

Worst-Case Expected Shortfall with Univariate and Bivariate Marginals

Anulekha Dhara*

Bikramjit Das[†]

Karthik Natarajan^{‡§}

Abstract

Worst-case bounds on the expected shortfall risk given only limited information on the distribution of the random variables has been studied extensively in the literature. In this paper, we develop a new worst-case bound on the expected shortfall when the univariate marginals are known exactly and additional expert information is available in terms of bivariate marginals. Such expert information allows for one to choose from among the many possible parametric families of bivariate copulas. By considering a neighborhood of distance ρ around the bivariate marginals with the Kullback-Leibler divergence measure, we model the trade-off between conservatism in the worst-case risk measure and confidence in the expert information. Our bound is developed when the only information available on the bivariate marginals forms a tree structure in which case it is efficiently computable using convex optimization. For consistent marginals,

*Engineering Systems and Design, Singapore University of Technology and Design, 8 Somapah Road, Singapore 487372. Email: anulekha@sutd.edu.sg

[†]Engineering Systems and Design, Singapore University of Technology and Design, 8 Somapah Road, Singapore 487372. Email: bikram@sutd.edu.sg

[‡]Engineering Systems and Design, Singapore University of Technology and Design, 8 Somapah Road, Singapore 487372. Email: karthik_natarajan@sutd.edu.sg

[§]This research was partly funded by the SUTD-MIT International Design Center grant number IDG31300105 on ‘Optimization for Complex Discrete Choice’ and the MOE Tier 2 grant number MOE2013-T2-2-168 on ‘Distributional Robust Optimization for Consumer Choice in Transportation Systems’.

as ρ approaches ∞ , the bound reduces to the comonotonic upper bound and as ρ approaches 0, the bound reduces to the worst-case bound with bivariate known exactly. We also discuss extensions to inconsistent marginals and instances where the expert information which might be captured using other parameters such as correlations.

Keyword: expected shortfall (CVaR), distributionally robust bound, marginals, KL-divergence.

1 Introduction

In recent years, there has been a growing interest in estimating worst-case bounds on the joint risk measure of a portfolio when the probability distribution of the underlying risks is uncertain. Such bounds are useful in evaluating the maximum amount of risk that a portfolio is exposed to given model uncertainty. Of particular relevance to this paper is the expected shortfall risk measure (also referred to as the Conditional Value-at-Risk (CVaR) measure) that has been proposed in the Basel Accord III as an alternative to the Value-at-Risk measure for the banking industry (see Embrechts et. al. 2014). For a continuous random variable \tilde{c} with a finite expected absolute value $\mathbb{E}_\theta(|\tilde{c}|) < \infty$ where θ is the probability measure and $F(\cdot)$ is the cumulative distribution function, the expected shortfall at confidence level $\alpha \in (0, 1)$ is defined as:

$$\text{ES}_\alpha^\theta(\tilde{c}) = \frac{1}{1-\alpha} \int_\alpha^1 \text{VaR}_\gamma^\theta(\tilde{c}) \, d\gamma,$$

where the Value-at-Risk at the probability level $\gamma \in (0, 1)$ is defined as:

$$\text{VaR}_\gamma^\theta(\tilde{c}) = \inf \{c \in \mathbb{R} \mid F(c) \geq \gamma\}.$$

An alternative representation of the expected shortfall that was popularized by Rockafellar and Uryasev (2002) which is valid for any random variable \tilde{c} with a finite expected absolute value is

given by the optimal objective value to the following convex minimization problem:

$$\text{ES}_\alpha^\theta(\tilde{c}) = \inf_{\beta \in \mathbb{R}} \left(\beta + \frac{1}{1-\alpha} \mathbb{E}_\theta [\tilde{c} - \beta]^+ \right).$$

Consider the portfolio optimization problem:

$$\min_{\mathbf{x} \in \mathcal{X}} \text{ES}_\alpha^\theta(\tilde{\mathbf{c}}^T \mathbf{x}) \tag{1.1}$$

where \mathbf{x} is a n -dimensional decision vector (portfolio allocations) that is chosen in the convex set \mathcal{X} and $\tilde{\mathbf{c}}$ is a n -dimensional random vector (loss of assets) with a joint distribution $\boldsymbol{\theta}$. The loss of the portfolio is given by the random variable $\tilde{\mathbf{c}}^T \mathbf{x} = \sum_i \tilde{c}_i x_i$ where the expected shortfall captures the risk of the portfolio. The expected shortfall risk measure has several attractive mathematical properties. Firstly, the portfolio optimization problem with an expected shortfall risk measure is a convex optimization problem (unlike the VaR measure) and formulated as follows:

$$\min_{\mathbf{x} \in \mathcal{X}, \beta \in \mathbb{R}} \left(\beta + \frac{1}{1-\alpha} \mathbb{E}_\theta [\tilde{\mathbf{c}}^T \mathbf{x} - \beta]^+ \right). \tag{1.2}$$

Moreover, expected shortfall is a coherent risk measure (see Artzner et. al. 1999) and encourages risk diversification. However, there remains challenges in using this risk measure in portfolio optimization. Lim et. al. (2011) have shown that the estimation error with the expected shortfall risk measure tends to be magnified since the model is very sensitive to the tail of the return distribution. Hanasusanto et. al. (2016) have shown that computing the expected value of the non-negative part of a linear combination of random variables of the form

$$\mathbb{E}_\theta [\tilde{\mathbf{c}}^T \mathbf{x} - \beta]^+,$$

for fixed values of \mathbf{x} and β is #P-hard, even for uniformly distributed independent random variables.

One approach to tackle these difficulties is to allow for the joint probability distribution $\boldsymbol{\theta}$ to lie in a possible set of distributions Θ . Given this set of distributions, the distributionally robust

portfolio optimization problem is formulated as

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\theta \in \Theta} \text{ES}_\alpha^\theta(\tilde{\mathbf{c}}^T \mathbf{x}), \quad (1.3)$$

where the portfolio allocation vector is chosen to minimize the worst-case expected shortfall risk over all possible distributions in the set Θ . Under specific assumptions on the set Θ , formulation (1.3) has been shown to be efficiently solvable. For discrete distributions, examples of the set Θ for which the worst-case expected shortfall and the corresponding distributionally robust optimization problem are efficiently solvable are:

- (a) Sets of distributions with fixed univariate marginals (see Rüschendorf 1983),
- (b) Sets of distributions with fixed nonoverlapping multivariate marginals (see Doan and Natarajan 2012, Embrechts and Pucetti 2006),
- (c) Sets of distributions with fixed overlapping multivariate marginals where the structure of the overlapping marginals satisfies a regularity property (see Doan, Li and Natarajan 2015),
- (d) Sets of distributions where the probabilities of the scenarios are assumed to lie in a box uncertainty set or an ellipsoidal uncertainty set (see Zhu and Fukushima 2009),
- (e) Sets of distributions where the probabilities of the scenarios are assumed to lie in a uncertainty set defined using the ϕ -divergence measure around a nominal probability distribution (see Ben-Tal et. al. 2013). This includes Kullback-Leibler divergence, Burg entropy and χ^2 -distance as special cases.

Note that in cases (b), (c), (d) and (e), the representation of the set Θ might itself be exponential in the number of random variables. For example, if each of the n random variables takes m possible values, the joint distribution can take up to m^n possible values. Hence, efficient solvability in cases (b)-(e) typically refers to an algorithm that solves the problem in polynomial time in the number of support points of the joint discrete distribution.

1.1 Motivation

In the literature, several approaches have been proposed to model distributions where information on the marginals is available and only limited dependence information is known. One popular approach is to find the maximum entropy distribution (see Jaynes 1957) where the joint distribution is chosen to maximize the Shannon entropy measure (see Shannon 1948) satisfying the given information. For example, given univariate marginals, the maximum entropy distribution is the independent distribution. Similarly, given bivariate information which forms a tree structure, the maximum entropy distribution is a conditionally independent distribution on a tree, referred to as a Markov (Chow-Liu 1968) tree distribution. An alternate approach is to use the theory of copulas in which the joint distribution is modeled in terms of the univariate marginals and a copula (see Sklar 1959). Specifying and estimating the copula in high dimensions however still remains challenging. A closely related model for distributions in high dimensions that builds on bivariate copulas is a vine copula (see Bedford and Cooke 2002). Unlike Chow-Liu tree distributions on n random variables where only dependence on $n - 1$ edges can be specified, in vine copulas it is possible to specify additional dependence information among the variables using bivariate copula. The original motivation for using vine copula stemmed from the need to incorporate expert opinions on dependencies into the specification of the distribution. In our work, instead of describing the joint distribution that is consistent with univariate and bivariate information by specifying a copula structure, we focus on finding the worst-case joint distribution for the expected shortfall risk measure.

In this paper, we assume that the univariate marginal distributions are precisely known but allow for the information on the bivariate distributions to be generated in a variety of ways including historical data which might involve noisy, imprecise or incomplete data observations and expert

opinions. Hence it is quite possible that there exist no distribution which is consistent with the given univariate and bivariate marginal information. For example in portfolio optimization, correlation between the stocks is often used to model dependencies. These correlations may be measured between pairs of stocks over different periods of time. Furthermore, experts often provide estimates on correlations between pairs of stocks based on their forecasts. Combining different expert views can lead to a negative semidefinite or indefinite correlation matrix. This has led to an interest in the “closest correlation matrix problem” where the objective is to find a positive semidefinite correlation matrix that is closest to the given matrix. While much work has been done towards finding and constructing the closest correlation matrix (see Higham 2002, Qi and Sun 2010), far fewer work involve studying its effect on portfolio optimization.

Our approach is closely related to four streams of work. Glasserman and Yang (2016) recently proposed an optimization formulation to compute the worst-case credit valuation adjustment risk by combining models of market and credit risk. In their approach, the marginal distributions are held fixed while the worst-case risk is computed by penalizing for deviation with respect to a bivariate reference model of the dependence between the market and credit risk models. While their optimization formulations are developed for $n = 2$, we generalize this to arbitrary n while preserving polynomial time solvability. The second stream of work is the model proposed by Ben-Tal et. al (2013) who studied robust optimization problems where the set of joint distributions lies in a neighborhood of a reference probability distribution with distance that is defined using the notion of ϕ -divergence. Their convex optimization formulation is polynomial in the number of support points of the joint distribution but can be exponential in the number of random variables. On the other hand in our work by considering neighborhoods of lower dimensional probability distributions, specifically bivariate marginals and assuming a tree structure, we develop a convex reformulation that grows polynomially in the number of random variables. The third stream of results is by

Roughgarden and Kearns (2013) who discuss the hardness of verifying if a joint distribution exists with the given univariate and bivariate marginals and the corresponding problem of finding the closest consistent marginals to inconsistent marginals. While their focus is on showing the NP-hardness of these problems on a complete graph, we show that these problems are easy to solve on a tree and furthermore it is possible to evaluate robust bounds. Finally, our results build on the earlier work of Doan, Li and Natarajan (2015) who develop bounds on expected shortfall using univariate and bivariate marginal information. In contrast to their work, we allow for the marginals to be inconsistent and even when consistent, we provide flexibility in modeling the confidence in the expert opinions by considering a neighborhood around the marginals. Thus, we explicitly model the trade-off between conservatism in the worst-case risk measure and confidence in the expert information.

The outline of the paper is as follows. In Section 2, we provide a formal description of the set of distributions Θ with given univariate marginals μ_i and expert information in terms of bivariate marginals μ_{ij} . We study the problem of finding the closest consistent marginals by finding the smallest ρ -neighborhood using the Kullback-Leibler divergence around the given bivariate marginals such that Θ is nonempty. We also provide a polynomial-sized convex reformulation for this problem by assuming a tree structure. We also discuss several extensions in this section. In Section 3, we define the uncertainty set of distributions using a ρ -neighborhood of the bivariate distributions such that Θ is nonempty. We provide a convex optimization problem to find the worst-case upper bound on the expected shortfall. We show that the bound is efficiently computable when the bivariate marginal information forms a tree structure. In Section 4, we present numerical experiments. All proofs are provided in the Appendix.

