Università
Ca' Foscari
Venezia

**Department
of Management**

Working Paper Series

**Lorenzo Bastianello and
Marco LiCalzi**

**Target-based solutions for
Nash bargaining**

**Working Paper n. 5/2015
July 2015**

ISSN: 2239-2734

This Working Paper is published under the auspices of the Department of Management at Università Ca' Foscari Venezia. Opinions expressed herein are those of the authors and not those of the Department or the University. The Working Paper series is designed to divulge preliminary or incomplete work, circulated to favour discussion and comments. Citation of this paper should consider its provisional nature.

Target-based solutions for Nash bargaining*

LORENZO BASTIANELLO

[Lorenzo.Bastianello@malix.univ-paris1.fr]

EDE-EM

Université Paris 1-Panthéon-Sorbonne

MARCO LICALZI

[licalzi@unive.it]

Dept. of Management

Università Ca' Foscari Venezia

(July 2015)

Abstract. We revisit the Nash model for two-person bargaining. A mediator knows agents' *ordinal* preferences over feasible proposals, but has incomplete information about their acceptance thresholds. We provide a behavioural characterisation under which the mediator recommends a proposal that maximises the probability that bargainers strike an agreement. Some major solutions are recovered as special cases; in particular, we offer a straightforward interpretation for the product operator underlying the Nash solution.

Keywords: cooperative bargaining, target-based preferences, Nash solution, egalitarian and utilitarian solutions, mediation, copulas.

JEL Classification Numbers: C78, D81, D74.

Correspondence to:

Marco LiCalzi	Department of Management Università Ca' Foscari Venezia S. Giobbe, Cannaregio 873 30121 Venice, Italy
Phone:	[+39] 041-234-6925
Fax:	[+39] 041-234-7444
E-mail:	licalzi@unive.it
Skype:	mlicalzi

* We are grateful to F. Dietrich, F. Durante, M. Marinacci, W. Trockel, and to audiences at RUD 2014, FUR 2014, SING 2015, and at the universities of Bielefeld, Padua, Paris 1, Salento, Salerno, and Turin for their comments. Financial support under PRIN 20103S5RN3 and FP7 n. 318723 (MATHEMACS) is acknowledged.

1 Introduction

The Nash model for two-person bargaining pivots on the assumption that agents are expected utility maximisers. The underlying feasible alternatives are abstracted away by mapping any proposal into a pair of von Neumann-Morgenstern utilities (u_1, u_2) . A Nash bargaining problem (S, d) consists of a compact and convex set $S \subset \mathbb{R}^2$ of feasible utility pairs, and a disagreement point d in S representing the utilities associated with bargaining breakdown. Moreover, the convexity of S is frequently justified by including lotteries among the feasible proposals.

A bargaining solution assigns to any Nash problem (S, d) a single pair of utility values in S . Nash (1950) proved that the unique solution satisfying four axioms is defined as the maximiser of the product $(u_1 - d_1)(u_2 - d_2)$ for (u_1, u_2) in S and $u_i \geq d_i$ for $i = 1, 2$. These four axioms are usually known as symmetry, Pareto optimality, invariance to positive affine transformations, and independence of irrelevant alternatives.

The Nash model is a cornerstone of two-person bargaining theory. Its simplicity and robustness have fostered both its widespread application and its theoretical prominence. A vast body of literature has adopted it, proposing different axiomatizations for the Nash solution as well as several alternative solutions; see f.i. Thomson (1994). Along its many glories, however, not all is well with the model: the Nash solution cannot stake a claim for being intuitively appealing. Rubinstein, Safra, and Thomson (1992) put it very sharply: “*the solution lacks a straightforward interpretation since the meaning of the product of two von Neumann–Morgenstern utility numbers is unclear*”; see Section 4.2.3.

Motivated by this, they review the foundations of the Nash model and offer a very lucid account of its interpretive limitations. Restating the classical utility-based Nash model in terms of agents’ preferences, they offer a more attractive definition of the Nash solution. This conceptual switch to a preference-based language is a key step for reinterpreting the logic underlying the axioms and the solution. However, while their focus on preferences brings substantial theoretical insights, it does not yet uncover an intuitive meaning neither for the product operator nor for the Nash solution.

This paper revisits the Nash model from a related viewpoint. We switch from a utility-based language to a probability-based language. (Specifically, we dispense with most of the formalities of expected utility.) This unlocks several theoretical dividends. We offer a behavioural characterisation for a general class of solutions, equivalent to maximising the probability that the bargainers strike an agreement. This provides a sound underpinning for giving prescriptive advice to a mediator. We also characterise a few major solutions as special cases of this approach, where the single feature separating them is the nature of the stochastic dependence between the bargainers’ stance.

Our probability-based approach suggests a straightforward interpretation for the *product of two von Neumann–Morgenstern utility numbers* advocated by the Nash solution. This is revealed as the product of two probabilities, and corresponds to an implicit assumption of stochastic independence between the bargainers’ positions. We then show how relaxing this assumption generates other well-known but less frequently used alternatives, namely the egalitarian and the (truncated) utilitarian solutions.

A simple example may be useful to elucidate our interpretation of the Nash solution,

leaving generalisations and details to the rest of the paper. Two agents are bargaining over a set A of feasible alternatives, described in physical terms. (The Nash model ignores A and focus on the space of utilities.) Assume that A is a nonempty, compact and connected subset of $\mathbb{R}^n$. Each agent $i = 1, 2$ has an ordinal continuous preference $\succsim_i$ over A . The two agents hire a mediator to suggest a solution and help them strike an agreement. The mediator knows agents' ordinal preferences, but she is not sure what it takes for an agent to accept a proposal x from A .

More formally, we postulate that i accepts a proposal x if and only if $x \succsim_i t_i$, where t_i in A is i 's acceptance threshold (for short, his *target*). The mediator has incomplete information about the bargainers' targets: she believes that each target is a random variable T_i , with a compact and convex support in A . Under her beliefs, she maps each proposal x to a pair of individual acceptance probabilities (p_1, p_2) in $[0, 1]^2$, where $p_i = P(x \succsim_i T_i)$. If the bargainers' targets are stochastically independent, the probability that both accept x is given by the product $p_1 \cdot p_2$ of the individual acceptance probabilities.

The mediator can recommend any feasible alternative, but she cannot impose it: if she suggests x , it is left to the bargainers to accept it. Her goal is to find a proposal x that maximises the probability that agents strike an agreement. Assuming that targets are stochastically independent, she should advance a proposal x that maximises the product $p_1 \cdot p_2$. Hence, the Nash solution may be interpreted as the rule that recommends to maximise the probability to strike an agreement when agents' targets are private information and independently distributed.

This is not merely an analogy. Section 4.2.2 shows that the distribution function $P(x \succsim_i T_i)$ for T_i is formally equivalent to the Bernoulli index function $U_i(x)$ and thus we can set $P(x \succsim_i T_i) = U_i(x)$. The Nash solution requires to maximise the product of two numbers, and they can be equivalently interpreted as utilities or probabilities. The existing axiomatizations are framed in a utility-based language for which the product operator is a puzzle. Switching to a probability-based language uncovers a straightforward interpretation.

The rest of the paper is organised as follows. Section 2 provides preliminary information on bivariate copulas. Section 3 describes our model, offers a general characterisation for preferences over bargaining solutions, as well as an axiomatization for the Nash solution, the egalitarian solution, and the (truncated) utilitarian solution. Section 4 reviews related literature, and offers a commentary on our results. Section 5 illustrates applications and extensions, as well as some testable restrictions on the model.

2 Preliminaries

Copulas are functions that link multivariate distributions to their one-dimensional marginal distributions; see Nelsen (2006). They are used to model different forms of statistical dependence and construct families of distributions exhibiting them. We focus on the prominent case of bivariate copulas.

A (bivariate) copula is a function $C : [0, 1]^2 \rightarrow [0, 1]$ that satisfies two properties:

- C1) for any p, q in $[0, 1]$, $C(p, 0) = C(0, q) = 0$, $C(p, 1) = p$, and $C(1, q) = q$;
- C2) for any $p_1 > q_1$ and $p_2 > q_2$ in $[0, 1]$, $C(p_1, p_2) + C(q_1, q_2) \geq C(p_1, q_2) + C(q_1, p_2)$.

Property C2 is usually called 2-increasingness in the literature, but is of course equivalent to supermodularity. We use this latter name because it is presumably more familiar for our readers. The combination of C1 and C2 implies that $C(p, q)$ is increasing in each argument; see Lemma 2.1.4 in Nelsen (2006). When the weak inequality in C2 is replaced by a strict one, the copula C is said to be strictly supermodular and it is strictly increasing in each argument.

