Internet Appendix

IA.A Inferring Dependence

Formal Description of BRA

We present hereafter the formal steps to infer dependence such that $\sum_{j=1}^{d} \omega_j X_j = S$ in each of the n states. To simplify exposition, it is convenient to define transformations $Y_j := \omega_j X_j$, so that $y_{ij} := \omega_j x_{ij}$. Let G_j denote the distribution of Y_j. We then introduce the following $n \times (d+1)$ matrix $\mathbf{M}$:

$$\mathbf{M} = \begin{bmatrix} y_{11} & y_{12} & \cdots & y_{1d} & -s_1 \\ y_{21} & y_{22} & \cdots & y_{2d} & -s_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ y_{n1} & y_{n2} & \cdots & y_{nd} & -s_n \end{bmatrix}, \tag{21}$$

where the last column consists of the elements $-s_i$, which are all possible discrete realizations of the negative sum $-S$; i.e., $-s_i := -F_S^{-1}(\frac{i-0.5}{n})$.

The objective is to find a rearrangement of the first d columns such that the row sums of the $d+1$ columns of $\mathbf{M}$ are all equal to zero. In other words, the opposite of the last column is the sum of the previous columns, i.e., $S = \sum_{j=1}^{d} Y_j$. Denote this last column by $Y_{d+1} = -S$. This procedure is equivalent to finding a rearrangement of the matrix $\mathbf{M}$ such that $Y_1 + \cdots + Y_{d+1}$ is identically equal to zero and thus such that $\text{var}(Y_1 + \cdots + Y_{d+1}) = 0$. We allow for rearrangements within columns, as doing so affects the dependence among Y_j, $j = 1, \ldots, d$ but not their respective marginal distributions. By contrast, swapping values among columns will affect the marginal distributions and is not allowed. Clearly, for $Y_1 + \cdots + Y_{d+1}$ to have the smallest possible variance, it must hold that for all $\ell = 1, \ldots, d+1$, Y_ℓ is as negatively correlated as possible with $\sum_{j=1, j\neq\ell}^{d+1} Y_j$ (Puccetti and Rüschendorf 2012, Theorem 2.1), i.e., is antimonotonic. This observation lies at the core of this rearrangement method.

In fact, it must actually hold that for any decomposition of $\{1, \ldots, d+1\} = I_1 \cup I_2$ into two disjoint sets I_1 and I_2, the sums $S_1 := \sum_{k \in I_1} Y_j$ and $S_2 := \sum_{k \in I_2} Y_j$ are antimonotonic and not only for singleton sets of the form $I_1 = \{j\}$. This observation makes it possible to generalize the standard RA by rearranging "blocks of columns" instead of one column at a time: the columns in the first set I_1 are stacked into a matrix (block) $\mathbf{Y}_1$, and its rows are rearranged (i.e., entire rows are swapped) such that the row sums of $\mathbf{Y}_1$ (reflecting S_1) are in increasing order. As for the matrix $\mathbf{Y}_2$ that is formed by stacking the remaining columns, the rows are rearranged such that the row sums (reflecting S_2) are in decreasing order.

Algorithm for inferring dependence

1. Select a random sample of n_{sim} possible partitions of the columns $\{1,\dots,d+1\}$ into two nonempty subsets $\{I_1, I_2\}$. In our case, $d = 9$ and it is feasible to consider all nontrivial partitions, thus we take $n_{sim} = 2^d - 1$.

2. For each of the n_{sim} partitions, create the matrices (blocks) $\mathbf{Y}_1$ and $\mathbf{Y}_2$ with corresponding row sums S_1 and S_2 and rearrange rows of $\mathbf{Y}_2$ so that S_2 is antimonotonic to S_1.

3. If there is no improvement in var $\left(\sum_{j=1}^{d+1} Y_j\right)$, output the current matrix $\mathbf{M}$; otherwise, return to step 1.

At each step of this algorithm, we ensure that the variance decreases or remains the same: the columns, say Y_j before rearranging and $\widetilde{Y}_j$ after rearranging, satisfy the inequality[38]

$$\text{var}\left(\sum_{j=1}^{d+1} Y_j\right) \geqslant \text{var}\left(\sum_{j=1}^{d+1} \widetilde{Y}_j\right).$$

The method we propose for inferring dependence is inspired by the so-called rearrangement algorithm (RA) of Puccetti and Rüschendorf (2012) and of Embrechts, Puccetti, and Rüschendorf (2013), which was originally introduced to deal with the assessment of model risk, adjusted by Bernard and McLeish (2016) and Bernard, Bondarenko, and Vanduffel (2018) to make it suitable for inferring dependence, and termed the block rearrangement algorithm (BRA).

Toy Example of BRA

To explain the method for constructing a joint dependence, we provide here an oversimplified example. There are $d = 3$ assets and $n = 5$ states of the world. Therefore, Y_1, Y_2, Y_3 and $-S$ all take five values with probability $1/5$, which are collected in a matrix:

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & 0 & -19 \\ 2 & 2 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 5 & 5 & 5 & -8 \\ 6 & 7 & 9 & -6 \end{bmatrix}. \tag{22}$$

The first three columns of matrix $\mathbf{M}$ depict the random vector (Y_1, Y_2, Y_3). Its joint outcomes are displayed in the five rows, each row reflecting one of the five states of the world. The random vector (Y_1, Y_2, Y_3) does not yet describe a compatible dependence, because the five row sums

[38] Indeed, var $\left(\sum_{k=1}^{d+1} Y_k\right)$ = var $\left(Y_j + \sum_{k\neq j} Y_k\right)$, and a necessary condition for var$\left(\sum_{k=1}^{d+1} Y_k\right)$ to become minimum is that each Y_j is antimonotonic with $\sum_{k\neq j} Y_k$.

(taken over all four columns) do not equal to zero; i.e., we do not yet meet the constraint that $Y_1 + Y_2 + Y_3 - S = 0$. However, permuting elements within a column is allowed, as it does not affect the marginal distributions. We thus aim at permuting elements within columns to satisfy the condition that $Y_1 + Y_2 + Y_3 - S = 0$ at each of the five states.

