


# Option-Implied Dependence and Correlation Risk Premium

Oleg Bondarenko 

University of Illinois at Chicago Department of Finance  
olegb@uic.edu (corresponding author)

Carole Bernard 

Grenoble Ecole de Management Department of Accounting, Law and Finance, and Vrije Universiteit Brussel (VUB)  
Department of Economics and Political Sciences  
carole.bernard@grenoble-em.com

## Abstract

We propose a novel model-free approach to obtain the joint risk-neutral distribution among several assets that is consistent with options on these assets and their weighted index. We implement this approach for the nine industry sectors comprising the S&P 500 index and find that their option-implied dependence is highly asymmetric and time-varying. We then study two conditional correlations: when the market moves down or up. The risk premium is strongly negative for the down correlation but positive for the up correlation. Intuitively, investors dislike the loss of diversification when markets fall, but they actually prefer high correlation when markets rally.

## 1. Introduction

Option markets provide rich information about assets future returns. There exists a no-arbitrage relationship that links prices of options to the risk-neutral density (RND) of an asset future return. First discovered by Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978), this fundamental relationship is one of the most useful in financial economics. Building on the relationship, researchers have developed effective techniques to estimate the risk-neutral

---

We thank Jennifer Conrad (the editor) and Lorenzo Schoenleber (the referee) for many constructive suggestions that helped to improve the article. For valuable comments, we also thank Torben G. Andersen (discussant), Nicole Branger, Adrian Buss (discussant), Ian Dew-Becker, Hamed Ghanbari, Jens Jackwerth, Philippe Mueller (discussant), Paul Schneider (discussant), Fabio Trojani, Steven Vanduffel, Grigoriy Vilkov, Alex Weissensteiner, and seminar participants at the 2018 CDI Conference in Montreal, 2018 Conference in Honor of Bruno Dupire, 2019 Conference on Risk Management and Financial Innovation, 2019 ITAM Finance Conference, 2019 EFMA Meeting, 2019 Risk Day organized by ETH Zurich, 2018 CUNEF Madrid workshop on “New Frontiers in Financial Markets,” 2021 Vienna workshop on “Econometrics of Option Market,” 2020 WFA meeting, the Collegio Carlo Alberto, Northwestern University, NYU, the University of Lugano, and the University of Illinois at Chicago. An earlier version of this article was titled “Option-Implied Dependence” (with Steven Vanduffel). We also gratefully acknowledge funding from the Canadian Derivatives Institute (formerly IFSID), the FWO research grant FWO G015320N at the Vrije Universiteit Brussel and the Global Risk Institute (GRI).

distributions in a model-free way (see, e.g., Jackwerth and Rubinstein (1996), Aït-Sahalia and Lo (2000), and Bondarenko (2003)). Option-implied distributions have since then been used in numerous applications. However, the fundamental relationship only works in one dimension. It allows researchers to obtain the *individual* risk-neutral distributions for future returns of stock *A* and stock *B*, but not their *joint* distribution. In this article, a novel extension to higher dimensions is proposed. Specifically, the options written on individual assets *and* on their index are used to fully describe the forward-looking risk-neutral dependence among the assets.

We refer to our approach as MFDR, or model-free dependence recovery. The approach consists of 3 main steps. First, we estimate risk-neutral marginal distributions of individual assets and also of their weighted sum (the index). As inputs, this step requires traded options on the individual assets and the index. Second, we frame the problem of finding a joint distribution among the assets as an integer optimization problem, in which a matrix of asset returns must be arranged in a suitable manner. Third, to solve the resulting large-scale optimization problem, we rely on a combinatorial technique, termed the block rearrangement algorithm (BRA) (see Bernard and McLeish (2016)). The BRA is the key to our approach, as alternative solution techniques are simply not feasible for the problem at hand.

Our MFDR methodology derives a complete description of the implied dependence and, as such, compares favorably with existing methods that primarily focus on the average pairwise correlation. As Buss, Schönleber, and Vilkov (2019a) point out, “[c]omputing the historical pairwise correlation among any two stocks is rather easy; however, computing an expected pairwise correlation from option data is, in practice, not possible.” Thus, additional simplifying assumptions must be imposed to obtain an estimate. A common assumption is that of a constant pairwise correlation, which makes the recovery of the single parameter possible by comparing the variances of the index and individual components. This approach has been employed by the Chicago Board Options Exchange (CBOE) that has been disseminating its S&P 500 Implied Correlation Indices since July 2009. These correlation indices are now widely accepted dependence measures. In academic literature, related approaches are followed among others by Driessen, Maenhout, and Vilkov (2009), (2013), Buraschi, Kosowski, and Trojani (2013), and Faria, Kosowski, and Wang (2018), who also assume constant pairwise correlations.

Compared to the existing methods, the MFDR approach offers two critical advantages. First, as the name suggests, the approach is completely model-free and requires no parametric assumptions. Second and most importantly, it yields a full dependence structure, not just a partial dependence measure. While the existing methods equate the index variance with that of the weighted sum of its components (one moment condition), MFDR matches their whole distributions (theoretically, infinitely many conditions). This results in an essentially perfect fit of the index implied volatility curve, instead of only matching one of its statistics. Since MFDR yields a feasible dependence structure, a proper correlation matrix is assured by construction, without the need for any additional assumptions. Finally, while there might be many compatible dependence structures, MFDR is shown to maximize entropy and thus it yields the “most likely” implied dependence among asset returns given the information contained in available options.

We implement MFDR using ETF options on S&P 500 index and its nine industry sectors. The industry sectors provide an ideal setting for our methodology.

On the one hand, the joint distribution of the nine sectors is rich and economically interesting. On the other hand, the dimensionality of the problem is not excessively large and, computationally, it can be readily handled by our methodology.<sup>1</sup> Importantly, options on the industry sectors have become sufficiently liquid in recent years, making it feasible to accurately estimate their marginal distributions. Starting from Jan. 2007, we are able to estimate option-implied dependence daily. To assess accuracy and stability of MFDR, we conduct an extensive Monte Carlo experiment. This experiment confirms the viability of the new methodology in realistic applications.

Empirically, we make several contributions. Our first contribution is to document empirical properties of option-implied dependence from Jan. 2007 to the end of 2020. We find that the dependence between the nine sectors is time-varying and highly nonnormal. In particular, the dependence was much stronger in the later part of the financial crisis than in the earlier part. Similarly, the strongest dependence was observed in the midst of the COVID-19 crisis. For most trading days, the option-implied dependence is grossly inconsistent with the assumption of multinormality.<sup>2</sup> The formal tests of Mardia (1970) demonstrate that violations of multinormality are predominantly due to nonzero skewness, whereas violations due to excess kurtosis are less prevalent. Overall, we find that for most days the option-implied dependence is highly asymmetric, with large negative returns being much more correlated than large positive returns.

Our second contribution is to present novel evidence regarding the correlation risk premium (CRP), defined as the difference between the correlation under the real-world and risk-neutral probability measures. The realized correlations are computed from historical returns, while the risk-neutral ones are obtained from the option-implied joint distribution. We focus on three types of average correlations computed for the nine sectors: *global*, *down*, and *up*. The first type is the standard unconditional correlation, the other two are correlations conditional on the S&P 500 return being below or above its median value. While the global correlation has been studied extensively in the literature, the other two types cannot be studied without our methodology.

The early articles on the global CRP include Driessen et al. (2009), (2013), which document a strong negative CRP for stocks in the S&P 100, S&P 500 and the DJIA and link it to diversification. During turbulent times, correlations tend to increase, making diversification less effective. Therefore, the index options are

<sup>1</sup>It is worth mentioning that even for this problem of moderate dimensionality, the computational demands are considerable. We estimate option-implied dependence for over 3,500 trading days. For each day, the first step of MFDR calls for 10 RND estimations (or multivariate quadratic optimizations). The third step of MFDR involves finding the optimal perturbation of a  $1,000 \times 10$  matrix. The latter is an NP-complete problem with the total of  $(1,000!)^9$  possible permutations. However, the BRA is able to find an approximate, but very accurate solution in a reasonable amount of time. Limitations of MFDR are discussed in Section II.E.

<sup>2</sup>Because the risk-neutral margins of the nine sectors are nonnormal (they are highly skewed and leptokurtic), their *joint distribution* cannot possibly be multinormal. However, the question remains whether or not their *dependence* can be represented by a Gaussian copula, as is commonly assumed in applications. For the first time, our methodology allows a direct examination of the option-implied dependence in a model-free fashion. We find that the dependence is highly skewed, but its kurtosis is not overly excessive.

expensive because they allow investors to hedge against the risk of reduced diversification. In our work, we also find that the global CRP is negative. However, our most intriguing empirical results pertain to the down and up correlations. Here, we document that the risk premium is significantly negative for the down correlation but positive for the up correlation. The magnitudes of the risk premium for the down and up correlations are much larger than for the global correlation. These findings are consistent with the economic intuition that investors are mainly concerned with the loss of diversification when the market falls. As a result, they are willing to pay a considerable premium to hedge against increases in the down correlation. On the other hand, investors actually *prefer* high correlation when the market rallies. That is, investors view the down correlation as “bad” and the up correlation as “good.” The net effect of the negative risk premium for the down correlation and the positive risk premium for the up correlation is a negative risk premium for the global correlation.

The signs of the down and up CRP mirror those for the down and up variance risk premium, as reported in Feunou, Jahan-Parvar, and Okou (2018) and Kilic and Shaliastovich (2019). However, it is important to emphasize that our results neither follow from nor imply the latter findings. The previous articles study the market variance risk premium and are based exclusively on index options. Intuitively, these articles demonstrate that index OTM puts are expensive and index OTM calls are cheap when compared to the historical distribution of the index returns. Our results, on the other hand, make a statement on the *relative* pricing of individual options compared to the index options. In this regard, we find that sector OTM puts are cheap compared to index OTM puts (thus, implying too high down correlations between the sectors), whereas sector OTM calls are expensive compared to index OTM calls (implying too low up correlations). Again, a priori, there is no theoretical reason why the sign of the down (up) CRP should match that for the down (up) variance risk premium. A toy model in Section IA.D of the Supplementary Material formalizes this point.

Perhaps the negative risk premium for the down correlation is not completely unexpected. If investors only care about volatility when it leads to losses, they will dislike the down correlation and its risk premium will be negative. As argued by Ang, Chen, and Xing (2006), similar conclusions can be obtained in an equilibrium with disappointment aversion preferences (see Gul (1991), Routledge and Zin (2010)). However, it might be more challenging to rationalize with standard preferences the *positive* risk premium for the up correlation. Furthermore, the *magnitudes* of the risk premia also appear quite remarkable. To put things in perspective, the magnitude of the risk premium for the down (respectively, up) correlation is approximately 2.8 (respectively, 2.5) times larger than for the global correlation. Motivated by these findings, we introduce a new derivative contract, the down minus up correlation (DUC) swap, which at maturity pays the difference between the realized down and up correlations. Historically, the strategy that sells the DUC swap would have been very profitable, as it takes advantage of both the “expensive” down correlation and the “cheap” up correlation. Thus, selling the swap earns the risk premium, which is about 5.3 times larger than selling the global correlation.

The regular (Pearson) correlation is jointly affected by the margins (such as volatility, skewness, and heavy tails) and the dependence (copula). Thus, an

important but difficult question is whether the CRP is driven by the priced components of the margins or from the priced dependence. Therefore, as our third contribution, we disentangle the respective roles on the CRP of the changes in the two components. To remove the effects of the margins in the computation of correlation, we consider Spearman (instead of Pearson) correlations. Spearman's is a rank correlation and it has the advantage of being unaffected by the margins. To compute Spearman correlation under the risk-neutral measure, however, the complete joint distribution is required, even for the case of the global correlation. Thus, our MFDR methodology is crucial. By contrasting the results for Pearson and Spearman correlations, we conclude that the CRP is mainly driven by the dependence and not by the margins. Along the same lines, when we investigate what causes the enormous spread between the risk-neutral down and up correlations, we find that only about 11% of the spread can be attributed to nonnormality of the margins, whereas the rest is due to the non-normality (skewness) of the dependence. We also extend our analysis of the down and up correlations by constructing tail indices, which measure the dependence of extreme negative and positive returns. We find that the left-tail index is always much larger than the right-tail index, providing further evidence of the asymmetry of the dependence.

As our fourth contribution, we study whether option-implied correlations could predict future market returns. Consistent with the prior literature, the implied global correlation has a strong predictive power (see Buss and Vilkov (2012), Buss, Schönleber, and Vilkov (2019b)). Interestingly enough, its predictive power comes mainly from the *up* correlation. The up correlation is a stronger predictor than the global correlation, which in turn is better than the down correlation. This holds true for all predictive horizons from 1 to 12 months. For the up correlation, the adjusted  $R^2$  increases with the horizon and reaches an impressive level of 19.8% for the 12-month horizon. We also find that, when the up correlation is high, cyclical stocks tend to outperform defensive stocks over next 6–12 months, whereas the opposite is true when the up correlation is low.

Our empirical results highlight the importance of proper modeling of the dependence, especially under the risk-neutral measure. To match the salient features of the option data, it is critical to allow for a highly asymmetric dependence. Standard models in the literature do not always have this property. Therefore, as our fifth contribution, we develop an alternative approach to model the multivariate joint distribution. We refer to it as the *hybrid* model because it combines i) fully nonparametric margins extracted from the individual options and ii) a parsimonious parametric copula. The proposed copula is based on the homogeneous multivariate skewed normal distribution driven by two parameters only. Despite its simplicity, the model captures reasonably well the most salient features of the option-implied dependence. Of course, in terms of fitting option prices the hybrid model cannot compete with MFDR, because the latter produces (essentially) a perfect fit. Instead, the hybrid model offers different advantages: it is transparent, intuitive, and easy to implement. Our primary motivation for developing this model is twofold. First, because the hybrid model does not rely on the somewhat opaque MFDR methodology, it provides an alternative confirmation of our key empirical results. Second, we believe that the hybrid model

could prove useful in other applications where MFDR cannot be implemented due to data limitations.

Our article is related to several important strands of the literature. A number of articles use traded options to infer various measures of dependence. Driessen et al. (2009), (2013), Buraschi et al. (2013), and Faria et al. (2018) estimate the average implied correlation assuming that pairwise correlations are all equal. Buss, Schönleber, and Vilkov (2017) estimate a block-diagonal correlation matrix with two possible values for pairwise correlations. Buss and Vilkov (2012) relax the assumption of equal pairwise correlations by assuming a linear transformation between the correlations under the real-world and risk-neutral measures. Kelly, Lustig, and Van Nieuwerburgh (2016) assess the dependence among several assets by computing the spread between a portfolio of individual puts and the put on the index. Dhaene, Linders, Schoutens, and Vyncke (2012) develop an option-implied measure of dependence, which is termed Herd Behavior Index (HIX).

The down and up CRP studied in this article extend the literature on the global CRP. The pioneering work on the global CRP includes Driessen et al. (2009), (2013), who document a strong negative CRP for stocks in the S&P 100, S&P 500, and the DJIA. Further investigation of the CRP and its link to macroeconomic variables can be found in Pollet and Wilson (2010), Buraschi et al. (2013), Engle and Figlewski (2014), Harvey, Liu, and Zhu (2016), Mueller, Stathopoulos, and Vedolin (2017), and Faria et al. (2018). The correlation risk is also closely related to the variance risk (Bondarenko (2004), (2014), Carr and Wu (2009), Bollerslev and Todorov (2011), and Schneider and Trojani (2015)) and to the disagreement risk (Buraschi, Trojani, and Vedolin (2014)). The down and up variance risk premium is studied by Kilic and Shaliastovich (2019) and Feunou et al. (2018).

The asymmetric behavior of the left and right tails of the asset returns has been widely recognized among both practitioners and academics following the path-breaking work of Longin and Solnik (2001). Ang and Chen (2002), Hong, Tu, and Zhou (2006), Alcock and Hatherley (2017), Jiang, Wu, and Zhou (2018), and Alcock and Sinagl (2022) provide evidence of asymmetric correlation and dependence between stocks and the market. Furthermore, Hong et al. (2006) assess the economic importance of asymmetric returns in the context of a portfolio choice problem. They find that investors can achieve over 2% annual certainty-equivalent gains when they account for asymmetric correlation. Longin and Solnik (2001) study the asymmetric tail dependence by computing the exceedance correlation, or the correlation between returns that are jointly above or below a given threshold. For international equity markets, they demonstrate that the correlation between large *negative* returns does not converge to 0 and instead tends to increase deeper in the left tail. On the other hand, the correlation between large *positive* returns does decrease to 0. More recently, Chabi-Yo, Ruenzi, and Weigert (2018) study the dependence of stock returns. Like us, they are motivated by the limitations of Pearson correlation. After correcting for the effect of margins, they find evidence of asymmetry between the left and right tails. MFDR allows us to explore related questions, but for option-implied counterparts. In the study of the nine sectors of the S&P 500 index, we find that the dependence is much more asymmetric under the

risk-neutral than under the real-world measure. Our work is also related to the literature on skewness (see Patton (2004), DeMiguel, Plyakha, Uppal, and Vilkov (2013), Amaya, Christoffersen, Jacobs, and Vasquez (2015), and Jondeau, Zhang, and Zhu (2019)) and on downside risk (see Bollerslev and Todorov (2011), Kelly and Jiang (2014), Kelly et al. (2016), Farago and Tédongap (2018), and Orłowski, Schneider, and Trojani (2020)).

The BRA used in the third step of MFDR generalizes the standard rearrangement algorithm (RA) introduced by Puccetti and Rüschendorf (2012). The RA has found important applications in various disciplines. Embrechts, Puccetti, and Rüschendorf (2013) use the RA in quantitative risk management to assess the impact of model uncertainty on Value-at-Risk estimates for portfolios. Other applications of the RA include operations research (fair allocation of goods, optimization) and engineering (image reconstruction). Bernard, Bondarenko, and Vanduffel (2018) study the theoretical properties of the BRA for the problem of finding a joint distribution given the marginal distribution of several random variables and of their sum. The first step of MFDR assumes that risk-neutral marginal distributions can be estimated from traded options in a model-free way (see, e.g., Jackwerth and Rubinstein (1996), Aït-Sahalia and Lo (2000), and Bondarenko (2003)). Figlewski (2018) reviews pros and cons of the various methods proposed for extracting risk-neutral densities.

The rest of the article is organized as follows: In Section II, we formulate the problem of inferring the dependence from option prices, discuss the existing approaches, and present our MFDR methodology. In Section III, we implement the approach using options on the S&P 500 index and its sectors and document the properties of the option-implied dependence. Section IV provides a detailed investigation of the CRP. Section V concludes.

## II. Extracting Information from Options

In this section, we recall how to extract the option-implied probability distribution for a single asset and discuss how index options can provide information about the dependence among several assets. We then review the existing approaches to constructing option-implied correlations and present our MFDR methodology.

### A. Univariate Case

For a given underlying asset, let  $X$  denote the return over a fixed time period  $[t, T]$ . Let  $C(K)$  and  $P(K)$  denote the time- $t$  price of the European-style call and put options with moneyness  $K$  and maturity  $T$  written on the asset's return  $X$ . For simplicity, we assume that the asset pays no dividends and that the risk-free rate is 0.<sup>3</sup> Under the standard assumptions, the option prices are equal to the

<sup>3</sup>In the empirical application, we use ETF options on the S&P 500 index and its industry sectors. These options are American-style and are written on an asset's price, not the return. Moreover, the underlying ETFs do pay quarterly dividends, and of course, the risk-free rate is not really 0. We address these real-world complications as follows: First, we convert *spot* prices of the options and the underlying asset into *forward* prices (for delivery at time- $T$ ). The forward prices account for dividends and the

expected value of their payoffs under a suitably chosen *risk-neutral* probability measure  $\mathbb{Q}$ :

$$C(K) = E^{\mathbb{Q}}[(X - K)^+] = \int_0^{\infty} (x - K)^+ f(x) dx,$$

$$P(K) = E^{\mathbb{Q}}[(K - X)^+] = \int_0^{\infty} (K - x)^+ f(x) dx,$$

where  $f(x)$  denotes the RND. The RND satisfies the relationship first established by Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978):

$$(1) \quad f(x) = \frac{\partial^2 C(K)}{\partial K^2} \Big|_{K=x} = \frac{\partial^2 P(K)}{\partial K^2} \Big|_{K=x}.$$

Similarly, the *risk-neutral cumulative distribution* (RNCD) satisfies

$$(2) \quad F(x) = \frac{\partial C(K)}{\partial K} \Big|_{K=x} = \frac{\partial P(K)}{\partial K} \Big|_{K=x}.$$

Although not directly observable, the RNCD can be recovered using the relationship in (2), provided that options with a continuum of strikes  $K$  are available. In practice, options are only available for a finite number of strikes. Nevertheless, a number of efficient nonparametric approaches have been proposed in the literature that make it possible to circumvent this shortcoming (see, e.g., Jackwerth and Rubinstein (1996), Ait-Sahalia and Lo (2000), and Bondarenko (2003)). Note that for our approach, we only need the information about the RNCD and not the RND. The former can be estimated considerably more accurately (this is because the so-called curse of differentiation is not as severe when estimating the first, rather than the second, derivative of a function).