The following result is in Sklar (1959) and characterizes how a copula links the bivariate distribution to its univariate marginals.

Theorem 1. *Let (X, Y) be a random vector with marginal distributions $F(x)$ and $G(y)$. The following are equivalent:*

- i) $H(x, y)$ is the joint distribution function of (X, Y) ;
- ii) there exists a copula $C(p, q)$ such that $H(x, y) = C[F(x), G(y)]$ for all x, y .

If $F(x)$ and $G(y)$ are continuous, then $C(p, q)$ is unique. Otherwise, $C(p, q)$ is uniquely defined on the cartesian product $\text{Ran}(F) \times \text{Ran}(G)$ of the ranges of the two marginal distributions. Conversely, if $C(p, q)$ is a copula and $F(x)$ and $G(y)$ are distribution functions, then the function $H(x, y)$ defined above is a joint distribution function with margins $F(x)$ and $G(y)$.

The best known example of a copula is the product $\Pi(p, q) = p \cdot q$, associated with stochastic independence. Two other important examples are $W(p, q) = \max(p + q - 1, 0)$ and $M(p, q) = \min(p, q)$. For any copula $C(p, q)$ and any (p, q) in $[0, 1]^2$, it is the case that $W(p, q) \leq C(p, q) \leq M(p, q)$. Intuitively, M is the copula associated with the strongest possible positive dependence between X and Y , given the marginal distributions F and G ; similarly, W describes the strongest possible negative dependence. The copulas W and M are known as the Fréchet lower and upper bound, respectively.

3 Model and results

In its simplest version, our model for two-person bargaining is as abstract as the Nash model. We postulate that each feasible proposal x is mapped to a pair of probabilities (p_1, p_2) . At this stage, it suffices to think of p_i as the individual acceptance probability that a third-party called the mediator attributes to Agent $i = 1, 2$ when he is offered the proposal x . Section 4.2 discusses two compatible interpretations for this mapping.

For our purposes, a *bargaining problem* is represented by a compact set B in $[0, 1]^2$ where each point $\mathbf{p}$ in B corresponds to a pair of (acceptance) probabilities. A *solution* is a map that for any problem B delivers (at least) one point in B .

We consider the preferences of the mediator over the set of lotteries on pairs of acceptance probabilities, and derive a behavioural characterisation under which she evaluates a proposal by the probability that both bargainers agree to it. More formally, we assume that the mediator has a preference relation $\succsim$ on the lotteries on $[0, 1]^2$, and provide a representation theorem under which the solution for B corresponds to a $\succsim$ -maximal point in B .

3.1 Assumptions on preferences

We denote an element (p_1, p_2) in $[0, 1]^2$ by $\mathbf{p}$. We view $[0, 1]^2$ as a mixture space for the $\oplus$ operation, under the standard interpretation where $\alpha\mathbf{p} \oplus (1 - \alpha)\mathbf{q}$ is a lottery that delivers $\mathbf{p}$ in $[0, 1]^2$ with probability α in $[0, 1]$ and $\mathbf{q}$ in $[0, 1]^2$ with the complementary probability $1 - \alpha$; see Herstein and Milnor (1953). Moreover, let $\mathbf{p} \vee \mathbf{q} = (\max(p_1, q_1), \max(p_2, q_2))$ and $\mathbf{p} \wedge \mathbf{q} = (\min(p_1, q_1), \min(p_2, q_2))$ denote the standard lattice-theoretical join and meet for the usual component-wise monotonic partial ordering $\geq$ in $\mathbb{R}^2$.

We make the following assumptions about the mediator's preference $\succsim$ over $[0, 1]^2$, where $\succ$ and $\sim$ have the usual meaning. For simplicity, we write “for sure” instead of the more accurate “with probability 1”.

A.1 (Regularity) $\succsim$ is a complete preorder, continuous and mixture independent.

This implies that there exists a real-valued function $V : [0, 1]^2 \rightarrow [0, 1]$, unique up to positive affine transformations, that represents $\succsim$ and is linear with respect to $\oplus$; that is, $V(\alpha\mathbf{p} \oplus (1 - \alpha)\mathbf{q}) = \alpha V(\mathbf{p}) + (1 - \alpha)V(\mathbf{q})$, for any α in $[0, 1]$ and any $\mathbf{p}, \mathbf{q}$ in $[0, 1]^2$. See Theorem 8.4 in Fishburn (1970). Quite interestingly, Nash (1950, p. 157) explicitly points out how an analog of A.1 is implied in his model by the assumption that both bargainers are expected utility maximizers.

A.2 (Non-triviality) $(1, 1) \succ (0, 0)$.

This rules out the trivial case where the mediator is indifferent between a proposal that is accepted for sure by both bargainers and another proposal that is refused for sure by both bargainers.

A.3 (Disagreement indifference) for any p, q in $[0, 1]$, $(p, 0) \sim (0, q)$.

This is named after Assumption DI in Border and Segal (1997), who also study a preference relation over solutions. Their paper is discussed in Section 5.2. Framed within the Nash model, Assumption DI states the following: a solution that assigns to either player the same utility he gets at the disagreement point is as good as the disagreement point itself. In simple words, a solution that gives one player the worst individually rational outcome is equivalent to a solution that gives both bargainers the same utility as the disagreement point. In our probability-based framework, it states that having one of the bargainers refusing for sure is equivalent to having both refusing for sure. A proposal is accepted if and only if both bargainers agree to it.

A.4 (Consistency over individual probabilities) for any p in $[0, 1]$,

$$p(1, 1) \oplus (1 - p)(0, 1) \sim (p, 1) \quad \text{and} \quad p(1, 1) \oplus (1 - p)(1, 0) \sim (1, p).$$

This states the following. Assume that one bargainer is known to accept for sure. Then the mediator is indifferent between a lottery that has the second bargainer accepting for sure with probability p and refusing for sure with probability $(1 - p)$, or a proposal where the second bargainer accepts with probability p . Intuitively, the first lottery has an “objective” probability p of success, while the second proposal has a “subjective” probability with the same value p . We assume that the mediator is indifferent between the two modalities.

Technically speaking, the assumption aligns the marginal acceptance probabilities with the corresponding joint acceptance probability if one of the two bargainers accepts for sure.

A.5 (Weak complementarity) for any $\mathbf{p}, \mathbf{q}$ in $[0, 1]^2$,

$$(1/2)(\mathbf{p} \vee \mathbf{q}) \oplus (1/2)(\mathbf{p} \wedge \mathbf{q}) \succsim (1/2)\mathbf{p} \oplus (1/2)\mathbf{q}$$

This is named after Axiom S in Francetich (2013). It states that a fifty-fifty lottery between two pairs of acceptance probabilities $\mathbf{p}$ and $\mathbf{q}$ is weakly inferior to a fifty-fifty lottery between their extremes (under the component-wise ordering). The interpretation is the following. Suppose $p_1 \geq q_1$ and $q_2 \geq p_2$. When the individual acceptance probabilities improve from (q_1, p_2) to (p_1, p_2) , the increase in the probability of success for a proposal cannot be greater than when they change from (q_1, q_2) to (p_1, q_2) . Whatever advantage is gained when the first bargainer's acceptance probability increases by $p_1 - q_1$, it adds more to the probability of success when the second bargainer is more likely to accept. In simple words, the individual acceptance probabilities are (weakly) complementary towards getting to an agreement. For a particularly sharp illustration, let $p_1 = q_2 = 1$ and $p_2 = q_1 = 0$: clearly, joint acceptance occurs only at $(1, 1)$, and a fifty-fifty lottery between $(1, 1)$ and $(0, 0)$ is strictly better than a fifty-fifty lottery between $(1, 0)$ and $(0, 1)$.

We show by example in Section 3.6 that, under A.1, the four assumptions A.2–A.5 are logically independent.

3.2 A general characterisation

Our first result gives a behavioural characterisation for the preferences of the mediator. Under A1–A5, there exists a unique copula that represents $\succsim$. Given the marginal distributions, any copula identifies a joint probability distribution consistent with them; see Section 2 and Nelsen (2006). For any possible dependence structure linking the marginals, there is a copula that describes it.

