Vanderbilt University Department of Economics Working Paper Series
Nashville, TN 37235-1819

Vanderbilt University Department of Economics Working Papers 18-00002

Feasible Shared Destiny Risk Distributions

Thibault Gajdos
CNRS, Aix-Marseille University

John A Weymark
Vanderbilt University

Claudio Zoli
University of Verona

Abstract

Social risk equity is concerned with the comparative evaluation of social risk distributions, which are probability distributions over the potential sets of fatalities. In the approach to the evaluation of social risk equity introduced by Gajdos, Weymark, and Zoli (Shared destinies and the measurement of social risk equity, Annals of Operations Research 176:409-424, 2010), the only information about such a distribution that is used in the evaluation is that contained in a shared destiny risk matrix whose entry in the kth row and ith column is the probability that person i dies in a group containing k individuals. Such a matrix is admissible if it satisfies a set of restrictions implied by its definition. It is feasible if it can be generated by a social risk distribution. It is shown that admissibility is equivalent to feasibility. Admissibility is much easier to directly verify than feasibility, so this result provides a simply way to identify which matrices to consider when the objective is to socially rank the feasible shared destiny risk matrices.

This paper has been prepared for Indraneel Dasgupta and Manipushpak Mitra, Eds., Deprivation, Inequality and Poverty: Essays in Honour of Satya Ranjan Chakravarty, to be published by Springer.
Citation: Thibault Gajdos and John A Weymark and Claudio Zoli, (2018) "Feasible Shared Destiny Risk Distributions", *Vanderbilt University Department of Economics Working Papers*, VUECON-18-00002.
Contact: Thibault Gajdos - thibault.gajdos@univ-amu.fr, John A Weymark - john.weymark@vanderbilt.edu, Claudio Zoli - claudio.zoli@univr.it.
Submitted: January 15, 2018.   **Published:** January 16, 2018.
URL:http://www.accessecon.com/Pubs/VUECON/VUECON-18-00002.pdf

Feasible Shared Destiny Risk Distributions

Thibault Gajdos, John A. Weymark, and Claudio Zoli

January 2018

Abstract Social risk equity is concerned with the comparative evaluation of social risk distributions, which are probability distributions over the potential sets of fatalities. In the approach to the evaluation of social risk equity introduced by Gajdos, Weymark, and Zoli (Shared destinies and the measurement of social risk equity, Annals of Operations Research 176:409–424, 2010), the only information about such a distribution that is used in the evaluation is that contained in a shared destiny risk matrix whose entry in the kth row and ith column is the probability that person i dies in a group containing k individuals. Such a matrix is admissible if it satisfies a set of restrictions implied by its definition. It is feasible if it can be generated by a social risk distribution. It is shown that admissibility is equivalent to feasibility. Admissibility is much easier to directly verify than feasibility, so this result provides a simply way to identify which matrices to consider when the objective is to socially rank the feasible shared destiny risk matrices.

Keywords social risk evaluation, social risk equity, public risk, shared destinies

JEL classification numbers D63, D81, H43.

Thibault Gajdos
CNRS and Laboratoire de Psychologie Cognitive, Aix-Marseille University, Bâtiment 9 Case D, 3 place Victor Hugo, 13331 Marseille Cedex 3, France
e-mail: thibault.gajdos@univ-amu.fr

John A. Weymark ✉
Department of Economics, Vanderbilt University, VU Station B #35189, 2301 Vanderbilt Place, Nashville, TN 37235-1819, USA
e-mail: john.weymark@vanderbilt.edu

Claudio Zoli
Department of Economics, University of Verona, Via Cantarane 24, 37129 Verona, Italy
e-mail: claudio.zoli@univr.it

1 Introduction

Governments routinely implement policies that affect the risks that a society faces. For example, barriers are installed to lessen the risk of a terrorist driving a vehicle into pedestrians, dikes are built to reduce the risk of flooding, and carbon taxes are imposed to slow down the rise in the temperature of the Earth's atmosphere so as to reduce the likelihood of the serious harms that result from climate change. Policies differ in the degree to which they change the expected aggregate amount of a harm and how it is distributed across the population. A consequentialist approach to evaluating the relative desirability of different policies that affect these kinds of social risks does so by ranking the possible distributions of the resulting harms. If this ranking, or an index representing it, takes account of the equity of the resulting distribution of risks, it is a measure of *social risk equity*.

The measurement of social risk equity has its origins in the work of Keeney (1980a,b,c). The analysis of social risk equity has been further developed by Broome (1982), Fishburn (1984), Fishburn and Sarin (1991), Fishburn and Straffin (1989), Gajdos et al (2010), Harvey (1985), Keeney and Winkler (1985), and Sarin (1985), among others. While these analyses apply to any kind of social harm, for the most part, the harm that they consider is death. Our analysis also applies to any socially risky situation in which a harm may affect some, but not necessarily all, of the society in question but, for concreteness, we, too, suppose that this harm is death.