## B. Implied Dependence of Several Assets

To obtain the joint distribution of several assets, the knowledge of individual marginal distributions is not enough. We also need their dependence. Dependence is implicit in the prices of multivariate options, such as index options. Consider an index comprising  $d$  assets. Let  $X_{j,t+\tau}$  denote the return of the  $j$ th asset for the period  $[t, t + \tau]$  for some fixed horizon  $\tau$  (e.g., 3 months). The return of the index can be written as

$$S_{t+\tau} = \sum_{j=1}^d \omega_j X_{j,t+\tau}, \quad \sum_{j=1}^d \omega_j = 1,$$

---

nonzero risk-free rate. Second, we adjust option prices for the early exercise feature by following the approach of Barone-Adesi and Whaley (1987). This gives us the prices of equivalent European-style options. Third, we rescale option prices and their strikes by the asset's forward price. This gives us options on the asset's *return*.



where  $\omega_1, \dots, \omega_d$  are the weights. When there is no confusion, we simplify the notation and drop the time index. For example, we write  $X_j$  instead of  $X_{j,t+\tau}$  and  $S$  instead of  $S_{t+\tau}$ . We assume that the options market offers a sufficient range of strikes so that the risk-neutral distributions  $F_S, F_1, \dots, F_d$  for the returns  $S, X_1, \dots, X_d$  can be estimated accurately.

Our goal is to find a dependence structure (or copula)  $C$  that can explain the distribution of the index. Specifically, for a given copula  $C$ , we can define a new random variable:

$$(3) \quad Z := Z(C) = Z(C; F_1, \dots, F_d, \omega_1, \dots, \omega_d) := \sum_j \omega_j X_j,$$

where  $(X_1, X_2, \dots, X_d)$  has a joint distribution fully described by the copula  $C$  and the respective margins  $F_j$ . We can think about the random variable  $Z$  as the weighted return of the  $d$  components, or replicated index return. Ideally, we would like to find a copula  $C$  such that the weighted return  $Z$  is equal in distribution to the observed index return  $S$ :

$$(4) \quad Z \stackrel{d}{=} S.$$

This is a difficult problem, as an equality in distribution imposes infinitely many constraints on the choice of the dependence structure  $C$ . We could start by matching various moments of  $Z$  and  $S$ . We note that the risk-neutral densities have a mean of 1 (when  $X_j$  and  $S$  are defined as gross returns); thus, the first central moment of  $Z$  and of  $S$  are matched automatically:

$$E[Z] = \sum_j \omega_j E[X_j] = \sum_j \omega_j = 1 = E[S].$$

Matching the second central moment leads to the standard identity between the variance of the index and the variances of its components:

$$(5) \quad \text{var}(S) = \sum_{j=1}^d \omega_j^2 \text{var}(X_j) + 2 \sum_{j=1}^{d-1} \sum_{j < k} \omega_j \omega_k \sqrt{\text{var}(X_j)} \sqrt{\text{var}(X_k)} \rho_{jk},$$

where  $\rho_{jk}$  is the correlation between  $X_j$  and  $X_k$ . Most existing approaches rely on the identity in (5) and additional auxiliary assumptions, e.g., the assumption of equal pairwise correlations,  $\rho_{jk} = \rho$ . We review them in the next subsection and then present our approach in [Section II.D](#).

### C. Existing Approaches

The CBOE Implied Correlation Index is an attempt to estimate the average pairwise correlation among the stocks in the S&P 500 index.<sup>4</sup> Its basic idea is to use the condition in (5) but to approximate the variances by the squares of the ATM

<sup>4</sup>The underlying methodology was recently updated by CBOE and is detailed in the white paper, Chicago Board Options Exchange (2022).

Black–Scholes implied volatilities (BSIV, or simply IV) and to replace the different  $\rho_{jk}$  with a single correlation parameter that we denote as  $\rho_{\text{CBOE}}$ . Then

$$(6) \quad \rho_{\text{CBOE}} = \frac{\sigma_S^2 - \sum_{j=1}^d \omega_j^2 \sigma_j^2}{2 \sum_{j=1}^{d-1} \sum_{k>j} \omega_j \omega_k \sigma_j \sigma_k}.$$

Assuming that this approximation is exact, the expression of  $\rho_{\text{CBOE}}$  in (6) can be rewritten as the weighted average pairwise correlation:

$$(7) \quad \rho_{\text{CBOE}} = \frac{\sum_{j<k} \omega_j \omega_k \sigma_j \sigma_k \rho_{jk}}{\sum_{j<k} \omega_j \omega_k \sigma_j \sigma_k}.$$

However, since this approximation is in fact *not* exact, this interpretation no longer holds. The CBOE index cannot be viewed as a genuine correlation, it may be very different from the true average correlation, and it can potentially take values that are greater than 1.<sup>5</sup>

The academic literature has proposed several modifications to the CBOE methodology. Most notably, Driessen et al. (2009), (2013) improve the CBOE approach in two ways. First, they apply the condition in (5) to all components and not a small subset. Second, for the standard deviations of  $X_j$  and  $S$  they use the model-free implied volatility (MFIV; see Britten-Jones and Neuberger (2000), Carr and Madan (2001), and Bakshi, Kapadia, and Madan (2003)). The variances are now driven not by the ATM options only, but are based on the whole cross section of options. For MFIVs, the condition in (5) holds exactly under the assumption of diffusion, but only approximately if prices could jump (which is likely in practice).

More recently, Buss et al. (2017), (2019a) use standard deviations for RNDs,  $\sigma^Q$ , for which the condition in (5) now holds identically.<sup>6</sup> Furthermore, Buss et al. (2017) relax the equal pairwise correlation constraint by estimating a block diagonal correlation matrix with two correlation parameters (i.e., pairwise correlations are constant for any two stocks within the same economic sector of the S&P 500 but take another value when two stocks belong to different sectors). Buss and Vilkov (2012) replace the assumption of constant pairwise correlations  $\rho_{jk} = \rho$  with a linear

<sup>5</sup>For example, the KCJ index was 100.8 on Nov. 6, 2008; 105.93 on Nov. 13, 2008; and 103.4 on Nov. 20, 2008. We observe that, strictly speaking, the condition in (5) does not hold for ATM IVs. Implied volatilities are standard deviations of log-returns under the assumption that these are normally distributed. Therefore, using implied volatilities in (5) would be justified when  $\log(S) = \sum_j \omega_j \log(X_j)$ . In reality, the index is an arithmetic, not a geometric, average. Hence, some bias is introduced and we can only state an approximate relation:  $\sigma_S^2 \approx \sum_{j=1}^d \omega_j^2 \sigma_j^2 + 2 \sum_{j=1}^{d-1} \sum_{k>j} \omega_j \omega_k \sigma_j \sigma_k \hat{\rho}_{jk}$  where  $\sigma_S$  and  $\sigma_j$  are the ATM IVs for the index and its components, and  $\hat{\rho}_{jk}$  is the pairwise correlation of  $\log(X_j)$  and  $\log(X_k)$ . In addition, CBOE uses a subset of  $d = 50$  largest components of the index, which introduces another bias: since the selected components are generally less volatile, the average implied correlation tends to be overstated.

<sup>6</sup>Note that  $\sigma^Q$  is also related to the concept of the simple model-free implied volatility (SMFIV) of Martin (2017) and can be computed from the *simple variance swap*.

specification, which ensures a negative CRP ( $\rho^Q > \rho^P$ ) that is higher in magnitude for stocks with low or negative correlation consistently with the empirical evidence in Mueller et al. (2017) for the FX market.

A rather different approach to assess the risk-neutral dependence among several assets is proposed by Kelly et al. (2016). It is based on the idea that a portfolio of individual options is always more expensive than an option on the portfolio, where the strikes are chosen appropriately. The spread between the two is larger when the correlation among the assets is lower. They find that during the 2007–2009 financial crisis, the OTM put options for financial firms were extraordinarily expensive relative to the matched OTM put options for the financial sector index. Therefore, they conclude that a large amount of aggregate tail risk was missing from the cost of the financial sector crash insurance, likely due to a perceived sector-wide government bailout guarantee. Overall, their approach uses options to infer a particular indicator of tail dependence, but it does not provide the entire joint distribution.

#### D. A More General Approach: Model-Free Dependence Recovery

One common limitation of the existing approaches is that only partial dependence information is obtained. Furthermore, strong assumptions are typically imposed on the correlation matrix under the risk-neutral measure. The identifying restriction in (5), which equates the index variance to the variance of the portfolio of the components, is clearly insufficient to recover the entire correlation structure – there are many possible ways to satisfy this single restriction.

As stated previously, our approach is more ambitious, as it attempts to find a full dependence (copula) and to do so in a completely model-free fashion. How is this even possible? The key to our approach is that it matches not just one moment in (5) but a *continuum* of moments. Specifically, our method attempts to construct the random variable  $Z$  such that it is equal to  $S$  almost surely, which is even stronger than the condition in (4) and which implies that for *any* function  $g(z, s)$ ,

$$(8) \quad E[g(Z, S)] = E[g(S, S)].$$

In this respect, our approach generalizes the existing approaches. Our solution satisfies the second moment condition in (5) as a special case, meaning that our solution will yield exactly the same average global pairwise correlation  $\rho$  as the existing approaches. The existing approaches use just one summary statistic from each option-implied distribution  $F_j$  (either BSIV, or MFIV, or  $\sigma^Q$ ) and match just one restriction in (5).<sup>7</sup> In contrast, our approach uses the complete information contained in each distribution  $F_j$  and attempts to satisfy a continuum of restrictions implied by (8). We discuss in more detail various implications of (8) in [Appendix A](#).

<sup>7</sup>BSIV depends on the price of the ATM option only. MFIV and  $\sigma^Q$  use information from the entire cross section of options with different strikes, but that information is still reduced to a single statistic. Clearly, there are many distributions, otherwise very different, but with identical MFIV (or  $\sigma^Q$ ). The existing approaches are unable to distinguish among them, yielding the same measure of dependence. In [Appendix C](#), we provide numerical examples, which illustrate that satisfying only condition (5) is not enough to reproduce the observed implied volatility curve of the index (see [Tables C.1](#) and [C.2](#)).

Let us briefly describe our method. Recall that  $S = \sum_{j=1}^d \omega_j X_j$  for some weights  $\omega_j$  that sum to 1. Assume that we observe the  $(d+1)$  risk-neutral distributions ( $F_j$  for the return  $X_j$  and  $F_S$  for the index return) and that we want to construct a *discrete* multivariate distribution among the  $d$  returns  $X_1, \dots, X_d$ . To do so, we first approximate  $F_j$  for each asset  $j \in \{1, \dots, d\}$  with a discrete distribution (which can be done to any degree of accuracy) as follows: There are  $n$  equiprobable states in which  $X_j$  takes the values  $x_{ij}$ ,  $i = 1, \dots, n$ , where the elements  $x_{ij}$  are defined as realizations  $x_{ij} := F_j^{-1}\left(\frac{i-0.5}{n}\right)$  ( $i = 1, \dots, n$ ), and  $S$  takes the values  $s_i$ ,  $i = 1, \dots, n$  defined as  $F_S^{-1}\left(\frac{i-0.5}{n}\right)$ . Using this discretization, we represent the multivariate vector of asset returns  $(X_1, X_2, \dots, X_d)$  by an  $n \times d$  matrix:

$$(9) \quad \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1d} \\ x_{21} & x_{22} & \dots & x_{2d} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n1} & x_{n2} & \dots & x_{nd} \end{bmatrix},$$

where the  $j$ th column corresponds to the  $j$ th asset return  $X_j$  and the  $i$ th row represents a state of the world in which a joint outcome  $(X_1 = x_{i1}, \dots, X_d = x_{id})$  occurs with probability  $1/n$ . If one permutes the elements in the  $j$ th column, the marginal distribution of  $X_j$  remains unchanged because all realizations are equally likely. In contrast, the dependence of  $X_j$  with the other variables  $X_k$  is affected, because permutations result in different states.

We then aim at rearranging the matrix in (9) by permuting elements within a column to satisfy the restriction that  $S = \sum_{j=1}^d \omega_j X_j$  in each of the  $n$  states. The solution technique is called BRA and its formal exposition is relegated to Section IA.A of the Supplementary Material. This appendix also presents a toy model with  $n = 5$  and  $d = 3$  for which it is possible to trace every step of the BRA procedure. However, this parsimonious example does not resemble anything close to what we would like to do in real applications. Therefore, to further develop intuition, Section IA.A of the Supplementary Material provides another illustration of the algorithm, now using continuous margins that are discretized into a large number of states. This is still an unrealistic, but pedagogical example. It uses only two assets and allows us to easily visualize the impact on the dependence of altering the distribution of their weighted sum. Specifically, we assume that the two returns  $X_1$  and  $X_2$  are normally distributed with standard deviations of 0.2 and 0.4, respectively. Their margins are thus fixed. We then consider 3 cases for the distribution of their weighted sum,  $S = \frac{1}{2}X_1 + \frac{1}{2}X_2$ . In the first 2 cases,  $S$  is also normally distributed with a variance chosen such that the implied correlation is equal to either 0 (no dependence, the first and second graphs in Figure IA.1 in the Supplementary Material) or 0.97 (strong dependence, the first and second graphs in Figure IA.2 in the Supplementary Material), respectively. In the last case,  $S$  has a skewed distribution with a heavy left tail (asymmetric dependence, the first and second graphs in Figure IA.3 in the Supplementary Material). To run the BRA, we discretize the marginal distributions of  $X_1$ ,  $X_2$  and  $S$  into  $n = 1,000$  equiprobable states. The resulting joint distribution is represented by 1,000 dots corresponding to pairs  $(X_1, X_2)$  in the

fourth graph of each figure (which in turn correspond to the 1,000 rows of the output matrix from the BRA). In Figure IA.1 in the Supplementary Material, the two variables  $X_1$  and  $X_2$  appear independent, which is expected when an implied correlation is equal to 0. In Figure IA.2 in the Supplementary Material, the implied dependence is strong and symmetric. This is also intuitive because the distribution of the sum is also symmetric. Finally, when the sum has a skewed distribution, the joint distribution is highly asymmetric, with a very pronounced left tail dependence, as can be seen in the third and fourth graphs of Figure IA.3 in the Supplementary Material.

## E. Discussion and Limitations of MFDR

Intuitively, our approach can be compared to GMM. Asset returns are permuted in such a way as to keep the margins fixed and to enforce the “moment” condition  $S = \sum_{j=1}^d Y_j$  state-by-state. Effectively, there are as many restrictions as states  $n$ , which could be chosen as large as desired. The approach finds a compatible joint distribution, which matches every available option on the assets and the index. Although there are many restrictions being enforced, the solution is, typically, not unique. This situation is not uncommon (when the market is incomplete, there are many risk-neutral measures, which correctly price available options). However, out of many possible ones, the BRA finds an economically sensible solution. As shown in Bernard et al. (2018), the BRA procedure has an important property: The obtained multivariate model for  $(Y_1, Y_2, \dots, Y_d)$  exhibits maximum entropy. This means that the procedure yields the “most likely” configuration given the information available and given that no additional information is used.<sup>8</sup>

To make clear, what is meant by “most likely,” let us take a step back and assume for a moment that we only have information about the marginal distributions of the assets, and not the sum. That is, we only agree on the values that appear in the first  $d$  columns but not on the order in which they appear. Consequently, all permutations within columns are equally plausible, and there is no reason to prefer one permutation over another. Hence, randomizing the assignment of realizations to the different states leads to marginal distributions (reflected by the columns) that are most likely to be independent, which corresponds precisely to the maximum-entropy case. Suppose now that the additional information is known, namely, the marginal distribution of the sum. In this case, the set of admissible permutations is simply reduced to those that yield row sums that are 0. The BRA method implements the idea of randomizing the assignment of realizations to the different states but now under the additional constraint provided by the knowledge of the distribution for the index.<sup>9</sup>

<sup>8</sup>In this regard, it is also worth citing Jaynes ((2003), p. 370), who developed the principle of maximum entropy in its modern form and who stated the following: “In summary, the principle of maximum entropy is not an oracle telling which predictions must be right; it is a rule for inductive reasoning that tells us which predictions are most strongly indicated by our present information.”

<sup>9</sup>In the context of extracting RNDs, the principle of maximum entropy has been explored in Rubinstein (1994), Jackwerth and Rubinstein (1996), and Stutzer (1996).

In sum, our approach amounts to constructing a numerical dependence and we refer to it as MFDR. The approach consists of main 3 steps. First, from traded options on the individual assets and the index, we estimate the risk-neutral marginal distributions. Second, we discretize the estimated marginal distributions into  $n$  states and reduce the problem of finding a joint distribution among  $d$  assets to finding a suitable permutation of an  $n \times (d+1)$  matrix. Finally, we solve the resulting integer optimization problem with the BRA.

In Section III, we implement MFDR using ETF options for the  $d=9$  industry sectors comprising the S&P 500 index. This application provides an ideal setting to showcase MFDR, which is best suited for small to moderate  $d$ . Generally, there are three potential limitations to our methodology. The first limitation is related to the availability of high-quality option data. Recall that in the first step of MFDR, we estimate risk-neutral marginal distributions from prices of traded options. For each asset, we need options with a wide range of strikes, densely covering the whole support of the RND. This requirement is not satisfied for many assets, including some large companies in the DJIA or the S&P 100. Without such a detailed information, accurate estimation of option-implied marginal distribution becomes infeasible, at least in a fully model-free manner. The second limitation is related to the computational time required to solve the optimization problem in the third step of MFDR. Although the BRA method is much faster than the direct search, its computational demands increase exponentially with  $d$ . In Section IA.E of the Supplementary Material, we implement MFDR for the case of  $d=30$  stocks in the DJIA, but the computational burden of the BRA becomes considerable. To speed up convergence, we can reduce the number of states  $n$ , and use the “randomized” BRA (which cycles through a random subset of partitions of the columns, instead of all possible partitions) or the standard RA. These steps enable us to find a solution for larger  $d$ , but with some loss of accuracy. Finally, as  $d$  increases, the weighted sum of the  $d$  components (the index) becomes less informative about the joint distribution. Intuitively, with more degrees of freedom, it becomes easier to find a copula that fits the distribution of the index. The BRA will still find a valid joint distribution, which perfectly fits all options on individual components and their index (provided that there are no arbitrage violations). However, the solution will be determined to a larger degree by the principle of maximum entropy and to a lesser degree by the constraints imposed by observed option prices.

### III. Empirical Application

We use daily closing prices of options on the nine SPDR Select Sector funds and on the SPDR S&P 500 Trust. These nine sector ETFs are capitalization-weighted portfolios comprising all stocks in the S&P 500 index. Option data are obtained directly from CBOE. Table 1 lists abbreviated names for the ETFs. Our sample covers the period from Jan. 1, 2007 to Dec. 31, 2020. Although options on sector ETFs were also traded in prior years, the availability of strikes and maturities was more limited. By 2007, the liquidity of sector options has improved considerably, making it feasible to accurately estimate their RNDs on a day-to-day basis. We focus on horizon  $\tau = 3$  months and infer the option-

TABLE 1  
S&P 500 Sectors

Table 1 lists the underlying assets used in this study. The sector ETFs appear in alphabetical order of the ticker.