Theorem 2. *The preference relation $\succsim$ satisfies A.1–A.5 if and only if there exists a unique copula $C : [0, 1]^2 \rightarrow [0, 1]$ that represents $\succsim$, in the sense that*

$$\mathbf{p} \succsim \mathbf{q} \quad \text{if and only if} \quad C(\mathbf{p}) \geq C(\mathbf{q}).$$

Proof. Necessity being obvious, we prove only sufficiency. By A.1, the Mixture Space Theorem (Herstein and Milnor, 1953) implies that there exists a unique (up to positive affine transformations) function $V : [0, 1]^2 \rightarrow \mathbb{R}$ that represents $\succsim$ and is linear with respect to $\oplus$. By A.2, $V(1, 1) - V(0, 0) > 0$. Apply the appropriate positive affine transformation and consider the (unique) function $C : [0, 1]^2 \rightarrow [0, 1]$ defined by

$$C(p, q) = \frac{V(p, q) - V(0, 0)}{V(1, 1) - V(0, 0)}.$$

We show that C satisfies the two defining properties C1–C2 of a copula given in Section 2.

By A.3, we have $C(p, 0) = C(0, p) = C(0, 0) = 0$, for any p in $[0, 1]$. Moreover, clearly $C(1, 1) = 1$. By A.4 and linearity, for any p in $[0, 1]$, we get $C(p, 1) = C(p(1, 1) \oplus (1 -$

$p)(0, 1)) = pC(1, 1) + (1 - p)C(0, 1) = p$; a similar argument shows that $C(1, p) = p$. This proves C1.

By A.5 and the linearity of C , for all $\mathbf{p}, \mathbf{q}$ in $[0, 1]^2$, it follows that

$$\frac{1}{2}C(\mathbf{p} \vee \mathbf{q}) + \frac{1}{2}C(\mathbf{p} \wedge \mathbf{q}) = C\left(\frac{1}{2}(\mathbf{p} \vee \mathbf{q}) \oplus \frac{1}{2}(\mathbf{p} \wedge \mathbf{q})\right) \geq C\left(\frac{1}{2}\mathbf{p} \oplus \frac{1}{2}\mathbf{q}\right) = \frac{1}{2}C(\mathbf{p}) + \frac{1}{2}C(\mathbf{q}),$$

so that C is supermodular, and C2 holds. $\square$

In our setup, this result has a straightforward interpretation. A pair (p, q) of acceptance probabilities in B represents the individual probability that each bargainer accepts the underlying proposal. When the mediator's preferences satisfy axioms A.1–A.5, she behaves as if she aggregates these individual probabilities by consistently using a (unique) copula C and computes the joint probability $C(p, q)$ that the bargainers strike an agreement. Since the copula is arbitrary, the mediator may entertain any subjective opinion regarding the dependence structure (as embedded in the copula) between the individual acceptance probabilities. The proposals in B are ranked accordingly to the resulting joint acceptance probability.

A $\succsim$ -maximal element in B is a choice that maximises the probability that the bargainers accept the underlying proposal and strike an agreement. Since any copula is necessarily Lipschitz continuous and B is compact, the set of $\succsim$ -maximal elements in B is not empty and a solution exists. On the other hand, our assumptions do not imply its uniqueness. In general, the solution in B corresponds to an equivalence class of pairs of individual acceptance probabilities (including lotteries over those) for which the mediator assesses the same joint acceptance probability. It is worth noting that B is not required to be convex or even connected.

The copula representing the preferences of the mediator in Theorem 2 is linear with respect to $\oplus$, in the sense that $C(\alpha\mathbf{p} \oplus (1 - \alpha)\mathbf{q}) = \alpha C(\mathbf{p}) + (1 - \alpha)C(\mathbf{q})$. When the mediator evaluates the probability of success for a lottery, she assesses first the probabilities of success for $\mathbf{p}$ and $\mathbf{q}$ through $C(\mathbf{p})$ and $C(\mathbf{q})$, and then she mixes them with the same weights defining the lottery. On the contrary, our assumptions do not imply linearity with respect to convex combinations of points in $[0, 1]^2$ based on the standard $+$ operator.

Two simple variations on Theorem 2 are worth mentioning. Recall that a copula $C(p, q)$ is (weakly) increasing in each argument. Therefore, the preference relation $\succsim$ does not violate the (weak) Pareto ordering: if $p_1 \geq q_1$ and $p_2 \geq q_2$, then $(p_1, p_2) \succsim (q_1, q_2)$. But it may not satisfy the (strong) Pareto ordering, under which the condition $(p_1 - q_1)(p_2 - q_2) > 0$ implies $(p_1, p_2) \succ (q_1, q_2)$. Consider a mild strengthening of A.5, where we write $\mathbf{p} \bowtie \mathbf{q}$ to indicate that $\mathbf{p} \neq \mathbf{q}$ and that $\mathbf{p}$ and $\mathbf{q}$ are not comparable with respect to the usual partial ordering $\geq$.

A.5* (Complementarity) for any $\mathbf{p}, \mathbf{q}$ in $[0, 1]^2$,

$$(1/2)(\mathbf{p} \vee \mathbf{q}) \oplus (1/2)(\mathbf{p} \wedge \mathbf{q}) \succsim (1/2)\mathbf{p} \oplus (1/2)\mathbf{q}.$$

Moreover, if $\mathbf{p} \bowtie \mathbf{q}$,

$$(1/2)(\mathbf{p} \vee \mathbf{q}) \oplus (1/2)(\mathbf{p} \wedge \mathbf{q}) \succ (1/2)\mathbf{p} \oplus (1/2)\mathbf{q}.$$

In combination with the other assumptions, this rules out “thick” indifference curves and makes $\succsim$ consistent with the (strong) Pareto ordering: if $p_1 \geq q_1$ and $p_2 \geq q_2$, then $(p_1, p_2) \succsim (q_1, q_2)$ and, moreover, $(p_1, p_2) \succ (q_1, q_2)$ if $(p_1 - q_1)(p_2 - q_2) > 0$. This follows from the next result, because any strictly supermodular copula is strictly increasing in each argument.

Theorem 3. *The preference relation $\succsim$ satisfies A.1–A.5* if and only if there exists a unique strictly supermodular copula $C : [0, 1]^2 \rightarrow [0, 1]$ that represents $\succsim$.*

A second variation embodies an elementary notion of fairness.

A.6 (Anonymity) for any p, q in $[0, 1]$, $(p, q) \sim (q, p)$.

This states that the evaluation for any pair (p, q) of individual acceptance probabilities is unaffected by permutations, and hence is anonymous with respect to the bargainers’ identity. The following result is immediate.

Theorem 4. *The preference relation $\succsim$ satisfies A.1–A.5 and A.6 if and only if there exists a unique symmetric copula $C : [0, 1]^2 \rightarrow [0, 1]$ that represents $\succsim$.*

Symmetric copulas are the most frequently studied class, and Anonymity seems to be a natural requirement. However, it remains a special case. The analog of some solutions (including the Nash solution) require Complementarity and imply Anonymity, and hence are associated with strictly supermodular and symmetric copulas.

Theorem 2 shows that we can characterise different solution concepts for cooperative bargaining as the aggregation (via copula) of two individual acceptance probabilities into a joint probability of success. As a special case, the following Theorem 5 proves that stochastic independence is the key assumption for deriving the product operator underlying the Nash solution.

Since each copula models a different dependence structure, other solutions for cooperative bargaining may be recovered under alternative assumptions. In particular, the two extreme assumptions of maximal positive dependence and maximal negative dependence between the individual acceptance probabilities bring about the egalitarian solution and a variant of the (truncated) utilitarian solution, discussed respectively in Sections 3.4 and 3.5.

3.3 The Nash solution

The Nash solution is a special case of Theorem 2 when the copula chosen by the mediator presumes stochastic independence among the individual acceptance probabilities. That is, the Nash solution emerges whenever we assume that these individual probabilities are independent. This seems by far a very natural requirement, and in our view it gives the Nash solution a central position among the special cases.

Under A.5*, we need only one additional assumption to characterise the Nash solution.

A.7 (Rescaling indifference) for any α, p, q in $[0, 1]$, $(\alpha p, q) \sim (p, \alpha q)$.

This states that the mediator is indifferent whether the same proportional reduction in the acceptance probability is applied to one bargainer or to the other one. The probability

of joint acceptance is equally affected when downsizing (by the same factor) the individual propensity to accept of either bargainer. Clearly, A.7 implies A.3 (Disagreement indifference) and A.6 (Anonymity).

Theorem 5. *The preference relation $\succsim$ satisfies A.1–A.2, A.4–A.5*, and A.7 if and only if it is represented by the copula $\Pi(p, q) = p \cdot q$.*

Proof. Necessity is obvious. Sufficiency follows if we show that all the indifference curves for the representing copula are hyperbolas. First, suppose that $\mathbf{p} = (p_1, p_2)$ and $\mathbf{q} = (q_1, q_2)$ satisfy $p_1 p_2 = q_1 q_2$; without loss of generality, assume $p_1 > q_1$ and $p_2 < q_2$. Consider $\mathbf{p} \vee \mathbf{q} = (p_1, q_2)$ and let $\alpha = q_1/p_1 = q_2/p_2 < 1$. By A.7, we obtain $(\alpha p_1, q_2) \sim (p_1, \alpha q_2)$ and, substituting for α , we find $\mathbf{q} = (q_1, q_2) \sim (p_1, p_2) = \mathbf{p}$. Therefore, two points on the same hyperbola are indifferent.