2 Consistent Marginals in Θ

In this section, we provide a formal description of the set of distributions Θ . We restrict our attention to univariate and bivariate marginal information. Denote the random vector by $\tilde{\mathbf{c}} = (\tilde{c}_1, \tilde{c}_2, \dots, \tilde{c}_n)$ and define the index set $\mathcal{N} = \{1, 2, \dots, n\}$. Let \mathcal{C}_i denote the set of values taken by \tilde{c}_i for $i \in \mathcal{N}$. For $i \in \mathcal{N}$, let $\boldsymbol{\mu}_i$ denote the given univariate distribution where $\mu_i(c_i) = \mathbb{P}_{\boldsymbol{\theta}}(\tilde{c}_i = c_i)$ for $c_i \in \mathcal{C}_i$. Let $\mathcal{N}' \subseteq \mathcal{N} \times \mathcal{N}$ denote the index set for the bivariate marginal information. Then for $(i, j) \in \mathcal{N}'$ with $i < j$, let $\boldsymbol{\mu}_{ij}$ denote the bivariate marginal where $\mu_{ij}(\mathbf{c}_{ij}) = \mathbb{P}_{\boldsymbol{\theta}}(\tilde{\mathbf{c}}_{ij} = \mathbf{c}_{ij}) = \mathbb{P}_{\boldsymbol{\theta}}(\tilde{c}_i = c_i, \tilde{c}_j = c_j)$ with $\tilde{\mathbf{c}}_{ij} = (\tilde{c}_i, \tilde{c}_j)$ and $\mathbf{c}_{ij} = (c_i, c_j)$ for $c_i \in \mathcal{C}_i$ and $c_j \in \mathcal{C}_j$. We assume that the bivariate marginal information is obtained from expert information. Such bivariate marginals might be chosen using one of many possible parametric families of bivariate copula. We also discuss extensions when such expert information on bivariates might be captured using simpler parameters such as correlations. The associated graph with this set of marginals consist of the vertex set \mathcal{N} where each vertex i corresponds to the random variable \tilde{c}_i and the arc set \mathcal{N}' where each undirected edge (i, j) between vertex i and j corresponds to a pair of random variables \tilde{c}_i and \tilde{c}_j for which the bivariate marginal is provided. Let $(\mathcal{N}, \mathcal{N}')$ denote the corresponding graph.

Let Θ denote the Fréchet class of joint distributions with the given univariate and bivariate marginals defined as

$$\Theta = \{\boldsymbol{\theta} : \text{proj}_i(\boldsymbol{\theta}) = \boldsymbol{\mu}_i, \forall i \in \mathcal{N}, \text{proj}_{ij}(\boldsymbol{\theta}) = \boldsymbol{\mu}_{ij}, \forall (i, j) \in \mathcal{N}'\},$$

where $\text{proj}_i(\cdot)$ and $\text{proj}_{ij}(\cdot)$ denote the projection of the joint distribution to the i -th and (i, j) -th variables. An important question is to identify conditions on the univariate and bivariate marginals which would guarantee that the set Θ is nonempty. A necessary set of conditions to guarantee the set is nonempty is the consistency of the univariate and bivariate marginals which is given by the

equality conditions:

$$\begin{aligned} \sum_{c_j \in \mathcal{C}_j} \mu_{ij}(\mathbf{c}_{ij}) &= \mu_i(c_i), \quad \forall c_i \in \mathcal{C}_i, \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \\ \sum_{c_i \in \mathcal{C}_i} \mu_{ij}(\mathbf{c}_{ij}) &= \mu_j(c_j), \quad \forall c_j \in \mathcal{C}_j, \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}'. \end{aligned} \tag{2.1}$$

However, the consistency of univariate and bivariate marginals is not a sufficient condition to guarantee the nonemptiness of Θ (see Vorob'ev 1962). For example, consider $\mathcal{N} = \{1, 2, 3\}$ with binary random variables. The univariate marginals are given as $\mu_1(0) = \mu_1(1) = 1/2$, $\mu_2(0) = \mu_2(1) = 1/2$, $\mu_3(0) = \mu_3(1) = 1/2$ and the bivariate marginals are given as $\mu_{12}(0, 1) = \mu_{12}(1, 0) = 1/2$, $\mu_{23}(0, 1) = \mu_{23}(1, 0) = 1/2$, $\mu_{13}(0, 1) = \mu_{13}(1, 0) = 1/2$. Though the univariate and bivariate marginals are consistent in this example, the set Θ is empty. The nonemptiness of the Fréchet class of distributions Θ is ensured under the consistency of multivariate marginals when the graph satisfies a graph theoretic condition known as the “running intersection property”. One such graph structure is a tree. The consistency of the univariate and bivariate marginals in this case will ensure that Θ is nonempty. The tree structure plays an important role in the present work as it simplifies the formulations as discussed next.

2.1 Closest Consistent Marginals

In this section, we consider the problem of finding the closest consistent marginals when the original set of distribution Θ is empty. To do so, we find the minimal perturbation ρ for all the given bivariate marginals such that the new bivariate marginals $\{\boldsymbol{\theta}_{ij}\}_{(i,j) \in \mathcal{N}'}$ lie within a ρ -neighborhood of $\{\boldsymbol{\mu}_{ij}\}_{(i,j) \in \mathcal{N}'}$ and there exists a joint distribution consistent with the given univariate marginals $\{\boldsymbol{\mu}_i\}_{i \in \mathcal{N}}$ and the new bivariate marginals $\{\boldsymbol{\theta}_{ij}\}_{(i,j) \in \mathcal{N}'}$. To define the ρ -neighborhood, we make use of the notion of Kullback-Leibler (KL)-divergence. For any two non-negative vectors $\mathbf{p} = (p_1, \dots, p_m)^T$ and $\mathbf{q} = (q_1, \dots, q_m)^T$, the KL-divergence is defined by:

$$I_{KL}(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^m p_i \log \left(\frac{p_i}{q_i} \right).$$

The KL-divergence is an example of ϕ -divergence (see Pardo 2006) which has recently gained popularity in the area of distributionally robust optimization (see Ben-Tal et al. 2013) where the set of probability distributions is defined to be the ρ -neighborhood of a reference probability distribution.

For verifying the nonemptiness of Θ , we need to check for the existence of a joint distribution θ with the given univariate marginals $\{\mu_i\}_{i \in \mathcal{N}}$ and bivariate marginals $\{\mu_{ij}\}_{(i,j) \in \mathcal{N}'}$. If such a θ exists, the minimum deviation is given by $\rho = 0$ else $\rho > 0$. We next formulate an optimization problem to find the smallest neighborhood of the given bivariate marginals which guarantees the nonemptiness of Θ . Let $\mathcal{C} = \mathcal{C}_1 \times \mathcal{C}_2 \times \dots \times \mathcal{C}_n$ be the set of all possible realizations of $\tilde{\mathbf{c}}$ and $\theta(\mathbf{c}) = \mathbb{P}(\tilde{\mathbf{c}} = \mathbf{c})$ for $\mathbf{c} \in \mathcal{C}$. The problem of finding the closest consistent marginals using the KL-divergence measure is formulated as the solution to a convex optimization problem:

$$\begin{aligned} & \min_{\rho, \theta, \theta_{ij}} \quad \rho \\ \text{subject to} & \quad \sum_{\mathbf{c} \in \mathcal{C}|c_i} \theta(\mathbf{c}) = \mu_i(c_i), \quad \forall i \in \mathcal{N}, \forall c_i \in \mathcal{C}_i, \\ & \quad \sum_{\mathbf{c} \in \mathcal{C}|c_{ij}} \theta(\mathbf{c}) = \theta_{ij}(c_{ij}), \quad \forall (i, j) \in \mathcal{N}', \forall c_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\ & \quad \sum_{\mathbf{c} \in \mathcal{C}} \theta(\mathbf{c}) = 1, \\ & \quad \theta(\mathbf{c}) \geq 0, \quad \forall \mathbf{c} \in \mathcal{C}, \\ & \quad I_{KL}(\theta_{ij}, \mu_{ij}) \leq \rho, \quad \forall (i, j) \in \mathcal{N}'. \end{aligned} \tag{2.2}$$

In formulation (2.2), the objective is to minimize the distance of the neighborhood given by ρ . The first four constraints ensure that there exists a joint distribution with univariate μ_i and bivariate θ_{ij} while the last constraint ensures that the bivariate θ_{ij} lies in a ρ -neighborhood of μ_{ij} . Note that for consistent marginals μ_i and μ_{ij} such that a joint distribution θ exists, the optimal solution of (2.2) is $\rho = 0$.

The number of variables in the convex optimization problem in (2.2) is however exponential in the number of variables n . We next identify an instance where the closest consistency problem can be solved as a polynomial sized convex optimization problem. The result is provided next and the proof is provided in the Appendix.

Theorem 2.1 *Consider the given univariate and bivariate marginals $\{\mu_i\}_{i \in \mathcal{N}}$ and $\{\mu_{ij}\}_{(i,j) \in \mathcal{N}'}$ where the graph associated with $(\mathcal{N}, \mathcal{N}')$ is a tree. Then the solution to the following convex optimization problem finds the minimal ρ -neighborhood for which there exists a joint distribution:*

$$\begin{aligned}
& \min_{\rho, \theta_{ij}} \rho \\
\text{subject to} & \sum_{c_j \in \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = \mu_i(c_i), \quad \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_i \in \mathcal{C}_i, \\
& \sum_{c_i \in \mathcal{C}_i} \theta_{ij}(\mathbf{c}_{ij}) = \mu_j(c_j), \quad \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_j \in \mathcal{C}_j, \\
& \sum_{\mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = 1, \quad \forall (i, j) \in \mathcal{N}', \\
& \theta_{ij}(\mathbf{c}_{ij}) \geq 0, \quad \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\
& I_{KL}(\theta_{ij}, \mu_{ij}) \leq \rho, \quad \forall (i, j) \in \mathcal{N}'.
\end{aligned} \tag{2.3}$$

In formulation (2.2), the number of decision variables is $O(m^n)$ where m is the number of possible values that each of the n random variables takes. On the other hand in formulation (2.3), the number of decision variables is only $O(nm^2)$.

2.2 Extensions

In this section, we relate Theorem 2.1 to three existing models and discuss possible extensions.