Description	Ticker	Abbreviation
Materials sector SPDR fund	XLB	MAT
Energy sector SPDR fund	XLE	ENE
Financial sector SPDR fund	XLF	FIN
Industrial sector SPDR fund	XLI	IND
Technology sector SPDR fund	XLK	TEC
Consumer staples sector SPDR fund	XLP	CST
Utilities sector SPDR fund	XLU	UTI
Health care sector SPDR fund	XLV	HEA
Consumer discretionary sector SPDR fund	XLY	CDI
SPDR S&P 500 ETF trust	SPY	SPX

implied joint distribution of the nine sectors.<sup>10</sup> The details of the estimation procedure are provided in [Appendix B](#). We conduct an extensive Monte Carlo experiment, which confirms accuracy and stability of MFDR in our setting, see Section IA.C of the Supplementary Material. The realized volatilities and correlations are computed using daily returns of the ETFs, which are obtained from CRSP. The daily sector weights  $\omega_j$  are obtained from Bloomberg. The risk-free rate is approximated by the rate of Treasury bills.

We first briefly discuss the main characteristics of the sectors and of the aggregate market over the studied period. [Figure 1](#) plots the cumulative returns and the 3-month ATM implied volatilities for the S&P 500 and for 3 specific sectors (FIN, TEC, and ENE). The performance of individual sectors often differs considerably. During the financial crisis (defined in this article as the period from Aug. 1, 2007 to Apr. 1, 2009), the financial sector starts to decline much earlier than other sectors, but from the last quarter of 2008 onward the crisis spreads to the rest of the economy. During the 2015 energy crisis, it is the energy sector that experiences a considerable decline, whereas the other sectors are mostly unaffected. The COVID-19 crisis in Mar. 2020 affects all sectors. As expected, the ATM implied volatilities (IVs) of the sectors are highly correlated over time. The IV of the financial sector is very high in 2008–2009, whereas the IV of the energy sector stands out in 2015–2016 and during the COVID-19 crisis.

[Table 2](#) provides descriptive statistics for the nine sectors and the S&P 500 index. In particular, it shows that the largest sectors of the S&P 500 index are TEC, FIN, and HEA, with the median portfolio weights of 24.1%, 16.0%, and 13.3%, respectively. The smallest sectors are MAT and UTI, with the median weights of 3.2% and 3.4%, respectively. Some sectors experience considerable variation in the weight  $\omega_j$ . For example, the weight for FIN ranges from 8.6% to 22.4% over the sample period, whereas the weight for ENE ranges from 1.9% to 16.2%. [Table 2](#) also reveals that the ATM IVs are consistently smaller than the standard deviations of the estimated risk-neutral densities ( $\sigma^Q$ ). This pattern reflects

<sup>10</sup>The choice of  $\tau = 3$  months represents a practical compromise. The horizon is long enough to allow accurate estimation of the realized correlations. At the same time, it is short enough so that sector options are still liquid.

FIGURE 1  
Cumulative Returns and ATM Implied Volatilities

Figure 1 shows implied volatilities corresponding to a maturity of 3 months. The gray-shaded areas indicate the financial crisis (Aug. 1, 2007 to Apr. 1, 2009) and the COVID-19 crisis (Feb. 21, 2020 to Mar. 23, 2020). Shown are 3 sectors (FIN, TEC, and ENE) as well as the S&P 500 index (the black line).

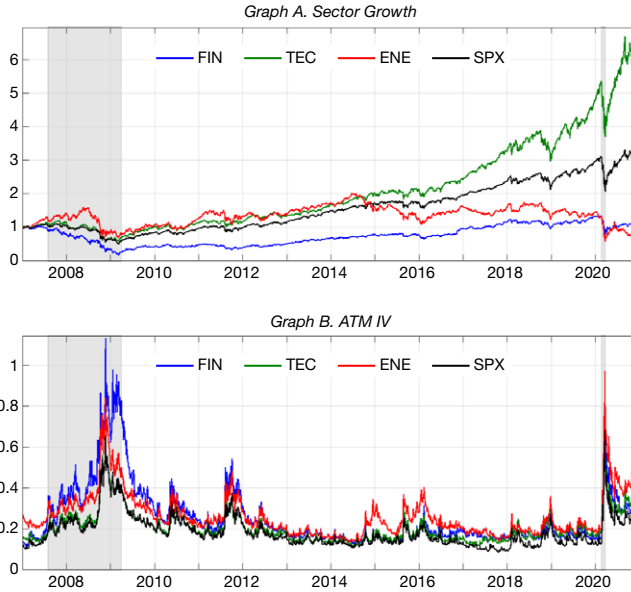


TABLE 2  
Summary Statistics

Table 2 reports time series averages over our sample for the nine sectors and for the aggregate market (S&P 500). The first 7 columns are computed from daily returns and include the mean, standard deviation  $\sigma^P$ , skewness, kurtosis, correlation with the market (Corr), beta with the market, and the Sharpe ratio (SR). The next 3 columns are the minimum, median, and maximum for sector weight  $\omega_j$ . The last 2 columns report Q statistics computed from 91-day options. They include the option ATM implied volatility (IV) and standard deviation of the RND. When applicable, statistics are reported in annualized form and as decimals.

	Mean	$\sigma^P$	Skew	Kurt	Corr	Beta	SR	Weight			IV	$\sigma^Q$
								Min	Med	Max		
MAT	0.086	0.255	-0.17	11.20	0.88	1.08	0.30	0.023	0.032	0.039	0.232	0.246
ENE	0.022	0.317	-0.29	16.50	0.81	1.24	0.04	0.019	0.091	0.162	0.263	0.272
FIN	0.046	0.336	0.27	16.51	0.86	1.40	0.11	0.086	0.160	0.224	0.256	0.272
IND	0.094	0.230	-0.19	11.82	0.92	1.02	0.37	0.075	0.102	0.122	0.210	0.223
TEC	0.150	0.226	0.01	15.03	0.92	1.01	0.62	0.184	0.241	0.402	0.206	0.219
CST	0.080	0.151	-0.19	14.63	0.80	0.59	0.47	0.065	0.098	0.141	0.149	0.164
UTI	0.059	0.199	0.41	19.79	0.70	0.67	0.25	0.026	0.034	0.047	0.178	0.190
HEA	0.106	0.181	-0.11	14.45	0.84	0.74	0.54	0.106	0.133	0.170	0.173	0.188
CDI	0.125	0.224	-0.40	12.00	0.91	0.99	0.52	0.076	0.110	0.133	0.208	0.222
SPX	0.092	0.207	-0.06	17.89	1.00	1.00	0.40				0.185	0.200


nonnormality of risk-neutral densities and manifests itself in the pronounced volatility skews.

To provide initial intuition, Figure 2 plots the average implied volatility smile for the nine sectors and for the index. The implied volatility smile corresponds to  $\tau = 3$  months and is shown as a function of the normalized moneyness:



FIGURE 2  
Implied Volatility

Figure 2 shows the average implied volatilities of the nine sectors (solid lines) and the aggregate market (dashed lines) for the normalized moneyness  $m$ . Time to maturity  $\tau = 3$  months. For the sectors, implied volatility smiles are value-weighted for each day. Graph A shows the averages computed separately for the financial crisis (FC, Aug. 1, 2007 to Apr. 1, 2009, red) and for the period excluding the crisis (xFC, blue). Graph B shows the averages computed for trading days when the realized correlation is low (blue) and high (red). Specifically, we first compute the weighted average pairwise correlation for the nine sectors over the 3-month trailing window. We then select the bottom and top quartiles and compute average implied volatilities for both groups.



$$(10) \quad m := \frac{\log \frac{K}{Z}}{\sigma^{\text{ATM}} \sqrt{\tau}}$$

where  $\sigma^{\text{ATM}}$  is the ATM IV and  $Z$  is the forward price. To construct the figure, we compute the daily volatility smiles for the nine sectors and value-weight them. The average sector volatility smile and the index volatility smile are then averaged across various subsamples. In Graph A, we isolate the financial crisis (defined as the period from Aug. 1, 2007 to Apr. 1, 2009). Generally, the volatility smile for the sectors is higher than for the index, although the gap between the two is relatively narrow, indicating a high correlation between the sectors. During the financial crisis, as expected, the implied volatilities are much higher. It is also noteworthy that the gap between the two smiles is wider for positive  $m$ , implying a potentially lower correlation on the upside. In Graph B, we distinguish between days when the realized correlation is low or high. Specifically, we first compute the average pairwise realized correlation for the nine sectors over the trailing 3-month window. We then select the bottom and top quartiles and compute average implied volatilities for each group. For the high correlation quartile, the two volatility smiles are much steeper and the gap between the two smiles is more narrow, especially, for the middle and high values of  $m$ .

### A. Three Types of Correlations

To describe the option-implied dependence of the nine sectors, we compute three types of pairwise correlations. The first type is the standard correlation:

$$(11) \quad \rho_{j,k}^{\text{Q}} = \text{corr}^{\text{Q}}(X_j, X_k),$$

which is the Pearson correlation coefficient between the returns for sectors  $X_j$  and  $X_k$  over the period  $[0, \tau]$ . The other two types are *down* and *up* correlations, or correlations conditional on the S&P 500 having a low or high return, respectively. Specifically,

$$(12) \quad \rho_{j,k}^{d,Q} = \text{corr}^Q(X_j, X_k | S \leq S^m),$$

$$(13) \quad \rho_{j,k}^{u,Q} = \text{corr}^Q(X_j, X_k | S > S^m),$$

where  $S^m$  denotes the median return of the S&P 500 index. In Section IV, we use the down and up correlations to better understand the nature of the CRP. We refer to the standard correlation as *global*, to distinguish it from the two truncated correlations.<sup>11</sup> This should not cause any confusion, as we do not work with international assets.

Since there are many sector pairs ( $\frac{1}{2}d(d-1) = 36$ ), it is often convenient to work with weighted average correlations. Specifically, for positive weights  $\pi_j$ , we define

$$(14) \quad \rho^{a,Q} = \frac{\sum_{j < k} \pi_j \pi_k \rho_{j,k}^{a,Q}}{\sum_{j < k} \pi_j \pi_k},$$

where the superscript  $a \in \{g, d, u\}$  indicates the type of the correlation (global, down, or up). However, if there is no confusion, we often drop the superscript  $g$  when discussing the usual, global correlation. There are several sensible choices for the weights  $\pi_j$ , including  $\pi_j = 1/d$  (equal-weighted) or  $\pi_j = \omega_j$  (value-weighted). Here, we focus on the “risk-weighted” averaging:

$$(15) \quad \pi_j = \omega_j \sigma_j^Q,$$

where  $\sigma_j^Q$  is the standard deviation of the RND. This case has been used in most existing approaches. In particular, assuming constant pairwise correlations, the average global correlation  $\rho^Q$  with the weights in (15) can be computed from the volatilities of the components and the index without our MFDR but using (6) instead. Thus, focusing on this case allows for direct comparison with the

<sup>11</sup>Intuitively, the down correlation is simply a correlation coefficient computed under the conditional probability with respect to the event  $\{S \leq S^m\}$  (i.e., the index return being below its median). Mathematically, the down correlation is defined as  $\text{corr}(X_j, X_k | S \leq S^m) = \frac{\text{cov}(X_j, X_k | S \leq S^m)}{\sqrt{\text{var}(X_j | S \leq S^m)} \cdot \sqrt{\text{var}(X_k | S \leq S^m)}}$ . A related conditional correlation is studied in Ang and Chen (2002) and Campbell, Forbes, Koedijk, and Kofman (2008), who also refers to it as a truncated correlation or correlation conditional on a partitioning event. See equation (2) in the latter article. In our empirical application, the joint distribution of the nine sectors on any day is represented by  $n = 1,000$  equiprobable states, where each state is given by the vector of the nine returns. In particular, these states are shown as  $n$  dots in Figures 3 and 4. Half of the states correspond to the index being below its median ( $S \leq S^m$ ) and we use these 500 draws from the joint distribution to compute the down correlation  $\rho_{j,k}^{d,Q}$ . The other 500 states are used to compute the up correlation  $\rho_{j,k}^{u,Q}$ .

existing literature.<sup>12</sup> However, the findings on the CRP that we present in Section IV.B are robust to the choice of the weights. When studying the CRP we need not only the risk-neutral correlations, but also their real-world counterparts,  $\rho^{g,\mathbb{P}}$ ,  $\rho^{d,\mathbb{P}}$ , and  $\rho^{u,\mathbb{P}}$ . The latter are computed from the same equations (11)–(15) but using the realized correlations instead of the option-implied correlations.<sup>13</sup> The realized correlations are obtained from the daily ex-dividend returns of the underlying ETFs over the trailing 3-month window.<sup>14</sup> More precisely, the window is equal to 63 trading days, which approximately corresponds to 3 calendar months, or 91 calendar days.

## B. Case Studies

We now illustrate our approach on two dates in the midst of the financial crisis: Sept. 8, 2008 and Nov. 17, 2008. The first date represents a relatively calm period, whereas the second represents an extremely turbulent one. For each day, we proceed as follows: We use 91-day options to obtain the risk-neutral marginal distributions  $F_j$  and  $F_S$  for the nine sectors  $X_j$  and the index  $S$ . We discretize  $(d+1) = 10$  distributions into  $n = 1,000$  states, collect them in an  $n \times 10$  matrix, and apply the method described in Section II.D. The output of the BRA is another  $n \times 10$  matrix that describes a joint model compatible with all marginal distributions. Since the full joint distribution is a 10-dimensional object, we need to make some choices on how to display it. Hence, from the output matrix, we extract triplets  $(x_i, y_i, z_i)$  for  $i = 1, \dots, n$  to examine the dependence among the S&P 500 index ( $x_i$ ), Financial sector ( $y_i$ ), and Utilities sector ( $z_i$ ). (Of the nine sectors, we arbitrarily pick two, the most and least dramatic ones.)

We present the results in Figures 3 and 4, where Graph A (in each) displays dependence for the pair  $(S, \text{FIN})$  and Graph B for the pair  $(S, \text{UTI})$ . We remove the effect of the marginal distributions on the joint distribution and display in the first column of the scatterplots  $(F_S(x_i), F_{\text{FIN}}(y_i))$  and  $(F_S(x_i), F_{\text{UTI}}(z_i))$ . By doing so, we bring all returns to the same (uniform) scale and obtain a visualization of the true copula. However, it is typically easier to interpret the dependence between *normally* distributed variables instead of *uniformly* distributed ones. Therefore, in the second column, we show the scatterplots of transformed variables that are now standard normal (normalized dependence). Specifically, we use the quantile function  $\Phi^{-1}$  of

<sup>12</sup>To obtain the correlation risk premium, Driessen et al. (2009), Buss et al. (2017), (2019a) have to rely on the relationship in (6) as their approach cannot produce pairwise correlations under  $\mathbb{Q}$ . A notable exception is Buss and Vilkov (2012) who impose a linear relationship between pairwise implied and realized correlations. They are then free to use an approach based on either (7) or (14).

<sup>13</sup>We also considered alternative definitions for the down and up correlations, where the mean return or 0 is used as the cutoff instead of the median  $S^m$ . The empirical results were very similar. However, there are important theoretical and practical advantages to defining the conditional correlations with respect to a specific quantile of the index return. In particular, using the median guarantees that the calculation of down and up correlations under  $\mathbb{P}$  is always based on an equal number of realizations. Specifically,  $S^m$  is computed every day of our sample based on the last 63 daily returns. This ensures that the down and up realized correlations can be computed with at least 31 observations. Under  $\mathbb{Q}$ , the down and up correlations can be readily computed because the full joint distribution is known for each trading day.

<sup>14</sup>See Jackwerth and Vilkov (2019) for a discussion of the impact of frequency on the estimation of real-world correlations.

FIGURE 3  
Implied Dependence on Sept. 8, 2008

Graphs A and D of Figure 3 show the dependence of the Financial and Utilities sectors relative to the S&P 500 index. Graphs B and E show the same dependence but after transformation to normally distributed variables. Graphs C and F display the corresponding contour plots.

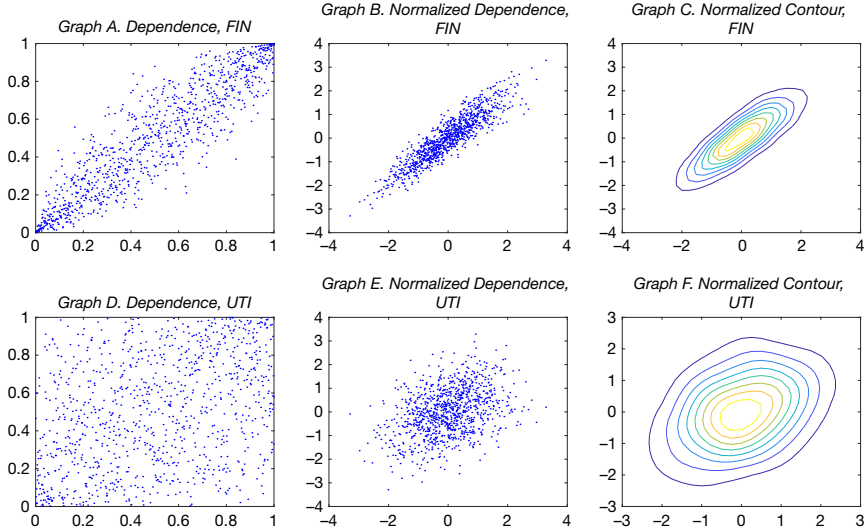
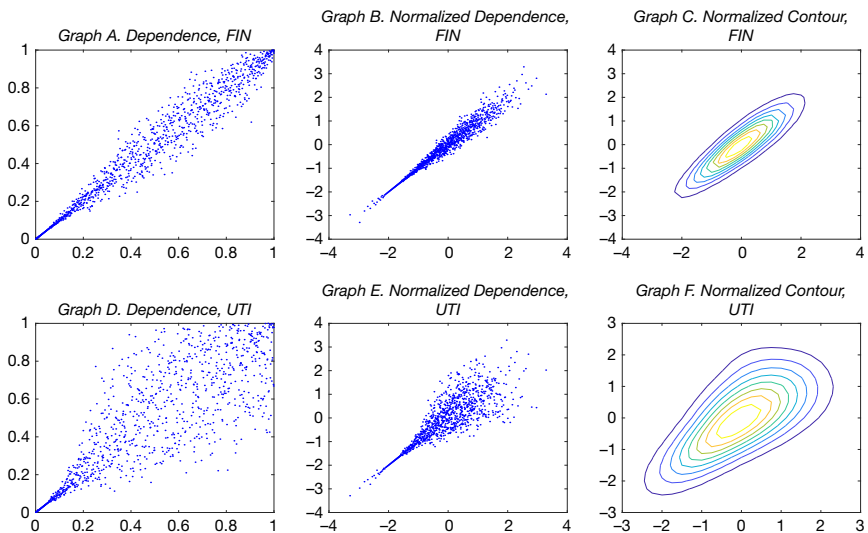


FIGURE 4  
Implied Dependence on Nov. 17, 2008

Graphs A and D of Figure 4 show the dependence of the Financial Utilities sectors relative to the S&P 500 index. Graphs B and E show the same dependence but for normally distributed variables. Graphs C and F display the corresponding contour plots.



the standard normal distribution to define transformation  $G_S(x) := \Phi^{-1}(F_S(x))$  for the S&P 500 index;  $G_{FIN}(x)$  and  $G_{UTI}(x)$  are defined similarly for FIN and UTI. In the second column, we then show scatterplots for the couples  $(G_S(x_i), G_{FIN}(y_i))$  and for the couples  $(G_S(x_i), G_{UTI}(z_i))$ . In the third column, we display the corresponding contour plots. When the dependence is normal, these contours must be perfect ellipsoids. On Sept. 8, 2008, we find positive dependences for both sectors, but the one for the financial sector is much stronger. Both dependences appear slightly asymmetric, with the left tail being stronger than the right tail. On Nov. 17, 2008, the same trends become much stronger for both sectors: the dependence is now noticeably more pronounced, the asymmetry is very obvious, and the left tail dependence is extreme, even for the “calm” Utility sector.

