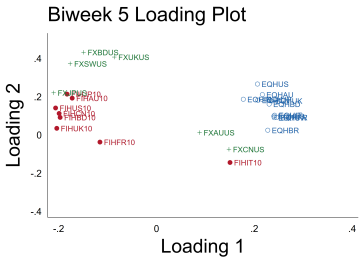
(4) 2020 1/29 – 2/11



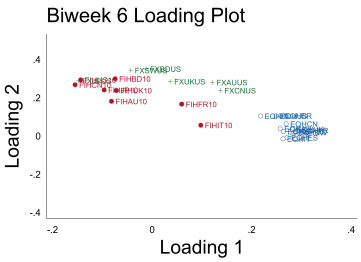
(5) 2020 2/12 – 2/25



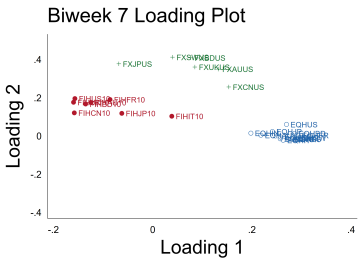
(6) 2020 2/26 – 3/10



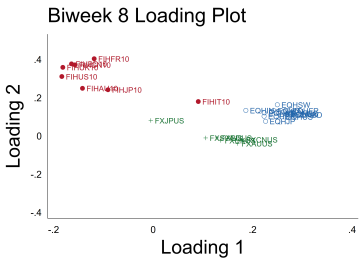
(7) 2020 3/11 – 3/24

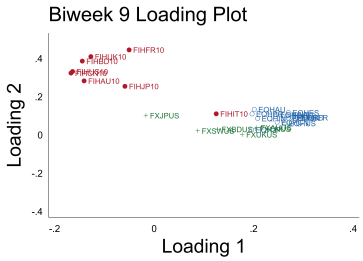


(8) 2020 3/25 – 4/7

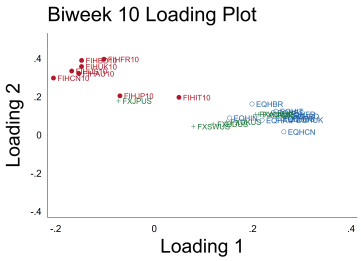


(9) 2020 4/8 – 4/21

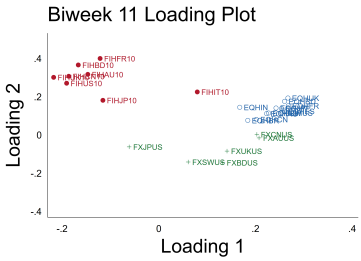




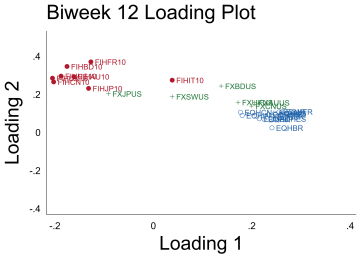
(10) 2020 4/22 – 5/5



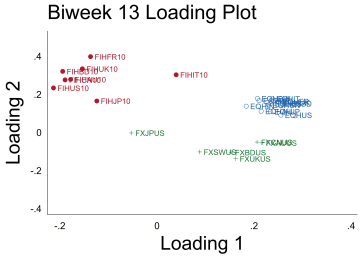
(11) 2020 5/6 – 5/19



(12) 2020 5/20 – 6/2

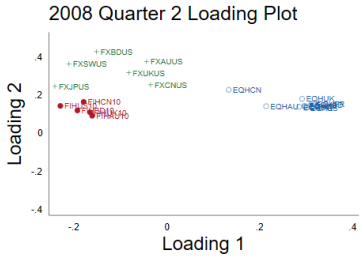


(13) 2020 6/3 – 6/16

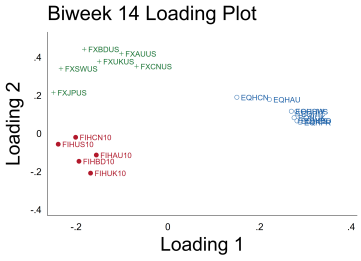


(14) 2020 6/17 – 6/30

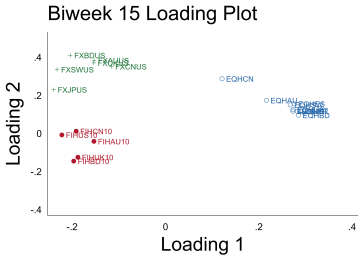
Panel B: GFC Period



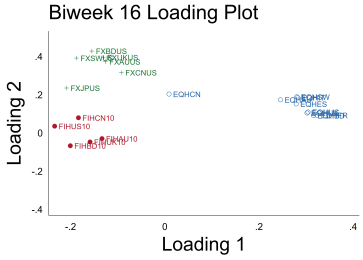
(15) 2008 4/1 – 6/31



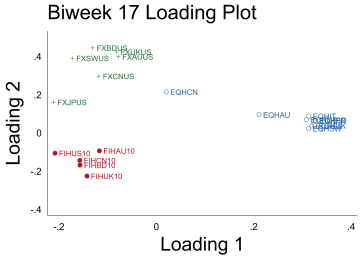
(16) 2008 7/1 – 7/14



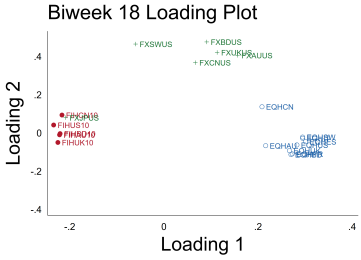
(17) 2008 7/15 – 7/28



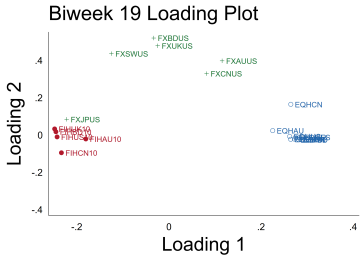
(18) 2008 7/29 – 8/11



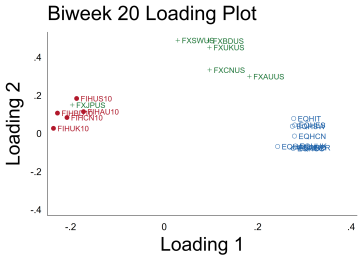
(19) 2008 8/12 – 8/25