Roughgarden and Kearns (2013) considered the problem of finding a minimal ρ -neighborhood over which a consistent joint distribution θ exists which is obtained by perturbing both the univariate marginals $\{\mu_i\}_{i \in \mathcal{N}}$ and the bivariate marginals $\{\mu_{ij}\}_{(i,j) \in \mathcal{N}'}$. In their work, the authors

formulated a linear program by minimizing the L_1 -norm distance instead of the KL-divergence measure that we use in (2.2). However since they consider a general graph structure for the bivariate, their problem is NP-hard to solve. We can build on their approach in cases when the univariate marginal information is not reliable by also perturbing the univariate marginals using the KL-divergence measure. The convex optimization problem (2.3) for the tree structure in this case can be extended to the following formulation:

$$\begin{aligned}
& \min_{\rho, \boldsymbol{\theta}_i, \boldsymbol{\theta}_{ij}} \quad \rho \\
\text{subject to} \quad & \sum_{c_j \in \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = \theta_i(c_i), \quad \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_i \in \mathcal{C}_i, \\
& \sum_{c_i \in \mathcal{C}_i} \theta_{ij}(\mathbf{c}_{ij}) = \theta_j(c_j), \quad \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_j \in \mathcal{C}_j, \\
& \sum_{\mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = 1, \quad \forall (i, j) \in \mathcal{N}', \\
& \theta_{ij}(\mathbf{c}_{ij}) \geq 0, \quad \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\
& I_{KL}(\boldsymbol{\theta}_i, \boldsymbol{\mu}_i) \leq \rho, \quad \forall i \in \mathcal{N}, \\
& I_{KL}(\boldsymbol{\theta}_{ij}, \boldsymbol{\mu}_{ij}) \leq \rho, \quad \forall (i, j) \in \mathcal{N}'.
\end{aligned} \tag{2.4}$$

Secondly, the optimization problems (2.2) and (2.3) is closely related to the formulation of the iterative proportional fitting (IPF) procedure (see Deming and Stephan 1940, Fienberg 1970, Glasserman and Yang 2016, Ireland and Kullback 1968) which is also known as biproportional fitting or raking. In the IPF procedure, the bivariate or higher-dimension multivariate marginal information is adjusted keeping the given univariate or lower-dimension marginals respectively fixed, thereby maximizing the bivariate (or multivariate as the case may be) entropy or relative entropy which is equivalent to minimizing the KL-divergence. The optimization problem associated with

the IPF method for given univariate and bivariate marginals can as follows:

$$\begin{aligned}
& \max_{\boldsymbol{\theta}_{ij}} && - \sum_{(i,j) \in \mathcal{N}'} \sum_{\mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) \log \left(\frac{\theta_{ij}(\mathbf{c}_{ij})}{\mu_{ij}(\mathbf{c}_{ij})} \right) \\
\text{subject to} &&& \sum_{c_j \in \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = \mu_i(c_i), && \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_i \in \mathcal{C}_i, \\
&&& \sum_{c_i \in \mathcal{C}_i} \theta_{ij}(\mathbf{c}_{ij}) = \mu_j(c_j), && \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_j \in \mathcal{C}_j, \\
&&& \theta_{ij}(\mathbf{c}_{ij}) \geq 0, && \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j.
\end{aligned} \tag{2.5}$$

This problem can be reformulated in terms of the ρ -neighborhood formulation as follows:

$$\begin{aligned}
& \min_{\rho, \boldsymbol{\theta}_{ij}} && \sum_{(i,j) \in \mathcal{N}'} \rho_{ij} \\
\text{subject to} &&& \sum_{c_j \in \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = \mu_i(c_i), \quad \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_i \in \mathcal{C}_i, \\
&&& \sum_{c_i \in \mathcal{C}_i} \theta_{ij}(\mathbf{c}_{ij}) = \mu_j(c_j), \quad \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_j \in \mathcal{C}_j, \\
&&& \theta_{ij}(\mathbf{c}_{ij}) \geq 0, \quad \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\
&&& I_{KL}(\boldsymbol{\theta}_{ij}, \boldsymbol{\mu}_{ij}) \leq \rho_{ij}, \quad \forall (i, j) \in \mathcal{N}'.
\end{aligned} \tag{2.6}$$

Observe that the above formulation is similar to that of (2.3) except that here the objective is defined using the L_1 -norm while in (2.3) the objective is defined using the L_∞ -norm.

Finally, in our model, we focus on the expert information given in terms of bivariate marginals. A related problem that has been studied is to generate distributions with given univariate marginals together with expert information specified in terms of a covariance matrix. One such method is the NORTA (NORmal To Anything) method (see Cario and Nelson 1997). While popular, it is known that the NORTA method might not be able to generate joint distributions in all instances where there exists sets of univariate marginals with a feasible covariance matrix (see Ghosh and Henderson (2002) for counterexamples). Ghosh and Henderson (2002) proposed a general linear optimization formulation to construct a joint distribution with given univariate marginals and a covariance matrix if one exists or else show that the problem is infeasible. However the size of their

linear program grows exponentially in the number of random variables and is hence easy to solve only in low dimensions. Using our previous result, we can extend their result as follows. Assume that, for each $(i, j) \in \mathcal{N}'$, we are given an estimate of the covariance between i and j represented by Σ_{ij} . Define the mean of the random variables as $E(\tilde{c}_i) = \sum_{c_i} c_i \mu_i(c_i)$. We can then generalize the formulation in (2.3) by using covariance information as follows:

$$\begin{aligned}
& \min_{\rho, \theta_{ij}} \quad \rho \\
\text{subject to} \quad & \sum_{c_j \in \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = \mu_i(c_i), & \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_i \in \mathcal{C}_i, \\
& \sum_{c_i \in \mathcal{C}_i} \theta_{ij}(\mathbf{c}_{ij}) = \mu_j(c_j), & \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_j \in \mathcal{C}_j, \\
& \sum_{\mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = 1, & \forall (i, j) \in \mathcal{N}', \\
& \theta_{ij}(\mathbf{c}_{ij}) \geq 0, & \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\
& \left(\sum_{\mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j} c_i c_j \theta_{ij}(\mathbf{c}_{ij}) - E(\tilde{c}_i)E(\tilde{c}_j) \right) - \Sigma_{ij} \leq \rho, & \forall (i, j) \in \mathcal{N}', \\
& - \left(\sum_{\mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j} c_i c_j \theta_{ij}(\mathbf{c}_{ij}) - E(\tilde{c}_i)E(\tilde{c}_j) \right) + \Sigma_{ij} \leq \rho, & \forall (i, j) \in \mathcal{N}'.
\end{aligned} \tag{2.7}$$

Note that at optimality $\rho = \max_{(i,j) \in \mathcal{N}'} |E(\tilde{c}_i \tilde{c}_j) - E(\tilde{c}_i)E(\tilde{c}_j) - \Sigma_{ij}|$. If the optimal objective value in (2.7) is $\rho = 0$, then by using a proof construction as in Theorem 2.1, there exists a joint distribution with the given univariate marginals and covariance matrix, else there exists no such distribution. Note that unlike the original linear program in Ghosh and Henderson (2002), the number of variables and constraints in the linear program (2.7) is polynomial in the number of random variables. This is because we make the assumption of a tree structure on the known covariance information.

3 Worst-Case Expected Shortfall

The distributionally robust CVaR (or expected shortfall) problem for portfolio optimization in (1.3) is typically formulated as:

$$\min_{\mathbf{x} \in \mathcal{X}, \beta \in \mathbb{R}} \left(\beta + \frac{1}{1 - \alpha} \max_{\boldsymbol{\theta} \in \Theta} \mathbb{E}_{\boldsymbol{\theta}}[\tilde{\mathbf{c}}^T \mathbf{x} - \beta]^+ \right)$$

In this section, we focus on solving the inner maximization problem for fixed \mathbf{x} and β , and finding the worst-case distribution corresponding to it. We consider a more general expected value of functions of the form $\mathbb{E}_{\boldsymbol{\theta}} [\max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k(\mathbf{x}) + b_k(\mathbf{x}))]$ where $\mathcal{K} = \{1, 2, \dots, K\}$ and $\mathbf{a}_k(\mathbf{x})$, $b_k(\mathbf{x})$ are assumed to be affine functions of the decision vector $\mathbf{x} \in \mathcal{X}$. Observe that for $\mathcal{K} = \{1, 2\}$, defining $\mathbf{a}_1(\mathbf{x}) = \mathbf{x}$, $b_1(\mathbf{x}) = -\beta$ and $\mathbf{a}_2(\mathbf{x}) = \mathbf{0}$, $b_2(\mathbf{x}) = 0$, the problem reduces to the inner maximization problem for the robust expected shortfall problem. For given univariate $\boldsymbol{\mu}_i$ and bivariate $\boldsymbol{\mu}_{ij}$, the set of distributions Θ_{ρ} is defined as:

$$\begin{aligned} \Theta_{\rho} = \left\{ \boldsymbol{\theta} : \right. & \text{proj}_i(\boldsymbol{\theta}) = \boldsymbol{\mu}_i, \forall i \in \mathcal{N}, \quad \text{proj}_{ij}(\boldsymbol{\theta}) = \boldsymbol{\mu}_{ij}, \forall (i, j) \in \mathcal{N}', \\ & \sum_{\mathbf{c} \in \mathcal{C}|c_i} \boldsymbol{\theta}(\mathbf{c}) = \boldsymbol{\mu}_i(c_i), \forall i \in \mathcal{N}, \forall c_i \in \mathcal{C}_i, \\ & \sum_{\mathbf{c} \in \mathcal{C}|c_{ij}} \boldsymbol{\theta}(\mathbf{c}) = \boldsymbol{\mu}_{ij}(c_{ij}), \forall (i, j) \in \mathcal{N}', \forall c_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\ & \sum_{\mathbf{c} \in \mathcal{C}} \boldsymbol{\theta}(\mathbf{c}) = 1, \quad \boldsymbol{\theta}(\mathbf{c}) \geq 0, \forall \mathbf{c} \in \mathcal{C}, \\ & \left. I_{KL}(\boldsymbol{\theta}_{ij}, \boldsymbol{\mu}_{ij}) \leq \rho, \forall (i, j) \in \mathcal{N}' \right\}, \end{aligned} \tag{3.1}$$

for a given $\rho \geq 0$. When no joint distribution exists corresponding to the given marginal information, ρ must be at least the minimal perturbation for which the set is nonempty obtained by solving (2.2). For the case where the marginals have a consistent joint distribution, the set of distributions is nonempty for any non-negative ρ . The choice of ρ in defining the set of distributions thus captures the confidence in the expert information. We next provide an equivalent formulation to compute the tight upper bound. The result is an extension of the linear programming problem

formulation provided in Doan and Natarajan (2012) where we consider a ρ -neighborhood around the bivariate marginals thus providing for a tradeoff between conservatism in the risk measure and confidence in the expert information. The proof is provided in the appendix.

Theorem 3.1 *Consider the univariate and bivariate marginals $\{\boldsymbol{\mu}_i\}_{i \in \mathcal{N}}$ and $\{\boldsymbol{\mu}_{ij}\}_{(i,j) \in \mathcal{N}'}$ with the set Θ_ρ defined in (3.1). Let \mathcal{V} denote the optimal value of the following convex program:*

$$\begin{aligned}
& \max_{v^k, v_i^k, v_{ij}^k, w_k, \boldsymbol{\theta}, \boldsymbol{\theta}_{ij}} \sum_{k \in \mathcal{K}} \sum_{\mathbf{c} \in \mathcal{C}} \mathbf{c}^T \mathbf{a}_k v^k(\mathbf{c}) + \sum_{k \in \mathcal{K}} b_k w_k \\
& \text{subject to} \quad \sum_{\mathbf{c} \in \mathcal{C}|c_{ij}} v^k(\mathbf{c}) = v_{ij}^k(c_{ij}), \quad \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \forall k \in \mathcal{K}, \\
& \quad \sum_{\mathbf{c} \in \mathcal{C}|c_i} v^k(\mathbf{c}) = v_i^k(c_i), \quad \forall i \in \mathcal{N}, \forall c_i \in \mathcal{C}_i, \forall k \in \mathcal{K}, \\
& \quad \sum_{\mathbf{c} \in \mathcal{C}} v^k(\mathbf{c}) = w_k, \quad \forall k \in \mathcal{K}, \\
& \quad \sum_{k \in \mathcal{K}} w_k = 1, \\
& \quad \sum_{k \in \mathcal{K}} v_i^k(c_i) = \mu_i(c_i), \quad \forall i \in \mathcal{N}, \forall c_i \in \mathcal{C}_i, \\
& \quad \sum_{k \in \mathcal{K}} v_{ij}^k(c_{ij}) = \theta_{ij}(c_{ij}), \quad \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\
& \quad \sum_{k \in \mathcal{K}} v^k(\mathbf{c}) = \boldsymbol{\theta}(\mathbf{c}), \quad \forall \mathbf{c} \in \mathcal{C}, \\
& \quad I_{KL}(\boldsymbol{\theta}_{ij}, \boldsymbol{\mu}_{ij}) \leq \rho, \quad \forall (i, j) \in \mathcal{N}', \\
& \quad v^k(\mathbf{c}) \geq 0, \quad \forall \mathbf{c} \in \mathcal{C}, \forall k \in \mathcal{K} \\
& \quad w_k \geq 0, \quad \forall k \in \mathcal{K}.
\end{aligned} \tag{3.2}$$

Then \mathcal{V} coincides with the Fréchet bound $\mathcal{M} = \max_{\boldsymbol{\theta} \in \Theta_\rho} \mathbb{E}_{\boldsymbol{\theta}} [\max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k)]$.