For the two dates, we show all pairwise correlations  $\rho_{j,k}^Q$  on Graph A of Figures 5 and 6. We observe that Financial, Energy, and Technology sectors are

FIGURE 5  
Implied Correlations for the Nine Sectors on Sept. 8, 2008

Graph A of Figure 5 shows the correlation matrix. Graph B shows the implied down correlation  $\rho_{j,S}^{d,Q}$  (y-axis) versus the implied up correlation  $\rho_{j,S}^{u,Q}$  (x-axis). Also shown is the 45-degree line.

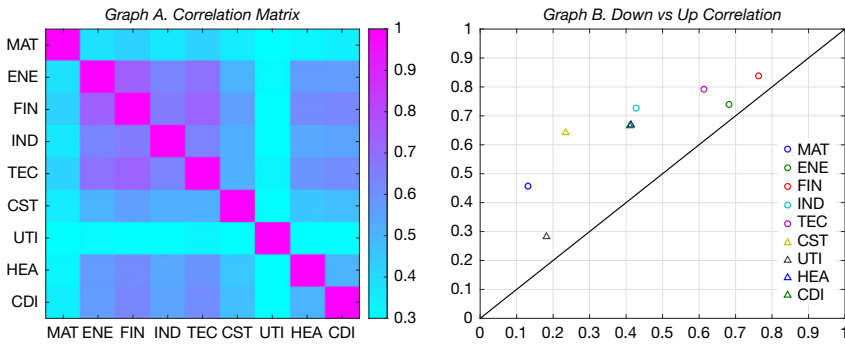
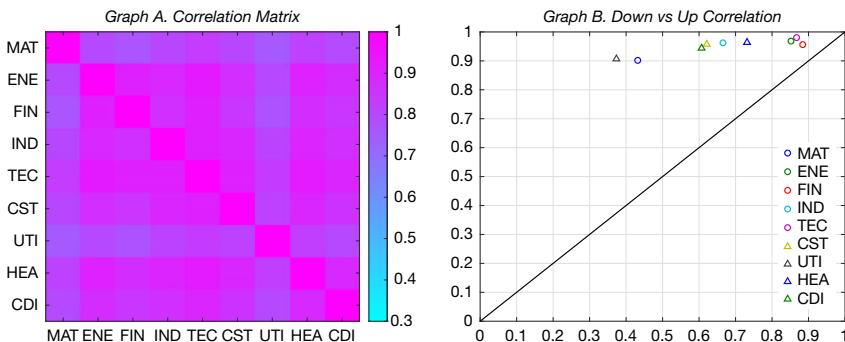


FIGURE 6  
Implied Correlations for the Nine Sectors on Nov. 17, 2008

Graph A of Figure 6 shows the correlation matrix. Graph B shows the implied down correlation  $\rho_{j,S}^{d,Q}$  (y-axis) versus the implied up correlation  $\rho_{j,S}^{u,Q}$  (x-axis). Also shown is the 45-degree line.



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  $\rho^Q$  is 0.58 for Sept. 8, 2008, and 0.88 for Nov. 17, 2008.

If  $X_k$  in definitions (12) and (13) is replaced with the index  $S$ , we obtain the global, down, and up correlations  $\rho_{j,S}^Q$ ,  $\rho_{j,S}^{d,Q}$ , and  $\rho_{j,S}^{u,Q}$  of sector  $j$  with the market. We display the scatter plot of the last two correlations in Graph B of Figures 5 and 6. There are nine points corresponding to the nine sectors. Also shown is the first bisectrix  $\rho_{j,S}^{d,Q} = \rho_{j,S}^{u,Q}$ . From Figures 5 and 6, it is clear that the down correlations tend to be much higher than the up correlations, i.e.,  $\rho_{j,S}^{d,Q} > \rho_{j,S}^{u,Q}$ . In fact, the average down and up correlations are 0.44 and 0.17 for Sept. 8, 2008, but 0.91 and 0.51 for Nov. 17, 2008.<sup>15</sup> On both days, the correlation conditional on the market going down is thus considerably stronger than the correlation conditional on the market going up. This feature is not unique for the two selected days and provides a strong indication that the implied dependence for the nine sectors is asymmetric and thus nonnormal. In Section III.C, we formally assess the extent to which the dependence deviates from normality.

In Section IA.B of the Supplementary Material, we report the option-implied dependence for two more recent dates: Oct. 20, 2017 and Mar. 23, 2020. The first one represents a very calm period, when the market implied volatility was unprecedentedly low. The second one is a very turbulent period at the peak of the COVID-19 crisis. For both dates, the implied dependence remains highly skewed, with the down correlations being much higher than the up correlations. 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). This suggests that, in the midst of the COVID-19 crisis, prices of sector options reflected a distinct possibility of a strong market rally, with most sectors advancing simultaneously. The market indeed had a stunning recovery, increasing by more than 35% over the next 3 months.

### C. Asymmetry of Risk-Neutral Dependence

To formally test whether the option-implied dependence is consistent with a multivariate Gaussian copula, we rely on the classical multinormality tests of Mardia (1970). He constructs two statistics for measuring multivariate skewness (MS) and kurtosis (MK), which can be used to test the hypothesis of normality (Mardia (1974), (1975), Mardia, Kent, and Bibby (1980)). Usually, the test for whether skewness and kurtosis are consistent with a normal model are performed separately; however, so-called omnibus tests can assess them simultaneously. We perform these two tests on the normalized dependence defined as in Section III.B. That is, we use option-implied dependence with attached normal margins and assess whether multivariate normality holds. The two tests are performed separately

<sup>15</sup>The reported average down and up correlations may seem too low when compared to correlations  $\rho_{j,S}^{d,Q}$  and  $\rho_{j,S}^{u,Q}$  shown in Figures 5 and 6. Keep in mind, however, that the average correlations are computed from *pairwise* correlations  $\rho_{j,k}^{d,Q}$  and  $\rho_{j,k}^{u,Q}$ , which are lower than  $\rho_{j,S}^{d,Q}$  and  $\rho_{j,S}^{u,Q}$ .

TABLE 3  
Test for Normal Dependence

Table 3 reports the results of two multivariate normal tests, MS and MK, which are run for all trading days. Shown are the average  $p$ -values and the fraction of the days when the  $p$ -value exceeds a threshold of 0.01, 0.05, or 0.10 (i.e., when the null hypothesis of normal dependence is not rejected).

	<u>No. of Obs.</u>	<u>Mean</u>	<u><math>p &gt; 0.01</math></u>	<u><math>p &gt; 0.05</math></u>	<u><math>p &gt; 0.10</math></u>
MS	3,513	0.0003	0.0011	0.0009	0.0009
MK	3,513	0.1169	0.3251	0.2747	0.2442

on each trading day in the sample, with the results reported in Table 3. Both tests provide strong evidence that the risk-neutral dependence is not of a normal nature, although the evidence for nonzero skewness is more pronounced. Moreover, given the  $p$ -values and the proportion of rejection for the respective tests MS and MK, it is clear that the asymmetry (skewness) is a key feature to reject the normal dependence hypothesis.

## IV. Correlation Risk

### A. Implied and Realized Correlations over Time

Figure 7 plots the average implied correlations for each trading day in our sample, where Graph A displays the time series of  $\rho^Q$ , whereas Graph B displays

FIGURE 7  
Average Implied Correlations

Graph A of Figure 7 shows the average implied global correlation  $\rho^Q$ . Graph B shows the average implied down (blue) and up (green) correlations  $\rho^{d,Q}$  and  $\rho^{u,Q}$ . The gray-shaded areas indicate the financial crisis and COVID-19 crisis.




TABLE 4  
Correlation Risk Premium

Table 4 reports statistics for the risk premia  $\theta$ ,  $\theta^d$ , and  $\theta^u$  computed for the average global, down, and up correlations. 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.

	No. of Obs.	Under P	Under Q	Premium	$t$ -Stat
Global	3,513	0.678	0.748	-0.069	-6.4
Down	3,513	0.527	0.719	-0.193	-10.9
Up	3,513	0.444	0.271	0.173	11.5
Down–up	3,513	0.083	0.449	-0.366	-19.8

both  $\rho^{d,Q}$  and  $\rho^{u,Q}$ . In particular, this figure makes it clear that the gap between the down and up correlations is always positive and consistently very wide. We observe that the sample average for the global correlation (0.748) is larger than for the down correlation (0.719) and is much larger than for the up correlation (0.271) (see Table 4). It is important to stress that the truncated correlations are not directly comparable to the global correlation. The conditioning bias shrinks correlations to 0, which can be derived analytically in the multivariate normal setting.<sup>16</sup> Assuming normality and accounting for the conditioning bias as in (16), the above average down and up correlations translate into equivalent global correlations of 0.864 and 0.423, respectively.

It is also worth stressing that the implied correlations computed for sectors are much higher than those for individual stocks reported in the prior literature. When 500 stocks are aggregated into nine sectors, the idiosyncratic risk gets largely diversified away and the correlation increases.<sup>17</sup> Buss et al. (2017) estimate the implied global correlation for the S&P 500 index in two ways: using the 500 individual stocks and the 9 sectors. Over the sample from 1996 to 2015, they find that the 91-day correlation for the stocks (0.423) is much lower for the sectors (0.700).<sup>18</sup>

<sup>16</sup>Under multivariate normality, the following relation between the down (up) correlation  $\rho_{j,S}^{d,Q}$  ( $\rho_{j,S}^{u,Q}$ ) and the global correlation  $\rho_{j,S}^Q$  holds (for convenience, we suppress the reference to the measure  $Q$  and assets  $j$  and  $S$ ):

$$(16) \quad \rho^d = \rho^u = \rho \sqrt{\frac{1 - \frac{2}{\pi}}{1 - 2\rho^2}}$$

which immediately implies that  $|\rho^d| = |\rho^u| < |\rho|$ . This formula is consistent with derivation in Appendix B of Ang and Chen (2002). See also equation (3) in Campbell et al. (2008).

<sup>17</sup>The effect can be intuitively explained in the homogeneous model, where the pairwise correlation among sectors is  $\rho_{sectors}$  and the pairwise correlation of stocks within each sector is  $\rho_{intra}$ . Under the assumption that there are many sectors and many stocks within each sector, one obtains the approximate relationship  $\rho_{all} \sim \rho_{sectors} \cdot \rho_{intra}$ , where  $\rho_{all}$  is the average pairwise correlation across all stocks. If  $\rho_{intra} = 0.70$  and the average correlation among the nine sectors  $\rho_{sect} = 0.75$ , then  $\rho_{all} = 0.53$ , which agrees with estimates in the literature. For example, the average realized pairwise correlations between sectors is around 0.7 versus around 0.4 between stocks (see Table A105 in Buss et al. (2017)).

<sup>18</sup>To further illustrate this effect, we have applied the MFDR methodology to estimate the dependence for the 30 stocks in DJIA in Section IA.E of the Supplementary Material. 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 opposed to 0.678, 0.527, and 0.444 for the nine sectors in S&P 500, as shown, respectively, in Table 4 and Table IA.4 in the Supplementary Material. This indicates that the correlations for individual stocks are much lower than their counterparts for the sectors. The implied dependence for



FIGURE 8

## Average Implied Correlation, Cumulative Returns, and ATM Implied Volatilities

Figure 8 focuses on the period of the financial crisis (Aug. 1, 2007 to Apr. 1, 2009). The red solid lines show the dates that separate Period I, Period II, and Period III of the financial crisis (Mar. 17, 2008 and Sept. 15, 2008). The black dashed lines indicate several extreme trading days (July 14, 2008, Oct. 9, 2008, Nov. 20, 2008, and Jan. 20, 2009). Graphs B and C show the financial sector (blue line) and the S&P 500 index (black line).




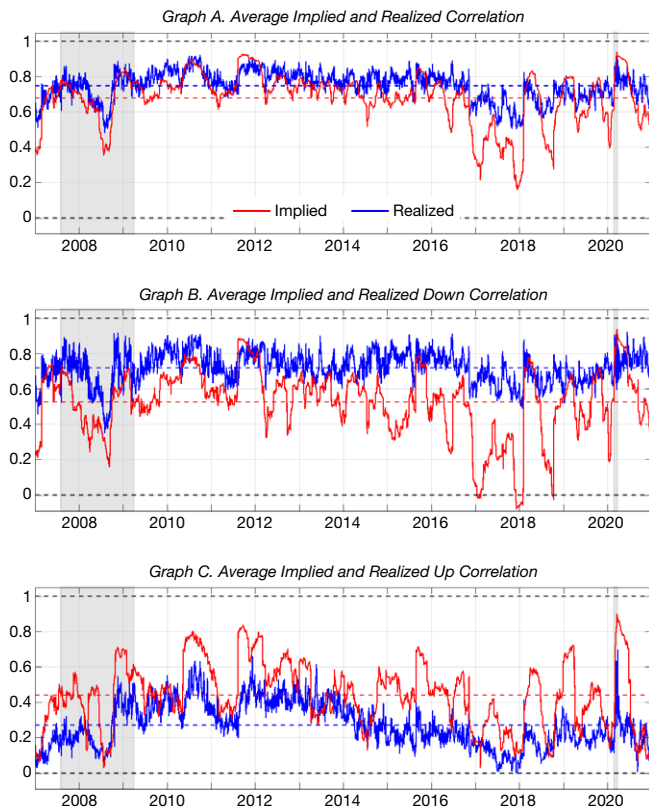
Figure 7 reveals that all three correlations drop sharply during the second half of 2008, which might seem counterintuitive, as this period is in the midst of the financial crisis. To better understand this finding, we split the time period Aug. 2007–July 2009 into three subperiods as shown in Figure 8: Period I runs until Mar. 2008 (the bailout of Bear Sterns), Period II is from Mar. 2008 until Sept. 2008 (the failure of Lehman Brothers), and Period III is from Sept. 2008 until July 2009.

In 2006, the U.S. financial industry was highly exposed to subprime mortgages. When house prices started to fall in July 2006, this had an immediate impact on banks. The financial sector declined, and its implied volatility increased, as confirmed by the blue lines in Figure 1. By Aug. 2007, correlations had also increased considerably, as confirmed by the red line in Figure 9. However, the crisis had yet not spread to other sectors, and the market implied volatility remained fairly stable, as evidenced by the black line in Graph C of Figure 8. At that point, there was a widespread belief among market participants that by lowering interest rates, the Federal Reserve could boost market liquidity and restore confidence. In Mar. 2008,

individual stocks still exhibits high asymmetry, with the down correlation being much larger than the up correlation.

FIGURE 9  
Implied and Realized Correlations

Figure 9 shows implied correlations (blue) 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 gray-shaded areas indicate the financial crisis and COVID-19 crisis.



Bear Stearns became the first of several financial institutions to be bailed out by the government. This event created an expectation that if needed, the government would rescue any other banks and thus that the crisis would not spill over to other sectors. Correlations and volatilities temporarily decreased during Period II, returning to the pre-crisis levels as evidenced by Graphs A and C of Figure 8. However, the situation changed drastically after the bankruptcy of Fannie Mae and Freddie Mac on Sept. 7, followed by that of Lehman Brothers a week later. As the government decided to let Lehman fail, the market stress increased sharply and investors started to panic. A new crisis period started as investors rushed to the safest investments, such as cash or government securities. During the last quarter of 2008, markets fell worldwide, and volatilities and correlation peaked and remained at high levels afterward, as confirmed by all 3 graphs of Figure 8 during Period III.

## B. Risk Premia

We use the average correlations defined in Section III.A to study down and up CRP. Formally, we define the global, down, and up correlation premia as

$$(17) \quad \theta^a = \rho^{a,\mathbb{P}} - \rho^{a,\mathbb{Q}},$$

where, as before,  $a \in \{g, d, u\}$  corresponds to the global, down, and up correlations. Recall that the correlations under  $\mathbb{P}$  are computed over the 3-month trailing window, so that the risk premia are defined in the ex ante fashion. The same approach has been used, among others, in Driessen et al. (2009), Buss et al. (2017), (2019a), and Bollerslev, Tauchen, and Zhou (2009).<sup>19</sup>

Figure 9 plots across time the average implied correlation (blue line) and realized correlation (red line). Graphs A, B, and C correspond to the three types of correlations: global (A), down (B), and up (C). We observe that implied and realized correlations generally behave very similarly. This is reassuring: the forward-looking implied correlation extracted from options using our MFDR methodology appears to be closely related to the real-world correlation computed from asset historical returns. Second, we observe that both implied and realized correlations display considerable variation over time and that the magnitude of the fall in the correlations that we observe during 2008 is not that unusual. For example, between Mar. 2018 and Oct. 2018, the realized global correlation first dropped from 60% to 10% and then rebounded to nearly 85% at the height of the COVID-19 crisis in Mar. 2020.

Graph A of Figure 9 documents that the global CRP  $\theta$ , which appears as the difference between the blue and the red lines, is mostly negative. Across the sample period, it has an average of  $-0.069$  and is statistically highly significant with a  $t$ -statistic of  $-6.4$  (see Table 4). Buss et al. (2017) also estimate the CRP for the nine sectors of the S&P 500 at the 3-month horizon, but over a different sample period from 1996 to 2015. Their estimate ( $-0.059$ ) is of similar magnitude as ours.<sup>20</sup>

Graphs B and C of Figure 9 illustrate our most intriguing contribution to the CRP literature. Graph B visually demonstrates that the average realized down correlations are systematically lower than their implied counterparts (the red line is consistently below the blue line). The opposite is true for Graph C: The average realized up correlations are systematically higher than their implied counterparts. Table 4 reveals that on average the up correlations are lower than the down correlations under  $\mathbb{P}$  and that this asymmetry is even more pronounced under  $\mathbb{Q}$ , with  $\rho^{u,\mathbb{Q}} < \rho^{u,\mathbb{P}} < \rho^{d,\mathbb{P}} < \rho^{d,\mathbb{Q}}$ . As the result, the average down CRP  $\theta^d$  is negative

<sup>19</sup>It is common to estimate the realized and implied variances at a given time- $t$  from the information available at that time. That is, the implied variance is estimated from options available at time- $t$ , while the realized variance is estimated from past historical returns up to time- $t$ . Specifically, Bollerslev et al. (2009) define the variance risk premium as “the difference between this ex ante risk-neutral expectation of the future return variation over the  $[t, t + 1]$  time interval and the ex post realized return variation over the  $[t - 1, t]$ .”

<sup>20</sup>Note that the magnitude of the CRP for *sectors* is generally smaller than that for *individual stocks*. In particular, Driessen et al. (2013) and Buss et al. (2017), (2019a) report the CRP for stocks in the S&P 500 index of  $-0.106$ ,  $-0.100$ , and  $-0.099$ , respectively. However, as mentioned earlier, the correlations for sectors are not directly comparable to those for stocks.

(−0.193), whereas the average up CRP  $\theta^u$  is positive (0.173). Both risk premia are highly significant and economically very large. They are consistent with the economic intuition that investors are mainly concerned with the loss of diversification when the market falls. Consequently, they are willing to pay a considerable premium to hedge against increases in the down correlation. However, investors actually *prefer* high correlation when the market rallies. That is, investors view the down correlation as “bad” and the up correlation as “good.” The net effect of the negative premium for the down correlation and the positive premium for the up correlation is a negative premium for the global correlation.

The last row of Table 4 focuses on the correlation spread (the difference between the down and up correlations),  $\Delta\rho^{\mathbb{M}} = \rho^{d,\mathbb{M}} - \rho^{u,\mathbb{M}}$ , where  $\mathbb{M} \in \{\mathbb{Q}, \mathbb{P}\}$  indicates the probability measure under which the expectations are evaluated. Unsurprisingly, the corresponding risk premium,  $\Delta\rho^{\mathbb{P}} - \Delta\rho^{\mathbb{Q}}$ , is negative and highly significant (−0.366,  $t$ -stat = −19.8). Its magnitude is about 5.3 times larger than the CRP for the global correlation: the magnitude of risk premium for the down (up) correlation is approximately 2.8 (2.5) times larger than for the global correlation.