Conversely, suppose without loss of generality that $p_1 p_2 > q_1 q_2$. We show that $(p_1, p_2) \succ (q_1, q_2)$. Let $\alpha = (q_1 q_2)/(p_1 p_2) < 1$. By strong Pareto dominance, $(p_1, p_2) \succ (\alpha p_1, p_2)$. Since the point $(\alpha p_1, p_2)$ lies on the same hyperbola as (q_1, q_2) , by the first part of this proof $(\alpha p_1, p_2) \sim (q_1, q_2)$. Hence, $(p_1, p_2) \succ (\alpha p_1, p_2) \sim (q_1, q_2)$. $\square$

This result uncovers an appealing interpretation for the product operator underlying the Nash solution for cooperative bargaining. When the bargaining problem is framed with respect to pairs of acceptance probabilities (instead of utilities), the product operator is the natural consequence of the assumption that the contribution of these individual probabilities to the joint acceptance probability satisfies stochastic independence.

The key behavioural implication of Theorem 5 is the following. Suppose that the mediator evaluates that the probability of joint acceptance for a proposal (p, q) is π . For any α in $[0, 1]$, by the linearity of C , the mediator attributes a probability of success $\alpha\pi$ to the lottery $\alpha(p, q) \oplus (1 - \alpha)\mathbf{0}$. Intuitively, the introduction of the α -randomisation reduces the probability π of success by a factor α . The reduction is multiplicative because the randomising device is (tacitly assumed as) stochastically independent. Under the Nash copula, any α -reduction to either of the individual acceptance probabilities has the same effect on the mediator's evaluation as an α -randomisation: $\alpha(p, q) \oplus (1 - \alpha)\mathbf{0} \sim (\alpha p, q) \sim (p, \alpha q)$. It is immaterial whether the reduction comes from an (objective) lottery or from a (subjective) assessment.

3.4 The egalitarian solution

In the utility-based Nash model, the egalitarian solution (Kalai, 1977) recommends the maximal point at which utility gains from the disagreement point d are equal. More simply, for a Nash problem (S, d) , the egalitarian solution selects the maximiser of the function $\min\{(u_1 - d_1), (u_2 - d_2)\}$ for (u_1, u_2) in S and $u_i \geq d_i$ for $i = 1, 2$.

In our probability-based formulation, this translates to the requirement that the solution for a bargaining problem B is the (set of) maximiser(s) of the function $\min(p_1, p_2)$ for $\mathbf{p}$ in B . Consider the following assumption.

A.8 (Meet indifference) for any p, q in $[0, 1]$, $(p, p \wedge q) \sim (p \wedge q, q)$.

This states that the mediator is indifferent between two pairs of acceptance probabilities as far as they have the same meet. Intuitively, preferences over (p, q) depend only on the

smallest value between p and q . Clearly, A.8 implies A.3 (Disagreement indifference) and A.6 (Anonymity). The manifest analogies between A.7 and A.8 reappear in the formulation of the following result, whose proof is similar and thus can be omitted.

Theorem 6. *The preference relation $\succsim$ satisfies A.1–A.2, A.4–A.5*, and A.8 if and only if it is represented by the copula $M(p, q) = \min(p, q)$.*

This characterises the representing copula under Meet indifference as the Fréchet upper bound $M(p, q) = \min(p, q)$, that provides the strongest possible positive dependence between two marginal distributions. Therefore, we can reinterpret the egalitarian solution as the recommendation that maximises the probability of joint acceptance when the mediator assumes that the individual acceptance probabilities are maximally positively dependent.

3.5 The utilitarian solution

There exist alternative formulations for the utilitarian solutions in the Nash model. They share the general principle that the solution should recommend an alternative that maximises the sum of utilities, or of utility increments over the disagreement point. The main obstacle impeding a unified definition is that utilities are defined only up to positive affine transformations. We consider relative utilitarianism, that normalises the individual utilities to have infimum zero and supremum one before considering their sum. This reasonable normalisation choice was first considered in Arrow (1963); see Dhillon and Mertens (1999).

In our probability-based framework, the normalisation issue is immaterial and we can simply map the utilitarian precept into the (still) generic recommendation of maximising the sum of individual acceptance probabilities. Clearly, if we are to reinterpret this sum as a probability of joint success, this generic recommendation needs to be suitably qualified. Consider the following assumption.

A.9 (Average indifference) for any p, q in $[0, 1]$, $(p, q) \sim (\frac{p+q}{2}, \frac{p+q}{2})$.

Consider all pairs of acceptance probabilities on the segment between (p, q) and $(\frac{p+q}{2}, \frac{p+q}{2})$. As we move inward towards the bisector, the components are “less spread out” and one individual acceptance probability decreases at the expense of the other. Assumption A.9 states that the mediator is indifferent among all these pairs of acceptance probabilities, because the increase of one exactly compensates the diminution of the other. Intuitively, the two individual probabilities behave as substitutes towards the joint probability of acceptance. Clearly, this is at odds with A.5*, but it is compatible with A.5. Moreover, A.9 implies A.6 (Anonymity) but not A.3 (Disagreement indifference).

Theorem 7. *The preference relation $\succsim$ satisfies A.1–A.5 and A.9 if and only if it is represented by the copula $W(p, q) = \max(p + q - 1, 0)$.*

Proof. Necessity is obvious. As for sufficiency, by Theorem 2 there exists a copula C representing $\succsim$. It is known that W is the only quasi-convex copula; see Example 3.27 in Nelsen (2006). Hence, it suffices to show that A.9 implies that C is quasi-convex; that is, for any α in $(0, 1)$ and $\mathbf{p}, \mathbf{q}$ in $[0, 1]^2$, we have $C(\alpha\mathbf{p} + (1 - \alpha)\mathbf{q}) \leq \max\{C(\mathbf{p}), C(\mathbf{q})\}$.

For ease of notation, given $\mathbf{p} = (p_1, p_2)$, let $\bar{\mathbf{p}} = (\frac{p_1+p_2}{2}, \frac{p_1+p_2}{2})$ denote its symmetrised counterpart lying on the main diagonal. Then A.9 states that $\mathbf{p} \sim \bar{\mathbf{p}}$ and thus $C(\mathbf{p}) = C(\bar{\mathbf{p}})$. Assume without loss of generality that $q_1 + q_2 \leq p_1 + p_2$. Since C is increasing over the main diagonal, we obtain

$$C(\alpha\mathbf{p} + (1-\alpha)\mathbf{q}) = C(\overline{\alpha\mathbf{p} + (1-\alpha)\mathbf{q}}) \leq C(\bar{\mathbf{p}}) = C(\mathbf{p}),$$

and thus C is quasi-convex. $\square$

This characterises the representing copula under Average indifference as the Fréchet lower bound $W(p, q) = \max(p+q-1, 0)$, that provides the strongest possible negative dependence between two marginal distributions. Therefore, we can reinterpret this form of (truncated) utilitarian solution as the recommendation that maximises the probability of joint acceptance when the mediator assumes that the individual acceptance probabilities are maximally negatively correlated.

This interpretation requires a comment. The copula $W(p, q)$ is strongly Pareto increasing on the triangle above the diagonal from $(0, 1)$ to $(1, 0)$, and is zero on the rest of its domain. If the mediator's preferences are represented by this copula, she behaves as an utilitarian for all pairs above the diagonal, and is indifferent for the pairs below the diagonal, suggesting that she may violate (strong) Pareto dominance. The reason for the mediator's indifference is that, under the assumptions of Theorem 7, she opines that any feasible proposal mapping to a pair (p, q) below the diagonal will be refused for sure. Hence, from her viewpoint, it is neither better nor worse than $\mathbf{0}$.

3.6 Logical independence of the assumptions

This section provides simple examples to show that, under A.1, the four assumptions A.2–A.5 used in Theorem 2 are logically independent. Recall that A.1 implies the existence of a real-valued function $V : [0, 1]^2 \rightarrow [0, 1]$, unique up to positive affine transformations, that represents $\succsim$ and is linear with respect to $\oplus$. We recall each assumption and list the associated counterexample immediately after. We omit quantifiers when they are obvious.