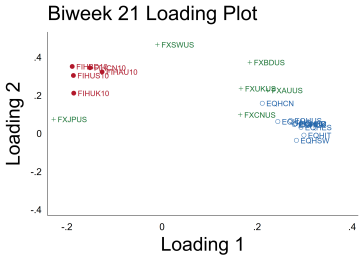
(20) 2008 8/26 – 9/8



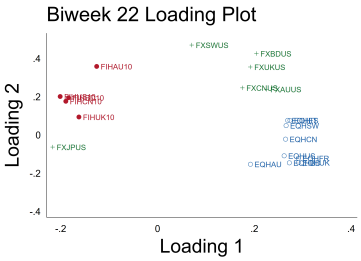
(21) 2008 9/9 – 9/22



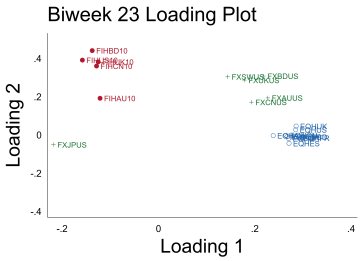
(22) 2008 9/23 – 10/6



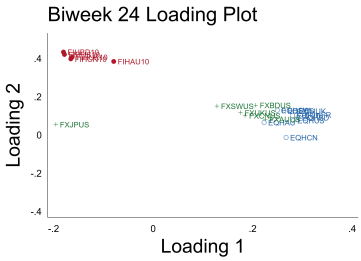
(23) 2008 10/7 – 10/20



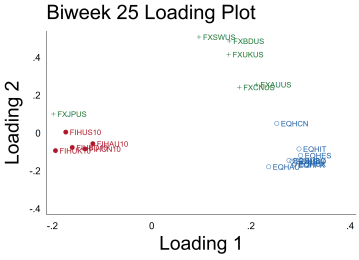
(24) 2008 10/21 – 11/3



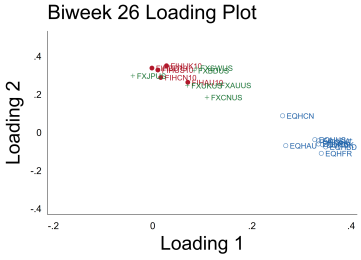
(25) 2008 11/4 – 11/17



(26) 2008 11/18 – 12/1

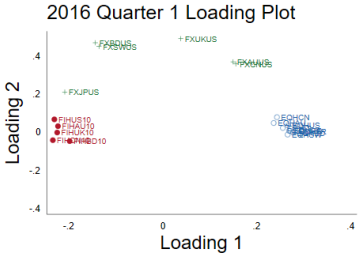


(27) 2008 12/2 – 12/15

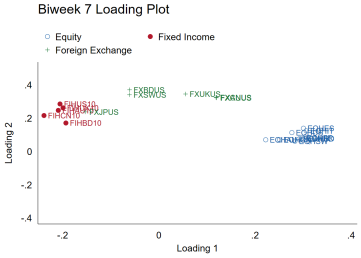


(28) 2008 12/16 – 12/31

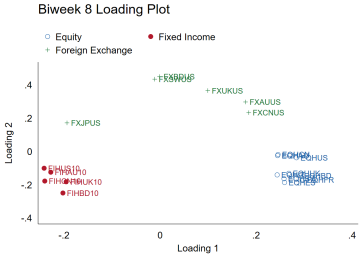
Panel B: Brexit Period



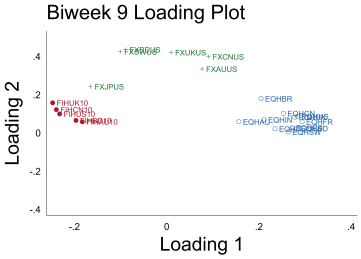
(29) 2016 1/1 – 3/31



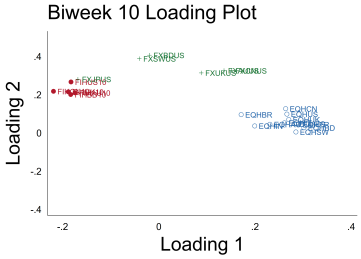
(30) 2016 3/25 – 4/7



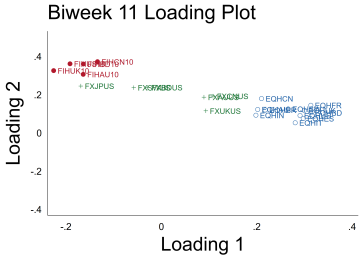
(31) 2016 4/8 – 4/21



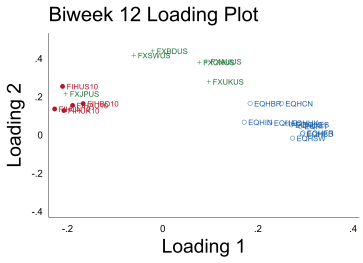
(32) 2016 4/22 – 5/5



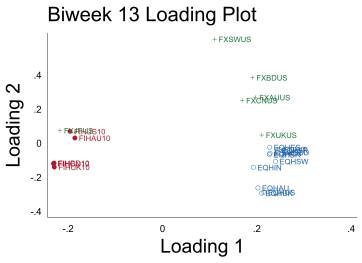
(33) 2016 5/6 – 5/19



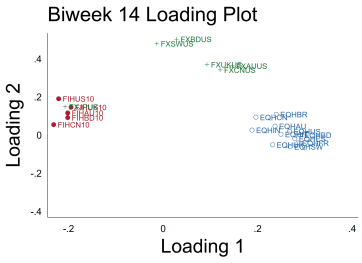
(34) 2016 5/20 – 6/2



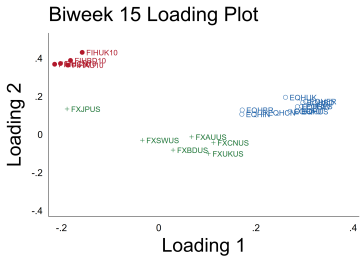
(35) 2016 6/3 – 6/16



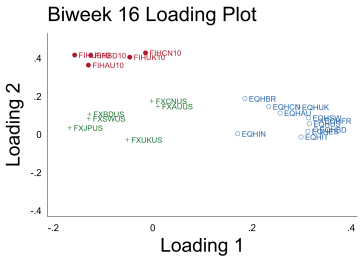