These observations motivate us to define the DUC swap. At time- $T$ , this swap has a payoff (the variable leg) equal to the difference between the *realized* down and up correlations, or  $\Delta\rho^{\mathbb{P}}$ . That is, the daily returns realized over the life of the swap are split into two equal groups, when the market return is low and high. The average weighted correlations are computed for each group and their difference gives the payoff of the swap. At time-0, the DUC swap has the initial cost (the fixed leg) equal to the difference between the *risk-neutral* down and up correlations, or  $\Delta\rho^{\mathbb{Q}}$ . Historically, selling the DUC swap would have been very profitable, as it takes advantage of both selling the “expensive” down correlation and buying the “cheap” up correlation. However, two caveats are in order. The DUC swap is not yet tradable. We assume that historically it would have been traded at the price implied by our BRA approach. Furthermore, the transaction costs, likely to be substantial, are not accounted for.<sup>21</sup>

We observe that the signs of the down and up CRP mirror those for the down and up variance risk premium, as recently reported in Feunou et al. (2018) and Kilic and Shaliastovich (2019). However, it is important to emphasize that our results neither follow from nor imply the latter findings. A priori, there is no theoretical reason why the sign of the down (up) CRP should match that for the down (up) variance risk premium. To make this point precise, we consider a simple model in which the  $d$  sectors have returns distributed according to a mixture of two regimes. In each regime, returns are jointly normally distributed with homogeneous parameters. The details of this model are relegated to Section IA.D of the Supplementary Material.

We focus on the equal-weighted index  $S = \sum_{j=1}^d \omega_j X_j$  with  $\omega_j = 1/d$  and  $d = 9$ . We keep the parameters of the model fixed under  $\mathbb{P}$ , but vary them under  $\mathbb{Q}$  to investigate possible signs of the variance and CRP. Let  $\text{VRP}^a$  and  $\text{CRP}^a$

<sup>21</sup>We leave for future research development of a tradable strategy, which could proxy the DUC swap.

TABLE 5  
Signs of the Variance and Correlation Risk Premia

The parameters for the 3 scenarios in Table 5 are given in Table IA.2 in the Supplementary Material.

	$VRP^g$	$VRP^d$	$VRP^u$	$CRP^g$	$CRP^d$	$CRP^u$
Scenario 1	-	-	+	-	-	-
Scenario 2	-	-	+	-	-	+
Scenario 3	-	-	+	-	+	-

denote the risk premium for the variance and correlation, respectively, and where  $a \in \{g, d, u\}$  indicates global, down, or up premia. VRP is computed from the index returns, as the difference between the variance under  $\mathbb{P}$  and under  $\mathbb{Q}$ , e.g.,  $VRP^g = \text{var}^{g, \mathbb{P}}(S) - \text{var}^{g, \mathbb{Q}}(S)$ . In Table 5, we report the signs of risk premia for 3 different scenarios. For all scenarios, we choose the parameters of the mixture model under  $\mathbb{Q}$  to be consistent with the prior empirical evidence on the risk premia. Therefore, the first 4 columns in the table are the same across the 3 scenarios reflecting the fact that the global VRP, down VRP, and global CRP are negative, whereas the up VRP is positive.<sup>22</sup> The last two columns, however, show that, depending on parameterization, the down and up CRP could be both negative, the first negative while the other positive, and vice versa. In general, the signs of the VRP do not determine the signs of the CRP. In Scenario 3, for instance, the signs of the down and up CRP are *opposite* to the signs of the down and up VRP.

The previous studies of the market variance risk premium (Feunou et al. (2018) and Kilic and Shaliastovich (2019)) are based exclusively on index options and, intuitively, they demonstrate that index OTM puts are expensive and index OTM calls are cheap when compared to the historical distribution of the index returns. Our results, on the other hand, make a statement on the *relative* pricing of individual options compared to the index options. That is, we evaluate the *joint* pricing of the index *and* individual options. In this regard, we find that sector OTM puts are *cheap* compared to index OTM puts (thus, implying too high down correlations between the sectors), whereas sector OTM calls are *expensive* compared to index OTM calls (implying too low up correlations).

### C. Marginal Distributions or Dependence?

As shown in Section IV.B, time-varying correlation is priced and its risk premium changes the sign between the down to up correlations. Since correlations are affected jointly by the margins and the dependence, the risk premium could potentially stem from either one. Disentangling the respective roles of the margins and the dependence is an important but hard question. Ideally, for each trading day we need to estimate the two components under both the risk-neutral and physical measures. Under  $\mathbb{Q}$ , this is feasible, as MFDR yields the full joint distribution. However, under  $\mathbb{P}$ , we only observe one draw from the unobserved (and time-varying) distribution.

<sup>22</sup>The literature documenting the negative global VRP includes Bondarenko (2004), (2014), Bollerslev et al. (2009), and Carr and Wu (2009).

To better understand the roles of the margins and the dependence, we perform two exercises. In the first exercise, we focus on  $\mathbb{Q}$  measure only, for which we have complete information. We formally investigate the causes of the enormous asymmetry between the down and up correlations under  $\mathbb{Q}$  in Table 4 (i.e., a large positive correlation spread)  $\Delta\rho^{\mathbb{Q}} = \rho^{d,\mathbb{Q}} - \rho^{u,\mathbb{Q}}$ . As argued earlier, this asymmetry is inconsistent with a multivariate normal (MVN) distribution, for which the down and up correlations must be the same. The observed asymmetry may be caused by non-normality of the margins or nonnormality of the dependence. For example, the observed asymmetry could at least partially be due to the heavier left tails of the margins compared to the right tails. To disentangle the effects of the margins and of the dependence in explaining the correlation spread  $\Delta\rho^{\mathbb{Q}}$ , we contrast 4 cases, labeled NN, EN, NE, and EE, where the first letter denotes the type of the margins (Normal or Empirical) and the second letter denotes the type of the copula (again, Normal or Empirical). The EE case corresponds to the output distribution obtained using our model-free BRA algorithm (Empirical margins and Empirical copula). The NN case corresponds to Normal margins and Normal copula (i.e., an MVN distribution). Specifically, the standard deviations for the normal margins match those of the empirical margins, and the normal copula is calibrated to a constant correlation matrix in such a way that the model preserves the average pairwise correlation among the nine sectors. The NE case uses Normal margins joined with the Empirical dependence from the BRA algorithm (i.e., misspecified margins and correct dependence). Finally, the EN case uses the Empirical margins joined with a Normal copula (i.e., correct margins and misspecified dependence).

For all 4 cases, we compute the average pairwise global, down, and up correlations and display them as time series in the 3 graphs of Figure 10: NN (blue), EN (green), NE (red), and EE (black). In Table 6, we report the time-series averages of each quantity. In the first graph of Figure 10 (the average global correlation), the 4 lines are all close to each other. In fact, the average global correlations for EE and NN are identical by construction. In the second and third graphs (the average down and up correlations), the 4 lines are farther apart. Furthermore, for the up correlation the EN and NN lines are close to each other, as are the NE and EE lines. For the down correlation, the pattern is similar, although there are generally wider gaps between the corresponding pairs. We thus conclude that the correlation spread  $\Delta\rho^{\mathbb{Q}}$  is mainly driven by the type of copula and not by the potential nonnormality of the margins. This conclusion is reinforced by the last row of Table 6. The correlation spread is 0 for the multivariate normal case (NN), 0.048 for the EN model, and 0.449 for the EE model, which fits perfectly options on individual sectors and the index. Even though the EN model has symmetric copula, its down and up correlations are not the same due to the skewed margins. When comparing the NN and EN models, we can intuitively see that adopting the nonnormal margins explains only about 11% of the true correlation spread (0.048/0.449). The rest can be attributed to the nonnormality of the dependence.

In the second exercise, we focus on both  $\mathbb{Q}$  and  $\mathbb{P}$  measures and attempt to disentangle the respective roles of the margins and the dependence on the CRP. We repeat the analysis of Table 4 but now using Spearman (instead of Pearson) correlations. Spearman's is a rank correlation and it has the advantage of being unaffected by the margins. To compute Spearman correlation under  $\mathbb{Q}$ , however, the

FIGURE 10  
Implied Correlations

In Figure 10, we compute average implied global, down, and up correlations for the 4 cases (NN, EN, NE, and EE), where the first letter denotes the type of margins (Normal or Empirical) and the second letter denotes the type of copula (Normal or Empirical). For example, EN denotes Empirical margins joined with a Normal copula. Statistics are plotted as 1-month moving averages.

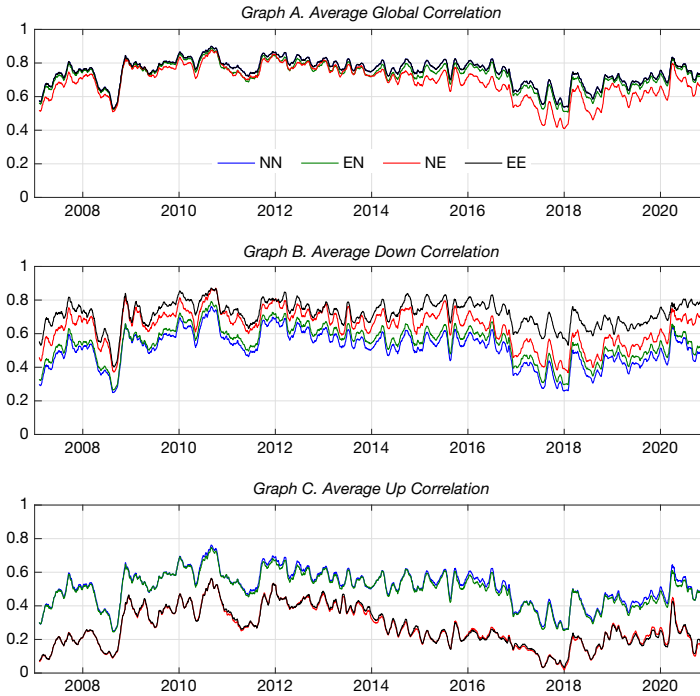


TABLE 6  
Implied Correlations

In Table 6, we compute average implied global, down, and up correlations for the 4 cases (NN, EN, NE, and EE), where the first letter denotes the type of margins (Normal or Empirical) and the second letter denotes the type of copula (Normal or Empirical). For example, EN denotes Empirical margins joined with a Normal copula. All statistics are under the  $\mathbb{Q}$  measure. The last row is the correlation spread,  $\Delta\rho^{\mathbb{Q}} = \rho^{\text{d},\mathbb{Q}} - \rho^{\text{u},\mathbb{Q}}$ .

	NN	EN	NE	EE
Global	0.748	0.728	0.689	0.748
Down	0.512	0.548	0.640	0.719
Up	0.511	0.500	0.269	0.271
Down-up	0.000	0.048	0.372	0.449

complete joint distribution is required, even for the case of the global correlation. Thus, our BRA methodology is crucial. The results are reported in Table 7. We observe that, compared to Table 4, all correlations are now slightly lower under both measures. However, the correlation spread  $\Delta\rho$  remains positive under both  $\mathbb{P}$  and  $\mathbb{Q}$ . Importantly, the risk premia for the three types of correlations (global, down, and up) remain quantitatively very similar, but now they are not confounded by the

TABLE 7  
Correlation Risk Premium, Spearman

Table 7 reports statistics for the risk premia  $\theta$ ,  $\theta^d$ , and  $\theta^u$  computed for the average global, down, and up correlations. It is the same as Table 3, except now the computations are based on Spearman (instead of Pearson) correlations. 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.

	No. of Obs.	Under P	Under Q	Premium	$t$ -Stat
Global	3,513	0.635	0.664	-0.029	-3.2
Down	3,513	0.447	0.622	-0.175	-12.9
Up	3,513	0.394	0.228	0.166	13.0
Down–up	3,513	0.053	0.394	-0.341	-21.4

dynamics of the marginal distributions. Therefore, by contrasting the results for Pearson and Spearman correlations, we conclude that the CRP is mainly driven by the dependence and not by the margins.

#### D. Tail Indices

In this article, we have studied two new dependence indicators, the down and up correlations, or correlations conditional on the index return being below or above its median. One natural generalization is to consider more detailed indicators that condition on other quantiles of index returns. This would permit a better characterization of the dependence deeper in the tails.<sup>23</sup> In fact, we can take this idea to the limit and estimate the left and right tail indices. Given a joint distribution of two random variables  $(X, Z)$ , let  $X_q$  and  $Z_q$  denote  $q$ -quantiles for  $X$  and  $Z$ . The left-tail index is defined as

$$LT(q) := \frac{2 \log(q)}{\log(\text{Prob}(X \leq X_q, Z \leq Z_q))} - 1.$$

This index is equal to 0 when the two variables are independent and 1 when they are comonotonic (have perfect positive dependence).<sup>24</sup> The right-tail index  $RT(q)$  is defined similarly. To reduce the noisiness of the estimator in finite samples, we choose  $q$  to be not too small, such as  $q=0.1$ . We compute the left-(right-) tail index between each sector and S&P 500 and value-weight them across the nine sectors. In Figure 11, we show  $LT(q)$  and  $RT(q)$  over time and observe that their empirical behavior is qualitatively similar to that of the down and up correlations. The left-tail index is always larger than the right-tail index. Quantitatively, both tail indices are higher than the corresponding correlations. Similarly to the analysis of Section IV.C, a careful check must be performed to better understand whether the asymmetry is driven by the properties of the marginal distributions or of the dependence.


<sup>23</sup>Longin and Solnik (2001) perform a study in this spirit under the real-world measure. We are now in position to extend their work to option-implied forward-looking measures.

<sup>24</sup>An alternative left tail index is considered in Chabi-Yo et al. (2018) to study the dependence between individual stock returns and the market return under the real-world probability.



FIGURE 11  
Implied Left- and Right-Tail Indices

In Figure 11, tail indices  $LT(q)$  (blue) and  $RT(q)$  (green) are computed for  $q=0.1$ . The gray-shaded areas indicate the financial crisis and COVID-19 crisis.



## E. Implications for Literature and Hybrid Model

The results in Section IV.C highlight the importance of proper modeling of the dependence, especially under  $\mathbb{Q}$ . The large positive correlation spread  $\Delta\rho^{\mathbb{Q}}$  is mainly driven by the nonnormality of the dependence and the CRP is mainly a compensation for the dependence risk. Therefore, a financial economist who wants to build a multi-asset model that matches the salient features of the option data must pay a special attention to the dependence. She could pursue two distinct strategies: a) to model the joint distribution directly and b) to model the margins and the dependence separately and then combine the two pieces to build the joint distribution.

The first strategy is more common because of its tractability (e.g., Chang, Christoffersen, Jacobs, and Vainberg (2012), Schreindorfer (2020)). However, it suffers from the fact that the number of parsimonious joint models is limited.

Here, we follow an alternative approach to model the multivariate joint distribution, which is based on previous strategy (b). It consists of estimating the marginal distributions from a model-free approach as discussed in the article but to use a copula based on the homogeneous multivariate skewed normal distribution driven by two parameters  $\rho$  and  $\delta$ , denoted by  $SN(\rho, \delta)$  (see Appendix C1 for the formal definition).

We refer to this model as *hybrid* because it combines i) fully nonparametric margins extracted from the individual options and ii) a parsimonious parametric copula. It offers a number of advantages. First, it fits perfectly all sector marginal distributions. That is, the model by construction matches exactly the prices of all sector options. Second, it automatically enforces any additional consistency condition (e.g., that the average beta of the nine sectors (computed for simple returns) must be equal to 1). Third, the model provides a very tight overall fit to option data.

To understand the appropriateness of the hybrid model, we compare it with some alternative approaches. Specifically, recall that the definition in (3) says that the weighted sum  $Z$  can be constructed from the known margins  $F_j$  and the copula  $C$ . By changing the copula  $C$ , we obtain different distributions of  $Z$  and, ideally, want to match the known distribution of  $S$  as in (4). To perform a comparison, we need some measure of goodness-of-fit. One intuitive approach is to compare the

two implied curves  $IV^Z$  and  $IV^S$ , which correspond to  $Z$  and  $S$ . Specifically, we define the distance as the root mean squared relative error (RMSRE):

$$(18) \quad D := D(Z, S) = \sqrt{\frac{1}{L} \sum_{k=1}^L \left( \frac{IV_k^Z - IV_k^S}{IV_k^S} \right)^2},$$

where we use the  $L = 20$  moneyness levels  $k = 1, \dots, L$ , equally spaced between 1% and 99%.

We consider several choices for the copula  $C$  and thus for the corresponding weighted sum  $Z$ . One choice corresponds to the output of our model-free approach, which we, as before, denote as EE. Alternatively, we can construct the copula from some parametric model. A common choice in the literature is the normal copula, but this model is not likely to perform well, given our finding on the very strong asymmetry of the implied dependence. Therefore, we consider a more flexible generalization, a multivariate skewed normal (SN) copula with two parameters  $\rho$  and  $\delta$ , denoted by  $SN(\rho, \delta)$  (see [Appendix C1](#) for the formal definition). We focus on the homogeneous case, where the pairwise correlations are identical and the skewness parameter is also identical across all assets.

In [Table C.1](#), we contrast the different models for Sept. 8, 2008, one of the two dates studied in [Section III](#). Initially, we choose the parameters of the SN copula to match the standard deviation of the index  $\sigma_S^Q$ . Specifically, for three (arbitrary) choices of the skewness parameter  $\delta$ , we vary the remaining free parameter  $\rho$  to match  $\sigma_S^Q$ . These 3 cases, denoted as  $SN_1$ ,  $SN_2$ , and  $SN_3$ , are presented in the first 3 rows of [Table C.1](#). The second and third columns in [Table C.1](#) report the parameters of the SN copulas. Next, we estimate the “optimal” SN copula, for which both parameters are chosen to minimize the distance  $D$  in (18). That model, denoted  $SN^*$ , is presented in the fourth row. Finally, the EE model from our MFDR approach is presented in the last row.

For each of the five models, [Table C.1](#) also reports the average correlations (global, down, and up). Since models  $SN_1$ ,  $SN_2$ , and  $SN_3$  all match  $\sigma_S^Q$ , the global correlation is forced to be the same by construction. That is, any approach based solely on matching  $\sigma_S^Q$ , as is typically done in the literature, would yield the exact same global correlation if the index distribution were given by any of the three models  $SN_1$ – $SN_3$ . In other words, the approach would not have been able to distinguish between the 3 cases even though their down and up correlations are drastically different.

Despite its simplicity, the model captures reasonably well the most salient features of the option-implied dependence.<sup>25</sup> As already mentioned, the hybrid model fits perfectly the sector options, but it could potentially misprice the index options. The mispricings, however, are much smaller than when a multivariate normal copula is used. Of course, in terms of fitting option prices the hybrid model cannot compete with the BRA approach, because the latter produces (essentially) a perfect fit. Instead, the hybrid model offers different advantages: it is transparent, intuitive, and easy to implement. Our primary motivation for developing this model

<sup>25</sup>The hybrid model has only 2 free parameters and, thus, is relatively easy to estimate.

is twofold. First, because the hybrid model does not rely on the somewhat opaque BRA methodology, it provides an alternative confirmation of our key empirical results, including the analysis of the CRP. Second, we believe that the hybrid model could prove useful in other applications where the BRA methodology cannot be implemented due to data limitations.<sup>26</sup> An additional analysis of the implied volatilities in the various models shows that the hybrid model may fit very well on most days, such as Sept. 8, 2008 but that it may also fail to reproduce the implied volatility pattern that only the EE approach succeeds in. Full details for the hybrid model are provided in [Appendix C](#).

## F. Return Predictability

A large body of the literature uses various option-implied factors to predict future market returns. In particular, these forward-looking factors include variance, semi-variance, correlation, or skewness (see, e.g., Buss and Vilkov (2012), Amaya et al. (2015), Stilger, Kostakis, and Poon (2017), Buss et al. (2019b), Jondeau et al. (2019), and Martin and Wagner (2019)). The MFDR methodology could provide new insights as to which specific option-implied factors drive return predictability. In this subsection, we explore whether the down and up correlations could predict future returns better than the standard global correlation.