The set of individuals who die as a result of their exposure to the risk is a fatality set and a social risk distribution is a probability distribution over all of the possible fatality sets. Social risk distributions are ranked using a social risk equity preference ordering. Not all of the information about a social risk distribution may be regarded as being relevant when determining the social preference relation. For example, Fishburn and Straffin (1989), Keeney and Winkler (1985), and Sarin (1985) only take account of the risk profiles for individuals and for fatalities. The former lists the likelihoods of each person dying, whereas the latter is the probability distribution over the number of fatalities. These statistics can be computed from a social risk distribution, but in doing so, some information is lost.

Gajdos et al (2010) propose also taking into account a concern for shared destinies; specifically, with the number of other individuals with whom someone perishes. Chew and Sagi (2012) describe this concern as being one of *ex post* fairness. For example, for a given probability of there being k fatalities, it might be socially desirable to have this risk spread more evenly over the individuals. As Example 3 in Gajdos et al (2010) demonstrates, it is possible for the distribution of how many people someone dies with to differ in two social risk distributions even though the risk profiles for individuals and for fatalities are the same in both distributions. As a consequence, a concern for shared destinies cannot be fully captured if one restricts attention to the information provided by the likelihoods of each person dying and the probability distribution over the number of fatalities.[1]

[1] There are other dimensions of social risk equity that may be of concern, such as dispersive equity and catastrophe avoidance. There is a concern for dispersive equity if account is taken of individual

The distribution of shared destiny risks can be expressed using a shared destiny risk matrix whose entry in the kth row and ith column is the probability that person i dies in a group containing k individuals. In the approach developed by Gajdos et al (2010), if two social risk distributions result in the same shared destiny risk matrix, they are regarded as being socially indifferent. Because the risk profiles for individuals and for fatalities can be computed from the information contained in a shared destiny risk matrix, their approach to evaluating social risks can take account of these two risk profiles, not just a concern for shared destinies. In effect, in their approach to social risk evaluation, ranking social risk distributions is equivalent to ranking shared destiny risk matrices.

The entries of a social destiny risk matrix are probabilities, and so all lie in the interval $[0,1]$. There are three other independent properties that such a matrix must necessarily satisfy as a matter of definition: (i) nobody's probability of dying can exceed 1, (ii) the probability that there are a positive number of fatalities cannot exceed 1, and (iii) nobody can have a probability of dying in a group of size k that exceeds the probability that there are k fatalities. A social destiny risk matrix that satisfies these properties is said to be admissible. Starting with a social risk distribution, we can compute the entries in the corresponding shared destiny risk matrix. A shared destiny risk matrix that can be generated in this way from a social risk distribution is said to be feasible. A feasible shared destiny risk matrix is necessarily admissible. The question we address is whether there are admissible shared destiny risk matrices that are not feasible. We show that there are not. Thus, a shared destiny risk matrix is admissible if and only if it is feasible.

In order to establish this result, we develop an algorithm that shows how to construct a social risk distribution from a shared destiny risk matrix in such a way that the resulting distribution can be used to generate the matrix. It is easy to determine if a shared destiny risk matrix is admissible but, as our algorithm makes clear, confirming that it is also feasible by finding a social risk distribution that generates it may be a formidable undertaking. However, if our objective is only to socially rank the feasible shared destiny risk matrices, our result tells us that this is equivalent to socially ranking the admissible shared destiny risk matrices. We do not need to know how to generate these matrices from social risk distributions in order to know that they are feasible; we only need to know that they are admissible.

In Section 2, we introduce the formal framework used in our analysis. The algorithm employed to determine a social risk distribution that generates a given admissible shared destiny risk matrix is presented in Section 3. We illustrate the operation of this algorithm in Section 4. We prove that a shared destiny risk matrix is admissible if and only if it is feasible in Section 5.

characteristics such as gender, race, or geographic location in addition to the individuals' exposures to social risks. See Fishburn and Sarin (1991) for an analysis of the evaluation of social risks that allows for dispersive equity. Bommier and Zuber (2008), Fishburn (1984), Fishburn and Straffin (1989), Harvey (1985), and Keeney (1980a) consider social preferences for catastrophe avoidance. We do not examine dispersive equity or catastrophe avoidance here.

2 Shared Destiny Risk Matrices

There is a society of $n \geq 2$ individuals individuals who face a social risk. Let $N = \{1, \ldots, n\}$ be the set of these individuals. A *fatality set* is a subset $S \subseteq N$ consisting of the set of individuals who *ex post* die as a consequence of the risk that this society faces. There are 2^n possible fatality sets, including $\varnothing$ (nobody dies) and N (everybody dies). A *social risk distribution* is a probability distribution p on 2^n, with $p(S)$ denoting the *ex ante* probability that the fatality set is S. We suppose that only this probability distribution is relevant for the purpose of social risk evaluation. The set of all such probability distributions is $\mathscr{P}$.

For each $k \in N$, let $M(k) \in [0,1]^n$ denote the vector whose ith component $M(k,i)$ is the *ex ante* probability that person i will die when there are exactly k fatalities. A *shared destiny risk matrix* is an $n \times n$ matrix M whose kth row is $M(k)$. Let $\bar{n}(k)$ be the number of positive entries in $M(k)$. The *risk profile for individuals* is the vector $\alpha \in [0,1]^n$, where

$$\alpha(i) = \sum_{k=1}^{n} M(k,i), \quad \forall i \in N, \tag{1}$$

which is the *ex ante* probability that person i will die. The *risk profile for fatalities* is the vector $\beta \in [0,1]^n$, where

$$\beta(k) = \frac{1}{k} \sum_{i=1}^{n} M(k,i), \quad \forall k \in N, \tag{2}$$

which is the *ex ante* probability that there will be exactly k fatalities. Note that a risk profile for fatalities does not explicitly specify the probability that nobody dies. The probability that there are no fatalities is $1 - \sum_{k=1}^{n} \beta(k)$.