(36) 2016 6/17 – 6/30



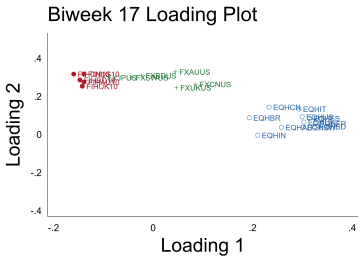
(37) 2016 7/1 – 7/14



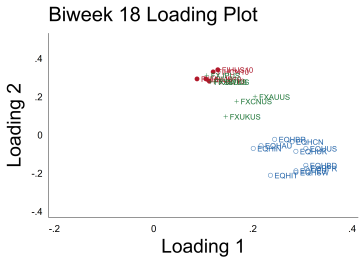
(38) 2016 7/15 – 7/28



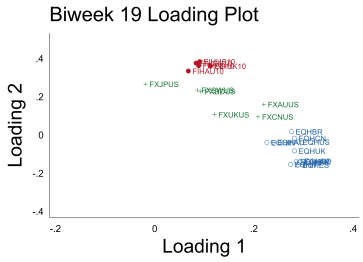
(39) 2016 7/29 – 8/11



(40) 2016 8/12 – 8/25



(41) 2016 8/26 – 9/8



(42) 2016 9/9 – 9/22

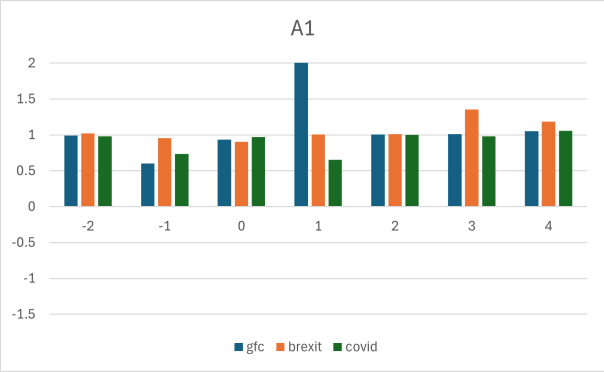
Figure 9: Changing PC Variances

Figure 9 reports results for decomposing changing PC variances into changing PC weights versus the changing variance-covariance matrix of global asset return. The sample covers the onset biweekly periods from January to June 2020, with the defining event happened in biweek. Changes of PC variances can be decomposed as:

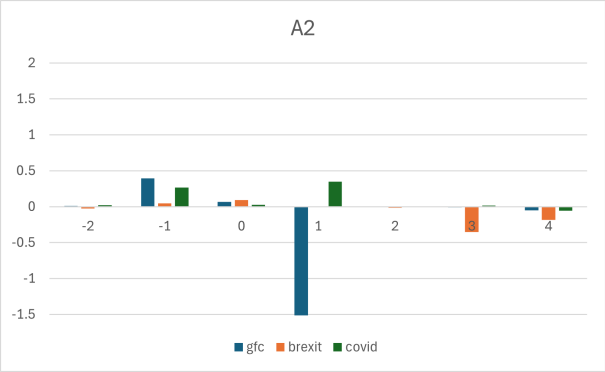
$$Var\left(PCi_{t+1}\right) - Var\left(PCi_t\right) = \underbrace{\left(w_t'^{i}\Sigma_{t+1}w_t^i - w_t'^{i}\Sigma_t w_t^i\right)}_{\text{A: Due to changing var-cov}} + \underbrace{\left(w_{t+1}'^{i}\Sigma_{t+1}w_{t+1}^i - w_t'^{i}\Sigma_{t+1}w_t^i\right)}_{\text{B: Due to changing PC weights}}$$

where w are the PC weights and $\sum$ is the variance-covariance matrix of global asset returns.