Our sample is relatively short and we start by focusing on the canonical in-sample univariate predictive regression:

$$(19) \quad r_{t,t+h}^e = \alpha + \beta \cdot x_t + \varepsilon_t,$$

where  $r_{t,t+h}^e$  denotes the cumulative excess return for the index over the period  $[t, t+h]$ ,  $h$  is the forecasting horizon, and  $x_t$  is the predictor known at time- $t$ . Using a similar regression, Buss et al. (2019b) study predictive power of several option-implied measures (correlation, variance, and semi-variance), as well as their realized counterparts and risk premia. They conclude that the implied correlation is by far the best univariate predictor. Therefore, for the predictor  $x_t$ , we use the three types of implied correlation: global, down, and up. We set the forecasting horizon  $h$  to 1, 3, 6, 9, or 12 months (i.e., 21, 63, 126, 189, and 252 trading days) and estimate the regression monthly. When using horizons longer than 1 month, we account for autocorrelation induced by overlapping returns by computing Newey–West  $t$ -statistics with the number of lags equal to the horizon.

Panel A of [Table 8](#) presents the results for the baseline specification. Using the OLS regression for the S&P 500 return, the table reports the beta coefficient, its Newey–West  $t$ -statistics, and the adjusted  $R^2$ . The results for the implied global correlation are consistent with the prior literature:  $\rho^{g,Q}$  is a strong predictor of the market return; the adjusted  $R^2$  statistics increases from 2.09% at the 1-month horizon to its maximum of 11.20% at the 6-month horizon and falls off after that.

<sup>26</sup>Our approach can be compared to Jackwerth and Vilkov (2019) who estimate the joint distribution of the S&P 500 return and its volatility. They also estimate a hybrid model, where model-free margins are combined with a parametric copula. Specifically, Jackwerth and Vilkov (2019) find that Frank copula works best for their application. Our problem is quite different and of higher dimension. The skewed normal copula provides a parsimonious way to model the strong asymmetric dependence observed in our application.

TABLE 8  
 Predictive Regressions

Table 8 reports the results of univariate predictive regressions over five horizons  $h$  (1, 3, 6, 9, and 12 months). The independent variable is the excess return for S&P 500. The predictor variable is the implied global, down, or up correlation. For each horizon and predictor, the table shows the beta coefficient, its Newey–West  $t$ -statistics computed with the number of lags equal to the horizon, and the adjusted  $R^2$  (in percentage).

$h$ Months	Global			Down			Up		
	$\beta$	$t$ -Stat	$R^2$	$\beta$	$t$ -Stat	$R^2$	$\beta$	$t$ -Stat	$R^2$
<i>Panel A. OLS Regression</i>									
1	0.10	2.02	2.09	0.07	1.23	1.06	0.08	2.34	3.60
3	0.29	2.15	6.66	0.23	1.69	5.26	0.19	2.74	7.96
6	0.51	2.22	11.20	0.36	1.39	6.34	0.38	3.47	16.67
9	0.58	2.41	9.45	0.34	1.42	3.62	0.50	3.34	19.22
12	0.64	2.87	9.09	0.31	1.57	2.19	0.59	2.98	19.80
<i>Panel B. Ex Financial Crisis</i>									
1	0.09	2.46	2.64	0.07	2.00	1.13	0.06	2.28	3.90
3	0.20	2.24	4.29	0.12	1.70	1.27	0.15	2.59	7.28
6	0.33	2.31	8.22	0.16	1.10	1.27	0.27	3.34	15.65
9	0.39	2.74	8.54	0.15	1.08	0.70	0.34	3.43	19.24
12	0.50	2.84	13.37	0.22	1.30	2.06	0.39	2.97	21.42
<i>Panel C. WLS Regression</i>									
1	0.08	1.63	1.59	0.07	1.23	1.28	0.06	1.85	2.51
3	0.22	1.65	4.79	0.20	1.40	4.17	0.16	2.22	6.11
6	0.40	1.74	8.95	0.32	1.24	6.03	0.32	2.86	13.54
9	0.45	1.90	7.47	0.32	1.31	3.73	0.40	2.68	14.63
12	0.59	2.68	9.64	0.36	1.83	3.48	0.53	2.69	18.68

The regression coefficient is always significantly positive, with  $t$ -statistics being 2.0 and higher. Turning to the other two correlations,  $\rho^{d,Q}$  and  $\rho^{u,Q}$ , we find that across all horizons the adjusted  $R^2$  is lower (higher) for the down (up) correlation. In particular, when compared to the global correlation, the predictive power for the up correlation is slightly higher for short horizons, but much stronger for longer horizons. Its  $R^2$  continues to increase all the way to the 12-month horizon, where it reaches an impressive level of 19.80%. The regression coefficient is positive and highly statistically significant. As for the down correlation, its  $R^2$  is much lower and its beta coefficient is never significant.

It is interesting to note that Feunou et al. (2018), Kilic and Shaliastovich (2019), and Buss et al. (2019b) find that the down semi-variance risk premium is a better predictor than the (global) variance risk premium, whereas the predictability by the up semi-variance risk premium is much weaker. In other words, it is the downside part of the variance that contains useful information about market future returns. In contrast, we find that the most valuable information for predicting market future returns is contained in the implied *up* correlation, whereas the predictability by the implied *down* correlation is weak.

We check the robustness of our predictability results to a number of alternative specifications. First, one might be concerned that the strong return predictability could be driven by a few influential outliers. Following Kilic and Shaliastovich (2019), we exclude from our analysis observations corresponding to the financial crisis, which was a period of high market volatility and extreme realized returns. Panel B of Table 8 reports the results for the sample excluding the financial crisis and they are similar to the results based on the entire sample.

The up correlation remains the strongest predictor. Although its betas are now slightly lower for all horizons, the adjusted  $R^2$ s are generally higher, ranging from 3.90% to 21.42%.<sup>27</sup>

Second, instead of completely removing influential observations, we can down-weight them using the weighted least squares (WLS) regression. Panel C of Table 8 reports the results for the WLS regression, where the weight at time- $t$  is equal to the inverse of the 3-month implied volatility. Thus, observations during volatile periods receive lower weights compared to those during calm periods. Intuitively, this corresponds to a market timing strategy, which uses implied correlation as a signal, but reduces the risk exposure during volatile periods. The main conclusions remain. Although the predictability for the up correlation now slightly decreases across all maturities, its  $R^2$  is as high as 18.68% for the 12-month horizon.

Third, a common concern with predictability regressions is that they often do not work well out-of-sample (see Campbell and Thompson (2008), Goyal and Welch (2008)). Therefore, Table 9 presents the results for the out-of-sample regressions. Specifically, the model in (19) is reestimated every month using 60-month rolling window. At any time- $t$ , the estimated coefficients are used to forecast the subsequent return  $r_{t,t+h}^e$  avoiding any look-ahead bias. Then we move the estimation window by 1 month, reestimate the model, and apply new coefficients to the next month.<sup>28</sup> In Table 9, we report the averages across all estimations for the beta,  $t$ -statistics, and in-sample adjusted  $R^2$ . Estimates for betas and  $R^2$ s are stable and their averages are similar to those in Table 8. Also reported is the out-of-sample  $R^2$  statistics, which compares the mean-squared error of a candidate predictive model to that of the benchmark model. The latter one uses the historical mean return as a

TABLE 9  
Out-of-Sample Regressions

Table 9 reports the results of univariate predictive regressions over five horizons  $h$  (1, 3, 6, 9, and 12 months). The independent variable is the excess return for S&P 500. The predictor variable is the implied global, down, or up correlation. For each horizon and predictor, the table shows the average beta coefficient, the average  $t$ -statistics (Newey–West adjusted with the number of lags equal to the horizon), the average adjusted in-sample  $R^2$  statistics, and the out-of-sample  $R^2$  statistics. Both  $R^2$  statistics are in percentage.

$h$	Global				Down				Up			
	$\beta$	$t$ -Stat	$R_{IS}^2$	$R_{OOS}^2$	$\beta$	$t$ -Stat	$R_{IS}^2$	$R_{OOS}^2$	$\beta$	$t$ -Stat	$R_{IS}^2$	$R_{OOS}^2$
1	0.11	1.32	1.37	1.08	0.05	0.82	-0.31	-0.73	0.09	1.63	2.88	8.76
3	0.28	1.40	5.76	1.93	0.12	0.72	1.97	-2.92	0.21	1.98	9.11	4.50
6	0.45	1.45	9.30	12.09	0.16	0.60	2.81	0.77	0.40	2.34	18.74	20.40
9	0.44	1.41	6.66	8.74	0.09	0.16	1.93	-4.49	0.42	2.34	17.91	15.73
12	0.47	1.91	6.83	15.75	0.07	0.32	2.17	-4.21	0.49	2.56	19.04	30.41

<sup>27</sup>If the COVID-19 crisis is also excluded, this has little effect on the predictability results. The peak of the COVID-19 was very short, affecting just one observation at monthly frequency.

<sup>28</sup>The choice of 60-month rolling window is a practical compromise. Ideally, we would like to have a sufficiently long estimation period to allow for accurate and stable in-sample fit of the model. However, we also want to have a long enough testing period to evaluate the forecasts.

forecast. Specifically, following Campbell and Thompson (2008), the out-of-sample  $R^2$  is computed as

$$R_{\text{OOS}}^2 = 1 - \frac{\text{MSE}_j}{\text{MSE}_0},$$

where subscripts  $j$  and  $0$  denote the candidate and benchmark models, respectively. The out-of-sample evidence is consistent with our in-sample results. The up correlation remains a strong predictor out-of-sample. Its  $R_{\text{OOS}}^2$  is high and increases with the horizon. As before, the global correlation is a useful, but weaker predictor. The down correlation cannot outperform the (naive) benchmark predictor, as evidenced by its negative out-of-sample  $R^2$ s.

Fourth, Table 10 presents the results of the OLS regression for the nine sectors. We find that, the predictive power of the up correlation continues to be strong. For all sectors and horizons, it delivers a higher  $R^2$  than the global correlation. The only exception is the Technology sector with  $h = 1$  and 3 months, where  $R^2$  for the up correlation is marginally lower. The difference in  $R^2$ s between the up and global correlations widens considerably for the longer horizons. For  $h = 9$  and 12 months,  $R^2$  for the up correlation typically exceeds 15%, often substantially. The two exceptions are Technology and Utilities sectors, for which predictability is relatively weak for all three predictors. For the down correlation, predictability is the weakest, with beta coefficients usually being insignificant.<sup>29</sup>

Overall, we conclude that the implied up correlation contains valuable information about future market returns. Future work and a longer sample will help us better understand economic mechanisms driving the up correlation predictability. Some initial intuition is provided by Figure 12, which investigates how different sectors behave when the up correlation is high or low. We group the nine sectors into three “Super Sectors”: Cyclical (MAT, FIN, and CDI), Sensitive (ENE, IND, and TEC), and Defensive (CST, UTI, and HEA). In defining Super Sectors, we follow Morningstar classification. Cyclical companies tend to have a high correlation with business cycles. Defensive companies are anti-cyclical and tend to stay stable whether the market is healthy or not. Sensitive companies fall between cyclical and defensive ones and tend to have moderate correlations with business cycles. For each Super Sector, we compute the average excess return as equal-weighted of its 3 industries. We sort all trading days into terciles based on the implied up correlation and then compute the average subsequent returns for three Super Sectors. The forecasting horizon  $h = 1, 3, 6$  and 12 months.

Figure 12 confirms nonparametrically a positive relationship between the up correlation and future returns. Across all Super Sectors, the average return monotonically increases from the bottom to the top tercile. When the up correlation is

<sup>29</sup>In unreported results, we repeated predictive regressions in (19), but now using the CBOE implied correlation index COR3M as the predictor. COR3M aims to measure the average correlation among the S&P 500 stocks for the 3-month tenor. The index has been recently redesigned by CBOE, but its backfilled values are available for our whole sample. We find that COR3M performs similarly to the global correlation  $\rho^{g,Q}$  as a predictor for the market and individual sectors and that both significantly underperform relative to the up correlation across all horizons. These results are available from the authors.

TABLE 10  
 Predictive Regressions for Individual Sectors

Table 10 reports the results of univariate predictive regressions over five horizons  $h$  (1, 3, 6, 9, and 12 months). The independent variable is the excess return for each of the nine sectors. The predictor variable is the implied global, down, or up correlation. For each sector, horizon, and predictor, the table shows the beta coefficient, its Newey–West  $t$ -statistics computed with the number of lags equal to the horizon, and the adjusted  $R^2$  (in percentage).

Sector	$h$ Months	Global			Down			Up		
		$\beta$	$t$ -Stat	$R^2$	$\beta$	$t$ -Stat	$R^2$	$\beta$	$t$ -Stat	$R^2$
MAT	1	0.12	2.02	1.77	0.08	1.03	0.56	0.09	2.36	2.70
	3	0.36	2.11	6.51	0.30	1.61	5.26	0.22	2.56	6.44
	6	0.67	2.23	11.56	0.48	1.50	6.94	0.46	2.96	14.55
	9	0.73	2.54	10.06	0.45	1.53	4.16	0.59	3.06	17.57
	12	0.72	2.95	8.91	0.35	1.45	2.08	0.65	2.98	18.10
ENE	1	0.17	2.27	2.15	0.09	1.13	0.31	0.13	2.16	3.84
	3	0.40	2.21	5.24	0.27	1.59	2.56	0.29	2.51	7.50
	6	0.61	1.91	8.12	0.41	1.45	3.95	0.47	2.53	12.46
	9	0.79	2.03	9.00	0.47	1.56	3.50	0.64	2.72	16.08
	12	0.91	2.21	10.46	0.51	1.59	3.61	0.76	2.93	18.21
FIN	1	0.09	1.24	0.21	0.06	0.57	-0.20	0.08	1.92	1.26
	3	0.33	1.75	3.15	0.29	1.34	2.85	0.23	2.26	4.00
	6	0.76	2.13	8.94	0.57	1.39	5.77	0.53	3.06	11.70
	9	0.95	2.35	9.55	0.64	1.72	4.86	0.72	2.77	14.83
	12	1.04	2.46	8.95	0.63	1.87	3.51	0.84	2.28	14.42
IND	1	0.13	2.14	2.00	0.08	1.13	0.78	0.09	2.56	2.90
	3	0.34	2.12	5.64	0.27	1.60	4.21	0.22	2.60	6.34
	6	0.63	2.31	10.62	0.43	1.37	5.47	0.46	3.47	14.70
	9	0.76	2.68	10.95	0.44	1.48	4.06	0.62	3.57	19.77
	12	0.87	3.47	11.35	0.44	1.64	3.01	0.75	3.32	20.90
TEC	1	0.10	1.70	1.69	0.10	1.44	1.94	0.06	1.87	1.51
	3	0.30	1.81	5.53	0.28	1.67	5.77	0.18	2.07	4.92
	6	0.41	1.43	5.30	0.32	1.08	3.69	0.30	2.04	7.63
	9	0.32	1.12	1.75	0.18	0.64	0.26	0.34	1.75	6.43
	12	0.32	1.15	1.18	0.11	0.44	-0.42	0.37	1.44	5.12
CST	1	0.06	1.82	1.26	0.04	1.11	0.36	0.05	1.84	2.80
	3	0.20	2.39	6.19	0.15	1.72	4.18	0.12	2.73	6.20
	6	0.43	3.30	16.79	0.27	1.61	7.84	0.29	5.14	20.40
	9	0.56	4.02	20.41	0.34	2.07	8.72	0.40	5.31	28.18
	12	0.62	4.55	20.10	0.35	2.43	7.18	0.49	5.15	31.18
UTI	1	0.07	1.20	0.49	0.06	0.93	0.35	0.05	1.22	0.90
	3	0.13	1.08	1.22	0.10	0.80	0.77	0.08	1.28	1.41
	6	0.24	1.39	3.31	0.14	0.70	1.00	0.18	2.18	5.49
	9	0.30	1.59	3.76	0.17	0.88	1.13	0.26	2.41	7.81
	12	0.34	1.85	3.76	0.18	1.10	0.94	0.31	2.28	8.46
HEA	1	0.08	1.73	1.40	0.05	1.16	0.36	0.08	1.91	4.33
	3	0.18	1.93	3.67	0.13	1.33	2.05	0.15	2.76	6.65
	6	0.34	2.39	7.20	0.17	1.02	1.78	0.32	3.96	17.31
	9	0.47	2.55	8.14	0.20	1.04	1.32	0.47	4.12	22.37
	12	0.63	2.79	10.94	0.25	1.33	1.52	0.64	3.59	28.58
CDI	1	0.12	2.09	2.06	0.08	1.06	0.66	0.10	2.90	4.32
	3	0.41	2.56	10.07	0.33	2.06	7.95	0.28	3.54	12.05
	6	0.72	2.51	14.30	0.52	1.76	8.64	0.54	3.86	21.06
	9	0.78	2.48	11.06	0.48	1.74	4.55	0.68	3.38	22.13
	12	0.83	2.70	9.40	0.40	1.67	2.11	0.80	2.94	22.17

high, all sectors perform well, with cyclical stocks outperforming defensive stocks. Conversely, when the up correlation is low, all sectors perform poorly, with defensive stocks outperforming cyclical stocks. Intuitively, when the up correlation is high, investors anticipate a broad market rally, with all sectors advancing simultaneously. As cyclical stocks have higher market betas, they tend to outperform defensive stocks during these days. Conversely, when the up correlation is low, the market is more fragile and sectors are pulling in different directions. During such times, defensive stocks tend to perform relatively better.<sup>30</sup>

<sup>30</sup>In unreported results, we estimated predictive regressions in (19), but when the dependent variable is the Super Sector's return in excess of the market. For all forecasting horizons, betas are positive for Cyclical and Sensitive and negative for Defensive sectors. However, betas are only significant for

FIGURE 12

## Subsequent Return Sorted by Implied Up Correlation

All trading days in Figure 12 are sorted into three terciles based on the implied up correlation  $\rho^{Q,u}$  and then the average subsequent returns are computed for three Super sectors (Cyclical, Sensitive, and Defensive). The forecasting horizon  $h = 1, 3, 6$  and 12 months.

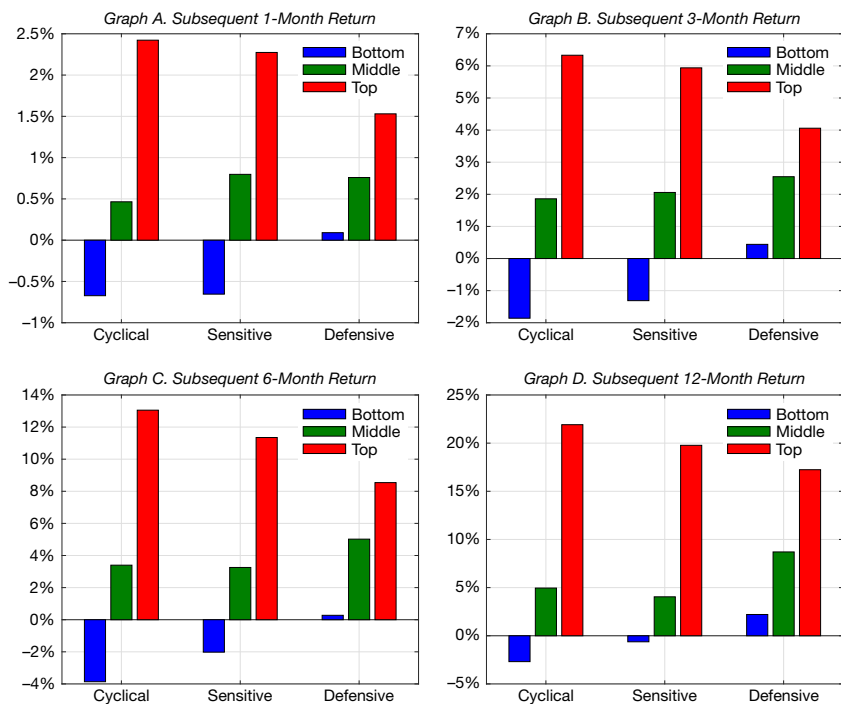


Table 11 reports the risk measures of several sector strategies, including the value-weighted portfolio of the nine sectors (SPX), the equal-weighted portfolio of the nine sectors (SPXEW), the three Super Sectors, and a simple sector rotation strategy (RS).<sup>31</sup> Our rotation strategy is rebalanced monthly based on the level of the implied up correlation. Specifically, RS invests in Defensive sectors when the up correlation is low (the bottom tercile), in all sectors when the up correlation is average (the middle tercile), and in both Cyclical and Sensitive sectors when the up correlation is high (the top tercile). Because their performances are similar across all three terciles, for simplicity, we treat Cyclical and Sensitive sectors symmetrically (and opposite to Defensive sectors).