3.1 Tree Structure

In Theorem 3.1, we considered the general set of distributions in (3.1) with no additional conditions on the structure of the bivariate marginal information. Building on Theorem 2.1, under the

assumption of a tree structure over the set of indices $(\mathcal{N}, \mathcal{N}')$, the set of distributions Θ_ρ can be defined in terms of univariate and bivariate marginals as:

$$\begin{aligned} \Theta_\rho = \left\{ \theta : \text{proj}_i(\theta) = \mu_i, \forall i \in \mathcal{N}, \quad \text{proj}_{ij}(\theta) = \theta_{ij}, \forall (i, j) \in \mathcal{N}', \right. \\ \sum_{c_j \in \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) = \mu_i(c_i), \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_i \in \mathcal{C}_i, \\ \sum_{c_i \in \mathcal{C}_i} \theta_{ij}(\mathbf{c}_{ij}) = \mu_j(c_j), \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall c_j \in \mathcal{C}_j, \\ \left. I_{KL}(\theta_{ij}, \mu_{ij}) \leq \rho, \forall (i, j) \in \mathcal{N}' \right\}. \end{aligned} \quad (3.3)$$

The formulation for the worst-case bound is provided next and the proof is provided in the appendix.

Theorem 3.2 *Consider the univariate and bivariate marginals $\{\mu_i\}_{i \in \mathcal{N}}$ and $\{\mu_{ij}\}_{(i,j) \in \mathcal{N}'}$ such that $(\mathcal{N}, \mathcal{N}')$ has a tree structure with the set Θ_ρ defined in (3.3). Let \mathcal{V} denote the optimal value of the following primal convex program*

$$\begin{aligned} \max_{v_i^k, v_{ij}^k, w_k, \theta_{ij}} \quad & \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}} \sum_{c_i \in \mathcal{C}_i} c_i a_i^k v_i^k(c_i) + \sum_{k \in \mathcal{K}} b_k w_k \\ \text{subject to} \quad & \sum_{k \in \mathcal{K}} v_{ij}^k(\mathbf{c}_{ij}) = \theta_{ij}(\mathbf{c}_{ij}), \quad \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\ & \sum_{k \in \mathcal{K}} v_i^k(c_i) = \mu_i(c_i), \quad \forall i \in \mathcal{N}, \forall c_i \in \mathcal{C}_i, \\ & \sum_{\mathbf{c}_{ij} \in \mathcal{C}_{ij}} v_{ij}^k(\mathbf{c}_{ij}) = w_k, \quad \forall (i, j) \in \mathcal{N}', \forall k \in \mathcal{K}, \\ & \sum_{c_i \in \mathcal{C}} v_i^k(c_i) = w_k, \quad \forall i \in \mathcal{N}, \forall k \in \mathcal{K}, \\ & \sum_{k \in \mathcal{K}} w_k = 1, \\ & \sum_{c_i \in \mathcal{C}_i} v_{ij}^k(\mathbf{c}_{ij}) = v_j^k(c_j), \quad \forall j \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall k \in \mathcal{K}, \forall c_j \in \mathcal{C}_j, \\ & \sum_{c_j \in \mathcal{C}_j} v_{ij}^k(\mathbf{c}_{ij}) = v_i^k(c_i), \quad \forall i \in \mathcal{N} : (i, j) \in \mathcal{N}', \forall k \in \mathcal{K}, \forall c_i \in \mathcal{C}_i, \end{aligned} \quad (3.4)$$

$$\begin{aligned}
I_{KL}(\boldsymbol{\theta}_{ij}, \boldsymbol{\mu}_{ij}) &\leq \rho, & \forall (i, j) \in \mathcal{N}', \\
v_{ij}^k(\mathbf{c}_{ij}) &\geq 0, & \forall (i, j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \forall k \in \mathcal{K}, \\
w_k &\geq 0, & \forall k \in \mathcal{K},
\end{aligned}$$

Then \mathcal{V} coincides with the Fréchet bound $\mathcal{M} = \max_{\boldsymbol{\theta} \in \Theta_\rho} \mathbb{E}_{\boldsymbol{\theta}} [\max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k)]$.

In the formulation (3.2), the number of decision variables is $O(Km^n)$ where m is the number of possible values that each of the n random variables takes whereas it is only $O(Knm^2)$ in formulation (3.4). Thus, under the tree structure formulation, the problem is polynomial time solvable.

In the distributionally robust CVaR problem (1.3) which has a minmax formulation, the aim is to find the optimal decision variable \mathbf{x} which minimizes the worst-case expected shortfall. Since the inner maximization problem is convex programming problem, it can be reformulated as the minimization problem using a dual formulation, thereby simplifying the distributionally robust CVaR problem to minimization problem. With the univariate and bivariate marginals $\{\boldsymbol{\mu}_i\}_{i \in \mathcal{N}}$ and $\{\boldsymbol{\mu}_{ij}\}_{(i,j) \in \mathcal{N}'}$ such that $(\mathcal{N}, \mathcal{N}')$ has a tree structure along with the uncertainty set Θ_ρ defined in (3.3) with $\rho > \rho^*$ where ρ^* is the optimal value of the problem (2.3), the worst-case CVaR problem (1.3) can be reformulated as follows

$$\begin{aligned}
\min_{\mathbf{x} \in \mathcal{X}, \beta, \lambda \geq 0, \xi, \zeta, \tau, \nu} & \left[\beta + \frac{1}{1 - \alpha} \left(\nu + \sum_{(i,j) \in \mathcal{N}'} \lambda_{ij} \rho + \sum_{i \in \mathcal{N}} \sum_{c_i \in \mathcal{C}_i} \xi_i(c_i) \mu_i(c_i) \right. \right. \\
& \left. \left. + \sum_{(i,j) \in \mathcal{N}'} \sum_{\mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j} \lambda_{ij} \mu_{ij}(\mathbf{c}_{ij}) \left(e^{\xi_{ij}(\mathbf{c}_{ij})/\lambda_{ij}} - 1 \right) \right) \right] \\
\text{subject to} & \quad \nu \geq \sum_{i \in \mathcal{N}} \tau_i^1 + \sum_{(i,j) \in \mathcal{N}'} \tau_{ij}^1 - \beta, \\
& \quad \nu \geq \sum_{i \in \mathcal{N}} \tau_i^2 + \sum_{(i,j) \in \mathcal{N}'} \tau_{ij}^2,
\end{aligned}$$

$$\begin{aligned}
\xi_i(c_i) &\geq c_i x_i - \tau_i^1 + \sum_{j \in \mathcal{N}: (i,j) \in \mathcal{N}'} \zeta_i^{1j}(c_i) + \sum_{l \in \mathcal{N}: (l,i) \in \mathcal{N}'} \zeta_i^{1l}(c_i), \quad \forall i \in \mathcal{N}, \forall c_i \in \mathcal{C}_i, \\
\xi_i(c_i) &\geq -\tau_i^2 + \sum_{j \in \mathcal{N}: (i,j) \in \mathcal{N}'} \zeta_i^{2j}(c_i) + \sum_{l \in \mathcal{N}: (l,i) \in \mathcal{N}'} \zeta_i^{2l}(c_i), \quad \forall i \in \mathcal{N}, \forall c_i \in \mathcal{C}_i, \\
\xi_{ij}(\mathbf{c}_{ij}) &\geq -\left(\tau_{ij}^1 + \zeta_j^{1i}(c_j) + \zeta_i^{1j}(c_i)\right), \quad \forall (i,j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j, \\
\xi_{ij}(\mathbf{c}_{ij}) &\geq -\left(\tau_{ij}^2 + \zeta_j^{2i}(c_j) + \zeta_i^{2j}(c_i)\right), \quad \forall (i,j) \in \mathcal{N}', \forall \mathbf{c}_{ij} \in \mathcal{C}_i \times \mathcal{C}_j.
\end{aligned}$$

4 Numerical Experiments

In this section, we present numerical results for the Fréchet upper bound for the inner maximization problem in the distributionally robust CVaR formulation. We consider examples wherein the complete univariate and partial bivariate information is provided with $(\mathcal{N}, \mathcal{N}')$ forming a tree structure.

Consider the worst-case upper bound:

$$\max_{\theta \in \Theta_\rho} \mathbb{E}_\theta \left[\sum_{i \in \mathcal{N}} \tilde{c}_i - \beta \right]^+.$$

As the parameter ρ becomes large, the bound converges to the worst-case bound obtained assuming only univariate distributions are known which is equivalent to the well-known comonotonic upper bound. For consistent marginals, as ρ tends to zero, the bound converges to the worst-case bound with bivariate distributions known exactly whereas for the inconsistent scenario, as ρ tends to ρ^* where ρ^* is the minimal perturbation obtained by solving the closest consistent marginal (2.3), the bound converges to the worst-case bound with optimal bivariate distributions for the maximum entropy problem (2.5). Thus as ρ is decreased and more confidence is attached to the expert information, the bound reduces from the comonotonic upper bound to the worst-case bound assuming the exact bivariate marginals known or that obtained from maximum entropy in the set \mathcal{N}' . In our computations, the bounds are estimated with the KNITRO solver that is accessed through the AMPL modeling language.

For the numerical examples we fix $\mathcal{N} = \{1, 2, 3, 4, 5\}$ with the random variable \tilde{c}_i taking values in $\mathcal{C}_i = \{1, 2, \dots, 10\}$ for $i \in \mathcal{N}$. Let the bivariate marginals be specified for $\mathcal{N}' = \{(1, 2), (2, 3), (3, 4), (4, 5)\}$ where the index set $(\mathcal{N}, \mathcal{N}')$ form a series graph.

We first consider an example with consistent marginal information. Corresponding to consistent univariate and bivariate marginals a joint distribution can be generated using the Chow-Liu tree distribution, thereby ensuring that the class of distributions Θ_ρ defined by (3.3) is nonempty for any $\rho \geq 0$. To compute the bound, we solve the convex optimization problem (3.4) with $\mathcal{K} = \{1, 2\}$, $a_i^1 = 1$ for $i \in \mathcal{N}$, $b_1 = -\beta$ and $a_i^2 = 0$ for $i \in \mathcal{N}$, $b_2 = 0$. Since $\sum_{i \in \mathcal{N}} \tilde{c}_i$ takes values in $\{5, 6, \dots, 50\}$, we can vary β in $[5, 50]$. We assume that each \tilde{c}_i is a discrete uniform random variable which takes values in \mathcal{C}_i with probability 0.1. We consider three different kinds of dependencies modeling expert information in the set \mathcal{N}' given by, (i) very positively correlated pairs of random variables, (ii) pairwise independent random variables and (iii) very negatively correlated pairs of random variables. For positively correlated random variables, the bivariate distributions in \mathcal{N}' are chosen with discrete uniform marginals and a Gaussian copula (discretized) with correlation parameter 0.69. Figure 1 represents the probability heat map of the optimal bivariate distributions for the random variables \tilde{c}_1 and \tilde{c}_2 obtained by solving the convex optimization problem (3.4) over $\beta = 15, 30, 45$ corresponding to different values of the parameter $\rho = 0.00001, 0.01, 0.1, 0.5$. We also provide the worst-case bounds in the figure (written as ‘Bound’ in the caption of each sub-figure). Optimal distributions for the other bivariate distributions in the set \mathcal{N}' are the same as observed in Figure 1 (by construction) and hence not shown here. In Figure 1, we observe that as ρ decreases, the bivariate distributions are positively correlated with small changes where the bound decreases slightly. In Figure 2, we consider the case where the expert information is pairwise independent bivariate distributions (can be thought of as Gaussian copula with correlation parameter 0) in the set \mathcal{N}' . As the figure indicates, the worst-case bivariate distributions as ρ

decrease, moves from positively correlated distributions to independent distributions. The bounds decrease more in these instances as compared to Figure 1. Lastly, in Figure 3, we consider the case where the bivariate distributions in \mathcal{N}' are chosen with uniform marginals and a Gaussian copula with correlation parameter -0.69 . In this case, the worst-case bivariate distributions as ρ decrease moves from positively correlated distributions to negatively correlated distributions. The bound reduces much more in this case since the expert information is very different from the worst-case bivariate distributions with comonotonic random variables.