With only five rows, it is feasible to try every possible permutation. However, in a realistic situation, this would be impractical, as the number of distinct configurations, $(n!)^d$, would be extremely large. In our empirical application, we employ our method using at least $n = 1,000$ states and, therefore, we need an efficient way to find a candidate solution. To achieve such efficiency, we observe that the condition $Y_1 + Y_2 + Y_3 - S = 0$ is equivalent to the condition that the random variable $Y_1 + Y_2 + Y_3 - S$ has zero variance. Clearly, to minimize the variance of $Y_1 + Y_2 + Y_3 - S$, it must hold that Y_1 is as negatively correlated as possible with $Y_2 + Y_3 - S$, which means that the elements of the first column of the matrix $\mathbf{M}$ in (22) should appear in opposite order (be antimonotonic) to those that correspond to $Y_2 + Y_3 - S$. Since permuting (rearranging) values within columns does not affect the marginal distributions, we rearrange the values in the first column to achieve this situation.

Let us illustrate this fundamental principle using the first column of the matrix in (22). We rearrange this column such that its realizations are placed in opposite order to the realizations of $Y_2 + Y_3 - S$. After this step, we obtain the matrix $\mathbf{M}^{(1)}$:

$$Y_2 + Y_3 - S = \begin{bmatrix} -18 \\ -8 \\ -3 \\ 2 \\ 10 \end{bmatrix} \qquad \mathbf{M}^{(1)} = \begin{bmatrix} 6 & 1 & 0 & -19 \\ 5 & 2 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 2 & 5 & 5 & -8 \\ 1 & 7 & 9 & -6 \end{bmatrix}. \tag{23}$$

For the starting configuration $\mathbf{M}$, the variance of row sums is $\text{var}(Y_1 + Y_2 + Y_3 - S) = 126$. After rearranging the first column, we obtain $\mathbf{M}^{(1)}$ and $\text{var}(Y_1 + Y_2 + Y_3 - S) = 58$ has been strictly decreased. We have not found a solution yet, but we are getting one step closer. We now repeat the same process for each of the subsequent columns of the matrix. We can further improve this procedure by noting that in order to yield zero row sums, we actually need $Y_1 + Y_2$ to be antimonotonic to $Y_3 - S$ and, likewise, $Y_1 + Y_3$ to be antimonotonic to $Y_2 - S$, and $Y_2 + Y_3$ to $Y_1 - S$. To simplify presentation, in what follows, we denote $Y_4 = -S$. For the original matrix, we have $V := \text{var}(Y_1 + Y_2 + Y_3 + Y_4) = 126$.

Step 1: Rearranging Y_1,

$$Y_2 + Y_3 + Y_4 = \begin{bmatrix} -18 \\ -8 \\ -3 \\ 2 \\ 10 \end{bmatrix} \qquad \mathbf{M}^{(1)} = \begin{bmatrix} 6 & 1 & 0 & -19 \\ 5 & 2 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 2 & 5 & 5 & -8 \\ 1 & 7 & 9 & -6 \end{bmatrix} \qquad V = 58.$$

Step 2: Rearranging Y_2,

$$Y_1 + Y_3 + Y_4 = \begin{bmatrix} -13 \\ -5 \\ -3 \\ -1 \\ 4 \end{bmatrix} \qquad \mathbf{M}^{(2)} = \begin{bmatrix} 6 & 7 & 0 & -19 \\ 5 & 5 & 3 & -13 \\ 3 & 3 & 4 & -10 \\ 2 & 2 & 5 & -8 \\ 1 & 1 & 9 & -6 \end{bmatrix} \qquad V = 12.4.$$

Step 3: Rearranging Y_3,

$$Y_1 + Y_2 + Y_4 = \begin{bmatrix} -6 \\ -3 \\ -6 \\ -4 \\ -4 \end{bmatrix} \qquad \mathbf{M}^{(3)} = \begin{bmatrix} 6 & 7 & 9 & -19 \\ 5 & 5 & 0 & -13 \\ 3 & 3 & 5 & -10 \\ 2 & 2 & 3 & -8 \\ 1 & 1 & 4 & -6 \end{bmatrix} \qquad V = 4.$$

In this case, the order of the fourth and fifth rows for Y_3 is arbitrary, and the rearrangement is not unique. Both lead to the same new variance of 4.

Step 4: Rearranging Y_4,

$$Y_1 + Y_2 + Y_3 = \begin{bmatrix} 22 \\ 10 \\ 11 \\ 7 \\ 6 \end{bmatrix} \qquad \mathbf{M}^{(4)} = \begin{bmatrix} 6 & 7 & 9 & -19 \\ 5 & 5 & 0 & -10 \\ 3 & 3 & 5 & -13 \\ 2 & 2 & 3 & -8 \\ 1 & 1 & 4 & -6 \end{bmatrix} \qquad V = 2.8.$$

Step 5: Rearranging the block $[Y_1 \;\; Y_2]$ does not reduce the variance, as it is already (weakly) antimonotonic. We can keep it unchanged or swap the first and second rows. We do the latter to

illustrate the algorithm:

$$Y_1 + Y_2 = \begin{bmatrix} 13 \\ 10 \\ 6 \\ 4 \\ 2 \end{bmatrix} \qquad Y_3 + Y_4 = \begin{bmatrix} -10 \\ -10 \\ -8 \\ -5 \\ -2 \end{bmatrix} \qquad \mathbf{M}^{(5)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 6 & 7 & 0 & -10 \\ 3 & 3 & 5 & -13 \\ 2 & 2 & 4 & -8 \\ 1 & 1 & 3 & -6 \end{bmatrix} \qquad V = 2.8.$$

Step 6: Rearranging the block $[Y_1 \;\; Y_3]$ does not help either, as it is already antimonotonic:

$$Y_1 + Y_3 = \begin{bmatrix} 14 \\ 6 \\ 8 \\ 6 \\ 4 \end{bmatrix} \qquad Y_2 + Y_4 = \begin{bmatrix} -14 \\ -3 \\ -10 \\ -6 \\ -5 \end{bmatrix} \qquad \mathbf{M}^{(6)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 6 & 7 & 0 & -10 \\ 3 & 3 & 5 & -13 \\ 2 & 2 & 4 & -8 \\ 1 & 1 & 3 & -6 \end{bmatrix} \qquad V = 2.8.$$