Cyclical for horizons of 3 months or more. While the implied up correlation is a strong predictor for the overall market, its marginal predictability for excess returns is weaker. Our sample period of 14 years is relatively short, resulting in a lack of statistical power. These results are available from the authors.

<sup>31</sup>Sector rotation is a popular investment strategy that exploits perceived differences in the relative performance of sectors at different stages of the business cycle (Beber, Brandt, and Kavajecz (2011)). We note that the success of this strategy critically depends on sectors not being overly correlated. In the presence of transaction costs, it would be harder to justify switching between sectors when they all perform similarly, even with the benefit of perfect foresight.



TABLE 11  
Risk Measures for Sector Strategies

Table 11 reports mean excess return, volatility, and Sharpe ratio for various sector strategies. The strategies are the value-weighted portfolio of the nine sectors (SPX), the equal-weighted portfolio of the nine sectors (SPXEW), Cyclical, Sensitive, Defensive, and the sector rotation strategy (RS). The rotation strategy is rebalanced monthly based on the tercile of the implied up correlation.

	<u>SPX</u>	<u>SPXEW</u>	<u>Cycl.</u>	<u>Sens.</u>	<u>Defens.</u>	<u>RS</u>
Mean return (%)	9.62	9.31	8.66	9.60	9.56	12.49
Std. Dev. (%)	16.81	17.17	21.30	20.38	13.52	16.54
Sharpe ratio	0.57	0.54	0.41	0.47	0.71	0.76

Since the rotation strategy equal-weights the selected sectors, SPXEW serves as the appropriate benchmark. Table 11 shows that RS outperforms SPXEW considerably: the average excess return is more than 3% higher (12.49% vs. 9.31%), the volatility is lower, resulting in a superior Sharpe ratio (0.76 vs. 0.54). RS also considerably outperforms all Super sectors in terms of the mean return, while it is only slightly worse than the Defensive sector in terms of risk.<sup>32</sup>

## V. Conclusions

We propose a novel methodology to estimate the risk-neutral dependence among several assets that is consistent with market prices of options on these assets and on their index. Termed MFDR, it offers two critical advantages compared to the existing methods. First, the methodology is completely model-free and requires no parametric assumptions. Second and most importantly, it yields a full dependence structure, not just the average correlation coefficient. To achieve so, the methodology matches a continuum of moments on the risk-neutral distributions, as opposed to satisfying just one restriction in the existing methods.

In the empirical application, we implement MFDR to the nine economic sectors comprising the S&P 500 index. We document that the option-implied dependence for the nine sectors is highly asymmetric and time-varying. We study the CRP and find that it is negative for the down correlation and positive for the up correlation. These findings are consistent with the economic intuition that investors are mainly concerned with the loss of diversification when the market falls and that they actually prefer high correlation when the market rallies. That is, investors view the down correlation as “bad” and the up correlation as “good.” While it might be possible to rationalize the negative risk premium for the down correlation with disappointment aversion preferences, the positive risk premium for the up correlation presents a bigger challenge for the theoretical literature.

We anticipate that our model-free methodology could be used in numerous additional applications and we now briefly discuss some of them. First, since our

<sup>32</sup>It is worth noting that, in our short and somewhat special sample period, the Defensive sector stands out on the risk-adjusted basis. It had about same mean return as the broad market, but with a much lower risk.

approach can estimate the full risk-neutral covariance matrix, it could be useful for optimal portfolio construction. The existing literature has already established that the moments of the univariate implied distribution can improve portfolio choice. Superior asset allocation can be achieved by using the risk-neutral moments estimated from options rather than their realized counterparts estimated from stock returns (e.g., Kostakis, Panigirtzoglou, and Skiadopoulos (2011), DeMiguel et al. (2013)). Therefore, a promising avenue for future research is to investigate whether these results can be extended to multidimensional portfolios, where the asset allocation is not limited to the risk-free assets and one risky asset but rather could include many risky assets.

Second, Bollerslev, Patton, and Quaadvlieg (2022) recently propose a novel extension of the CAPM model by decomposing the traditional market beta into four semibetas depending on the signed covariation between the market and individual asset returns (this extends the work of Ang et al. (2006) on the downside and upside betas). They show that the decomposition into the four semibetas offers superior cross-sectional predictions compared to those obtained by the traditional betas. The semibetas are estimated from historical returns, either daily or intraday. It would be of interest to use our methodology to estimate the forward-looking option-implied semibetas and to contrast them to the semibetas estimated under  $\mathbb{P}$ .

Third, MFDR is designed to find a joint distribution, which perfectly reproduces both the individual and index options. Importantly, this methodology can also inform us when no feasible joint distribution exists. This happens when the prices of individual options are inconsistent with the index options, implying that there exists an arbitrage opportunity. Typically, such a situation arises when a specific portfolio of individual options is too cheap compared to the index option. Our methodology can be used to detect potential arbitrage opportunities and to verify whether they can survive realistic trading costs.<sup>33</sup> This idea generalizes dispersion arbitrage, which is based solely on global correlations.

Finally, our methodology could shed more light on the phenomenon studied in Kelly et al. (2016). Specifically, they construct the so-called put spread for the financial industry as a measure of tail dependence. The spread is defined as the difference in costs between OTM puts for individual banks and OTM puts for the financial sector index. Kelly et al. (2016) find that during the 2007–2009 financial crisis, the put spread for the financial industry was extraordinary large, which implies a very low implied correlation among the banks and is consistent with the perceived sector-wide government bailout guarantee. When implemented at the industry level, our approach would potentially be able to recover the full joint distribution for the banks (not just one specific partial measure of tail behavior) and to disentangle the effects due to changes in the margins (e.g., the volatility of the banks) and the dependence (e.g., interaction between banks).

---

<sup>33</sup>The problem of detecting arbitrage opportunities is considered in Hobson, Laurence, and Wang (2005) and Chen, Deelstra, Dhaene, and Vanmaele (2008). These authors derive the highest possible price for a given basket option consistent with observed prices of individual options.

## Appendix A. Implications of the Restriction in (8)

It might be instructive to consider in more detail a few specific choices for function  $g(z, s)$ . When  $g(z, s) = z$  or  $g(z, s) = z^2$ , we recover conditions on the first and second central moments discussed previously. Similarly, the choice of  $g(z, s) = z^3$  delivers a condition on the average measure of coskewness. However, let us consider some simple “cross-moments” when  $g(z, s)$  depends on both  $z$  and  $s$ . When  $g(z, s) = zs$ , we obtain the condition similar to (5):

$$\sum_{j=1}^d \omega_j \sigma_j \rho_{jS} = \sigma_S.$$

With the additional assumption of *equal correlations with the index*,  $\rho_{jS} = \rho_S$ , it leads to a new identifying equation for the implied correlation:

$$\rho_S = \frac{\sigma_S}{\sum_{j=1}^d \omega_j \sigma_j}.$$

The above is *not* the same as the equation used in the existing literature. This is because the two auxiliary assumptions of constant correlations  $\rho_{jk} = \text{const}$  (between two assets) and  $\rho_{jS} = \text{const}$  (between an asset and the index) are not equivalent when  $\sigma_j$  and  $\omega_j$  are not constant.

Finally, consider expectations conditional on  $S$ . Let  $g(z, s) = z\mathbb{I}(s \leq K)$ , for some level  $K$ . We obtain the following restriction:

$$\sum_{j=1}^d \omega_j E[X_j | S \leq K] = E[S | S \leq K].$$

That is, conditional on  $S$  being below some critical level  $K$ , the average value of the portfolio of the  $d$  assets must be equal to that of the index itself. Generally, there will be many copulas  $C$  that are consistent with (4). This situation is not uncommon given that in incomplete markets, the risk-neutral measure is not unique. However, although they could differ in “micro” details, all solutions will agree on broad, “aggregate” features. In particular, all solutions will imply the same average pairwise correlation  $\rho$  obtained by the existing approaches. More generally, they will all agree on moments  $E[g(Z, S)]$  for any function  $g(z, s)$ . In this sense, any solution will provide very valuable information.

## Appendix B. CBOE Options

We use CBOE options on the SPDR ETFs for the nine Select Sectors and the S&P 500 itself. The ETFs are managed by State Street Global Advisors. The stocks in the S&P 500 index are divided into 11 industry sectors, but Information Technology and Telecommunications are combined in a single ETF (ticker XLK), while Financial and Real Estate are combined in another ETF (ticker XLF). The ETFs pay quarterly dividends. The ETF options are physically settled and have an American-style exercise. The contract size is 100 shares of the corresponding ETF. The minimum price movement is 0.05. The strikes are multiples of \$1. Sector options all expire on the same day.

We obtain option data directly from CBOE. The data includes bid and ask quotes recorded at both 14:45 and 15:15 CDT. Due to lower liquidity, spreads widen considerably right before the market closes at 15:15. Therefore, in our analysis we use quotes at 14:45. On any given trading day, we estimate option-implied RNDs with a constant time to maturity of 91 days. The details of the procedure are summarized below.

### Data Set Construction

1. We compute midpoint prices. In the data set, we match all puts and calls by trading date  $t$ , maturity  $T$ , and strike  $K$ . For each pair  $(t, T)$ , we drop very low (high) strikes with 0 bids. We approximate the risk-free rate  $r$  over  $[t, T]$  by the rate of 3-month Treasury bills.

2. Because sector options are American-style, their prices  $P_t^A(K)$  and  $C_t^A(K)$  could be slightly higher than the prices of the corresponding European options  $P_t(K)$  and  $C_t(K)$ . The difference, however, is small for the short maturities on which we focus. This is particularly true for OTM and ATM options. To infer the prices of European options  $P_t(K)$  and  $C_t(K)$  on a given underlying  $X_t$  and maturity  $\tau = T - t$ , we proceed as follows: First, we discard all ITM options. That is, we use put prices for  $K/Z_t \leq 1.00$  and call prices for  $K/Z_t \geq 1.00$ , where  $Z_t := X_t e^{(r-\delta)\tau}$  is the forward price. Prices of OTM and ATM options are both more reliable and less affected by the early exercise feature. Second, we correct American option prices  $P_t^A(K)$  and  $C_t^A(K)$  for the value of the early exercise feature by using Barone-Adesi and Whaley (1987) approximation. Third, we compute the prices of ITM options through the put-call parity relationship:

$$C_t(K) - P_t(K) = (Z_t - K)e^{-r\tau}.$$

3. We check option prices for violations of the no-arbitrage restrictions. To preclude arbitrage opportunities, European call and put prices must be monotonic and convex functions of the strike. In particular, the call pricing function  $C_t(K)$  must satisfy

$$(a) C_t(K) \geq (F_t - K)^+ e^{-r\tau}, \quad (b) -e^{-r\tau} \leq C_t'(K) \leq 0, \quad (c) C_t''(K) \geq 0.$$

In real data, however, restrictions (a)–(c) can sometimes be violated and we enforce them by running the *Constrained Convex Regression* (CCR) introduced in Bondarenko (2000). Intuitively, CCR searches for the smallest (in the sense of least squares) perturbation of option prices that restores the no-arbitrage restrictions. The procedure is also useful for identifying possible recording errors or typos.

4. We construct prices of synthetic options with constant time to maturity  $\tau_c = T_c - t = 91$  days. Specifically, we start with two available time to maturities  $\tau_1$  and  $\tau_2$  which bracket the target time to maturity  $\tau_c$ . Cleaned European options for  $\tau_1$  and  $\tau_2$  are converted into implied volatilities and linearly interpolated with respect to the normalized moneyness  $m$ , defined in (10). The interpolated implied volatilities for  $\tau_c$  are then converted back into option prices, which are used for the RND estimation.<sup>34</sup>

5. For each trading day  $t$ , we estimate the RND corresponding to time to maturity  $\tau_c$  using the method of *Positive Convolution Approximation* (PCA) developed in Bondarenko

<sup>34</sup>For most days, constructing options with constant time to maturity does not require extrapolation, as there exist two maturities such that  $\tau_1 \leq \tau_c \leq \tau_2$ . However, in the earlier part of our sample, there are a few days for which the shortest available maturity  $\tau_1$  exceeds  $\tau_c$  and extrapolation is unavoidable. We discard days which require extensive extrapolation, i.e., when the weight for the shortest maturity is negative and less than  $-0.5$ .

(2003). The method allows one to infer the RND  $f_t(x)$  and RNCD  $F_t(x)$  through the relationships in (1) and (2). The method directly addresses the important limitations of option data that (a) options are only traded for a discrete set of strikes, as opposed to a continuum of strikes, (b) very low and very high strikes are unavailable, and (c) option prices are recorded with substantial measurement errors, which arise from nonsynchronous trading, price discreteness, and the bid–ask bounce. The PCA method is fully nonparametric, always produces arbitrage-free estimators, and controls against overfitting while allowing for small samples. We implement a version of PCA for which extreme left and right tails of the RND are extended in accordance with power laws.

6. Given the estimated RNCDs for the nine sectors and the index, we discretize each of them into  $n = 1,000$  equally probable points, collect them into  $n \times (d + 1)$  matrix  $\mathbf{M}$ , and apply the BRA method described in Section II.D. The output of the BRA is given by another  $n \times (d + 1)$  matrix that describes the joint distribution between the nine sectors compatible with all marginal distributions. Armed with the full joint distribution, we are able to compute various statistics of interest, including global, down, and up pairwise correlation for sector returns.

## Appendix C. Hybrid Model

TABLE C.1  
Five Fitted Models

Table C.1 reports statistics for the five models estimated on Sept. 8, 2008.  $SN_1$ ,  $SN_2$ , and  $SN_3$  are the skewed normal copulas fit to match the standard deviation of the index  $\sigma_S^Q$ ;  $SN^*$  is the SN copula fit to match the whole IV curve by minimizing the distance  $D$  in (18); EE is the empirical copula obtained from our algorithm. Here,  $\delta$  and  $\rho$  are the parameters of the SN models;  $\rho^g$ ,  $\rho^d$ , and  $\rho^u$  are the average global, down, and up correlations;  $\Delta\rho$  is the correlation spread.

Model	$\delta$	$\rho$	$\sigma^Q$	$\rho^g$	$\rho^d$	$\rho^u$	$\Delta\rho$	$D \times 100$
$SN_1$	-0.87	0.140	0.456	0.585	0.437	0.130	0.307	2.28
$SN_2$	0.00	0.596	0.456	0.585	0.322	0.290	0.032	7.66
$SN_3$	0.80	0.366	0.456	0.585	0.244	0.394	-0.151	13.98
$SN^*$	-0.84	0.243	0.455	0.583	0.414	0.167	0.247	0.77
EE			0.456	0.584	0.437	0.172	0.265	0.19

Figure C.1 plots the implied volatility curves  $IV^Z$  for the five models  $SN_1$ ,  $SN_2$ ,  $SN_3$ ,  $SN^*$ , and EE. For reference, also shown is the implied volatility curve for the index  $IV^S$  (the black curve). Graph A of Figure C.1 shows the three models  $SN_1$ – $SN_3$ , which are fit solely to match  $\sigma^Q$  and thus (5). They are very far from the observed  $IV^S$  (the black curve). Graph B of Figure C.1 shows models  $SN^*$  and EE, with the latter one matching  $IV^S$  almost perfectly. Note that model  $SN^*$  does not match  $\sigma^Q$  perfectly, as it is fit by minimizing a different objective function (the distance  $D$ ). Even for the optimal SN model, the distance  $D$  is about 4 times larger than for the EE model, which does not even aim at minimizing the distance  $D$  as its objective. Visually, the fit with model  $SN^*$  seems quite good on Sept. 8, 2008, but this is not always the case.

We repeat here the study performed on Sept. 8, 2008, with another date to show that the hybrid model may not always be a good fit. Indeed, Table C.2 and Figure C.2 repeat the same analysis, but now for a different day, Oct. 20, 2017. For that day, even the optimal SN model is too inflexible to approximate the implied volatility curve well. The distance  $D$  is now about 20 times as large as that for the EE model.

FIGURE C.1

Implied Volatilities for the Five Models on Sept. 8, 2008

Graph A of Figure C.1 shows the skewed normal copulas  $SN_1$ – $SN_3$ . Graph B shows the skewed normal copula  $SN^*$  and empirical copula (from MFDR)  $EE$ . The black curves in both graphs show the true implied volatilities (for the index)  $IV^S$ .

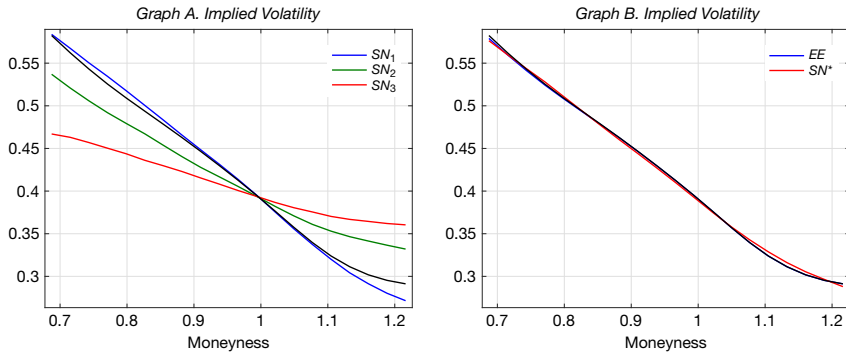


TABLE C.2  
Five Fitted Models

Table C.2 reports statistics for the five models estimated on Oct. 20, 2017.  $SN_1$ ,  $SN_2$ , and  $SN_3$  are the skewed normal copulas fit to match the standard deviation of the index  $\sigma_S^2$ ;  $SN^*$  is the SN copula fit to match the whole IV curve by minimizing the distance  $D$  in (18);  $EE$  is the empirical copula obtained from our algorithm. Here,  $\delta$  and  $\rho$  are the parameters of the SN models;  $\rho^g$ ,  $\rho^d$ , and  $\rho^u$  are the average global, down, and up correlations;  $\Delta\rho$  is the correlation spread.

Model	$\delta$	$\rho$	$\sigma^g$	$\rho^g$	$\rho^d$	$\rho^u$	$\Delta\rho$	$D \times 100$
$SN_1$	-0.87	-0.008	0.217	0.535	0.422	0.041	0.380	12.18
$SN_2$	-0.40	0.537	0.217	0.535	0.319	0.262	0.057	17.16
$SN_3$	0.80	0.368	0.217	0.535	0.246	0.402	-0.156	23.03
$SN^*$	-0.83	0.077	0.213	0.503	0.370	0.061	0.309	11.91
$EE$			0.217	0.516	0.538	0.060	0.477	0.59

FIGURE C.2

Implied Volatilities for the Five Models on Oct. 20, 2017

Graph A of Figure C.2 shows skewed normal copulas  $SN_1$ – $SN_3$ . Graph B shows skewed normal copula  $SN^*$  and empirical copula (from MFDR)  $EE$ . The black curves in both graphs show the true implied volatilities (for the index)  $IV^S$ .