In Figure 4, we compare the Fréchet bound for the perturbation parameter $\rho = 0.1$ with the bounds obtained from the bivariate distribution known exactly and from only the univariate distributions known for different β in all the three cases. Observe that depending on the expert information (that is the bivariate distribution) known, the Fréchet bound varies between the bounds obtained from univariate and the bivariate information. For the positively correlated scenario, the Fréchet bound for $\rho = 0.1$ is closer to the univariate bound while for the negatively correlated case, it is closer to the bivariate bound. Thus we observe that the model is able to capture the trade-off between conservatism in the worst-case risk measure and confidence in expert information.

For our next example, unlike the discrete uniform distribution above for the univariate marginals, we consider a non-uniform discrete univariate distribution for $i \in \mathcal{N}$ with the marginals given in Table 1, while generating three different bivariate distributions for $(i, j) \in \mathcal{N}'$ using the discrete Gaussian copula as in the preceding example, that is with correlation coefficients $0.69, 0$ and -0.69 .

c	1	2	3	4	5	6	7	8	9	10
$\mu_i(c)$	0.025	0.050	0.075	0.15	0.20	0.20	0.15	0.075	0.050	0.025

Table 1: Univariate marginals μ_i

Clearly the marginals are inconsistent in all the three cases (since the discretized Gaussian copulas

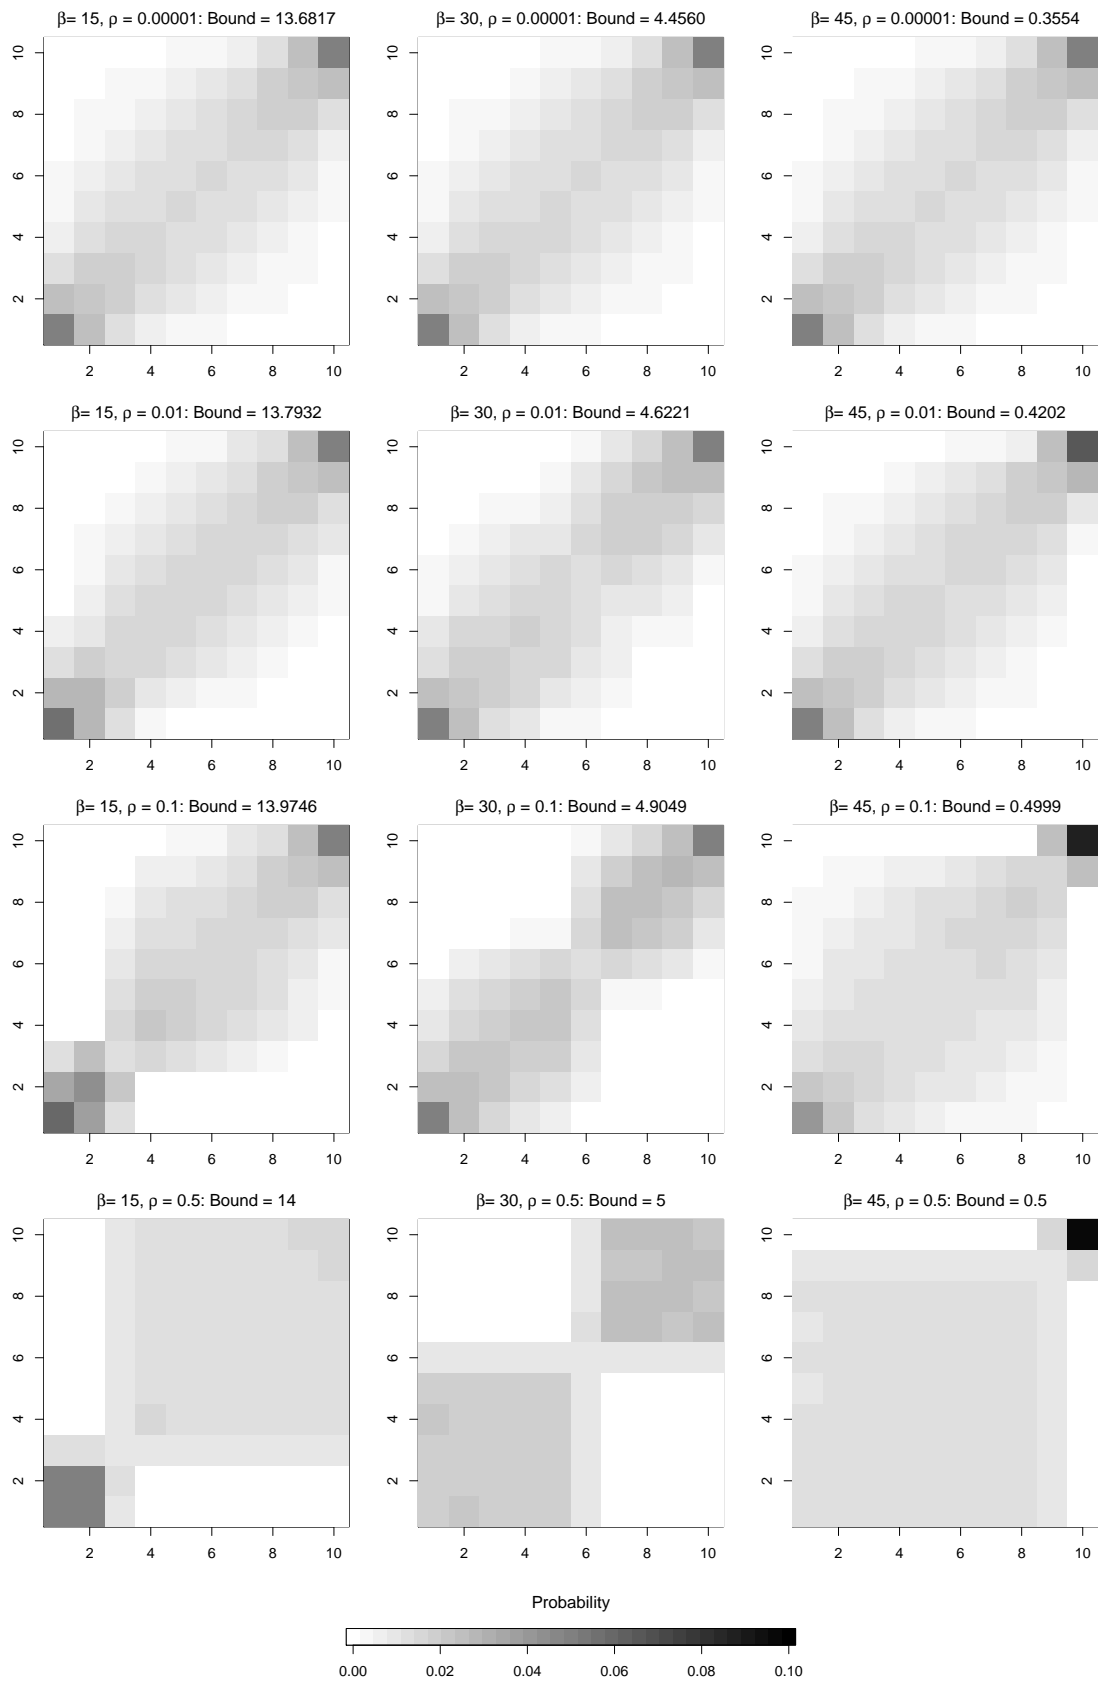


Figure 1: Heat map for θ_{12} for discrete uniform random variables with correlation coefficient 0.69

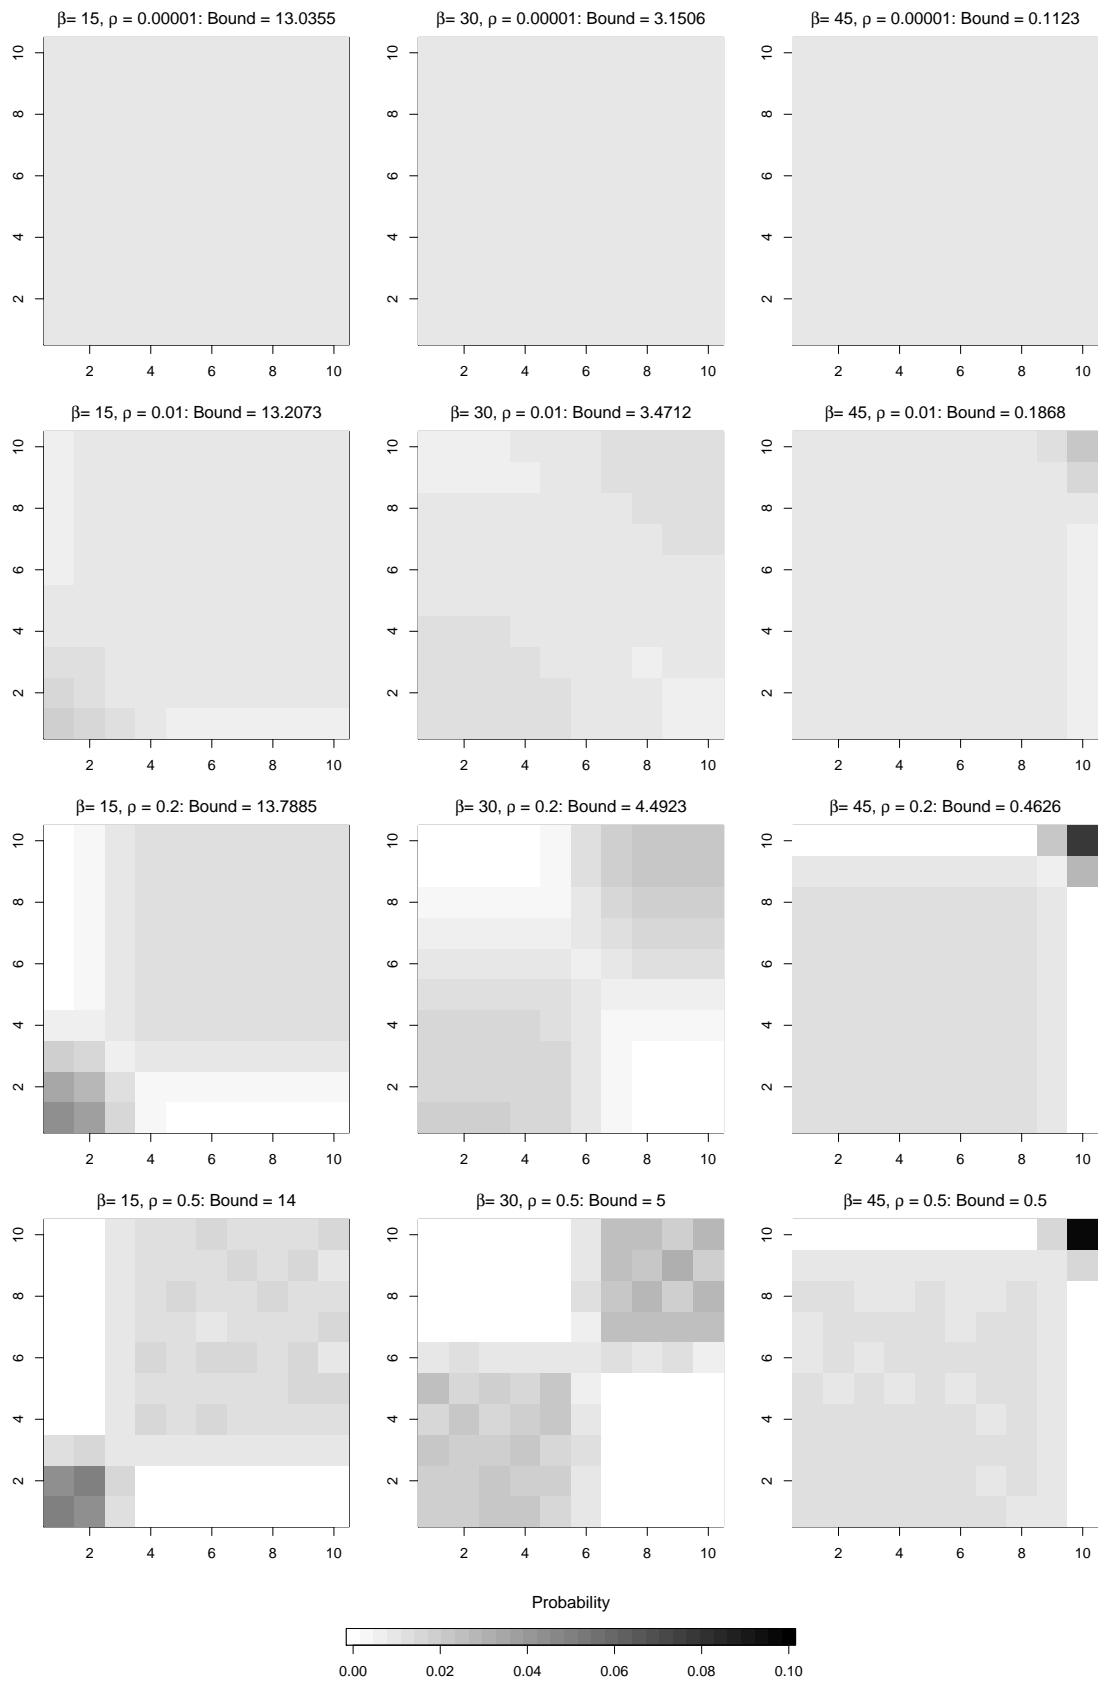


Figure 2: Heat map for θ_{12} for discrete uniform independent random variables




Figure 3: Heat map for θ_{12} for discrete uniform random variables with correlation coefficient -0.69