Step 7: Rearranging the block $[Y_1 \;\; Y_4]$, we need to swap the second, fourth, and fifth rows:

$$Y_1 + Y_4 = \begin{bmatrix} -14 \\ -4 \\ -10 \\ -6 \\ -5 \end{bmatrix} \qquad Y_2 + Y_3 = \begin{bmatrix} 14 \\ 7 \\ 8 \\ 6 \\ 4 \end{bmatrix} \qquad \mathbf{M}^{(7)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 1 & 7 & 0 & -6 \\ 3 & 3 & 5 & -13 \\ 6 & 2 & 4 & -10 \\ 2 & 1 & 3 & -8 \end{bmatrix} \qquad V = 2.$$

Step 8-11: Now, we go back to Step 1 and look again at columns Y_1, Y_2, Y_3, and Y_4. Columns Y_1, Y_3, and Y_4 are already antimonotonic, while rearranging column Y_2 decreases the variance to 1.6:

$$\mathbf{M}^{(11)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 1 & 7 & 0 & -6 \\ 3 & 3 & 5 & -13 \\ 6 & 1 & 4 & -10 \\ 2 & 2 & 3 & -8 \end{bmatrix} \qquad V = 1.6.$$

Step 12: Next, we apply again the rearrangement on block $[Y_1 \;\; Y_2]$. We switch rows 2 and 3 and find that the variance is equal to 0. The algorithm has converged and the final matrix is

$$\widetilde{\mathbf{M}} = \mathbf{M}^{(12)} = \begin{bmatrix} 5 & 5 & 9 & -19 \\ 3 & 3 & 0 & -6 \\ 1 & 7 & 5 & -13 \\ 6 & 1 & 4 & -10 \\ 2 & 2 & 3 & -8 \end{bmatrix} \qquad V = 0. \tag{24}$$

We have arrived to the ideal situation in which the row sums of the rearranged matrix are all equal to zero; i.e., we have found an admissible multivariate model for the assets (Y_1, Y_2, Y_3), which is consistent with the distribution of their sum S and which can now be used to compute various statistics of interest. For instance, we can find the conditional probability $P(Y_2 > 3|Y_1 > 3) = 0.5$. To specifically study the dependence among the three assets, we can remove the effect of the marginal distributions by applying the transformation $Y_j \to G_j(Y_j)$, $j = 1, 2, 3$. Then we obtain the following discrete dependence structure:

$$\widetilde{\mathbf{U}} = \begin{bmatrix} 0.8 & 0.8 & 1 \\ 0.6 & 0.6 & 0.2 \\ 0.2 & 1 & 0.8 \\ 1 & 0.2 & 0.4 \\ 0.4 & 0.4 & 0.6 \end{bmatrix}. \tag{25}$$

In practice, due to discretization errors and the fact that the algorithm is a heuristic, the variance of the row sums might not be exactly equal to zero. However, empirically the procedure usually performs extremely well. Although the final row sums do deviate from zero, deviations are trivial for all practical purposes and could be safely ignored.

Illustration of BRA for $d = 2$

There are two assets, whose returns X_1 and X_2 are normally distributed with zero mean and standard deviations $\sigma_1 = 0.2$ and $\sigma_2 = 0.4$. The two marginal distributions are fixed, but we vary the distribution for the weighted sum $S = \frac{1}{2}X_1 + \frac{1}{2}X_2$. We consider three cases:

(1) S is $N(0, \sigma_S^2)$, where σ_S is chosen such that the implied correlation is 0;

(2) S is $N(0, \sigma_S^2)$, where σ_S is chosen such that the implied correlation is 0.97;

(3) S follows a skewed distribution with a heavy left tail, modeled by a mixture of two normals.

The three cases (no dependence, strong dependence, and asymmetric dependence) are shown in Figures IA.1–IA.3. In each figure, the top panels show PDFs and CDFs of X_1, X_2 and S, which are all centered at zero. We discretize the three CDFs into $n = 1,000$ equiprobable states (only ten of them are shown on the plots) and run BRA to extract the dependence and joint distribution between X_1 and X_2. Those are represented by 1,000 dots and are shown in the two bottom panels.