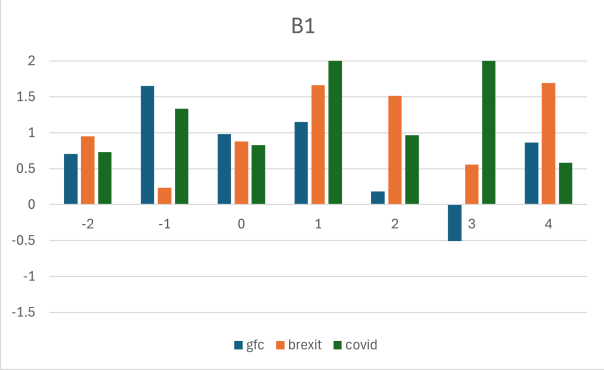
(1) Panel A: Fraction of PC1 changes from var-cov changes



(2) Panel B: Fraction of PC1 changes from changing PC weights



(3) Panel C: Fraction of PC2 changes from var-cov changes



(4) Panel D: Fraction of PC2 changes from changing PC weights

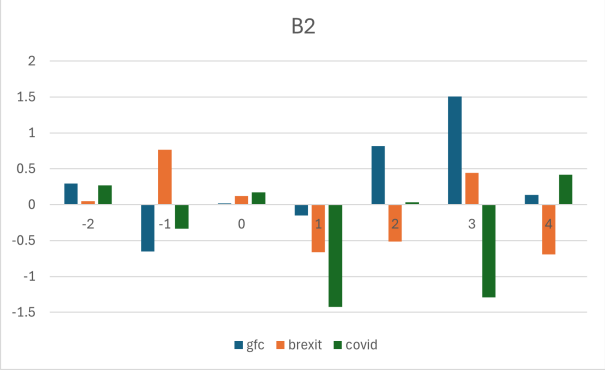


Figure 10: Correlation Between Crisis PC2 and Different Asset Returns Over Time

Figure 10 shows the correlation between crisis PC2 and different asset returns over time. Crisis PC2 uses the average PC2 of for each of the assets across the three crises periods. The assets include equal weighted equity, equal weighted fixed income, dollar basket, gold, and yen. The correlations are based on the high-frequency variance-covariance matrix.

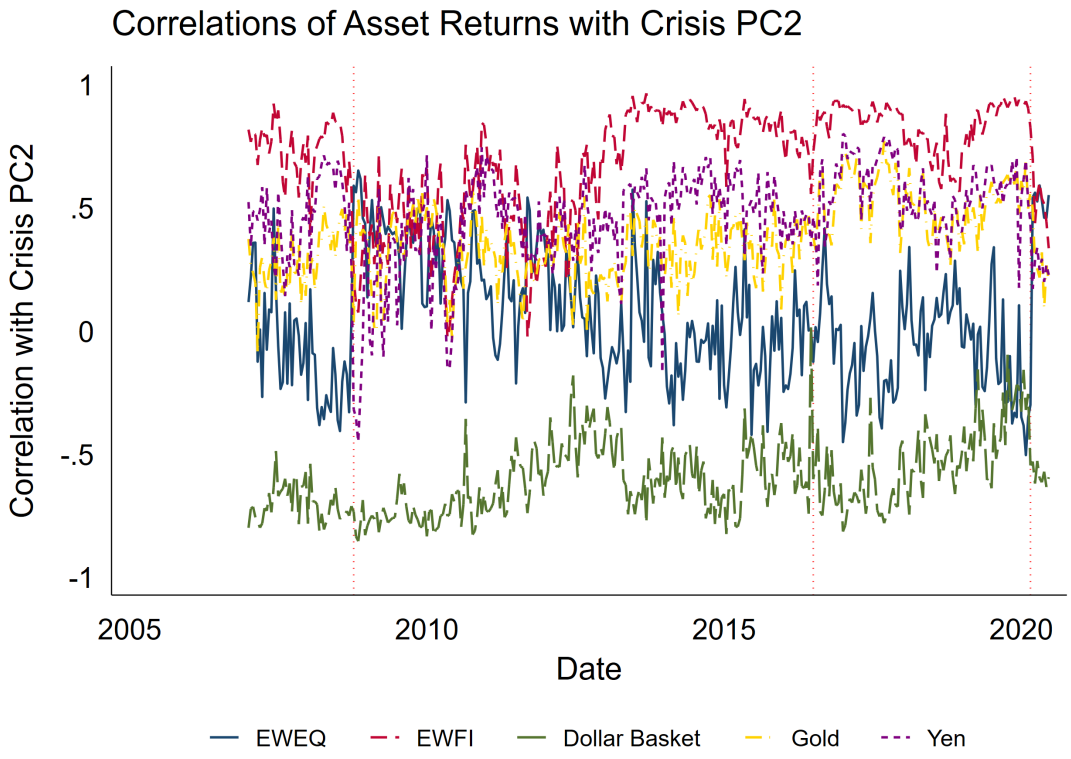
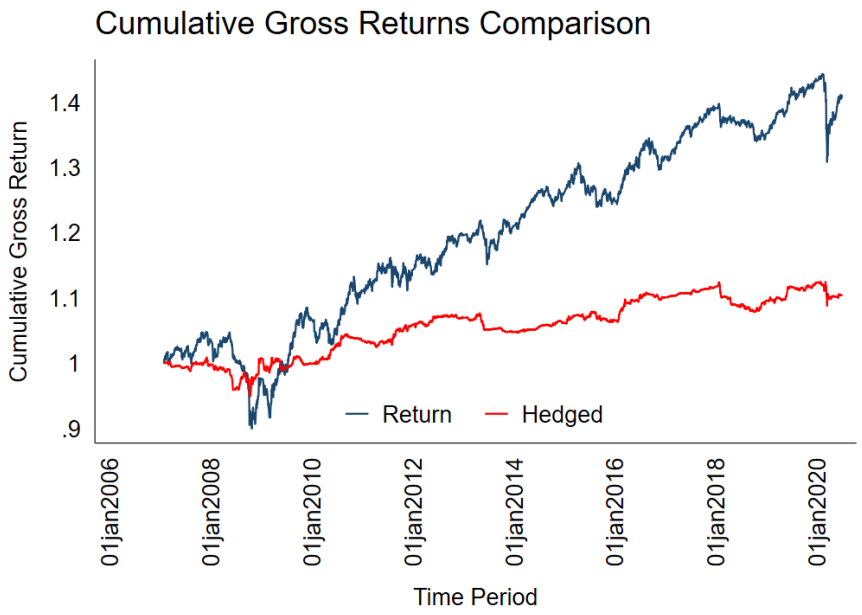