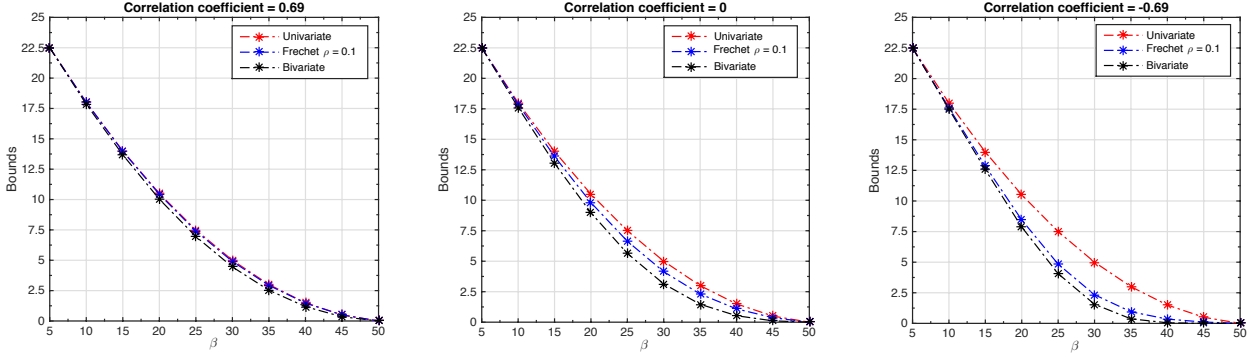


Figure 4: Comparison of Fréchet bound for $\rho = 0.1$ with univariate and bivariate bounds

has discrete Uniform marginals). Thus we solve the closest consistency problem (2.3) and consider the class of distribution Θ_ρ with $\rho \geq \rho^*$ where ρ^* is the minimal perturbation obtained. For the cases with correlation coefficient 0.69, 0 and -0.69 , the value of ρ^* as obtained by solving (2.3) are 0.342356, 0.434234 and 0.342356 respectively. Moreover, the optimal bivariate distribution θ_{ij} consistent with the given univariate as obtained by solving the relative maximum entropy problem (2.5) are feasible for Θ_ρ with $\rho \geq \rho^*$. For $\beta = 5, 10, \dots, 45, 50$ and different values of $\rho > \rho^*$, the Fréchet bounds were calculated using (3.4) which is always more than the bound calculated with given univariate and the bivariate calculated from maximum entropy. As ρ increases, the Fréchet bound converges to the univariate bound whereas as ρ decreases to ρ^* , the Fréchet bound converges to the maximum entropy bound. In Figure 5, we present the heat map of the optimal bivariate distributions for the random variables \tilde{c}_1 and \tilde{c}_2 obtained by solving (3.4) over $\beta = 15, 30, 45$ corresponding to different values of the parameter $\rho = 0.34236, 0.4, 0.75, 1.1$ for the case with correlation coefficient -0.69 while in Figure 6 we compare the Fréchet bound for $\rho = 0.5$ with the bound obtained with maximum entropy bivariate distribution and that obtained with univariate marginals only known for all three cases. Unlike the uniform univariate marginals case in the first example, the probabilities for the optimal bivariate distributions are concentrated towards the center; see Table 1. Comparing Figure 6 with Figure 4 (same bivariate information, but

different univariate information), we observe that although the concentration of mass is affected by the marginal univariate information provided; as ρ decreases we move from a more positively correlated structure to a more negatively correlated structure in both examples.

5 Conclusions

In this paper, we have developed a Fréchet upper bound for the distributionally robust CVaR problem under univariate and bivariate marginal information. By allowing for expert information in terms of bivariate marginals, we model the trade-off between conservatism in the worst-case risk measure and confidence in the expert information. Importantly, we show that as long as the bivariate marginals form a tree structure, the problem is efficiently solvable as a convex optimization problem. We end by discussing some possible directions for future research.

In our model, we considered a class of distributions in the ρ -neighborhood expressed in terms of KL-divergence. Under the KL-divergence measure, the new bivariate marginal has non-zero probability only on support points where the original bivariate marginal has non-zero probability. It will be interesting to extend our formulations to consider other ϕ -divergence measures which allows one to consider bivariate marginals in a neighborhood where non-zero probabilities might be assigned to support points with zero probabilities in the expert information (see Ben-Tal et. al. 2013). Secondly, while our study has been primarily restricted to discrete distributions, it would be interesting to study how well such a method would do for continuous distributions using discretization methods. Finally, in our results, the tree structure plays a pivotal role in simplifying the formulation in comparison to the general graph structure. A natural question is whether for arbitrary bivariate information, there are methods to find an optimal tree structure for the worst-case CVaR problem.

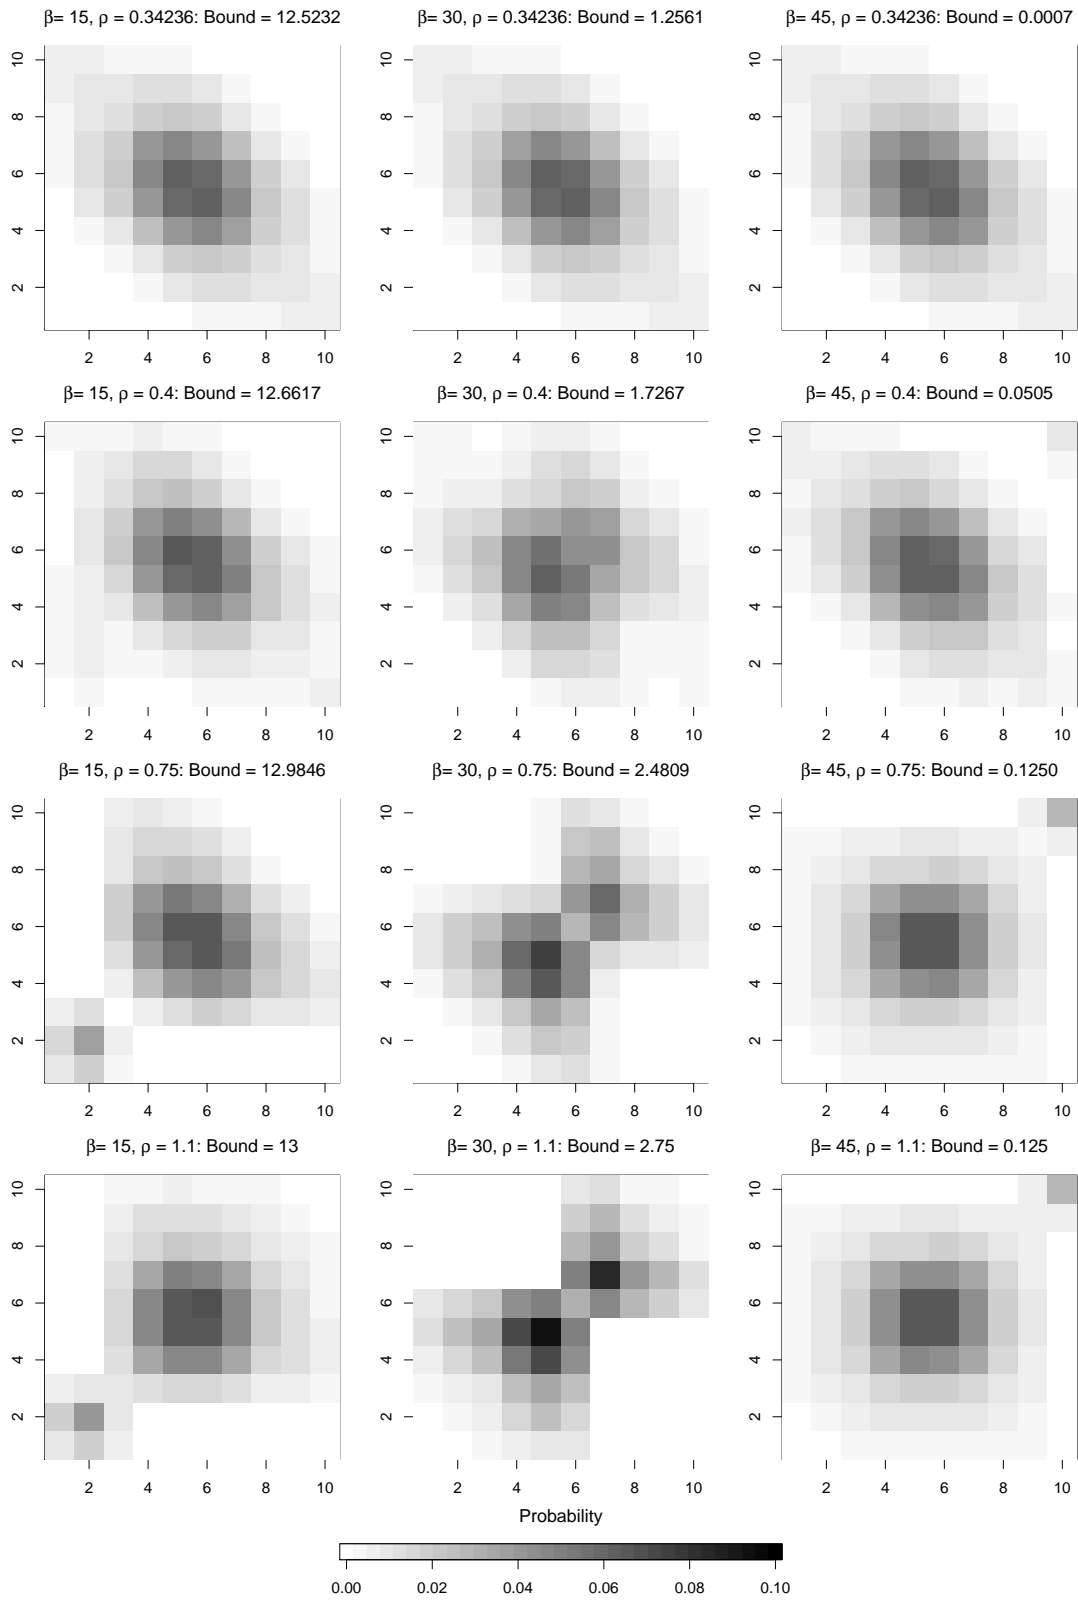


Figure 5: Heat map for θ_{12} for discrete random variables with correlation coefficient -0.69