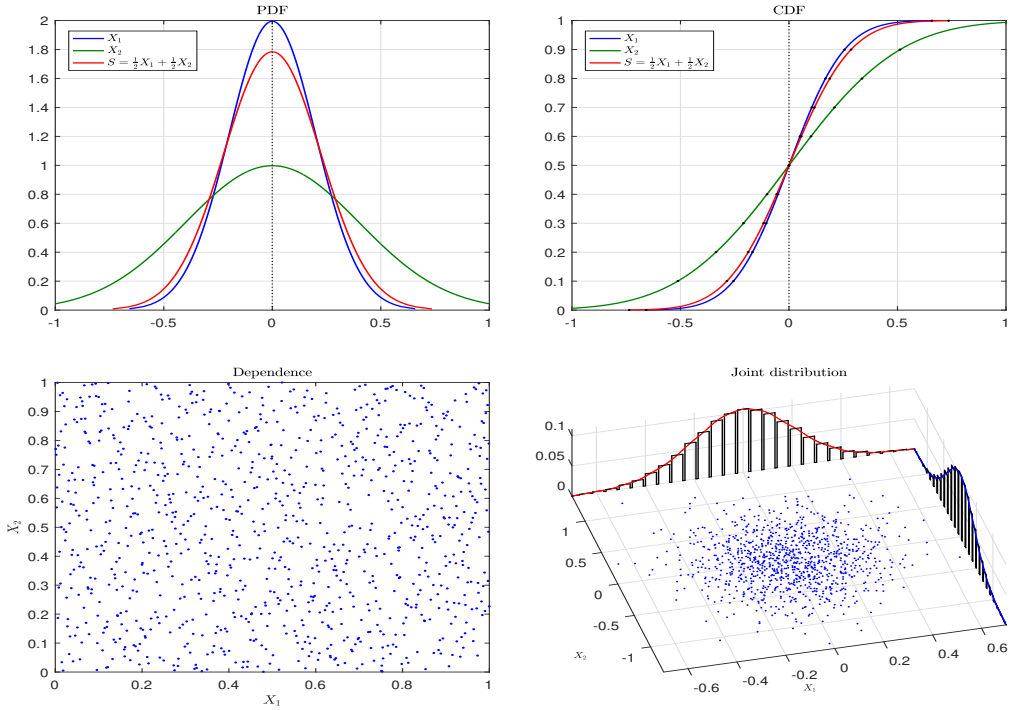

Figure IA.1: **No Dependence.**

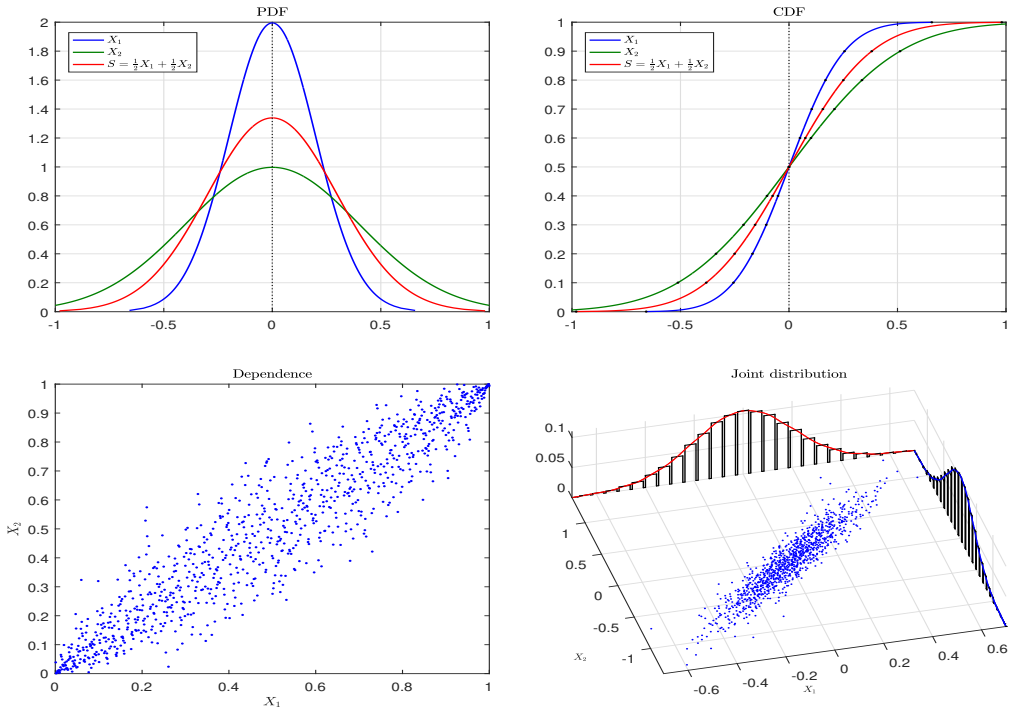

Figure IA.2: **Strong Dependence.**

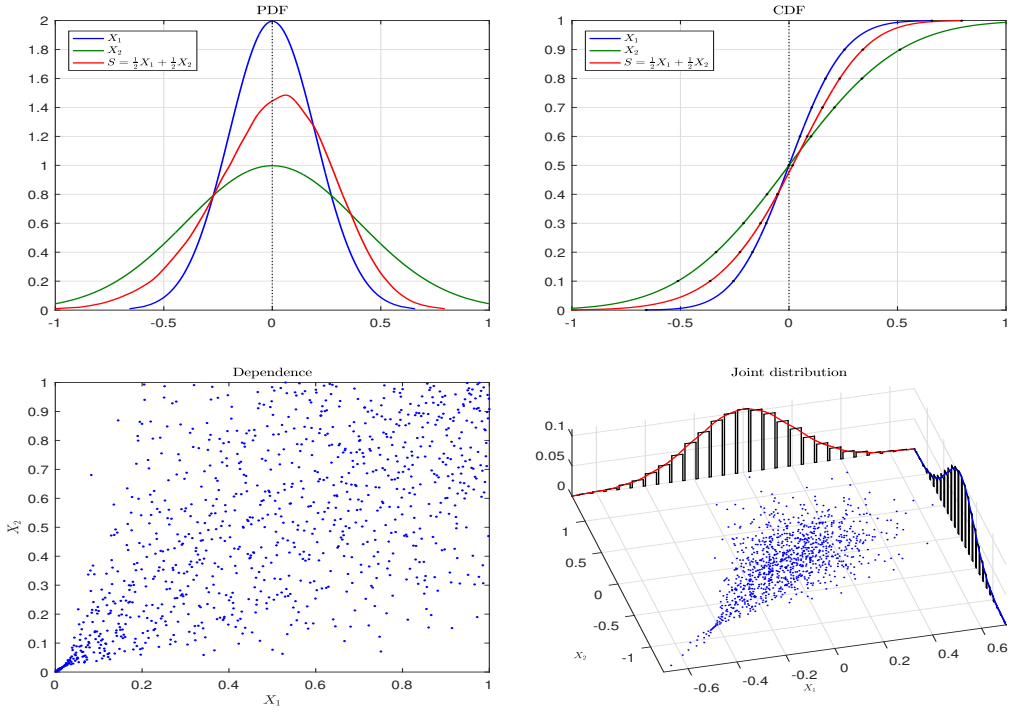

Figure IA.3: **Asymmetric Dependence.**

IA.B Additional Case Studies

We report the option-implied dependence for two additional selected dates: October 20, 2017 and March 23, 2020. The first date represents a very calm period in the midst of a historic bull market. On that day, the market 3-month implied volatility was only 0.084. The market return was 22.7% over the trailing 12-month period and 10.7% over the following 3 months. The second one represents a very turbulent period at the peak of COVID-19 crisis, which was followed by an extreme V-shaped recovery. We show dependences in Figures IA.4 and IA.5 and correlations in Figures IA.6 and IA.7. We observe that Financial, Energy, and Technology sectors are highly correlated with themselves and the other sectors. The best diversifiers are Materials and Utilities. The pairwise correlations are higher across the board for the second date compared to the first date. The average global correlation $\rho^{\mathbb{Q}}$ is 0.52 for October 20, 2017 and 0.74 for March 23, 2020. From Figures IA.6 and IA.7, it is clear that the down correlations again tend to be much higher than the up correlations. The average down and up correlations are 0.54 and 0.06 for October 