A.2 (Non-triviality) $(1, 1) \succ (0, 0)$.

Consider $V(p, q) = k$, for some constant k in $[0, 1]$. Then A.3 holds because $V(p, 0) = k = V(0, q)$. A.4 holds because $pV(1, 1) + (1-p)V(0, 1) = k = V(p, 1)$, and similarly for the second relation. And A.5 holds because $(1/2)V(\mathbf{p} \vee \mathbf{q}) + (1/2)V(\mathbf{p} \wedge \mathbf{q}) = k = (1/2)V(\mathbf{p}) + (1/2)V(\mathbf{q})$. However, A.2 does not hold because $V(1, 1) = k = V(0, 0)$.

A.3 (Disagreement indifference) for any p, q in $[0, 1]$, $(p, 0) \sim (0, q)$.

Consider $V(p, q) = p$. Then A.2 holds because $V(1, 1) = 1 > 0 = V(0, 0)$. A.4 holds because $pV(1, 1) + (1-p)V(0, 1) = p = V(p, 1)$, and $pV(1, 1) + (1-p)V(1, 0) = 1 = V(1, p)$. And A.5 holds because

$$\frac{1}{2}V(\mathbf{p} \vee \mathbf{q}) + \frac{1}{2}V(\mathbf{p} \wedge \mathbf{q}) = \frac{1}{2}(p_1 \vee q_1) + \frac{1}{2}(p_1 \wedge q_1) = \frac{1}{2}p_1 + \frac{1}{2}q_1 = \frac{1}{2}V(\mathbf{p}) + \frac{1}{2}V(\mathbf{q})$$

However, A.3 does not hold: for $p > 0$ and any q , we have $V(p, 0) = p > 0 = V(0, q)$.

A.4 (Consistency over individual probabilities) for any p in $[0, 1]$,

$$p(1, 1) \oplus (1 - p)(0, 1) \sim (p, 1) \quad \text{and} \quad p(1, 1) \oplus (1 - p)(1, 0) \sim (1, p).$$

Consider $V(p, q) = p^2q$. Then A.2 holds because $V(1, 1) = 1 > 0 = V(0, 0)$. A.3 holds because $V(p, 0) = 0 = V(0, q)$. A.5 holds because the first mixed derivative of V is positive, and hence V is supermodular. However, A.4 does not hold because $pV(1, 1) + (1 - p)V(0, 1) = p > p^2 = V(p, 1)$.

A.5 (Weak complementarity) for any $\mathbf{p}, \mathbf{q}$ in $[0, 1]^2$,

$$(1/2)(\mathbf{p} \vee \mathbf{q}) \oplus (1/2)(\mathbf{p} \wedge \mathbf{q}) \succsim (1/2)\mathbf{p} \oplus (1/2)\mathbf{q}$$

Consider the function

$$V(p, q) = \begin{cases} \min(p, q, \frac{1}{3}, p + q - \frac{2}{3}) & \text{if } \frac{2}{3} \leq p + q \leq \frac{4}{3} \\ \max(p + q - 1, 0) & \text{otherwise,} \end{cases}$$

borrowed from Exercise 2.11 in Nelsen (2006). Then A.2 holds because $V(1, 1) = 1 > 0 = V(0, 0)$. A.3 holds because $V(p, 0) = 0 = V(0, q)$. A.4 holds because $pV(1, 1) + (1 - p)V(0, 1) = p = V(p, 1)$, and similarly for the second relation. But A.5 does not hold: let $\mathbf{p} = (1/3, 2/3)$ and $\mathbf{q} = (2/3, 1/3)$, so that $\mathbf{p} \vee \mathbf{q} = (2/3, 2/3)$ and $\mathbf{p} \wedge \mathbf{q} = (1/3, 1/3)$. Then $V(\mathbf{p} \vee \mathbf{q}) + V(\mathbf{p} \wedge \mathbf{q}) - V(\mathbf{p}) - V(\mathbf{q}) = -1/3 < 0$, contradicting supermodularity.

4 Commentary

This section discusses the model and two interpretations of the results presented in Section 3, contrasting them with the related literature.

4.1 Fundamentals

The Nash model is an abstraction of real bargaining situations. As mentioned, Rubinstein, Safra, and Thomson (1992) — from now on, RST — have re-examined the Nash model, moving away from the utility-based language of the Nash model towards the fundamentals of a two-person bargaining problem. This work indirectly provides a foundation for the Nash model that is independent of the assumption that bargainers maximise expected utility. We provide a similar description of the fundamentals for our model, and argue that they nest RST's formulation.

There are two agents and a set A of feasible alternatives, described in physical terms and viewed as deterministic outcomes. The set A is a nonempty compact subset of a connected topological space X . Each agent $i = 1, 2$ has an ordinal preference order $\succsim_i$ over X , and hence over A . Preferences are continuous: therefore, there exist ordinal (i.e., unique up to increasing transformations) and continuous value functions $v_i(x)$ such that $x \succsim_i y$ if and only if $v_i(x) \geq v_i(y)$. Moreover, agents' preferences are (jointly) non-trivial: there are alternatives

x, y in A such that $x \succ_i y$ for both i . We add a final simplifying assumption: there are no alternatives x, y in A such that $x \sim_i y$ for both i . This avoids the need to rephrase definitions and results in terms of equivalence classes.

A *native bargaining problem* is a triple $(A, \succsim_1, \succsim_2)$ that satisfies the assumptions above. Compared to RST, we assume neither a commonly known disagreement point δ (in physical terms) nor that agents' preferences are defined over lotteries where the prizes are elements of X . In particular, we do not require that agents are expected utility maximisers or, more generally, that their preferences over A are representable by (cardinal) utility functions that are invariant only to positive affine transformation. Compared to the six restrictions made in RST (Section 2), we maintain (i)-(ii)-(v), weaken (iii) to non-triviality, and drop (iv)-(vi).

4.2 Interpretations

We provide two compatible interpretations for the bare-bones model of Section 3. The first one is behavioural and casts the mediator's problem as a decision problem with incomplete information. The second interpretation hijacks this setup, and shows that the classical Nash model (based on utilities) is mathematically equivalent to our probability-based formulation.

4.2.1 The target-based interpretation

The first interpretation takes the viewpoint of a mediator, hired by two bargainers to recommend them a feasible proposal over which they could strike an agreement. Given a native bargaining problem $(A, \succsim_1, \succsim_2)$, the mediator may suggest any feasible alternative in A , but cannot impose it. Her goal is to put on the table a proposal that each bargainer will individually evaluate and decide whether to accept. She may take into account issues of fairness or other considerations, but eventually her task is to select a proposal from A and her success is defined by its joint acceptance on the part of the bargainers. The mediator wants to maximise her probability of success.

We cast this situation as a decision problem under incomplete information. A bargainer $i = 1, 2$ accepts a proposal x when $x \succsim_i t_i$. We say that t_i is the minimum acceptance *target* for i . We may think of t_i as the minimal "fair" outcome that the bargainer has in mind, or as a proxy for his toughness (akin to his *type*), or as the outcome of a deliberative process over the risk of disagreement. The crucial assumption is that the mediator has incomplete information about t_i . She knows that $\succsim_i$ is continuous, and she believes that the type T_i of i has a random distribution with an (order-)convex support on $(X, \succsim_i)$.

Given that $\succsim_i$ is continuous on X , it admits a real-valued representation by a continuous function $v_i : X \rightarrow \mathbb{R}$. Therefore, $x \succsim_i t_i$ if and only if $v_i(x) \geq v_i(t_i)$ and we can equivalently reformulate the incomplete information about T_i as a random variable $v(T_i)$ on $(\mathbb{R}, \geq)$ with c.d.f. F_i and convex support. Consequently, $P(x \succsim_i T_i) = F_i \circ v_i(x)$, where F_i is a strictly increasing and continuous c.d.f. on $\mathbb{R}$. Note that v_i is unique only up to increasing transformations, so F_i is not uniquely defined; however, the function $F_i \circ v_i$ is unique. Therefore, we assume that the incomplete information of the mediator about each bargainer's type is summarised by $F_i \circ v_i$; in particular, given a proposal x , this is mapped into an individual acceptance probability $p_i = F_i \circ v_i(x)$.

The model in Section 3 takes this as its point of departure and provides a behavioural characterisation for the mediator's preferences over proposals in A . She ranks the feasible alternatives by their probability of being accepted by both bargainers, after combining their individual acceptance probabilities into a joint probability of success based on their dependence structure. For any native bargaining problem on A , the solution proposed by the mediator is a feasible alternative that maximises her induced preference order $\succsim$ over A .