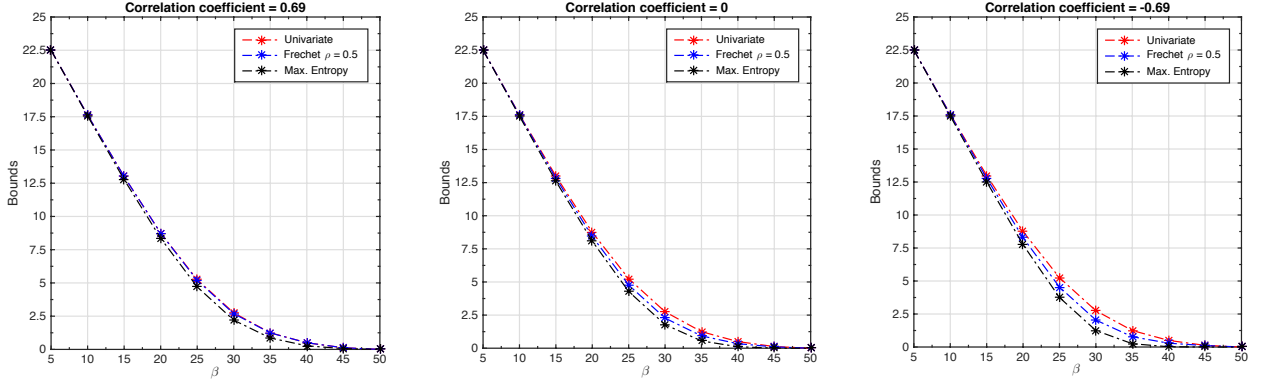


Figure 6: Comparison of Fréchet bound for $\rho = 0.5$ with univariate and max. entropy bounds

Appendix

Proof of Theorem 2.1

Proof. We show the equivalence of formulations (2.2) and (2.3) when the graph associated with the univariate and bivariate marginal index sets $(\mathcal{N}, \mathcal{N}')$ forms a tree structure. Given a feasible solution to (2.2) denoted by $(\rho, \boldsymbol{\theta}, \boldsymbol{\theta}_{ij})$, we consider the corresponding solution $(\rho, \boldsymbol{\theta}_{ij})$ to (2.3) with the same objective value. This solution is feasible to (2.3) since:

$$\begin{aligned} \sum_{c_j \in \mathcal{C}_j} \theta_{ij}(\mathbf{c}_{ij}) &= \sum_{c_j \in \mathcal{C}_j} \sum_{\mathbf{c} \in \mathcal{C} | \mathbf{c}_{ij}} \boldsymbol{\theta}(\mathbf{c}) = \sum_{\mathbf{c} \in \mathcal{C} | c_i} \boldsymbol{\theta}(\mathbf{c}) = \mu_i(c_i). \\ \sum_{c_i \in \mathcal{C}_i} \theta_{ij}(\mathbf{c}_{ij}) &= \sum_{c_i \in \mathcal{C}_i} \sum_{\mathbf{c} \in \mathcal{C} | \mathbf{c}_{ij}} \boldsymbol{\theta}(\mathbf{c}) = \sum_{\mathbf{c} \in \mathcal{C} | c_j} \boldsymbol{\theta}(\mathbf{c}) = \mu_j(c_j). \end{aligned}$$

By the non-negativity of $\boldsymbol{\theta}(\mathbf{c})$, the non-negativity condition of $\theta_{ij}(\mathbf{c}_{ij})$ holds in (2.3).

To show the converse, observe that given a feasible solution to (2.3), we can construct a joint distribution by using a conditionally independent distribution (Chow-Liu (1968) tree). For the tree structured index set $(\mathcal{N}, \mathcal{N}')$, the joint distribution is defined as

$$\boldsymbol{\theta}(\mathbf{c}) = \prod_{i \in \mathcal{N}} \mu_i(c_i) \prod_{(i,j) \in \mathcal{N}'} \frac{\theta_{ij}(\mathbf{c}_{ij})}{\mu_i(c_i) \mu_j(c_j)}$$

satisfying the univariate and bivariate marginal conditions

$$\sum_{\mathbf{c} \in \mathcal{C} | \mathbf{c}_{ij}} \boldsymbol{\theta}(\mathbf{c}) = \theta_{ij}(\mathbf{c}_{ij}) \quad \text{and} \quad \sum_{\mathbf{c} \in \mathcal{C} | c_i} \boldsymbol{\theta}(\mathbf{c}) = \mu_i(c_i),$$

thereby satisfying the conditions of (2.2). Hence, under the tree structure (2.2) reduces to (2.3). ■

Proof of Theorem 3.1

Proof. Define $\psi(\tilde{\mathbf{c}}) = \max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k)$. For any joint distribution $\boldsymbol{\theta} \in \Theta_\rho$, we have:

$$\begin{aligned} \mathbb{E}_\theta(\psi(\tilde{\mathbf{c}})) &= \mathbb{E}_\theta \left[\max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k) \right] \\ &= \sum_{k \in \mathcal{K}} \mathbb{E}_\theta \left(\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k \mid k\text{-th term is maximum} \right) \mathbb{P}_\theta(k\text{-th term is maximum}) \\ &= \sum_{k \in \mathcal{K}} \sum_{\mathbf{c} \in \mathcal{C}} \mathbf{c}^T \mathbf{a}_k \mathbb{P}_\theta(\tilde{\mathbf{c}} = \mathbf{c}, k\text{-th term is maximum}) \\ &\quad + \sum_{k \in \mathcal{K}} b_k \mathbb{P}_\theta(k\text{-th term is maximum}) \\ &= \sum_{k \in \mathcal{K}} \sum_{\mathbf{c} \in \mathcal{C}} \mathbf{c}^T \mathbf{a}_k v_k(\mathbf{c}) + \sum_{k \in \mathcal{K}} b_k w_k \end{aligned}$$

where the decision variables $v^k(\mathbf{c})$ and w_k denote

$$v^k(\mathbf{c}) = \mathbb{P}_\theta(\tilde{\mathbf{c}} = \mathbf{c}, k\text{-th term is maximum}) \quad \text{and} \quad w_k = \mathbb{P}_\theta(k\text{-th term is maximum}).$$

Since the decision variables $v^k(\mathbf{c})$ and w_k are probability measures, they are non-negative. By the definition of probability, observe that

$$\begin{aligned} \sum_{k \in \mathcal{K}} v^k(\mathbf{c}) &= \sum_{k \in \mathcal{K}} \mathbb{P}_\theta(\tilde{\mathbf{c}} = \mathbf{c}, k\text{-th term is maximum}) = \mathbb{P}_\theta(\tilde{\mathbf{c}} = \mathbf{c}) = \boldsymbol{\theta}(\mathbf{c}), \\ \sum_{\mathbf{c} \in \mathcal{C}} v^k(\mathbf{c}) &= \sum_{\mathbf{c} \in \mathcal{C}} \mathbb{P}_\theta(\tilde{\mathbf{c}} = \mathbf{c}, k\text{-th term is maximum}) = \mathbb{P}_\theta(k\text{-th term is maximum}) = w_k, \\ \sum_{k \in \mathcal{K}} w_k &= \sum_{k \in \mathcal{K}} \mathbb{P}_\theta(k\text{-th term is maximum}) = 1. \end{aligned}$$

Further we introduce the following decision variables

$$\begin{aligned} v_{ij}^k(\mathbf{c}_{ij}) &= \mathbb{P}_\theta(\tilde{\mathbf{c}}_{ij} = \mathbf{c}_{ij}, k\text{-th term is maximum}) \quad \text{and} \\ v_i^k(c_i) &= \mathbb{P}_\theta(\tilde{c}_i = c_i, k\text{-th term is maximum}). \end{aligned}$$

The variables $v^k(\mathbf{c})$, $v_{ij}^k(\mathbf{c}_{ij})$ and $v_i^k(c_i)$ are related as follows:

$$\begin{aligned}
\sum_{\mathbf{c} \in \mathcal{C} | \mathbf{c}_{ij}} v^k(\mathbf{c}) &= \sum_{\mathbf{c} \in \mathcal{C} | \mathbf{c}_{ij}} \mathbb{P}_{\boldsymbol{\theta}}(\tilde{\mathbf{c}} = \mathbf{c}, k\text{-th term is maximum}) \\
&= \mathbb{P}_{\boldsymbol{\theta}}(\tilde{\mathbf{c}}_{ij} = \mathbf{c}_{ij}, k\text{-th term is maximum}) = v_{ij}^k(\mathbf{c}_{ij}), \\
\sum_{\mathbf{c} \in \mathcal{C} | c_i} v^k(\mathbf{c}) &= \sum_{\mathbf{c} \in \mathcal{C} | c_i} \mathbb{P}_{\boldsymbol{\theta}}(\tilde{\mathbf{c}} = \mathbf{c}, k\text{-th term is maximum}) \\
&= \mathbb{P}_{\boldsymbol{\theta}}(\tilde{c}_i = c_i, k\text{-th term is maximum}) = v_i^k(c_i).
\end{aligned}$$

Moreover, $v_{ij}^k(\mathbf{c}_{ij})$ and $v_i^k(c_i)$ are related to $\theta_{ij}(\mathbf{c}_{ij})$ and $\mu_i(c_i)$ respectively as

$$\begin{aligned}
\sum_{k \in \mathcal{K}} v_{ij}^k(\mathbf{c}_{ij}) &= \sum_{k \in \mathcal{K}} \mathbb{P}_{\boldsymbol{\theta}}(\tilde{\mathbf{c}}_{ij} = \mathbf{c}_{ij}, k\text{-th term is maximum}) = \mathbb{P}_{\boldsymbol{\theta}}(\tilde{\mathbf{c}}_{ij} = \mathbf{c}_{ij}) = \theta_{ij}(\mathbf{c}_{ij}) \\
\sum_{k \in \mathcal{K}} v_i^k(c_i) &= \sum_{k \in \mathcal{K}} \mathbb{P}_{\boldsymbol{\theta}}(\tilde{c}_i = c_i, k\text{-th term is maximum}) = \mathbb{P}_{\boldsymbol{\theta}}(\tilde{c}_i = c_i) = \mu_i(c_i)
\end{aligned}$$

For the given univariate and bivariate marginals $\boldsymbol{\mu}_i$ and $\boldsymbol{\mu}_{ij}$, $\boldsymbol{\theta}_{ij}$ satisfy the following ρ -neighborhood condition

$$I_{KL}(\boldsymbol{\theta}_{ij}, \boldsymbol{\mu}_{ij}) \leq \rho, \quad \forall (i, j) \in \mathcal{N}'$$

such that it is consistent with $\boldsymbol{\mu}_i$. Thus for any $\boldsymbol{\theta} \in \boldsymbol{\Theta}_\rho$ along with the decision variables \mathbf{v}_k , \mathbf{v}_i^k , \mathbf{v}_{ij}^k , w_k satisfying the constraints,

$$\max_{\mathbf{v}_k, \mathbf{v}_{ij}^k, \mathbf{v}_i^k, w_k} \sum_{k \in \mathcal{K}} \sum_{\mathbf{c} \in \mathcal{C}} \mathbf{c}^T \mathbf{a}_k v_k(\mathbf{c}) + \sum_{k \in \mathcal{K}} b_k w_k \geq \mathbb{E}_{\boldsymbol{\theta}}(\psi(\tilde{\mathbf{c}})),$$

which implies

$$\mathcal{V} = \max_{\mathbf{v}_k, \mathbf{v}_{ij}^k, \mathbf{v}_i^k, w_k, \boldsymbol{\theta}} \sum_{k \in \mathcal{K}} \sum_{\mathbf{c} \in \mathcal{C}} \mathbf{c}^T \mathbf{a}_k v_k(\mathbf{c}) + \sum_{k \in \mathcal{K}} b_k w_k \geq \max_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \mathbb{E}_{\boldsymbol{\theta}}(\psi(\tilde{\mathbf{c}})) = \mathcal{M}.$$

To prove the result, we need to show that the bound is tight, that is $\mathcal{V} = \mathcal{M}$. Consider the optimal decision variables \mathbf{v}_k^* , \mathbf{v}_{ij}^{k*} , \mathbf{v}_i^{k*} , w_k^* , $\boldsymbol{\theta}^*$ and $\boldsymbol{\theta}_{ij}^*$ of the convex programming problem (3.2). Observe that w_k^* is a probability measure. Next for a fixed $k \in \mathcal{K}$, we construct distributions $\boldsymbol{\theta}^{k*}$ and $\boldsymbol{\theta}_{ij}^{k*}$ as follows:

- (a) Choose the k -th term $\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k$ with probability w_k^* .