By definition, each of the entries of M is a probability and so must lie in the interval $[0,1]$. Hence, each of components of α and β must be nonnegative as they are sums of entries in M. There are three other restrictions on M. They are

$$\alpha(i) \leq 1, \quad \forall i \in N, \tag{3}$$

$$\sum_{k=1}^{n} \beta(k) \leq 1, \quad \forall k \in N, \tag{4}$$

and

$$M(k,i) \leq \beta(k), \quad \forall (k,i) \in N^2. \tag{5}$$

The first of these requirements is that no person can die with a probability greater than 1. The second is that the probability that there are a positive number of fatalities cannot exceed 1. The third is that nobody's probability of dying in a group of size k can exceed the probability of there being k fatalities. Of course, it must also be the case that

$$\beta(k) \leq 1, \quad \forall k \in N. \tag{6}$$

That is, the probability that there are a particular number of fatalities cannot exceed 1. However, (6) follows from (4) because all probabilities are nonnegative. A social risk equity matrix M is *admissible* if it satisfies (3), (4), and (5).

It is obvious that M must satisfy (3) and (4), but the necessity of (5) is less so because there is more than one way that someone can die with $k-1$ other individuals when $k>1$. To see why (5) is required, suppose, on the contrary, that $M(k,i) > \beta(k)$ for some $i \in N$. Then, because $M(k,i)$ is the probability that person i perishes with $k-1$ other individuals, it must be the case that $\sum_{j\neq i} M(k,j) > (k-1)\beta(k))$. Hence, $\sum_{i=1}^{n} M(k,i) > k\beta(k)$. It then follows that $\beta(k) = \frac{1}{k}\sum_{i=1}^{n} M(k,i) > \beta(k)$, a contradiction.

For each $k \in N$, $\mathscr{T}(k) = \{S \in 2^n || S| = k\}$ is the set of subgroups of the society in which exactly k individuals die. For each $(k,i) \in N^2$, $\mathscr{S}(k,i) = \{S \in \mathscr{T}(k) | i \in S\}$ is the set of subgroups in which exactly k people die and i is one of them. A shared destiny risk matrix M is *feasible* if there exists a social risk distribution $p \in \mathscr{P}$ and an $n \times n$ matrix M_p such that $M(k,i) = M_p(k,i)$, where $M_p(k,i) = \sum_{S\in\mathscr{S}(k,i)} p(S)$. That is, $M_p(k,i)$ is the probability that there are k deaths and i is one of them when the social risk distribution is p.

3 The Decomposition Algorithm

By construction, for any $p \in \mathscr{P}$, M_p is an admissible shared destiny risk matrix. In other words, any feasible shared destiny risk matrix is admissible. The question then arises as to whether feasibility imposes any restrictions on M other than that it be admissible. We show that it does not.

For any admissible shared destiny risk matrix M, we need to show that there exists a social risk distribution $p \in \mathscr{P}$ such that $M_p = M$. This is done by considering each value of k separately. For each $k \in N$, we know that the probability of having this number of fatalities is $\beta(k)$. We need to distribute this probability among the subgroups in $\mathscr{T}(k)$ (the subgroups for which there are k fatalities) in such a way that the probability that person i dies in a group of size k is $M(k,i)$. The resulting probabilities for the subgroups in $\mathscr{T}(k)$ is called a *probability decomposition*. Put another way, for each $i \in N$, we need to distribute the probability $M(k,i)$ among the subgroups in $\mathscr{S}(k,i)$ (the subgroups containing person i for which there are k fatalities) in such a way that the amount π^S allocated to any $S \in \mathscr{S}(k,i)$ is the same for everybody in this group. The value π^S is then the probability that the set of individuals who perish is S.

If $M(k,i) = 0$ for all $i \in N$ (so $\bar{n}(k) = 0$), then $\beta(k) = 0$, so we assign probability 0 to each $S \in \mathscr{T}(k)$. If $k = n$, only N is in $\mathscr{T}(n)$, so no decomposition is needed; N is simply assigned the probability $\beta(n)$. When $\beta(k) \neq 0$ and $k < n$, we construct an algorithm that produces the requisite probability decomposition. The algorithm proceeds through a number of steps, which we denote by $t = 0, 1, 2, \ldots$. We show that the algorithm terminates in no more than $\bar{n}(k)$ steps. The relevant variables in each step are distinguished using a superscript whose value is the step number.

The vector $\hat{M}(k)$ is a *nonincreasing rearrangement* of $M(k)$ if $\hat{M}(k,i) \geq \hat{M}(k,i+1)$ for all $k = 1,\ldots,n-1$. Whenever a vector of probabilities for the n individuals is rearranged in this way, ties are broken in such a way that the original order of the individuals is preserved. For example, if $n = 3$, in the rearrangement $(2,1,1)$ of $(1,2,1)$, the first 1 is associated with person 1 and the second 1 with person 3. Without loss of generality, we suppose that $M(k)$ is initially ranked in nonincreasing order. We now describe our algorithm.