We note two advantages for this behavioural interpretation. It is compatible with (but does not require) the assumption that agents are expected utility maximisers. In fact, it is not even necessary to include lotteries over feasible alternatives among the objects of choice for the bargainers. A similar comment applies for the disagreement point: in the literature, it is customary to mention it and immediately dispatch it by normalising its value to zero. Our model does not presume that a disagreement point δ (in physical terms) is known; however, if one is given, then the agent's individual rationality implies that his target has zero probability to lie below δ , and thus $P(\delta \succ_i T_i) = 0$.

4.2.2 Nash bargaining redux

The original Nash model selects a solution by maximising the product of two von Neumann–Morgenstern utilities (from now on, NM). Our approach picks a solution by maximising a copula that aggregates two individual probabilities. The goal of this section is to show that these two approaches are consistent and strictly linked.

The key decision-theoretic observation is that the utility-based NM model may be recast in an exclusively probability-based language; see Castagnoli and LiCalzi (1996) and Bordley and LiCalzi (2000). For simplicity, consider preferences over a compact nonempty interval $B = [x_*, x^*]$ in $\mathbb{R}$. Suppose that the Bernoulli index U is strictly increasing, bounded, and continuous; see Grandmont (1972) for an axiomatization. Applying if necessary a positive affine transformation, let $U(x_*) = 0$ and $U(x^*) = 1$. Then $U(x)$ has the formal properties of a cumulative distribution function (c.d.f.), and on some appropriate probability space there exists a random variable T with c.d.f. $U(x) = P(T \leq x)$. We call T a (random) target.

Let $\mathcal{L}$ the space of lotteries over B . If a lottery X in $\mathcal{L}$ has c.d.f. F and is stochastically independent of T , the chain of equalities

$$EU(X) = \int U(x) dF(x) = \int P(T \leq x) dF(x) = P(X \geq T) \quad (1)$$

shows that the expected utility $EU(X)$ is formally equivalent to the probability that the lottery X scores better than the target T . Hence, a claim made in a utility-based language for $EU(X)$ maps to an equally valid statement in a target-based language for $P(X \geq T)$. In particular, we can replace the notion of a cardinal Bernoulli index $U(x)$ that is unique only up to positive affine transformations by the simpler concept of a c.d.f. $P(T \leq x)$ for the target T . The NM model for preferences under risk postulates that preferences are linear in probabilities. It can be equivalently interpreted as a procedure that ranks lotteries by the expected value of their Bernoulli index or by the probability that they score better than a target T ; see LiCalzi (1999).

We are ready to consider the Nash bargaining problem (S, d) . Recall that S is a compact

and convex subset of feasible NM utility pairs, while d represents bargainers' utilities in case of breakdown. The crucial, but often implicit, assumption of the Nash model is that the NM-utility functions U_1 and U_2 are commonly known. Using (1), this reads as the assumption that the distributions of the bargainers' targets T_1 and T_2 are *ex ante* commonly known, whereas the targets are private information.

In our former interpretation, given a proposal x , the mapping $p_i = F_i \circ v_i(x)$ is based on the mediator's beliefs. Under common knowledge, the alternative interpretation is that the bargainers themselves agree on their own individual acceptance probabilities and may directly use these as input in constructing a bargaining solution. The missing step for the two bargainers is how to aggregate the commonly known individual probabilities and evaluate the proposals. This aggregation problem may be attacked in different ways. Ours is a behavioural characterisation: if the agents have common knowledge of the joint distribution of their targets *ex ante*, they maximise the probability of success by settling on the commonly known copula. In particular, if it is common knowledge that their two targets are *ex ante* stochastically independent, they should settle for the Nash solution. A related normative approach is pursued by Border and Segal (1997), discussed in Section 5.2.

4.2.3 Related literature

A relevant byproduct of our approach is the interpretation of the product operator in the Nash solution as the consequence of an assumption of stochastic independence between individual acceptance probabilities. To the best of our knowledge, the utility-based literature offers two competing interpretations for the product operator. We recall them briefly, for comparison.

Roth (1979, Section I.C) shows that we can frame the bargaining model as a single-person decision problem, where each agent chooses how much to claim by maximising his expected utility under the assumption that the claim of the other bargainer i is uniformly distributed between his disagreement point d_i and his ideal point m_i (defined as i 's highest feasible utility). The Nash solution emerges from the independent choices of the two bargainers. A very similar interpretation is also in Glycopantis and Miur (1994), who make no reference to Roth (1979). As Roth himself acknowledges, this approach is outside the game-theoretic tradition, because the processes by which the two agents form their expectations are not mutually consistent. On the other hand, similarly to ours, this interpretation is grounded on an assumption of independence between the evaluations made by the two agents.

A second interpretation for the Nash product is proposed in Trockel (2008). He views the Nash product as a special case of a social welfare function that aggregates the admissible payoff pairs into a social ranking. It evaluates a recommendation $\mathbf{u}$ by the Lebesgue measure of the set of utility pairs that are Pareto-dominated by $\mathbf{u}$. Without advocating more than a formal analogy, we note that any copula in Theorem 2 may be interpreted as a social welfare function adopted by the mediator to select her recommendation.

4.3 Domain

The Nash model is framed in the space of utilities: it implicitly assumes that all native bargaining problems with the same utility representation are indistinguishable, and thus

must have the same solution. This hidden assumption is probably extreme. RST discuss at length its implications, and lay bare the tradeoff between the power of Nash’s axioms and the granularity of the domain. Switching to a preference-based language, they remould the assumption by keeping fixed the set A of alternatives and varying the bargainers’ preferences. In their approach, a bargaining solution is a function that assigns a unique element of a given set A to every pair of bargainers’ preferences over lotteries on $A \cup \{d\}$.

Consistent with Nash (1950), however, the standard way to specify the domain of the solution is to hold bargainers’ preferences fixed and let the set A of alternatives vary. The typical formulation considers all the Nash problems based on the same disagreement point (in utilities). For generality, we cast our presentation assuming that the domain of our model contains all compact (but not necessarily convex) subsets of $[0, 1]^2$. However, this domain may be considerably shrunk and, in practical applications, it is reasonable to do so. The domain must be rich enough to include enough problems and elicit preferences over $[0, 1]^2$. A smaller domain may be more appropriate to ensure that the dependence embedded in the copula refers to comparable situations: for instance, the two bargainers should be the same, and the problems submitted to the mediator should justify similar answers.

Here is an exemplification. Fix two bargainers and their preferences on X . For any feasible (compact) set A , let d_i be the acceptance probability for a $\succsim_i$ -minimal proposal in A ; similarly, let m_i be the acceptance probability for a $\succsim_i$ -maximal proposal. Borrowing language from the Nash model, let $\mathbf{d}$ and $\mathbf{m}$ be called the disagreement point and the ideal point (in probabilities). A rich domain for our model may consider only the compact subsets in $[0, 1]^2$ associated with the same $\mathbf{d}$, and for those we may characterise a preference relation $\succsim_{\mathbf{d}}$ represented by a copula C such that $C(\mathbf{d}) = 0$. (We assume that an agent rejects a minimal proposal for sure.) A smaller but still rich domain is formed by the bargaining problems with the same $\mathbf{d}$ and $\mathbf{m}$: if we assume that an agent accepts his ideal point for sure, the representing copula would have both $C(\mathbf{d}) = 0$ and $C(\mathbf{m}) = 1$; see Cao (1982) for a similar normalisation over utility functions in the Nash model.

5 Applications and extensions

This section illustrates the richness and robustness of the target-based approach. We present a few applications and extensions, including comparative statics, testable restrictions, and prescriptive advice.

5.1 Comparative statics

A small but elegant literature deals with the comparative statics of the Nash solution. For a typical result, consider Theorem 1 in Kihlstrom et al. (1981): “The utility which Nash’s solution assigns to a player increases as his opponent becomes more risk averse.” Let us consider the implications of this result for the target-based approach, when the Bernoulli index function $U_i(x) = F_i \circ v_i(x)$ is interpreted as the cumulative distribution function $P(x \succsim_i T_i)$.

An agent with a Bernoulli index function $V_1(x)$ is more risk averse than an agent characterised by $U_1(x)$ if and only if there exists an increasing concave transformation K such that

$V_1(x) = K \circ U_1(x)$. Under the target-based interpretation, both V_1 and U_1 are distribution functions over the same domain and with the same range in $[0, 1]$. Therefore, the function K is bounded in $[0, 1]$ with $\lim_{x \downarrow 0} K(x) \geq 0$ and $\lim_{x \uparrow 1} K(x) = 1$; by concavity, it follows that $K(x) \geq x$ for all x . Hence, $V_1(x) \geq U_1(x)$ for all x ; that is, viewed as c.d.f.'s, V_1 is stochastically dominated by U_1 . In simple words, the target associated with V_1 can be regarded as “less demanding” than the target associated with U_1 . Theorem 1 may be reformulated as follows. At the Nash solution, $i = 1, 2$ is offered a better proposal and his individual acceptance probability increases when the target of the other agent $j = 3 - i$ becomes less demanding. If an agent becomes more accommodating, the Nash solution ends up rewarding the other one.