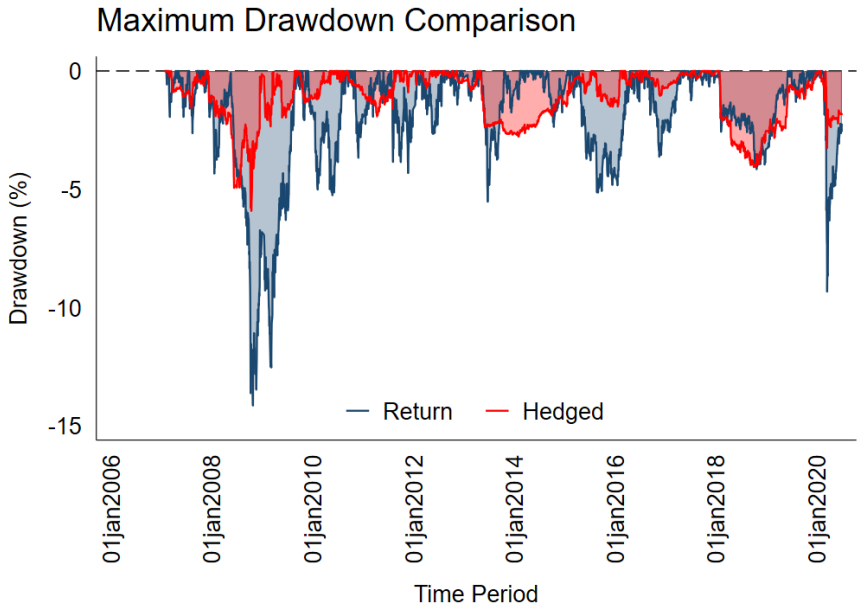
Figure 11: Cumulative Return and Max Drawdown of Risk Parity and Hedged Risk Parity

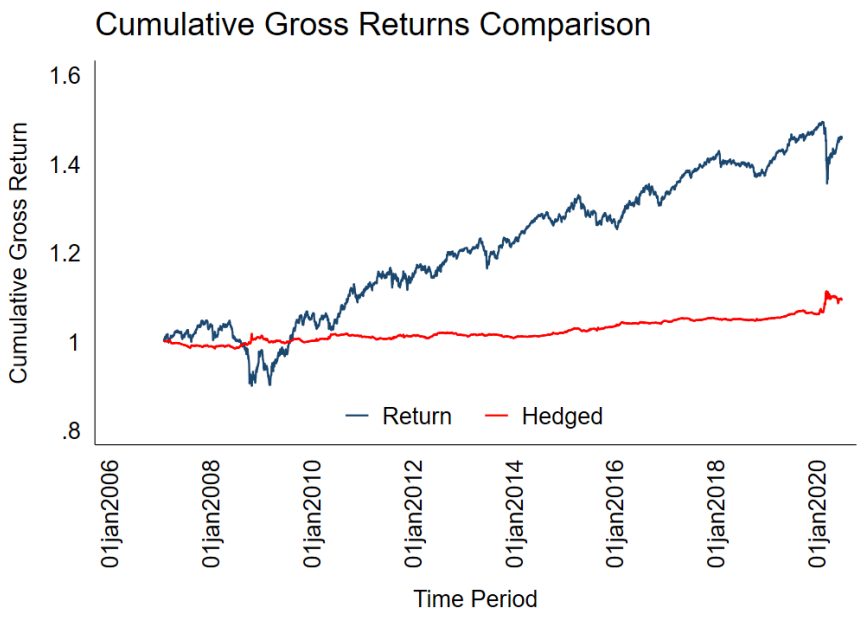
Figure 11 shows the cumulative returns and the max drawdowns of the raw unhedged risk parity portfolio and the hedged risk parity portfolio. Panels (1) and (2) show results using HF hedge and Panels (3) and (4) show results using the fixed crisis-weight hedge. The HF hedging portfolio PC loadings use the prior biweek loadings. The crisis hedging portfolio PC loadings use the average crisis weight of the other two crises for each crisis period. Both PC1 and PC2 are used for the hedging exercise. The betas are calculated each period prior to the hedging period using the high frequency covariance matrix.