20, 2017 and 0.64 and 0.35 for March 23, 2020. It is interesting to observe that, on the second date, there is evidence of a strong *right-tail* dependence (in addition to an even stronger *left-tail* dependence). In particular, the middle panels of Figure IA.5 display a strong dependence in the deep right tail for both Financial and Utilities sectors. This suggests that, in the midst of the COVID-19 crisis, prices of sector options reflected a distinct possibility of a large market rally, with most sectors increasing at the same time. The market indeed had a stunning recovery,

increasing by more than 35% over the next three months.

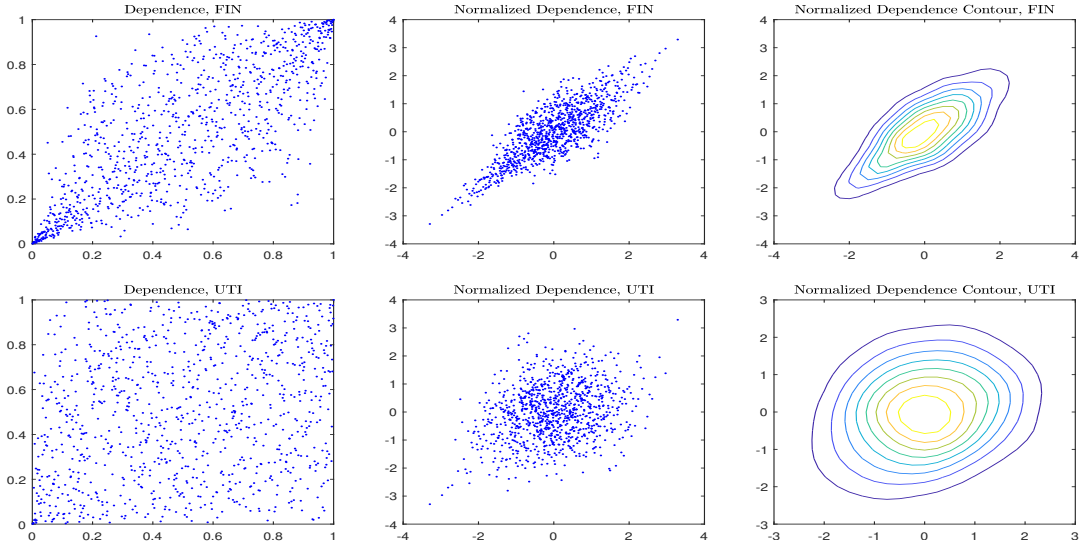

Figure IA.4: **Implied dependence on October 20, 2017.** The first column shows the dependence of the Financial sector (top panels) and the Utilities sector (bottom panels) relative to the S&P 500 index. The middle column shows the same dependence but for normally distributed variables. The third column displays the corresponding contour plots.

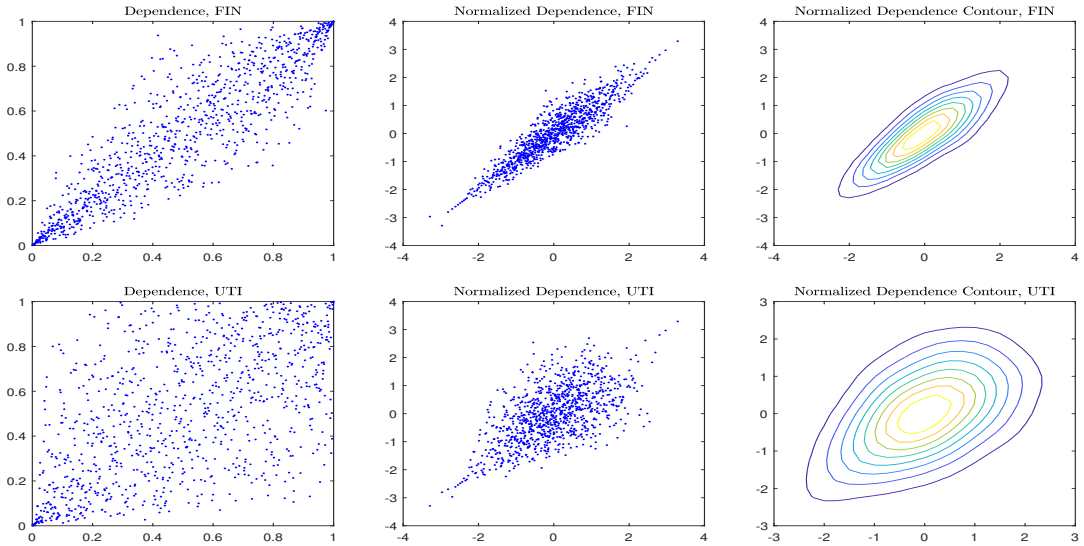

Figure IA.5: **Implied dependence on March 23, 2020.** The first column shows the dependence of the Financial sector (top panels) and the Utilities sector (bottom panels) relative to the S&P 500 index. The middle column shows the same dependence but for normally distributed variables. The third column displays the corresponding contour plots.