A related comment concerns the observation recently made in Alon and Lehrer (2014) that the Nash solution is not ordinally equivalent. Suppose that two agents are to divide a unit amount of money. Let x_i in $[0, 1]$ be the quantity attributed to Agent i , with $x_1 + x_2 = 1$. From an ordinal viewpoint, assume that each agent has increasing (and continuous) preferences in the amount x_i he secures. A model with ordinal preferences $U_1(x) = U_2(x) = x$ is equivalent to a model with ordinal preferences $U_1(x) = \sqrt{x}$ and $U_2(x) = 1 - \sqrt{1 - x}$. However, when applied literally, the Nash solution predicts $x_1 = x_2 = 1/2$ for the first model, and $x_1 = 1/4$ and $x_2 = 3/4$ for the second one. In our approach, the function $U_i(x) = F_i \circ v_i(x)$ embodies both the ordinal preference expressed by v_i and the individual acceptance probability represented by F_i . The two models are ordinally equivalent, but in the second model Agent 1 has a less demanding target while Agent 2 has a more demanding target. (In a utility-based language, what drives the difference is that in the second model Agent 1 is more risk averse while Agent 2 is less risk averse.)

These examples deal with comparative statics over the distribution of individual acceptance probabilities. The target-based approach, however, allows to study analogous results based on (partial) orderings for the copulas underlying the dependence that links individual probabilities to the joint probability of success. We consider only a simple example. Given two copulas C_1 and C_2 , the *concordance* ordering states that C_1 is more concordant than C_2 if $C_1(\mathbf{p}) \geq C_2(\mathbf{p})$ for all $\mathbf{p}$ in $[0, 1]^2$. Suppose this is the case for C_1 and C_2 , and let $\mathbf{p}_i^*$ be a solution under the copula C_i , for $i = 1, 2$ and the same bargaining problem B . Clearly,

$$C_1(\mathbf{p}_1^*) \geq C_1(\mathbf{p}_2^*) \geq C_2(\mathbf{p}_2^*)$$

so that the joint probability of success at the solution is increasing in the concordance ordering. Agents with concordant targets are more likely to strike a deal.

5.2 Social preferences and implementation

The model in Section 3 characterises a preference relation $\succsim$ and derives the bargaining solution through the maximisation of a copula. A similar approach is followed by Blackorby et al. (1994), who define a bargaining solution as the (set of) maximisers for a generalised Gini ordering, represented by a quasi-concave, increasing function that is linear on the rank-ordered subsets of $[0, 1]^2$. The class of generalised Gini orderings spans a continuum of solutions, exhibiting different levels of inequality aversion and including the egalitarian and

the utilitarian solutions as extreme cases.

We use their model as an example to illustrate the flexibility and limitations of our copula-based approach. For $0 \leq \alpha \leq 1/2$, the family of symmetric copulas

$$C_{\alpha,\alpha}(u, v) = \begin{cases} u + \frac{\alpha}{1-\alpha}v - \frac{\alpha}{1-\alpha} & \text{if } \alpha \leq u \leq v \leq 1 \\ v + \frac{\alpha}{1-\alpha}u - \frac{\alpha}{1-\alpha} & \text{if } \alpha \leq v \leq u \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

represents a continuum of affine (truncated) generalised Gini orderings. This family includes the egalitarian solution for $\alpha = 0$ and the utilitarian solution for $\alpha = 1/2$. Accordingly, the parameter α may be interpreted as an index of (increasing) inequality aversion. (This family is nested in a larger class of asymmetric copulas $C_{\alpha,\beta}$ with $\alpha, \beta > 0$ and $\alpha + \beta < 1$; see Exercise 3.8 in Nelsen, 2006.) We say that this family is truncated because $C(u, v) = 0$ if $\min\{u, v\} < \alpha$. The same comment made for the utilitarian solution in Section 3.5 applies here, and clearly these copulas do not satisfy A5* when $\alpha > 0$.

In our model, the copulas aggregate individual acceptance probabilities. One should be careful not to carry formal analogies too far, but the copulas $C_{\alpha,\alpha}$ may be used to model the mediator's aversion to inequalities among the individual acceptance probabilities. The alternative interpretation is that the copulas represent the stochastic dependence between agents' targets and that this dependence drives the choices. If the mediator's opinion is captured by $C_{\alpha,\alpha}$, then her recommended solution exhibits (some degree of) inequality aversion between the individual acceptance probabilities. If $C_{\alpha,\alpha}$ rationalises the mediator's recommendations, these two interpretations are behaviourally indistinguishable.

More generally, when $\succsim$ is viewed as a social preference relation, one may recast the representing copula as a social welfare function. This approach was pioneered by Kaneko and Nakamura (1979), who define and characterise a Nash social welfare function that evaluates the relative increases in individuals' welfare from a state δ (in physical terms) unanimously considered as the worst possible. In a subsequent paper, Kaneko (1980) takes care to point out the conceptual differences between a social welfare function and a bargaining solution. Put simply, however, the underlying idea is to pick a solution by maximising a function: this works also over non-convex (compact) sets, but it may generate set-valued solutions. The applicability of the Nash model is extended at the cost of forfeiting uniqueness.

Border and Segal (1997) provide an axiomatisation of the Nash (bargaining) solution that is very close in spirit both to Kaneko's insight and to our model. The natural interpretation for their setup is that "the two bargainers hire an arbitrator to make choices for them" (p. 1) and that the arbitrator has a preference order $\succsim$ over solutions. Border and Segal (1997) motivate these preferences by analogy to social choice, suggesting that the arbitrator relies on her notions of fairness to come up with a decision rule for all the bargaining problems. An ancillary interpretation views the axioms as guidelines that both bargainers should find acceptable before they agree to hire her. In their approach, the arbitrator's selection is justified by (normatively) binding axioms. Instead, we share with RST the search for a behavioural characterisation that downplays the normative undertones: the mediator in our model issues recommendations but cannot impose a solution.

A related but independent issue concerns the implementation problem. Our model assumes that the mediator knows the individual acceptance probabilities or, more precisely, that she holds subjective beliefs about the agents' targets. However, it may be in the agents' interest to misrepresent their objectives and manipulate the mediator's beliefs. Therefore, one should worry about devising mechanisms that help elicit correct information from the agents. An important step forward in this direction is made by Miyagawa (2002), who studies this problem in the context of two-person bargaining. He provides a simple four-stage sequential game that fully implements a reasonably large class of two-person bargaining solutions in subgame-perfect equilibrium.

Miyagawa (2002) relies on four restrictive assumptions to derive his results. First, the bargaining problem B must contain an alternative d that both agents consider least preferred. Second, any most preferred alternative for i is judged by $j = 3 - i$ as indifferent to d . Third, the Pareto frontier of B is strictly convex. Finally, the bargaining solution must be generated by the maximisation of a (component-wise) increasing and quasi-concave function. This generating function is an analog of the social welfare function just discussed. Coincidentally, Miyagawa (2002) also assumes that this function is normalized to $[0, 1]$ as in the copula-based model. Each assumption has technical implications; for instance, the combination of strict convexity for the Pareto frontier and quasi-concavity for the generating function entails the uniqueness of the solution.

Compared to our setup, the crucial restriction is that the generating function must be quasi-concave. A copula need not be quasi-concave; for instance, the function $C(u, v) = [M(u, v) + W(u, v)]/2$ is a copula, but it is not quasi-concave. Hence, the applicability of Miyagawa (2002) cannot generally be taken for granted. However, many common examples, and all the functional forms presented so far, are quasi-concave. Therefore, the class of solutions that are implementable includes the Nash, the egalitarian, and the (truncated) utilitarian solutions. This is duly noted in Miyagawa (2002, p. 293), although an unfortunate typo incorrectly mentions Kalai and Smorodinski (1975) instead of Kalai (1977).

5.3 Testable restrictions

An important requirement for a theory is the ability to generate falsifiable predictions. A direct test exists for the target-based approach. We first illustrate the idea, and then formalise it. As introduced in Section 3, a *bargaining problem* is represented by a compact set B in $[0, 1]^2$ where each point $\mathbf{p}$ in B corresponds to a pair of (acceptance) probabilities. A *solution* is a map that for any problem B selects (at least) one point in B . The target-based approach recommends a solution by maximising a suitable copula C over B .