(1) Cumulative Return Graph – HF Hedge

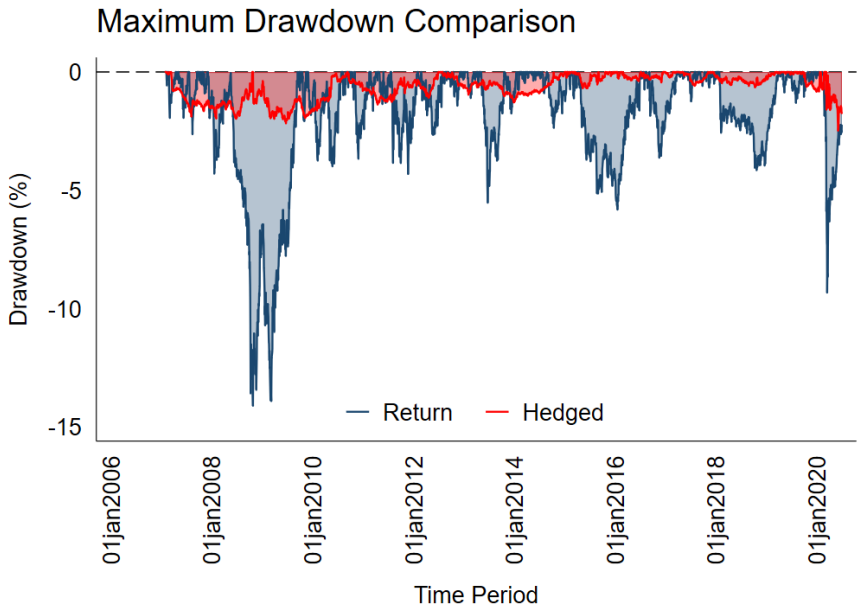


(2) Max Drawdown Graph – HF Hedge





(3) Cumulative Return Graph – Crisis Hedge



(4) Max Drawdown Graph – Crisis Hedge

A Online Appendix

A.1 Asset Loadings on PC Factors Prior to and During Crises

Table A1: PC1 and PC2 Loadings Prior and During Crises – Asset-by-Asset

This table shows the PC1 and PC2 loadings for each of the assets during each crisis and during the normal period prior to each of the crisis period. We define the crises periods for six biweeks starting from the defining events of each crisis. The defining events for the Great Financial Crisis period, and the Brexit period occurr in biweek 5, 19 and 13 for these three crises, respectively. The "Average" columns present the average PC1 and PC2 of for each of the assets across the three crises periods.

	COVID		GFC		Brexit		Average	
Panel A: Loadings for Each Crisis Period								
	PC1	PC2	PC1	PC2	PC1	PC2	PC1	PC2
EQHAU	0.223	0.096	0.227	-0.013	0.221	0.105	0.224	0.063
EQHBD	0.272	0.082	0.281	-0.020	0.271	0.056	0.275	0.039
EQHCN	0.247	0.073	0.258	0.044	0.206	0.108	0.237	0.075
EQHES	0.262	0.070	0.276	0.036	0.258	0.045	0.265	0.051
EQHFR	0.276	0.075	0.287	-0.012	0.272	0.042	0.278	0.035
EQHIT	0.254	0.070	0.280	0.033	0.239	0.047	0.258	0.050
EQHSW	0.257	0.077	0.267	0.024	0.254	0.040	0.259	0.047
EQHUK	0.268	0.079	0.282	-0.003	0.239	0.110	0.263	0.062
EQHUS	0.251	0.112	0.273	-0.005	0.257	0.100	0.260	0.069
FIHAU10	-0.136	0.233	-0.134	0.224	-0.127	0.221	-0.132	0.226
FIHBD10	-0.144	0.284	-0.194	0.257	-0.134	0.257	-0.157	0.266
FIHCN10	-0.174	0.248	-0.179	0.210	-0.122	0.243	-0.158	0.234
FIHUK10	-0.149	0.262	-0.187	0.192	-0.113	0.257	-0.150	0.237
FIHUS10	-0.170	0.267	-0.192	0.248	-0.134	0.244	-0.166	0.253
FXAUUS	0.154	0.124	0.199	0.239	0.097	0.115	0.150	0.159
FXBDUS	0.048	0.210	0.146	0.377	0.032	0.094	0.075	0.227
FXCNUS	0.172	0.089	0.149	0.212	0.096	0.105	0.139	0.135
FXJPUS	-0.087	0.207	-0.213	0.041	-0.129	0.109	-0.143	0.119
FXSWUS	0.015	0.196	0.039	0.385	-0.001	0.067	0.018	0.216
FXUKUS	0.082	0.182	0.131	0.320	0.082	0.059	0.098	0.187
Panel B: Loadings for Each Pre-Crisis Period								
	PC1	PC2	PC1	PC2	PC1	PC2	PC1	PC2
EQHAU	0.211	0.166	0.217	0.112	0.219	0.068	0.215	0.115
EQHBD	0.278	0.161	0.305	0.109	0.264	0.115	0.282	0.128
EQHCN	0.215	0.168	0.150	0.179	0.226	0.080	0.197	0.142
EQHES	0.253	0.112	0.297	0.096	0.258	0.139	0.269	0.116
EQHFR	0.274	0.175	0.307	0.110	0.267	0.123	0.283	0.136
EQHIT	0.242	0.118	0.295	0.108	0.248	0.119	0.262	0.115
EQHSW	0.243	0.170	0.281	0.111	0.250	0.117	0.258	0.133
EQHUK	0.207	0.216	0.290	0.127	0.254	0.110	0.250	0.151
EQHUS	0.266	0.128	0.286	0.092	0.255	0.086	0.269	0.102
FIHAU10	-0.184	0.258	-0.161	0.035	-0.193	0.222	-0.179	0.172
FIHBD10	-0.183	0.323	-0.194	0.046	-0.175	0.346	-0.184	0.238
FIHCN10	-0.194	0.268	-0.178	0.121	-0.204	0.243	-0.192	0.211
FIHUK10	-0.182	0.309	-0.176	0.037	-0.192	0.301	-0.183	0.216
FIHUS10	-0.215	0.228	-0.225	0.097	-0.203	0.242	-0.214	0.189
FXAUUS	0.094	-0.038	-0.013	0.373	0.137	0.016	0.072	0.117
FXBDUS	-0.006	-0.127	-0.137	0.435	-0.101	-0.123	-0.081	0.062
FXCNUS	0.070	-0.018	-0.014	0.277	0.142	0.008	0.066	0.089
FXJPUS	-0.210	0.044	-0.236	0.234	-0.189	-0.015	-0.212	0.088
FXSWUS	-0.113	-0.062	-0.203	0.359	-0.100	-0.115	-0.138	0.061
FXUKUS	0.045	-0.127	-0.067	0.342	0.046	-0.068	0.008	0.049