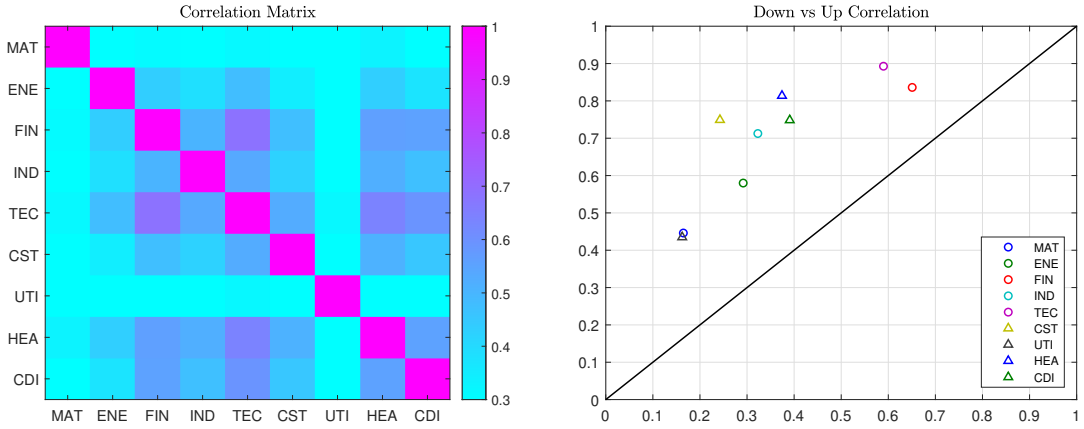

Figure IA.6: **Implied Correlations for the Nine Sectors on October 20, 2017.** The left panel shows the correlation matrix. The right panel shows the implied down correlation $\rho_{j,S}^{d,\mathbb{Q}}$ (y-axis) versus the implied up correlation $\rho_{j,S}^{u,\mathbb{Q}}$ (x-axis). Also shown is the 45-degree line.

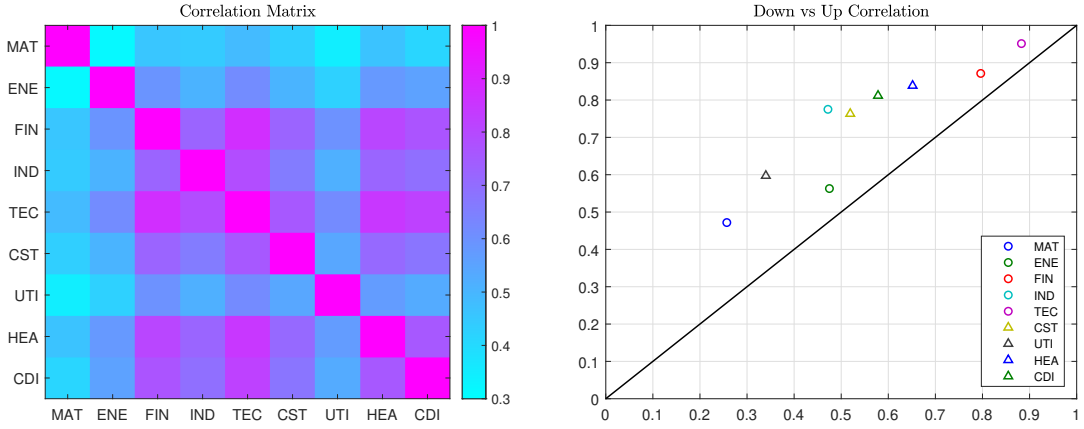

Figure IA.7: **Implied Correlations for the Nine Sectors on March 23, 2020.** The left panel shows the correlation matrix. The right panel shows the implied down correlation $\rho_{j,S}^{d,\mathbb{Q}}$ (y-axis) versus the implied up correlation $\rho_{j,S}^{u,\mathbb{Q}}$ (x-axis). Also shown is the 45-degree line.

IA.C Stability of the Algorithm

To assess the stability of our MFDR, we conduct a Monte-Carlo study. We distinguish between two separate effects: (i) instability of the optimization algorithm (BRA) itself, and (ii) instability of the inputs (i.e., measurement errors in the marginal distributions F_j and F_S estimated in the first step of MFDR). The first effect (unrealistically) assumes that the inputs are measured perfectly and focuses on how much noise is introduced by the optimization algorithm. Since the BRA is an heuristic approach, it finds an *approximate* solution. In addition, since the algorithm relies on

initial randomization, each run actually yields a *different* approximate solution. However, as we show next, the first effect is extremely small and can easily be ignored in practice. Moreover, the effect can be made arbitrarily small by increasing the number of states n. The second effect is not specific to our method only. The existing approaches need estimates of MFIV or $\sigma^{\mathbb{Q}}$, and those too require the knowledge of the whole distribution. Options prices are recorded with considerable measurement errors (due to illiquidity, non-synchronous trading, and other market imperfections) and these errors affect the estimation of the corresponding risk-neutral distributions and derived measures. The second effect is typically much larger than the first one.

To assess the first effect, we proceed as follows. For each day in our sample, we run the algorithm $N_{sim} = 10$ times for fixed inputs. As before, we use the distributions F_j and F_S estimated from midpoint prices of available options. This gives us N_{sim} joint distributions, which are not identical due to noise introduced by the BRA. From each joint distribution we compute the corresponding correlations ρ^g, ρ^d, and ρ^u and then bootstrap 50,000 times to obtain the confidence intervals for the time-series averages. The results for this experiment are reported in the first three columns in Table IA.1. The confidence intervals are extremely narrow. For example, the widest 1%–99% interval is for the up correlation and it is only 0.00017. As it turns out, the confidence intervals for a *single* day are already very narrow, but the intervals calculated for the *sample averages* become negligible.[39]

The last three columns in Table IA.1 report on the second effect. We conduct a similar experiment, except now joint distributions are estimated from *perturbed* option prices. We want to add noise to option prices and assess how this affects the estimates of the three types of correlations. It is not completely obvious how to model realistic measurement errors. Our approach is to assume that the true, but unobserved, option price is uniformly distributed between the bid and the ask. That is, instead of the midpoint prices, we use simulated prices. This leads to N_{sim} sets of perturbed marginal distributions F_j and F_S and, thus, perturbed joint distributions. The last three columns in Table IA.1 demonstrate that the second effect is 5 times or more larger than the first effect, but it is still very small in economic terms. Again, the confidence intervals for the up correlation are widest because the information on the right tail of RND is typically less precise than for the left tail (there are more OTM puts available than OTM calls). The 1%–99% interval for the up correlation is now 0.00086.