Probability Decomposition Algorithm. The initial values of the relevant variables are

$$M^0(k) = \hat{M}^0(k) = M(k) = \hat{M}(k)$$

and

$$\beta^0(k) = \beta(k).$$

Step 1. In Step 1, we assign a probability π^1 to the first k individuals, which is the set of individuals with the k highest probabilities in $\hat{M}^0(k)$. After π^1 is subtracted from each of the first k components of $\hat{M}^0(k)$, we are left with the fatality probability

$$\beta^1(k) = \beta^0(k) - \pi^1$$

to distribute among the groups of size k using the probabilities in

$$M^1(k) = \hat{M}^0(k) - (\pi^1,\ldots,\pi^1,0,\ldots,0).$$

Letting $\rho^0(k,i)$ denote the rank of individual i in $\hat{M}^0(k)$, we define the vector

$$\pi^1(k) = \pi^1 \cdot (I_1^0, I_2^0, \ldots, I_n^0),$$

where $I_i^0 = 1$ if $\rho^0(k,i) \leq k$ and $I_i^0 = 0$ otherwise. Using $\pi^1(k)$, $M^1(k)$ can be equivalently written as

$$M^1(k) = M^0(k) - \pi^1(k).$$

We need to ensure that each of the probabilities in $M^1(k)$ is nonnegative. Because $\hat{M}^0(k)$ is a nonincreasing rearrangement of $M^0(k)$ and $\hat{M}^1(k,i) = \hat{M}^0(k,i)$ for $i = k+1,\ldots,n$, it must therefore be the case that $\pi^1 \leq \hat{M}^0(k,k)$. We also need to ensure that none of these probabilities exceeds the fatality probability $\beta^1(k)$ left to distribute. This condition is satisfied by construction for the first k individuals. Hence, because $\hat{M}^0(k)$ is a nonincreasing rearrangement of $M^0(k)$, in order to satisfy this condition, it is only necessary that $\hat{M}^0(k,k+1) \leq \beta^1(k) = \beta^0(k) - \pi^1$. Both of these requirements are satisfied by setting

$$\pi^1 = \min\{\hat{M}^0(k,k), \beta^0(k) - \hat{M}^0(k,k+1)\}.$$

By (5), $\hat{M}^0(k,k) \leq \beta^0(k)$. Therefore, $\pi^1 \leq \beta^0(k)$ and, hence, $\beta^1(k) \leq \beta^0(k)$. Because $\beta^1(k) = \frac{1}{k}\sum_{i=1}^n M^1(k,i)$ and $M^1(k,i) \geq 0$ for all $i \in N$, $\beta^1(k) \geq 0$.

Let S^1 denote the first k individuals in $\hat{M}^0(k)$. We choose $p(S^1)$ to be π^1.

If $M^1(k) = (0,0,0,\dots,0)$, the algorithm terminates. Otherwise, it proceeds to the next step.

Step t ($t \geq 2$). The operation of the algorithm in this step follows the same basic logic as in Step 1. The value of π^t is chosen by setting

$$\pi^t = \min\{\hat{M}^{t-1}(k,k), \beta^{t-1}(k) - \hat{M}^{t-1}(k,k+1)\}. \tag{7}$$

Letting $\rho^{t-1}(k,i)$ denote the rank of individual i in $\hat{M}^{t-1}(k)$, we define the vector

$$\pi^t(k) = \pi^t \cdot (I_1^{t-1}, I_2^{t-1}, \dots, I_n^{t-1}), \tag{8}$$

where $I_i^{t-1} = 1$ if $\rho^{t-1}(k,i) \leq k$ and $I_i^{t-1} = 0$ otherwise.

We define $M^t(k)$ and $\beta^t(k)$ by setting

$$M^t(k) = M^{t-1}(k) - \pi^t(k) \tag{9}$$

and

$$\beta^t(k) = \beta^{t-1}(k) - \pi^t. \tag{10}$$

Analogous reasoning to that used in Step 1 shows that $0 \leq \beta^t(k) \leq \beta^{t-1}(k)$.

Let S^t denote the first k individuals in $\hat{M}^{t-1}(k)$. We choose $p(S^t)$ to be π^t.

If $M^t(k) = (0,0,0,\dots,0)$, the algorithm terminates. Otherwise, it proceeds to the next step.

If the algorithm terminates and a group S with k members has not been assigned a probability by the algorithm, we set $p(S) = 0$.

4 Examples of the Probability Decomposition Algorithm

The operation of the probability decomposition algorithm is illustrated with three examples. In each of these examples, it is assumed that M is admissible. In the first example, the algorithm is applied to the case in which nobody dies with anybody else.

Example 1. Let $k = 1$ with $M(1,i) > 0$ for some $i \in N$. If M is feasible, for each $i \in N$, we must have $p(\{i\}) = M(1,i)$. We show that the algorithm produces this result. We first consider the case in which $M(1,i) > 0$ for all $i \in N$.