(b) For a fixed $k \in \mathcal{K}$, define $\theta^{k*}(\mathbf{c}) = v_k^*(\mathbf{c})/w_k^*$. In addition, for each $(i, j) \in \mathcal{N}'$, define $\theta_{ij}^{k*}(\mathbf{c}_{ij}) = v_{ij}^{k*}(\mathbf{c}_{ij})/w_k^*$ and for $i \in \mathcal{N}$, $\theta_i^{k*}(c_i) = v_i^{k*}(c_i)/w_k^*$. Note that if $w_k^* = 0$, we simply drop that index. It is easy to observe that

$$\sum_{\mathbf{c} \in \mathcal{C} | \mathbf{c}_{ij}} \theta^{k*}(\mathbf{c}) = \theta_{ij}^{k*}(\mathbf{c}_{ij}) \quad \text{and} \quad \sum_{\mathbf{c} \in \mathcal{C} | c_i} \theta^{k*}(\mathbf{c}) = \theta_i^{k*}(c_i).$$

Hence the following inequality holds:

$$\mathbb{E}_{\theta^{k*}} \left[\max_{l \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_l + b_l) \right] \geq \mathbb{E}_{\theta^{k*}} [\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k] = \frac{1}{w_k^*} \sum_{\mathbf{c} \in \mathcal{C}} \mathbf{c}^T \mathbf{a}_k v^{k*}(\mathbf{c}) + b_k$$

where the first inequality is obtained by simply choosing the k -th term in the function for θ^{k*} .

Since

$$\begin{aligned} \mathbb{E}_{\theta^*} \left[\max_{l \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_l + b_l) \right] &= \sum_{k \in \mathcal{K}} w_k^* \mathbb{E}_{\theta^{k*}} \left[\max_{l \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_l + b_l) \right] \\ &\geq \sum_{k \in \mathcal{K}} w_k^* \left[\frac{1}{w_k^*} \sum_{\mathbf{c} \in \mathcal{C}} \tilde{\mathbf{c}}^T \mathbf{a}_k v^{k*}(\mathbf{c}) + b_k \right] = \sum_{k \in \mathcal{K}} \sum_{\mathbf{c} \in \mathcal{C}} \tilde{\mathbf{c}}^T \mathbf{a}_k v^{k*}(\mathbf{c}) + \sum_{k \in \mathcal{K}} b_k w_k^*. \end{aligned}$$

Therefore, $\mathcal{M} = \max_{\theta \in \Theta_\rho} \mathbb{E}_\theta \left[\max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k) \right] \geq \mathbb{E}_{\theta^*} \left[\max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k) \right] \geq \mathcal{V}$ which together with the fact that $\mathcal{V} \geq \mathcal{M}$ ensures that $\mathcal{V} = \mathcal{M}$. ■

Proof of Theorem 3.2

Proof. Define $\psi(\tilde{\mathbf{c}}) = \max_{k \in \mathcal{K}} (\tilde{\mathbf{c}}^T \mathbf{a}_k + b_k)$. For any joint distribution $\theta \in \Theta_\rho$, the expected value is expressed as:

$$\mathbb{E}_\theta(\psi(\tilde{\mathbf{c}})) = \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}} \sum_{c_i \in \mathcal{C}_i} c_i a_i^k v_i^k(c_i) + \sum_{k \in \mathcal{K}} b_k w_k$$

where the decision variable $v_i^k(c_i)$ and w_k denote

$$v_i^k(c_i) = \mathbb{P}_\theta(\tilde{c}_i = c_i, k\text{-th term is maximum}) \quad \text{and} \quad w_k = \mathbb{P}_\theta(k\text{-th term is maximum})$$

Further we introduce another decision variable $v_{ij}^k(\mathbf{c}_{ij}) = \mathbb{P}_{\boldsymbol{\theta}}(\tilde{\mathbf{c}}_{ij} = \mathbf{c}_{ij}, k\text{-th term is maximum})$ to incorporate the consistency relation between the given univariate $\boldsymbol{\mu}_i$ and the perturbed bivariate $\boldsymbol{\theta}_{ij}$. Using the basic properties of probability measures, the constraints are easily satisfied. Thus for any $\boldsymbol{\theta} \in \Theta_{\rho}$ along with the decision variables $\mathbf{v}_i^k, \mathbf{v}_{ij}^k, w_k$ satisfying the constraints,

$$\mathcal{V} = \max_{\mathbf{v}_{ij}^k, \mathbf{v}_i^k, w_k, \boldsymbol{\theta}} \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}} \sum_{c_i \in \mathcal{C}_i} c_i a_i^k v_i^k(c_i) + \sum_{k \in \mathcal{K}} b_k w_k \geq \max_{\boldsymbol{\theta} \in \Theta} \mathbb{E}_{\boldsymbol{\theta}}(\psi(\tilde{\mathbf{c}})) = \mathcal{M}.$$

For the reverse inequality $\mathcal{M} \geq \mathcal{V}$, consider the optimal decision variables $\mathbf{v}_{ij}^{k*}, \mathbf{v}_i^{k*}$ and w_k^* of the convex programming problem. Next for a fixed $k \in \mathcal{K}$ we construct a joint distribution $\boldsymbol{\theta}^{k*}$ as follows:

- (a) Choose the k -th term $\sum_{i \in \mathcal{N}} \sum_{c_i \in \mathcal{C}_i} c_i a_i^k + b_k$ with probability w_k^* .
- (b) For each $(i, j) \in \mathcal{N}'$, define $\theta_{ij}^{k*}(\mathbf{c}_{ij}) = v_{ij}^{k*}(\mathbf{c}_{ij})/w_k^*$ and for each $i \in \mathcal{N}$, define $\theta_i^{k*}(c_i) = v_i^{k*}(c_i)/w_k^*$. Note that if $w_k^* = 0$, we simply drop that index. Using the consistency conditions of $v_{ij}^{k*}(\mathbf{c}_{ij})$ and $v_i^{k*}(c_i), v_j^{k*}(c_j)$, it is easy to observe that $\theta_{ij}^{k*}(\mathbf{c}_{ij})$ and $\theta_i^{k*}(c_i), \theta_j^{k*}(c_j)$ are consistent.
- (c) For $k \in \mathcal{K}$, since $(\mathcal{N}, \mathcal{N}')$ forms a tree structure, there exists a consistent joint distribution $\boldsymbol{\theta}^{k*}$ corresponding to $\theta_{ij}^{k*}(\mathbf{c}_{ij})$ and $\theta_i^{k*}(c_i), \theta_j^{k*}(c_j)$.

Now working along the lines of the proof of Theorem 3.1, we obtain

$$\mathbb{E}_{\boldsymbol{\theta}^*} \left[\max_{l \in \mathcal{K}} \left(\tilde{\mathbf{c}}^T \mathbf{a}_l + b_l \right) \right] \geq \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}} \sum_{c_i \in \mathcal{C}_i} c_i a_i^k v_i^{k*}(c_i) + \sum_{k \in \mathcal{K}} b_k w_k^*.$$

Therefore, $\mathcal{M} \geq \mathcal{V}$ which together with the fact that $\mathcal{V} \geq \mathcal{M}$ ensures that $\mathcal{V} = \mathcal{M}$. ■

References

Artzner, P., F. Delbaen, J-M. Eber, and D. Heath (1999): Coherent measures of risk, *Mathematical Finance*, 9, 203-228.

- Bedford, T., and R. M. Cooke (2002): A new graphical model for dependent random variables, *Annals of Statistics*, 30, 1031-1068.
- Ben-Tal, A., D. den Hertog, A. De Waegenaere, B. Melenberg, and G. Rennen (2013): Robust solutions of optimization problems affected by uncertain probabilities, *Management Science*, 59, 341-357.
- Cario, M. C., and B. L. Nelson (1997): Modeling and generating random vectors with arbitrary marginal distributions and correlation matrix, *Technical Report*, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL.
- Chow, C. K., and C. N. Liu (1968): Approximating discrete probability distributions with dependence trees, *IEEE Transactions on Information Theory*, 14, 462-467.
- Deming, W. E., and F. F. Stephan (1940): On a least squares adjustment of a sampled frequency table when the expected marginal totals are known, *The Annals of Mathematical Statistics*, 11, 427-444.
- Doan, X. V., and K. Natarajan (2012): On the complexity of nonoverlapping multivariate marginal bounds for probabilistic combinatorial optimization problems, *Operations Research*, 60, 138-149.
- Doan, X. V., X. Li, and K. Natarajan (2015): Robustness to dependency in portfolio optimization using overlapping marginals, *Operations Research*, 63, 1468-1488.
- Embrechts, P., and G. Puccetti (2006): Bounds for functions of multivariate risks, *Journal of Multivariate Analysis*, 97, 526-547.
- Embrechts, P., G. Puccetti, L. Rüschendorf, R. Wang, and A. Beleraaj (2014): An academic response to Basel 3.5, *Risks*, 2, 25-58.
- Fienberg, S. E. (1970): An iterative procedure for estimation in contingency tables, *The Annals of Mathematical Statistics*, 41, 907-917.
- Ghosh, S., and S. G. Henderson (2002): Chessboard Distributions and Random Vectors with

- Specified Marginals and Covariance Matrix, *Operations Research*, 50, 820-834.
- Glasserman, P., and L. Yang (2016): Bounding wrong-way risk in CVA calculation, *Mathematical Finance*, doi:10.1111/mafi.12141.
- Hanasusanto, G. A., D. Kuhn, and W. Wiesemann (2016): A comment on “Computational complexity of stochastic programming problems”, *Mathematical Programming*, doi:10.1007/s10107-015-0958-2.
- Higham, N. J. (2002): Computing the nearest correlation matrix—a problem from finance, *IMA Journal of Numerical Analysis*, 22, 329-343.
- Ireland, C. T., and S. Kullback (1968): Contingency tables with given marginals, *Biometrika*, 55, 179-188.
- Jaynes, E. T. (1957): Information theory and statistical mechanics, *Physical Review*, 106, 620-630.
- Lim, A. E. B., J. G. Shanthikumar, and G.-Y. Vahn (2011): Conditional value-at-risk in portfolio optimization: Coherent but fragile, *Operations Research Letters*, 39, 163-171.
- Pardo, L. (2006): *Statistical inference based on divergence measures*, Statistics: Textbooks and Monographs, 185, Boca Raton: Chapman & Hall/CRC FL.
- Qi, H. D., and D. Sun (2010): Correlation stress testing for Value-at-Risk: An unconstrained convex optimization approach, *Computational Optimization and Applications*, 45, 427-462.
- Rockafellar, R. T., and S. Uryasev (2002): Conditional value-at-risk for general loss distributions, *Journal of Banking & Finance*, 26, 1443-1471.
- Roughgarden, T., and M. Kearns (2013): Marginals-to-models reducibility, *Proceedings of the 26th International Conference on Neural Information Processing Systems*, Lake Tahoe, Nevada, 1043-1051.
- Rüschendorf, L. (1983): Solution of a statistical optimization problem by rearrangement method, *Metrika*, 30, 55-61.

Shannon, C. E. (1948): A mathematical theory of communication, *Bell System Technical Journal*, 27, 379-423, 623-656.

Sklar, A. (1959): Fonctions de répartition à n dimensions et leurs marges, *Publications de l'Institut de Statistique de l'Université de Paris*, 8, 229-231.

Vorob'ev, N. N. (1962): Consistent families of measures and their extensions, *Theory of Probability and Its Applications*, 7, 147-163.

Zhu, S., and M. Fukushima (2009): Worst-case Conditional Value-at-Risk with application to robust portfolio management, *Operations Research*, 57, 1155-1168.