[39] As mentioned earlier, the accuracy of the BRA can be further improved by increasing the number of states, to, say, $n = 10,000$. This only increases the computation time but has no real effect on the empirical results.

	Fixed Inputs			Perturbed Inputs		
	$P_{0.90}$-$P_{0.10}$	$P_{0.95}$-$P_{0.05}$	$P_{0.99}$-$P_{0.01}$	$P_{0.90}$-$P_{0.10}$	$P_{0.95}$-$P_{0.05}$	$P_{0.99}$-$P_{0.01}$
Global	0.000001	0.000083	0.000098	0.00025	0.00050	0.00047
Down	0.000001	0.000106	0.000126	0.00033	0.00064	0.00061
Up	0.000002	0.000151	0.000178	0.00047	0.00092	0.00086

Table IA.1: **Stability of MFDR.** The table reports the size of the confidence intervals for the average global, down, and up correlations. The confidence intervals are obtained via bootstrap for the two cases: (1) when the BRA inputs are fixed (the first three columns), and (2) when the BRA inputs are perturbed (the last three columns). The confidence intervals are computed as the difference between percentiles at 10% and 90%, 5% and 95%, and 1% and 99%, respectively.

IA.D Down and Up Variance and Correlation Risk Premia

Let B denote a regime variable which takes on values of zero and one with probability of $p = 0.5$. This binary variable switches between two regimes, L and H ("Low" and "High"). As before, the vector $(X_1, ..., X_d)$ denotes returns for the d sectors, so that

$$(X_1, ..., X_d) = (X_1^L, ..., X_d^L)\mathbb{I}_{B=0} + (X_1^H, ..., X_d^H)\mathbb{I}_{B=1}.$$

In each regime ($B = 0$ or $B = 1$), the d returns are multivariate normal with homogeneous parameters. In particular, in regime L, all components have the same expected return $\mu_i^L = \mu^L$, same individual volatility $\sigma_i^L = \sigma^L$, and same pairwise correlation $\rho_{ij}^L = \rho^L$. Similar notation applies for regime H. Generally, the parameters differ under the physical probability $\mathbb{P}$ and the risk-neutral probability $\mathbb{Q}$. Under $\mathbb{Q}$, the expected return must be equal to the risk-free rate r assumed to be zero, i.e., $0.5\mu^{L,\mathbb{Q}} + 0.5\mu^{H,\mathbb{Q}} = 0$. We keep the parameters of the mixture model fixed under $\mathbb{P}$, but consider three different scenarios under $\mathbb{Q}$, as reported in Table IA.2. We focus on the equally weighted index $S = \sum_{j=1}^{d} \omega_j X_j$ with $\omega_j = 1/d$ and $d = 9$. In each scenario, we simulate $N_{sim} = 10,000,000$ returns, compute the corresponding VRP and CRP (global, down, and up), and report the results in Table IA.3.

	μ^L	μ^H	σ^L	σ^H	ρ^L	ρ^H
Under $\mathbb{P}$	-0.05	0.15	0.25	0.20	0.7	0.3
Under $\mathbb{Q}$, Scenario 1	-0.10	0.10	0.30	0.10	0.7	0.3
Under $\mathbb{Q}$, Scenario 2	-0.10	0.10	0.30	0.15	0.7	0.0
Under $\mathbb{Q}$, Scenario 3	-0.10	0.10	0.30	0.15	0.6	0.4

Table IA.2: **Parameters of the Mixture Model under $\mathbb{P}$ and $\mathbb{Q}$.**

	VRP^g	VRP^d	VRP^u	CRP^g	CRP^d	CRP^u
Scenario 1	-0.0044	-0.0140	0.0014	-0.098	-0.125	-0.138
Scenario 2	-0.0038	-0.0153	0.0017	-0.008	-0.089	0.033
Scenario 3	-0.0038	-0.0067	0.0005	-0.008	0.020	-0.047

Table IA.3: **Variance and Correlation Risk Premia in the Mixture Model.** The table reports global, down, and up VRP and CRP for the three scenarios, whose parameters are given in Table IA.2.

IA.E DJIA Options

As a robustness check, we applied our MFDR methodology to options on the Dow Jones Industrial Average (DJIA). To obtain the composition of the DJIA index, we used data from Compustat and merged it with data from CRSP using the CCM Linking Table. Specifically, we followed the approach of Dobelman, Kang, and Park (2014) and used GVKEY and IID to link to PERMNO. The data on returns and market prices are obtained from CRSP. To approximate the index weights on each day, we used the price of each stock in the index from the previous day as a proxy.

To match the historical data with options, we used the historical CUSIP link provided by OptionMetrics. We used the DJIA index directly as the underlying asset for options (ticker DIA). However, applying the MFDR methodology to DJIA options posed two challenges. Firstly, in the first step of MFDR, we needed to estimate the RND for each index component, which required options with a wide range of strikes, that densely covered the support of the RND. However, many stocks did not satisfy this requirement, especially in the early part of the sample. To overcome this challenge, we relied on the Volatility Surface File from OptionMetrics, and selected options with 91 days to maturity for each underlying asset. The Volatility Surface File contained the interpolated volatility surface for each security on each date for the set of standardized strikes with (absolute) delta from 0.1 to 0.9. The surface data involved extensive inter- and extrapolations of the market data with a smoothing kernel. As such, the surface data is not suitable for testing trading rules, but it proved to be a valuable source of information that could be used to generate signals for trading as demonstrated by previous studies such as DeMiguel, Plyakha, Uppal, and Vilkov (2013) and Driessen, Maenhout, and Vilkov (2009), among others.

The second challenge we encountered when applying the MFDR methodology to DJIA options was the increased computational demands required to solve the optimization problem in the third step of MFDR. The computational time of the BRA method increased exponentially with the number of components, making it infeasible for $d = 30$. To speed up the solution-finding process, we used a variant of the original method called the "randomized" BRA. This method cycles through a random subset of partitions of columns, instead of onsidering all possible partitions. Additionally, the method assigns higher probabilities to smaller blocks. While the randomized BRA method finds the solution much faster, there is a slight loss of accuracy. However, this has

a negligible effect on the accuracy of the aggregate measures (such as the global, down, and up correlations).[40]

We present the results in Table IA.4 and Figure IA.8. Over the 14-year sample period, the average implied global, down, and up correlations are 0.538, 0.507, and 0.144 for the 30 stocks in DJIA as compared to 0.678, 0.527, and 0.444 for the nine sectors in S&P 500 (Table 4). This suggests that the correlations for individual stocks are much lower than those for sectors. The implied dependence for individual stocks is still highly asymmetric, with the down correlation being much larger than the up correlation.