In Step 1, person 1 is the highest ranked individual in $\hat{M}^0(1)$. Therefore, we have $\beta^0(1) - \hat{M}^0(1,2) = \frac{1}{1}\sum_{i=1}^n \hat{M}^0(1,i) - \hat{M}^0(1,2) = \sum_{i\neq 2}^n \hat{M}^0(1,i) \geq \hat{M}^0(1,1)$. It then follows that $\pi^1 = \hat{M}^0(1,1)$ and, hence, $p(\{1\}) = \hat{M}^0(1,1) = M(1,1)$. We have $\pi^1(1) = (M(1,1),0,\dots,0)$ and so $M^0(1)$ is now replaced with $M^1(1) = (0, M^0(2),\dots,M^0(n))$. Because $M^1(1,1) = 0$, person 1 is never considered again by the algorithm. There is fatality probability $\beta^1(1) = \beta^0(1) - \pi^1 = \sum_{i=1}^n M(1,i) - M(1,1) = \sum_{i=2}^n M(1,i)$ left to allocate.

Step 2 uses the vector $\hat{M}^1(1)$ Person 2 is the second highest ranked individual in $M^0(1)$ and so is first ranked in $\hat{M}^1(1)$. As in Step 1, $\pi^2 = \hat{M}^1(1,1)$ and, hence,

$p(\{2\}) = \hat{M}^1(1,1) = M(1,2)$. We have $\pi^1(1) = (0, M(1,2), 0, \ldots, 0)$ and so $M^1(1)$ is replaced with $M^2(1) = (0, 0, M^0(3), \ldots, M^0(n))$ and person 2 is never considered again. There is fatality probability $\beta^2(1) = \beta^1(1) - \pi^2 = \sum_{i=2}^n M(1,i) - M(1,2) = \sum_{i=3}^n M(1,i)$ left to allocate.

More generally, person $i \in N$ is singled out in Step i and assigned the probability $p(\{i\}) = M(1,i)$. In Step n, $M^n(1) = (0,0,0,\ldots,0)$, and so the algorithm terminates.

If $M(1,i) = 0$ for some $i \in N$, then the algorithm proceeds as above but terminates in Step $\bar{n}(1) < n$, where it is recalled that $\bar{n}(1)$ is the number of individuals for whom $M(1,i)$ is positive. For any individual i for whom $M(1,i) = 0$, when the algorithm terminates, $p(\{i\})$ is set equal to 0.

In the next two examples, individuals do not die alone. In these examples, all probabilities are expressed in terms of percentages, so, for example, 5 is the probability 0.05.

Example 2. Let $n = 7$ and $k = 4$. We suppose that $M^0(4) = M(4) = \hat{M}(4) = \hat{M}^0(4) = (5,4,4,4,4,2,1)$. Consequently, $\beta^0(4) = \beta(4) = \frac{1}{4}\sum_{i=1}^n M(4,i) = 6$.

Step 1. We have $\pi^1 = \min\{\hat{M}^0(4,4), \beta^0(4) - \hat{M}^0(4,5)\} = 2$. Therefore, $\pi^1(4) = (2,2,2,2,0,0,0)$, $M^1(4) = M^0(4) - \pi^1(4) = (3,2,2,2,4,2,1)$, and $\beta^1(4) = \beta^0(4) - \pi^1 = 6 - 2 = 4$. Hence, $p(\{1,2,3,4\}) = \pi^1 = 2$.

Step 2. There are four individuals with the third highest probability in $M^1(4)$. Using our tie-breaking rule, individuals 1, 2, 3, and 5 have the first four probabilities in $\hat{M}^1(4)$. We thus have $\hat{M}^1(4) = (4,3,2,2,2,2,1)$, so $\pi^2 = \min\{\hat{M}^1(4,4), \beta^1(4) - \hat{M}^1(4,5)\} = 2$. Therefore, $\pi^2(4) = (2,2,2,0,2,0,0)$, $M^2(4) = M^1(4) - \pi^2(4) = (1,0,0,2,2,2,1)$, and $\beta^2(4) = \beta^1(4) - \pi^2 = 4 - 2 = 2$. Hence, $p(\{1,2,3,5\}) = \pi^2 = 2$.

Step 3. There are two individuals with the fourth highest probability in $M^2(4)$. The tie is broken in favour of person 1, so the first four individuals in $\hat{M}^2(4)$ are 1, 4, 5, and 6. We have $\hat{M}^2(4) = (2,2,2,1,1,0,0)$, so $\pi^3 = \min\{\hat{M}^2(4,4), \beta^2(4) - \hat{M}^2(4,5)\} = 1$. Therefore, $\pi^3(4) = (1,0,0,1,1,1,0)$, $M^3(4) = M^2(4) - \pi^3(4) = (0,0,0,1,1,1,1)$, and $\beta^3(4) = \beta^2(4) - \pi^3 = 2 - 1 = 1$. Hence, $p(\{1,4,5,6\}) = \pi^3 = 1$.