A.2 News Data

Ravenpack is a commercial vendor primarily focused on the asset management industry. They employ natural language processing (NLP) algorithms to measure specific events and sentiment generated from unstructured text contained in a variety of media sources and company reports. Their analysis is both at the firm level as well as at the aggregate (macro) level. As described in Section 3.3.1, in response to the coronavirus outbreak, Ravenpack produced a set of related Covid-19 news indices for every country, as well as a worldwide index on a daily basis (i) Coronavirus Panic Index, (ii) Coronavirus Media Hype Index, and (iii) Coronavirus Media Coverage Index for each country. All of these measures try to capture relevant Covid-19 news (such as references to panic/hysteria, percent of news coverage, *etc*...). Specifically, as described by Ravenpack, the Coronavirus Panic Index "measures the level of news chatter that makes reference to panic or hysteria and coronavirus;" the Coronavirus Media Hype Index "measures the percentage of news talking about the novel coronavirus;" and the Coronavirus Media Coverage Index "calculates the percentage of all news sources covering the topic of novel coronavirus." Changes in these indices can potentially be used to represent unexpected "news."

We also construct a composite index of these three Covid-19 news indices. Each index is standardized to have a zero mean and unit variance, and then the composite index is the average of the three standardized indices. Note that because news has a seasonal component (*e.g.*, weekday versus weekend, Monday versus Friday), we deseasonalize each index in the following manner. For the 6 months of data, we calculate the average ratio of each day's index level relative to a -3 day/+3 day window. The average for each day of the week is then calculated over the entire sample and then used to adjust the index level through time. Figure A1 graphs the deseasonalized composite index in our sample from January 1, 2020 to June 30, 2020.

A few observations can be made. First, the timeline of the indices is quite apparent. Early on in January, 2020, there is little news about Covid-19. Towards the end of January, the index spikes up as China locked down and many countries imposed bans on travel. For much of February, the index level plateaus, until a large spike in Italy (as people became aware that the coronavirus was spreading in northern Italy), followed by all other countries, in late February. The indices reached their peak throughout March into early April as it became clear the virus had spread widely, and economies began to lockdown globally. Thereafter, there is a slow downward trend as news about Covid-19 levels off. Finally, while the trend of the index is apparent through time, there is clearly time variation across the indices. Whether this time variation is simply noise or fundamentals is the topic of the empirical analysis in section 3.3.1.

A.3 Epidemiological Data

As described in Section 3.3.2, a particular focus of epidemiological groups, such as the one at Imperial College, was to forecast models of infection rates and deaths both within and across countries. They analyzed the impact of mitigation efforts such as sheltering-in-place, partial shutdowns, social distancing, mask wearing, etc. The Imperial College group's models are more in the vein of Susceptible Exposed Infectious Recovered (SEIR) epidemiological models.[13]

As one of their forecasting analyses, the Imperial College of London group produces one-week ahead forecasts of cumulated deaths for up to 69 countries as reported by the European Centre for Disease Prevention and Control (ECDC).[14] These one-week ahead forecasts start on March 15, 2020 for Italy and the U.S. with the other eight countries in our sample having

[13]For an excellent discussion of SEIR models in the recent economic literature, see, for example, Atkeson (2020) and Stock (2020), and Eichenbaum et al. (2021) and Jones et al. (2021) for theoretical economic models using SEIR models.

[14]The ECDC is considered one of the main reporters of Covid-19 data along with Johns Hopkins University, the World Health Organization (WHO) and the U.S. Center for Disease Control (CDC). All of the data produced by these agencies are highly correlated but nevertheless differ at any given point. The Imperial College research group calibrates its models to the ECDC.

forecasts initiated over the next few weeks. The forecast model is an ensemble average of four related models.[15] The first two models are based on estimating transmissibility of the coronavirus (the reproduction rate R_t) and use this rate to forecast the incidence of death. Both these models use back data to estimate the model's parameters and simulations to provide mean and median one week forecasts. The third model uses reported cases, probability of death given infection, and distribution of times from infection to death, to forecast the number of deaths over the next week. As with the first two models, this model is couched in an assumed distributional framework, resulting in mean and median forecasts from simulations of the model. The fourth model is an empirical Bayesian model that works backwards from deaths to estimate transmission rates and infections under different mitigation approaches such as social distancing, public event bans, school closures, quarantining, lockdowns, *etc*... An important aspect of these models is the adjustment for demographical differences in age distributions across countries and different fatality rates within these age distributions.