	N_{obs}	Under $\mathbb{P}$	Under $\mathbb{Q}$	Premium	t-stat
Global	3475	0.415	0.538	-0.123	-10.8
Down	3475	0.284	0.507	-0.222	-16.5
Up	3475	0.228	0.144	0.084	7.5
Down-Up	3475	0.056	0.363	-0.307	-21.2

Table IA.4: **Correlation Risk Premium for DJIA.** The table reports statistics for the risk premia θ, θ^d, and θ^u computed for the average global, down, and up correlations estimated for DJIA index. The last row is the correlation spread, $\Delta\rho = \rho^d - \rho^u$. The last column shows the Newey-West t-statistics computed with 63 lags.

Figure IA.8 plots across time the average implied correlation (blue line) and realized correlation (red line). There are three panels, corresponding to the three types of correlations (global, down, and up). The first panel documents that the global correlation risk premium (CRP) θ, which appears as the difference between the blue and the red lines, is mostly negative. Across the sample period, it has an average of -0.123 and with a t-statistic of -10.8, see Table IA.4. The second and third panels demonstrate that the average realized down correlations are systematically lower than their implied counterparts, while the opposite is true for the up correlations. The up correlations are lower than the down correlations under $\mathbb{P}$ and this asymmetry is more pronounced under $\mathbb{Q}$, with $\rho^{u,\mathbb{Q}} < \rho^{u,\mathbb{P}} < \rho^{d,\mathbb{P}} < \rho^{d,\mathbb{Q}}$. As a result, the average down CRP θ^d is negative (-0.222), while the average up CRP θ^u is positive (0.084). Both risk premia are highly significant with t-statistic of -16.5 and 7.5, respectively.

[40] On our computer cluster, the third step of MFDR can be computed for each day in approximately 2 seconds for the nine sectors of the S&P 500 and approximately 12 seconds for the 30 stocks of the DJIA. Thus, the randomized BRA for $d = 30$ takes about 6 times longer than the original BRA for $d = 9$. The total computational time for all 3500 days in the sample is greatly reduced by utilizing parallel computing.

We estimate that the third step of MFDR would take 10 to 40 times longer for the 100 stocks in the S&P 100 than for the DJIA. Without further steps to accelerate the procedure, an application to the 500 stocks in the S&P 500 would unlikely be feasible.

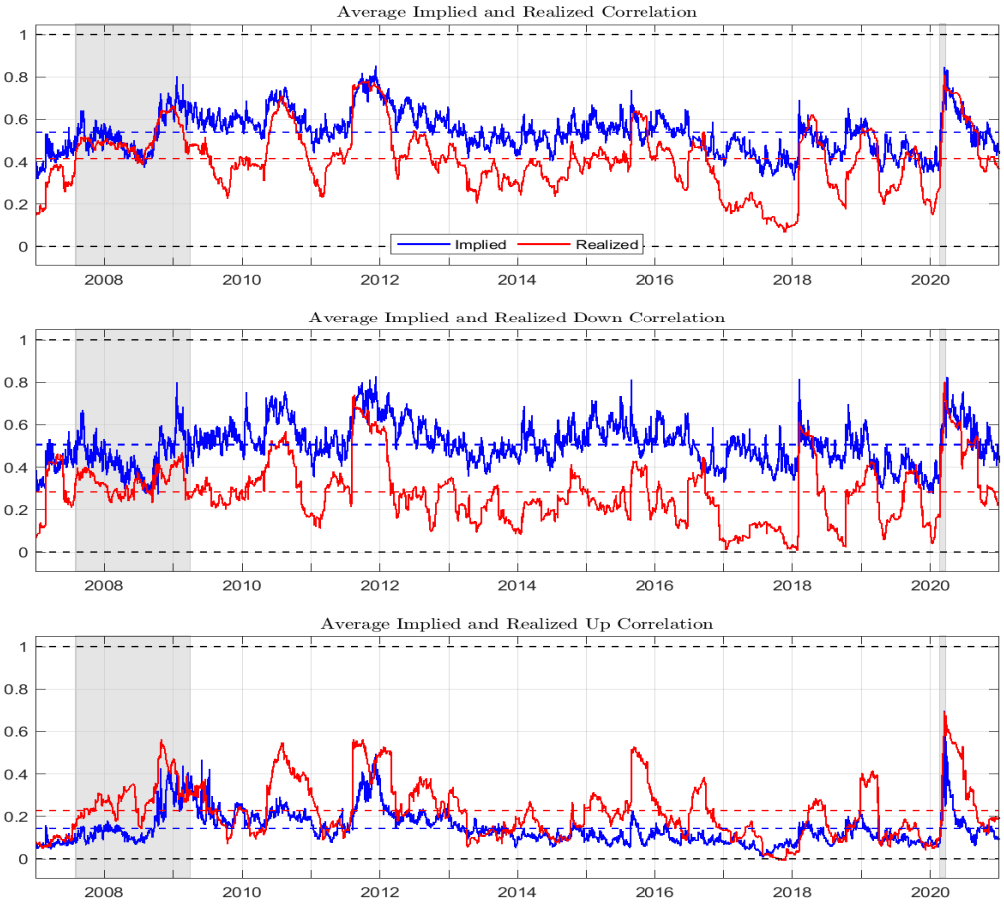

Figure IA.8: **Implied and Realized Correlations for DJIA**. Implied correlations (blue) are computed from option-implied dependence; realized correlations (red) are computed from sector index returns. The corresponding means of the two series are shown with the horizontal dashed lines. The grey shaded areas indicate the financial crisis and COVID-19 crisis.