Step 4. The four individuals with the highest probabilities in $M^3(4)$ are 4, 5, 6, and 7. We have $\hat{M}^3(4) = (1,1,1,1,0,0,0)$, so $\pi^4 = \min\{\hat{M}^3(4,4), \beta^3(4) - \hat{M}^3(4,5)\} = 1$. Therefore, $\pi^4(4) = (0,0,0,1,1,1,1)$, $M^4(4) = M^3(4) - \pi^4(4) = (0,0,0,0,0,0,0)$, and $\beta^4(4) = \beta^3(4) - \pi^4 = 1 - 1 = 0$. Hence, $p(\{4,5,6,7\}) = \pi^4 = 1$.

Because $M^4(4) = (0,0,0,0,0,0,0)$, the algorithm terminates in Step 4. The four groups identified in Steps 1–4 are assigned positive probability. For any other set of individuals S with four members, $p(S) = 0$. There are $\frac{n!}{k!(n-k)!} = \frac{7!}{4!3!} = 35$ possible groups of of this size, so 31 of them are assigned a zero probability.

In Examples 1 and 2, in Step t, π^t is set equal to $\hat{M}^t(k,k)$. In Example 3, it is instead sometimes set equal to $\beta^{t-1}(k) - \hat{M}^{t-1}(k,k+1)$.

Example 3. Let $n = 3$ and $k = 2$. We suppose that $M^0(2) = M(2) = \hat{M}(2) = \hat{M}^0(2) = (7,5,4)$. Consequently, $\beta^0(2) = \beta(2) = \frac{1}{2}\sum_{i=1}^n M(2,i) = 8$.

Step 1. We have $\pi^1 = \min\{\hat{M}^0(2,2), \beta^0(2) - \hat{M}^0(2,3)\} = 4$. Therefore, $\pi^1(2) = (4,4,0)$, $M^1(2) = M^0(2) - \pi^1(2) = (3,1,4)$, and $\beta^1(2) = \beta^0(2) - \pi^1 = 8 - 4 = 4$. Hence, $p(\{1,2\}) = \pi^1 = 4$.

Step 2. The two individuals with the highest probabilities in $M^1(2)$ are 1 and 3. We have $\hat{M}^1(2) = (4,3,1)$, so $\pi^2 = \min\{\hat{M}^1(2,2), \beta^1(2) - \hat{M}^1(2,3)\} = 3$. Therefore, $\pi^2(2) = (3,0,3)$, $M^2(2) = M^1(2) - \pi^2(2) = (0,1,1)$, and $\beta^2(2) = \beta^1(2) - \pi^2 = 4 - 3 = 1$. Hence, $p(\{1,3\}) = \pi^2 = 3$.

Step 3. There only two individuals (2 and 3) left with positive probabilities, and these probabilities are the same. Hence, this group of individuals must be assigned the unallocated fatality probability, so $p(\{2,3\}) = 1$. We confirm that the algorithm produces this result. We have $\hat{M}^2(2) = (1,1,0)$, so $\pi^3 = \min\{\hat{M}^2(1,2), \beta^2(2) - \hat{M}^2(2,2)\} = 1$. Therefore, $\pi^3(1) = (0,1,1)$, $M^3(2) = M^2(2) - \pi^3(2) = (0,0,0)$, and $\beta^3(2) = \beta^2(2) - \pi^3 = 1 - 1 = 0$. Hence, $p(\{2,3\}) = 1$, as was to be shown.

The algorithm terminates in Step 3. All subgroups of with two members are assigned positive probability.

As these examples illustrate, each step of the algorithm identifies a subgroup with k members and determines the probability that it is this group that perishes. For each individual i in this group, this probability must be subtracted from whatever part of the probability $M(k,i)$ that remains unallocated at the end of the previous step. In all three of the examples, at the end of the penultimate step of the algorithm, there is a group of size k whose members all have the same probability left to distribute. In the next section, we show that this is a general feature of the algorithm. When this amount has been allocated as the probability of this group perishing together, we have $M^t(k) = (0,\ldots,0)$, and so the algorithm terminates because, for each $i \in N$, the probability $M(k,i)$ that person i dies in a group of size k has been distributed among each of the groups of size k that include i.

The distribution of the probability in $\beta(k)$ across the groups with k members need not be unique. This is the case in Example 2 because there is more than one way to rearrange the vector of fatality probabilities being considered in a nonincreasing way in some of the steps. For example, if in Step 2 in this example, with the tie-breaking rule used in our algorithm, individuals 1, 2, 3, and 5 are regarded as having the four highest probabilities in $M^1(4)$. However, we could have used a tie-breaking rule that selects individuals 1, 2, 5, and 6 instead, in which case $p(\{(1,2,5,6)\}) > 0$, which is not the case with the tie-breaking rule used in the algorithm. Feasibility of a shared destiny risk matrix M only requires that there there exists a social risk distribution p such that $M_p = M$, not that this distribution be unique.

5 The Equivalence of Admissibility and Feasibility

In order to show that an admissible shared destiny risk matrix is feasible, we first establish a number of lemmas that identify some important properties of the probability decomposition algorithm. In each of our lemmas, we suppose that $k \neq n$ and that the probability decomposition algorithm is being applied to the kth row $M(k)$

of an admissible shared destiny risk matrix M for which the probability $\beta(k)$ that there are k fatalities is positive.

Lemma 1 shows that in each step of this algorithm, analogues of (2) and the admissibility restriction in (5) hold.