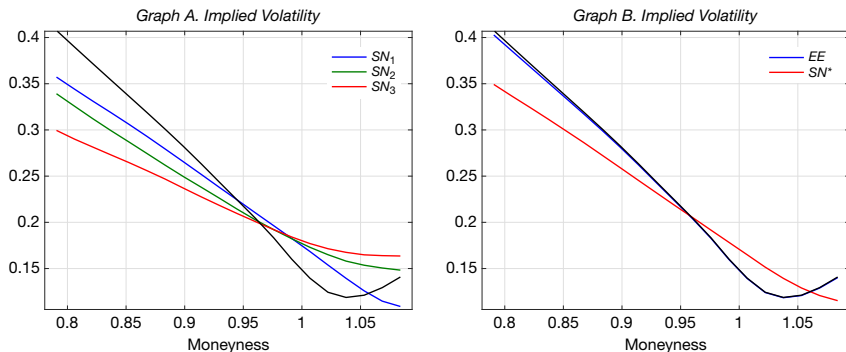


TABLE C.3  
Three Models over the Full Sample

Table C.3 reports time-series averages for three models estimated daily from Jan. 1, 2007 to Dec. 31, 2020.  $N^*$  is the normal copula with parameter  $\rho$ , and  $SN^*$  is the skewed normal copula with two parameters,  $\rho$  and  $\delta$ . They are fit to match the whole IV curve by minimizing the distance  $D$  in (18). EE is the model-free copula from our algorithm. The first 3 columns are the average global, down, and up correlations; the fourth column is the correlation spread  $\Delta\rho = \rho^d - \rho^u$ .

Model	$\rho^g$	$\rho^d$	$\rho^u$	$\Delta\rho$	$\rho$	$\delta$	$D \times 100$
$N^*$	0.724	0.538	0.490	0.047	0.748		8.96
$SN^*$	0.747	0.654	0.313	0.341	0.300	-0.905	4.01
EE	0.748	0.719	0.271	0.449			0.43

The previous illustrations focus on two specific dates. To generalize the analysis, we fit a multivariate normal copula (with one parameter  $\rho$ ) and a multivariate skewed normal copula (with two parameters  $\rho$  and  $\delta$ ) for all days in our sample. Specifically, we search for the normal copula  $N^*$  and the skewed normal copula  $SN^*$  that minimize the distance  $D$  in (18) on each day. In Table C.3, we report the time-series averages of the fitted parameters and distance  $D$ . It is clear that adding the skewness parameter greatly reduces that distance (by a factor of more than 2), but the distance for  $SN^*$  is still much larger than for the model-free EE (by a factor more than 9). Note also that the correlation coefficient in the normal copula  $\rho$  is not directly comparable to the coefficient  $\rho$  in the skewed normal copula. For  $SN^*$ , the average global correlation matches almost exactly that for EE (0.747 and 0.748, respectively), even though  $\rho$  is on average equal to only 0.300 (combined with a very negative skewness parameter  $\delta = -0.905$ ). Recall that the three models do not have to match the average global correlation, as the objective is not to match  $\sigma^Q$  but instead to minimize the distance  $D$ .<sup>35</sup>

Table C.3 shows that the normal copula is unable to reproduce the asymmetry between the down and up correlations that is exhibited by the model-free approach EE. Specifically, using the EE approach, the correlation spread  $\Delta\rho = \rho^d - \rho^u = 0.449$ . Even though a normal copula is symmetric, the down and up correlations are not identical because the margins are skewed (the heavy left tail). For model  $N^*$ , the correlation spread  $\Delta\rho = 0.047$ , or approximately 10.5% of that for EE. Intuitively, about 10.5% of the correlation spread can be attributed to nonnormality of the margins. This complements the results of Table 6 in Section IV.C. On the other hand, the skewed normal model with only one extra parameter can reproduce the down and up correlations considerably better, now accounting for approximately 76% of the correlation spread (0.341/0.449). This also means that many of our conclusions regarding the down and up CRP can be (approximately) confirmed by adopting the hybrid model  $SN^*$ , for which nonparametric margins are joined by the SN copula.

### Multivariate Skewed Normal Model

We follow the procedure of Azzalini and Valle (1996) to simulate a  $d$ -dimensional skewed normal copula. We are interested in the special case of constant pairwise correlation and constant skewness. Generally, the joint pdf of  $(X_1, X_2, \dots, X_d)$  is a skewed normal distribution with mean parameter  $\mathbf{0}$ , correlation matrix  $R$  and skewness parameter  $\lambda = (\lambda_1, \dots, \lambda_d)$  if

<sup>35</sup>Furthermore, recall that the parameter  $\rho$  of the normal copula can only be interpreted as a correlation coefficient if the margins are normal. Thus, the coefficient  $\rho$  does not match the average global correlation.

$$f(\mathbf{x}) = 2\phi_d(\mathbf{x}, \Sigma)\Phi(\boldsymbol{\alpha}'\mathbf{x}),$$

where  $\phi_d(\mathbf{x}, \Sigma)$  is the pdf of an MVN distribution with mean  $\mathbf{0}$  and covariance matrix  $\Sigma$  and

$$(20) \quad \boldsymbol{\alpha}' = \frac{\boldsymbol{\lambda}'R^{-1}\Delta^{-1}}{\sqrt{1 + \boldsymbol{\lambda}'R^{-1}\boldsymbol{\lambda}}}, \quad \Sigma = \Delta(R + \boldsymbol{\lambda}\boldsymbol{\lambda}')\Delta, \quad \Delta = \text{diag}(\sqrt{1 - \delta_1^2}, \sqrt{1 - \delta_2^2}, \dots, \sqrt{1 - \delta_d^2}),$$

$$\lambda_j = \frac{\delta_j}{\sqrt{1 - \delta_j^2}}, \text{ for some } \delta_j \in (-1, 1).$$

The simulation procedure can be summarized as follows:

- Simulate  $Z$  as a standard normal  $N(0, 1)$  and simulate  $(Y_1, \dots, Y_d)$  as an MVN vector with mean  $\mathbf{0}$  and correlation matrix  $R$ .
- Define  $X_j = \delta_j|Z| + \sqrt{1 - \delta_j^2}Y_j$ . Then each  $X_j$  is the standard skewed normal with parameter  $\lambda_j$ , and the  $d$ -dimensional vector  $(X_1, \dots, X_d)$  follows  $\text{SN}_d(R, \boldsymbol{\lambda})$ .
- Obtain the  $d$ -dimensional skewed normal copula by replacing simulated values of  $(X_1, \dots, X_d)$  with their ranks.

In our case,  $d=9$ , and we fit a skewed normal distribution with only two free parameters. One parameter,  $\rho$ , is the constant pairwise correlation in the correlation matrix  $R$ . The other parameter is a constant skewness coefficient (i.e.,  $\delta_1 = \dots = \delta_d = \delta$  where  $\delta$  is linked to  $\lambda$  by (20)).<sup>36</sup>

## Supplementary Material

To view supplementary material for this article, please visit <http://doi.org/10.1017/S0022109023000960>.

## References

- Ait-Sahalia, Y., and A. W. Lo. "Nonparametric Risk Management and Implied Risk Aversion." *Journal of Econometrics*, 94 (2000), 9–51.
- Alcock, J., and A. Hatherley. "Characterizing the Asymmetric Dependence Premium." *Review of Finance*, 21 (2017), 1701–1737.
- Alcock, J., and P. Sinagl. "International Determinants of Asymmetric Dependence in Investment Returns." *Journal of International Money and Finance*, 122 (2022), 102576.
- Amaya, D.; P. Christoffersen; K. Jacobs; and A. Vasquez. "Does Realized Skewness Predict the Cross-Section of Equity Returns?" *Journal of Financial Economics*, 118 (2015), 135–167.
- Ang, A., and J. Chen. "Asymmetric Correlations of Equity Portfolios." *Journal of Financial Economics*, 63 (2002), 443–494.
- Ang, A.; J. Chen; and Y. Xing. "Downside Risk." *Review of Financial Studies*, 19 (2006), 1191–1239.
- Azzalini, A., and A. D. Valle. "The Multivariate Skew-Normal Distribution." *Biometrika*, 83 (1996), 715–726.
- Bakshi, G.; N. Kapadia; and D. Madan. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies*, 16 (2003), 101–143.

<sup>36</sup>Schreindorfer (2020) has recently used another multivariate skewed distribution to model asymmetric tail dependence. Similarly to the skewed normal distribution, it is also constructed as a mixture, but in step b) of the above simulation an exponential distribution is used instead of the absolute value of a standard normal distribution ( $|Z|$ ) and  $\rho$  is set to 0.



- Banz, R., and M. Miller. "Prices for State-Contingent Claims: Some Estimates and Applications." *Journal of Business*, 51 (1978), 653–672.
- Barone-Adesi, G., and R. E. Whaley. "Efficient Analytic Approximation of American Option Values." *Journal of Finance*, 42 (1987), 301–320.
- Beber, A.; M. W. Brandt; and K. A. Kavajecz. "What Does Equity Sector Orderflow Tell Us About the Economy?" *Review of Financial Studies*, 24 (2011), 3688–3730.
- Bernard, C.; O. Bondarenko; and S. Vanduffel. "Rearrangement Algorithm and Maximum Entropy." *Annals of Operations Research*, 261 (2018), 107–134.
- Bernard, C., and D. McLeish. "Algorithms for Finding Copulas Minimizing Convex Functions of Sums." *Asia-Pacific Journal of Operational Research*, 33 (2016), 1650040.
- Bollerslev, T.; A. J. Patton; and R. Quaedvlieg. "Realized Semibetas: Disentangling "Good" and "Bad" Downside Risks." *Journal of Financial Economics*, 144 (2022), 227–246.
- Bollerslev, T.; G. Tauchen; and H. Zhou. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies*, 22 (2009), 4463–4492.
- Bollerslev, T., and V. Todorov. "Tails, Fears, and Risk Premia." *Journal of Finance*, 66 (2011), 2165–2211.
- Bondarenko, O. "Recovering Risk-Neutral Densities: A New Nonparametric Approach." SSRN Working Paper, <https://ssrn.com/abstract=246063> (2000).
- Bondarenko, O. "Estimation of Risk-Neutral Densities Using Positive Convolution Approximation." *Journal of Econometrics*, 116 (2003), 85–112.
- Bondarenko, O. "Market Price of Variance Risk and Performance of Hedge Funds." SSRN Working Paper, <https://ssrn.com/abstract=542182> (2004).
- Bondarenko, O. "Variance Trading and Market Price of Variance Risk." *Journal of Econometrics*, 180 (2014), 81–97.
- Breedon, D. T., and R. H. Litzenberger. "Prices of State-Contingent Claims Implicit in Option Prices." *Journal of Business*, 51 (1978), 621–651.
- Britten-Jones, M., and A. Neuberger. "Option Prices, Implied Price Processes, and Stochastic Volatility." *Journal of Finance*, 55 (2000), 839–866.
- Buraschi, A.; R. Kosowski; and F. Trojani. "When There Is No Place to Hide: Correlation Risk and the Cross-Section of Hedge Fund Returns." *Review of Financial Studies*, 27 (2013), 581–616.
- Buraschi, A.; F. Trojani; and A. Vedolin. "When Uncertainty Blows in the Orchard: Comovement and Equilibrium Volatility Risk Premia." *Journal of Finance*, 69 (2014), 101–137.
- Buss, A.; L. Schönleber; and G. Vilkov. "Option-Implied Correlations, Factor Models, and Market Risk." INSEAD Working Paper No. 2017/20/FIN, <https://ssrn.com/abstract=2906484> (2017).
- Buss, A.; L. Schönleber; and G. Vilkov. "Expected Correlation and Future Market Returns." CDI Working Paper, <https://ssrn.com/abstract=3114063> (2019a).
- Buss, A.; L. Schönleber; and G. Vilkov. "Expected Correlation and Future Market Returns: The Sum of Parts is More than the Whole." *Presentation at the 2019 Conference on Derivatives and Volatility, Financial Management Association* (2019b).
- Buss, A., and G. Vilkov. "Measuring Equity Risk with Option-Implied Correlations." *Review of Financial Studies*, 25 (2012), 3113–3140.
- Campbell, J. Y., and S. Thompson. "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21 (2008), 1509–1531.
- Campbell, R. A.; C. S. Forbes; K. G. Koedijk; and P. Kofman. "Increasing Correlations or Just Fat Tails?" *Journal of Empirical Finance*, 15 (2008), 287–309.
- Carr, P., and D. Madan. "Optimal Positioning in Derivative Securities." *Quantitative Finance*, 1 (2001), 19–37.
- Carr, P., and L. Wu. "Stock Options and Credit Default Swaps: A Joint Framework for Valuation and Estimation." *Journal of Financial Econometrics*, 8 (2009), 409–449.
- Chabi-Yo, F.; S. Ruenzi; and F. Weigert. "Crash Sensitivity and the Cross Section of Expected Stock Returns." *Journal of Financial and Quantitative Analysis*, 53 (2018), 1059–1100.
- Chang, B.-Y.; P. Christoffersen; K. Jacobs; and G. Vainberg. "Option-Implied Measures of Equity Risk." *Review of Finance*, 16 (2012), 385–428.
- Chen, X.; G. Deelstra; J. Dhaene; and M. Vanmaele. "Static Super-Replicating Strategies for a Class of Exotic Options." *Insurance: Mathematics and Economics*, 42 (2008), 1067–1085.
- Chicago Board Options Exchange. "CBOE S&P 500 Implied Correlation Index." White Paper (2022).
- DeMiguel, V.; Y. Plyakha; R. Uppal; and G. Vilkov. "Improving Portfolio Selection Using Option-Implied Volatility and Skewness." *Journal of Financial and Quantitative Analysis*, 48 (2013), 1813–1845.
- Dhaene, J.; D. Linders; W. Schoutens; and D. Vyncke. "The Herd Behavior Index: A New Measure for the Implied Degree of Co-Movement in Stock Markets." *Insurance: Mathematics and Economics*, 50 (2012), 357–370.

- Driessen, J.; P. J. Maenhout; and G. Vilkov. "The Price of Correlation Risk: Evidence from Equity Options." *Journal of Finance*, 64 (2009), 1377–1406.
- Driessen, J.; P. J. Maenhout; and G. Vilkov. "Option-Implied Correlations and the Price of Correlation Risk." Working Paper, INSEAD (2013).
- Embrechts, P.; G. Puccetti; and L. Rüschendorf. "Model Uncertainty and VaR Aggregation." *Journal of Banking & Finance*, 37 (2013), 2750–2764.
- Engle, R., and S. Figlewski. "Modeling the Dynamics of Correlations Among Implied Volatilities." *Review of Finance*, 19 (2014), 991–1018.
- Farago, A., and R. Tédongap. "Downside Risks and the Cross-Section of Asset Returns." *Journal of Financial Economics*, 129 (2018), 69–86.
- Faria, G.; R. Kosowski; and T. Wang. "The Correlation Risk Premium: International Evidence." Working Paper, Imperial College Business School (2018).
- Feunou, B.; M. R. Jahan-Parvar; and C. Okou. "Downside Variance Risk Premium." *Journal of Financial Econometrics*, 16 (2018), 341–383.
- Figlewski, S. "Risk-Neutral Densities: A Review." *Annual Review of Financial Economics*, 10 (2018), 329–359.
- Goyal, A., and I. Welch. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *Review of Financial Studies*, 21 (2008), 1455–1508.
- Gul, F. "A Theory of Disappointment Aversion." *Econometrica: Journal of the Econometric Society*, 59 (1991), 667–686.
- Harvey, C. R.; Y. Liu; and H. Zhu. "... and the Cross-Section of Expected Returns." *Review of Financial Studies*, 29 (2016), 5–68.
- Hobson, D.; P. Laurence; and T.-H. Wang. "Static-Arbitrage Upper Bounds for the Prices of Basket Options." *Quantitative Finance*, 5 (2005), 329–342.
- Hong, Y.; J. Tu; and G. Zhou. "Asymmetries in Stock Returns: Statistical Tests and Economic Evaluation." *Review of Financial Studies*, 20 (2006), 1547–1581.
- Jackwerth, J., and G. Vilkov. "Asymmetric Volatility Risk: Evidence from Option Markets." *Review of Finance*, 23 (2019), 777–799.
- Jackwerth, J. C., and M. Rubinstein. "Recovering Probability Distributions from Option Prices." *Journal of Finance*, 51 (1996), 1611–1631.
- Jaynes, E. T. *Probability Theory: The Logic of Science*. Cambridge: Cambridge University Press (2003).
- Jiang, L.; K. Wu; and G. Zhou. "Asymmetry in Stock Comovements: An Entropy Approach." *Journal of Financial and Quantitative Analysis*, 53 (2018), 1479–1507.
- Jondeau, E.; Q. Zhang; and X. Zhu. "Average Skewness Matters." *Journal of Financial Economics*, 134 (2019), 29–47.
- Kelly, B., and H. Jiang. "Tail Risk and Asset Prices." *Review of Financial Studies*, 27 (2014), 2841–2871.
- Kelly, B.; H. Lustig; and S. Van Nieuwerburgh. "Too-Systemic-to-Fail: What Option Markets Imply About Sector-Wide Government Guarantees." *American Economic Review*, 106 (2016), 1278–1319.
- Kilic, M., and I. Shaliastovich. "Good and Bad Variance Premia and Expected Returns." *Management Science*, 65 (2019), 2522–2544.
- Kostakis, A.; N. Panigirtzoglou; and G. Skiadopoulos. "Market Timing with Option-Implied Distributions: A Forward-Looking Approach." *Management Science*, 57 (2011), 1231–1249.
- Longin, F., and B. Solnik. "Extreme Correlation of International Equity Markets." *Journal of Finance*, 56 (2001), 649–676.
- Mardia, K. "Assessment of Multinormality and the Robustness of Hotelling's  $T^2$  Test." *Applied Statistics*, 24 (1975), 163–171.
- Mardia, K. V. "Measures of Multivariate Skewness and Kurtosis with Applications." *Biometrika*, 57 (1970), 519–530.
- Mardia, K. V. "Applications of Some Measures of Multivariate Skewness and Kurtosis in Testing Normality and Robustness Studies." *Sankhyā: Indian Journal of Statistics, Series B*, 36 (1974), 115–128.
- Mardia, K. V.; J. T. Kent; and J. M. Bibby. *Multivariate Analysis*. London: Academic Press (1980).
- Martin, I. "What is the Expected Return on the Market?" *Quarterly Journal of Economics*, 132 (2017), 367–433.
- Martin, I., and C. Wagner. "What is the Expected Return on a Stock?" *Journal of Finance*, 74 (2019), 1887–1929.
- Mueller, P.; A. Stathopoulos; and A. Vedolin. "International Correlation Risk." *Journal of Financial Economics*, 126 (2017), 270–299.
- Orłowski, P.; P. Schneider; and F. Trojani. "On the Nature of Jump Risk Premia." Research Paper, Swiss Finance Institute (2020).

- Patton, A. J. "On the Out-of-Sample Importance of Skewness and Asymmetric Dependence for Asset Allocation." *Journal of Financial Econometrics*, 2 (2004), 130–168.
- Pollet, J. M., and M. Wilson. "Average Correlation and Stock Market Returns." *Journal of Financial Economics*, 96 (2010), 364–380.
- Puccetti, G., and L. Rüschendorf. "Computation of Sharp Bounds on the Distribution of a Function of Dependent Risks." *Journal of Computational and Applied Mathematics*, 236 (2012), 1833–1840.
- Ross, S. "Options and Efficiency." *Quarterly Journal of Economic*, 90 (1976), 75–89.
- Routledge, B. R., and S. E. Zin. "Generalized Disappointment Aversion and Asset Prices." *Journal of Finance*, 65 (2010), 1303–1332.
- Rubinstein, M. "Implied Binomial Trees." *Journal of Finance*, 49 (1994), 771–818.
- Schneider, P., and F. Trojani. "Fear Trading." Swiss Finance Institute Research Paper 15–03 (2015).
- Schreindorfer, D. "Macroeconomic Tail Risks and Asset Prices." *Review of Financial Studies*, 33 (2020), 3541–3582.
- Stilger, P. S.; A. Kostakis; and S.-H. Poon. "What Does Risk-Neutral Skewness Tell Us About Future Stock Returns?" *Management Science*, 63 (2017), 1814–1834.
- Stutzer, M. "A Simple Nonparametric Approach to Derivative Security Valuation." *Journal of Finance*, 51 (1996), 1633–1652.