For four of the representative countries (in 3 continents, i.e., United States (North America), Japan (Asia), and Germany and Italy (Europe)), Figure A2 documents the total number of deaths each day (in solid blue) and the one to seven day ahead forecasts of deaths each week (given by red dots) from the start of the forecast period to June 31, 2020. The difference between the blue and red points therefore represents forecast errors. While some of the variation in ECDC deaths are related to random variation in reported deaths, there is nevertheless some clear patterns. Japan aside, the forecasts of the other representative countries' death rates increase sharply, then rates switch to declining ones thereafter. When deaths started to increase at a high rate, almost all the forecasts, irrespective of the country, overestimated deaths in the following weeks. Moreover, these forecast errors occur at different times for different countries. For example, the errors take place in late March to early April for Italy but early April to mid-April for Germany. Finally, the forecasts for the

[15]For a detailed description of the model, see "Short-term Forecasts of Covid-19 Deaths in Multiple Countries" developed by the Imperial College Covid-19 Response Team, https://mrc-ide.github.io/Covid19-short-term-forecasts/. We thank Sangeeta Bhatia of Imperial College for providing us with the relevant forecast data.

U.S. tend to be over pessimistic over most of the period. While the U.S.'s Covid-19 cases continued to climb over much of the period, fatality rates ex-post declined, possibly due to increasing testing of people who were less sick, different infection rates across demographics (*i.e.*, younger people getting sick) and better hospital care for those more seriously ill. These variations across countries and time allow an analysis between a country's systematic factor return component for equity and Covid-19 "news" as represented by forecast errors.

Figure A1: Covid-19 Composite New Index

Figure A1 shows a composite Covid-19 measure during the January-June 2020 Covid-19 period. The composite index is an average of three standardized Coronavirus indices provided by Ravenpack (Panic, Media Hype, and Media Coverage), who employ a natural language processing (NLP) algorithm to measure specific events and sentiment related to Covid-19 generated from unstructured text contained in a variety of media sources and company reports.

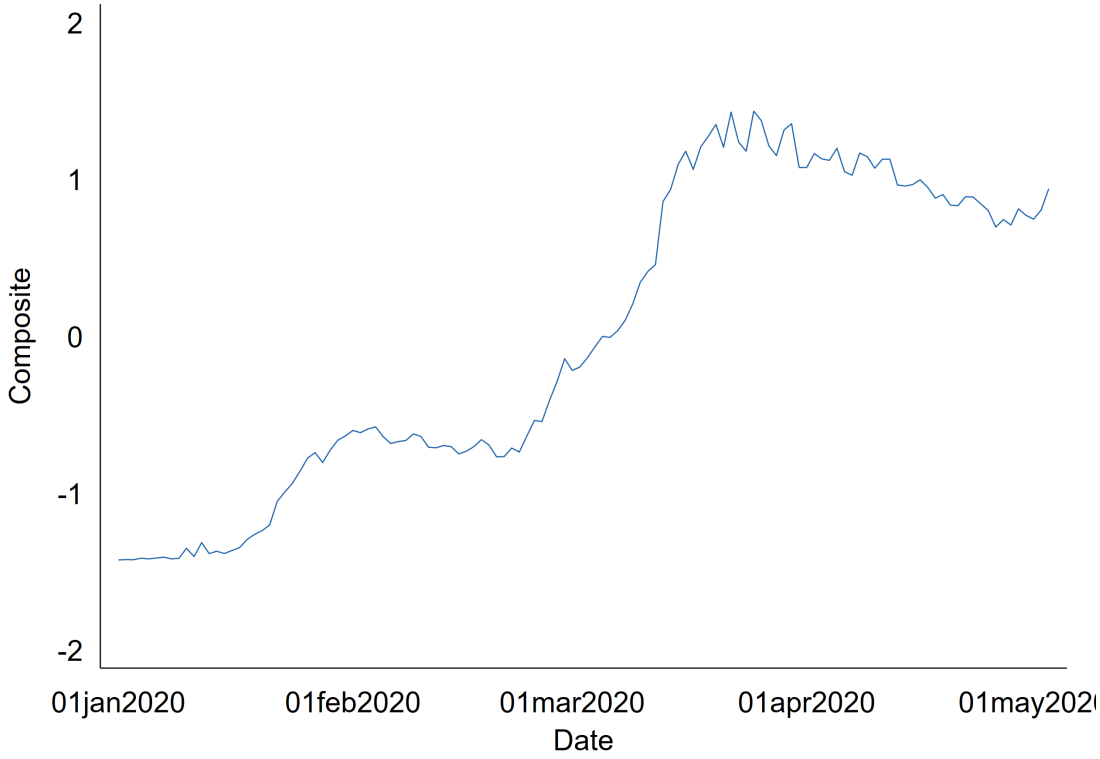


Figure A2: Epidemiological Model Forecast Errors

Figure A2 shows the actual Covid-19 deaths and epidemiological model forecasted deaths in four representative countries in 3 continents, i.e., the United States (North America), Japan (Asia), and Germany and Italy (Europe)) during the Covid-19 period. The forecasts are provided by the MRC Centre of Global Infectious Disease Analysis and Abdul Latif Jameel Institute for Disease and Emergency Analytics at Imperial College London. Specifically, the figure documents the total number of deaths each day (in solid blue) and the one to seven day-ahead forecasts of deaths each week (given by red dots) from the start of the forecast period to June 31, 2020. The difference between the blue and red points therefore represents forecast errors.