Lemma 1. *In any Step t of the algorithm,*

$$\beta^t(k) = \frac{1}{k}\sum_{i=1}^{n} M^t(k,i) \tag{11}$$

and

$$0 \leq M^t(k,i) \leq \beta^t(k), \quad \forall i \in N. \tag{12}$$

Proof. For any $k \neq n$, at the end of Step $t-1$ of the algorithm, from the probability $\beta(k)$ that there will be exactly k fatalities, there is still $\beta^{t-1}(k)$ left to allocate. In Step t, π^t is subtracted from the first k components of $\hat{M}^{t-1}(k)$ and 0 from the other $n-k$ components. Hence, by (2), (9) and (10), at the end of Step t, the amount from $\beta(k)$ left to allocate is (11).

Because $\pi^t \leq \hat{M}^{t-1}(k,k)$, $M^t(k,i) \geq 0$ for all $i \in N$. The argument used to show that $M^t(k,i) \leq \beta^t(k)$ for all $i \in N$ is the same as the argument used in Section 2 to show that (5) holds but with $M^t(k,i)$ substituting for $M(k,i)$ and $\beta^t(k)$ substituting for $\beta(k)$. □

In order for the probability decomposition algorithm to distribute *all* of the probability $\beta(k)$ that there are k fatalities among the subgroups of size k, the algorithm must terminate in a finite number of steps. Lemma 2 shows that this is the case if the algorithm reaches a step in which there are k positive entries left to distribute.

Lemma 2. *The algorithm terminates in Step $t+1$ if there are k positive entries in $M^t(k)$.*

Proof. By Lemma 1, (11) and (12) hold. If $M^t(k)$ contains k positive entries, (11) and (12) imply that they are all equal to $\beta^t(k)$. Thus, the algorithm terminates in the next step because $\pi^t = \beta^t(k)$ is subtracted from $\hat{M}^t(k,i)$ for each $i = 1,\ldots,k$, and so $M^{t+1}(k) = (0,\ldots,0)$. □

There are $\bar{n}(k)$ individuals who have a positive probability of dying in a group of size k. Lemma 3 shows that the probability decomposition algorithm terminates in a finite number of steps that does not exceed this value.

Lemma 3. *The algorithm terminates in at most $\bar{n}(k)$ steps.*

Proof. If $k = \bar{n}(k)$, then Lemma 2 applies with $t = 0$, so the algorithm terminates in Step 1.

Now, suppose that $k < \bar{n}(k)$. From (7), we know that in Step t of the algorithm, π^t is either $\hat{M}^{t-1}(k,k)$ or $\beta^{t-1}(k) - \hat{M}^{t-1}(k,k+1)$, whichever is smallest. We consider two cases distinguished by whether the first of these possibilities holds for all t or not.

Case 1. For each Step t of the algorithm, $\pi^t = \hat{M}^{t-1}(k,k)$. Then, by (7)–(9), $M^t(k)$ has at least one more 0 entry than $M^{t-1}(k)$. Thus, $M^t(k)$ has at least $n - \bar{n}(k) + t$ entries equal to 0 and, hence, has at most k positive entries in Step $\bar{n}(k) - k$. It follows from (11) and (12) (which hold by Lemma 1) that there is no Step t such that the number of positive entries in $M^t(k)$ is positive but less than k. Therefore, because the algorithm subtracts a common positive amount of probability from k individuals in each step, for some $t \leq \bar{n}(k) - k$, $M^t(k)$ has exactly k positive entries, which, by Lemma 2, implies that the algorithm terminates in at most $\bar{n}(k) - k + 1$ steps. Because $k < \bar{n}(k)$, this upper bound is at most $\bar{n}$.

Case 2. In some Step t of the algorithm, $\pi^t \neq \hat{M}^{t-1}(k,k)$. Let t^* be the first step for which this is the case. By (7), we then have that $\pi^{t^*} = \beta^{t^*-1}(k) - \hat{M}^{t^*-1}(k,k+1)$. Let i^* the individual for whom $\rho^{t^*-1}(k,i^*) = k+1$. That is, i^* is the individual for whom $M^{t^*-1}(k,i^*) = \hat{M}^{t^*-1}(k,k+1)$. Because $\pi^{t^*} = \beta^{t^*-1}(k) - \hat{M}^{t^*-1}(k,k+1)$, by (10), $\beta^{t^*}(k) = \hat{M}^{t^*-1}(k,k+1)$. Because $M^{t^*}(k,i^*) = M^{t^*-1}(k,i^*)$, it follows that $M^{t^*}(k,i^*) = \beta^{t^*}(k)$.

By (11) and (12), there cannot be more that k entries in $M^t(k)$ which are at least as large as $\beta^t(k)$. Hence, i^* must occupy one of the first k ranks in $M^{t^*}(k)$ and so i^*'s probability is reduced by π^{t^*} in Step t^*. By (10), for all t, $\pi^t = \beta^{t-1}(k) - \beta^t(k)$. Therefore, $M^{t^*+1}(k,i^*) = \beta^{t^*+1}(k)$. Iteratively applying the same reasoning in each of the subsequent non-terminal steps of the algorithm, we conclude that $M^\tau(k,i^*) = \beta^\tau(k)$ for any Step τ for which $\tau \geq t^*$ which is not a terminal step.

Because there cannot be more that k entries in $M^\tau(k)$ which are at least as large as $\beta^\tau(k)$, we now know that for each $\tau \geq t^*$, i^* has a rank not exceeding k in $M^\tau(k)$. Hence, in any Step t^{**} for which $t^{**} > t^*$, the individual who occupies rank $k+1$ in $\hat{M}^{t^{**}-1}(k)$ is someone, say i^{**}, who is different from i^*. Reasoning as above, if $\pi^{t^{**}} \neq \hat{M}^{t^{**}-1}(k,k)$, then $M^\tau(k,i^{**}) = \beta^\tau(k)$ for any Step τ for which $\tau \geq t^{**}$ which is not a terminal step. Furthermore, both i^* and i^{**} have ranks not exceeding k in $M^\tau(k)$ for any such τ.

By an iterative application of the preceding argument, we conclude that there can be at most k steps in which $\pi^t \neq \hat{M}^{t-1}(k,k)$. Because $M^t(k)$ has at least one more 0 entry than $M^{t-1}(k)$ in each Step t for which $\pi^t = \hat{M}^{t-1}(k,k)$, there are at most $\bar{n}(k) - k - 1$ values of t for which (i) $\pi^t = \hat{M}^{t-1}(k,k)$ and (ii) there are at least $k+1$ positive entries in $M^t(k)$. Thus, the algorithm terminates in at most $\bar{n}(k)$ steps. □

In each step of the algorithm, a group of size k is identified and assigned a probability. Lemma 4 shows that no group is considered in more than one step of the algorithm and, therefore, no group is assigned more than one probability.

Lemma 4. *No group of individuals with k members is assigned a probability in more than one step of the algorithm.*

Proof. We need to show that for all Steps t and t' of the algorithm for which $t \neq t'$, $(I_1^t, I_2^t, \ldots, I_n^t) \neq (I_1^{t'}, I_2^{t'}, \ldots, I_n^{t'})$. On the contrary, suppose that there exist $t < t'$ for which $(I_1^t, I_2^t, \ldots, I_n^t) = (I_1^{t'}, I_2^{t'}, \ldots, I_n^{t'})$. Let S be the set of individuals for whom the value of these indicator functions is 1. Because both π^t and $\pi^{t'}$ are positive, by (7), we must have $\pi^t \geq \pi^t + \pi^{t'}$, which is impossible. That is, both π^t and $\pi^{t'}$ must be subtracted in the same step from the probabilities of the members of S that have yet to be allocated when this group is the one being considered. □

With these lemmas in hand, we can now prove our equivalence theorem.

Theorem. *A shared destiny risk matrix M is admissible if and only it is is feasible.*

Proof. Because a feasible shared destiny risk matrix is necessarily admissible, we only need to show the reverse implication. Suppose that M is an admissible shared destiny risk matrix. For each $k \neq n$ for which $\beta(k) > 0$, Lemmas 3 and 4 imply that the probability decomposition algorithm assigns a probability $p(S) \in [0,1]$ to each group $S \in \mathscr{T}(k)$ (the set of groups with k members) in such a way that $\sum_{S\in\mathscr{T}(k)} p(S) = \beta(k)$. If $k \neq n$ and $\beta(k) = 0$, we let $p(S) = 0$ for all $S \in \mathscr{T}(k)$. Because $\mathscr{T}(n) = \{N\}$, we set $p(N) = \beta(n)$. Finally, we set $p(\varnothing) = 1 - \sum_{i=1}^{n} \beta(k)$. The function $p\colon 2^n \to [0,1]$ is therefore a social risk distribution. By construction, the corresponding shared destiny risk matrix M_p is the same as M because $M_p(k,i) = \sum_{S\in\mathscr{S}(k,i)} p(S) = M(k,i)$ for all (k,i). Hence, M is feasible. □

References

Bommier A, Zuber S (2008) Can preferences for catastrophe avoidance reconcile social discounting with intergenerational equity? Social Choice and Welfare 31:415–434

Broome J (1982) Equity in risk bearing. Operations Research 30:412–414

Chew SH, Sagi J (2012) An inequality measure for stochastic allocations. Journal of Economic Theory 147:1517–1544

Fishburn PC (1984) Equity axioms for public risks. Operations Research 32:901–908

Fishburn PC, Sarin RK (1991) Dispersive equity and social risk. Management Science 37:751–769

Fishburn PC, Straffin PD (1989) Equity considerations in public risks evaluation. Operations Research 37:229–239

Gajdos T, Weymark JA, Zoli C (2010) Shared destinies and the measurement of social risk equity. Annals of Operations Research 176:409–424

Harvey CM (1985) Preference functions for catastrophe and risk inequity. Large Scale Systems 8:131–146

Keeney RL (1980a) Equity and public risk. Operations Research 28:527–534

Keeney RL (1980b) Evaluating alternatives involving potential fatalities. Operations Research 28:188–205

Keeney RL (1980c) Utility functions for equity and public risk. Management Science 26:345–353

Keeney RL, Winkler RL (1985) Evaluating decision strategies for equity of public risks. Operations Research 33:955–970

Sarin RK (1985) Measuring equity in public risk. Operations Research 33:210–217