Suppose that B contains a point $\mathbf{p} = (p_1, p_2)$ with $p_1 + p_2 > 1$. The Fréchet lower bound implies $C(\mathbf{p}) \geq p_1 + p_2 - 1 > 0$ for any C . Consider another feasible point $\mathbf{q} = (q_1, q_2)$ in B . The Fréchet upper bound implies $C(\mathbf{q}) \leq \min(q_1, q_2)$. Therefore, if $\min(q_1, q_2) < p_1 + p_2 - 1$, the point $\mathbf{p}$ must be strictly preferred to $\mathbf{q}$, and $\mathbf{q}$ cannot be a solution for any copula C . More generally, by picking a point $\mathbf{p}^*$ that maximises $p_1 + p_2 - 1$ in B , we can formulate the following more stringent test.

Proposition 8. *Let $\mathbf{p}^* \in \arg \max_B (p_1 + p_2 - 1)$. Define the quadrant*

$$Q := \{\mathbf{q} \in [0, 1]^2 : \min(q_1, q_2) \geq p_1^* + p_2^* - 1\}.$$

Then the solution must belong to $B \cap Q$.

Clearly, under A.5*, a second immediate test is that a solution for B cannot be dominated (in the strong Pareto ordering) by another point available in B .

A different class of restrictions is the following. As discussed in Section 5.2, our model characterises bargaining solutions that can be rationalised by a copula. Roughly speaking, this implies that we can generate only solutions that satisfy independence of irrelevant alternatives. For instance, the well-known Kalai-Smorodinsky (1975) solution (for short, KS solution) cannot be derived in our model. Formally speaking, this solution violates A1 as can be shown by a simple example. For the bargaining problem B_1 represented by the convex hull of the three points $(0, 0)$, $(0, 1)$, and $(1, 0)$, the KS solution uniquely prescribes the point $\mathbf{p} = (1/2, 1/2)$. On the other hand, for the bargaining problem B_2 represented by the convex hull of the four points $(0, 0)$, $(0, 1)$, $(1/2, 1/2)$, and $(1/2, 0)$, the KS solution uniquely prescribes the point $\mathbf{q} = (1/3, 2/3)$. Since both $\mathbf{p}$ and $\mathbf{q}$ belong to $B_2 \subset B_1$, the first KS solution reveals $\mathbf{p} \succ \mathbf{q}$ while the second one reveals $\mathbf{q} \succ \mathbf{p}$, in violation of A1. Although it is formally possible to rationalise the KS solution as the outcome of a lexicographic maximisation, we find this approach interpretively unsatisfactory and thus we do not pursue it here.

5.4 Bargaining power

Since Binmore et al. (1986), the economic literature has made extensive use of an asymmetric version of the Nash Solution as a reduced form for the differences in the bargaining power of the agents. Formally, given a Nash bargaining problem (S, d) , the *asymmetric Nash solution* is defined as the maximiser of the product $(u_1 - d_1)^a (u_2 - d_2)^{1-a}$, for a in $[0, 1]$. As a increases from 0 to 1, the asymmetric solution increasingly favours the first agent, with symmetry holding at $a = 1/2$. For instance, consider the illustrative example from Section 5.1 where two agents are to divide a unit amount of money, and $U_1(x) = U_2(x) = x$. The asymmetric Nash solution gives $x_1 = a$ and $x_2 = 1 - a$.

It is natural to suggest using the function $N_a(p, q) = p^a q^{1-a}$ but this is not a copula, because it fails the second half of C1 in Section 2. Therefore, one cannot import the asymmetric Nash solution under the copula-based approach. On the other hand, since asymmetric copulas exist, we might search for some suitable (possibly parametric) substitute copula with similar properties. For instance, by Theorem 2.1 in Liebscher (2008), for any copula C the expression

$$C^a(p, q) = p^a q^{1-a} C(p^{1-a}, q^a) = N_a(p, q) \cdot C(p^{1-a}, q^a)$$

defines an asymmetric copula, bearing an obvious relationship with N_a . We argue that this approach cannot work either.

Given a bargaining problem B , let $\Delta_B(C) = \sup_{(p, q) \in B} |C(p, q) - C(q, p)|$ be the degree of asymmetry for a function C over B . In our illustrative example, we obtain $\Delta_B(N_a) = 1$ at $a = 0$ and $a = 1$; that is, the family of asymmetric Nash solutions includes the case of maximal

asymmetry. On the other hand, it is known that over $B = [0, 1]^2$ we have $\Delta_B(C) \leq 1/3$ for any copula C (Nelsen, 2007); moreover, $\Delta_B(C) \leq 1/5$ if C is also quasi-concave (Alvoni and Papini, 2007). Therefore, there exist no copula that can mimic the extreme asymmetry allowed by N_a . In our illustrative example, for instance, $\Delta_B(N_a)$ is symmetric around $a = 1/2$ and increasing in $|a - 1/2|$, with $\Delta_B(N_a) = 1/3$ at $|a - 1/2| \approx 0.219$; therefore, the copula-based approach cannot replicate N_a for a outside of the interval $[\cdot 281, \cdot 719]$.

In fact, a sound interpretation for modelling bargaining power requires more care. Consider again the standard utility-based approach. The widespread use of the asymmetric Nash solution is most likely due to its technical convenience, without much concern about the lack of any intuitive meaning for the weighted geometric mean of utilities represented by N_a . However, utility-based axiomatic foundations exist in the literature. Kalai (1977a) provides a characterisation for the asymmetric solution based on a replication argument for symmetric solutions. His approach is technically clean and simple, but provides little intuition.

A richer and more convincing approach is Harsanyi and Selten (1972), who study the case of fixed threats under incomplete information. This latter assumption, absent in the Nash model, explicitly recognises that each player has private information about some of his characteristics that are relevant for the bargaining process. Harsanyi and Selten (1972) show that differential information provides a foundation for asymmetric solutions. We follow in their footsteps and discuss a reduced form for modelling bargaining power in our illustrative example, based on the copula-based approach.

Recall the distinction between the native bargaining problem $(A, \succsim_1, \succsim_2)$ and the derived bargaining problem B ; see Section 4.1. Bargaining power must be traced back to the native problem. The individual acceptance probabilities provide the link between this and the derived problem, while the choice of a copula applies only to B as a tool to reconcile the individual acceptance probabilities into a joint probability of success. Therefore, bargaining power should affect the shape of the derived problem B rather than the functional form C .

In the illustrative example, the feasible set $A = \{(x_1, x_2) \in [0, 1]^2 : x_1 + x_2 \leq 1\}$ for the native problem is especially simple to describe. In particular, it can be reduced to a uni-dimensional problem $\{(x_1, 1 - x_1) \in [0, 1]^2 : 0 \leq x_1 \leq 1\}$, where each agent only cares about the money he receives. Assume Theorem 5; that is, the mediator aggregates agents' acceptance probabilities assuming stochastic independence and $C(p, q) = p \cdot q$. When Agent 1 has more bargaining power than Agent 2, he can aspire to higher targets. (As in Harsanyi and Selten (1972), bargaining power is only one among several issues that may cause this relationship to hold.) Everything else being equal, this is captured by the assumption that T_1 stochastically dominates T_2 .

For practical purposes, it is convenient to adopt a parametric formulation. For instance, let $P(x \succsim_i T_i) = x^{a_i}$ where $a_i \geq 0$ can be interpreted as the bargaining strength of Agent i , and assume $a_1 + a_2 > 0$. Clearly, T_1 stochastically dominates T_2 if and only if $a_1 \geq a_2$. The shape of the derived bargaining problem

$$B(a_1, a_2) = \{(p_1, p_2) \in [0, 1]^2 : p_1 = x^{a_1} \text{ and } p_2 = (1 - x)^{a_2}, \text{ for } 0 \leq x \leq 1\}$$

depends on how the bargaining strength a_i of each agent affects his own acceptance probability in the eyes of the mediator. Given the product copula, the acceptance probabilities aggregate

into the function $x^{a_1} \cdot (1 - x)^{a_2}$, and the corresponding recommendation is $x_1^* = a_1/(a_1 + a_2)$ and $x_2^* = a_2/(a_1 + a_2)$. We recover the same solution set associated with the asymmetric Nash solution.

From an interpretive viewpoint, we emphasise that the asymmetry associated with differences in bargaining power affects the individual acceptance probabilities. The approach described at the beginning of this section fails because it tries to impose the asymmetry on the copula, that instead aggregates those into the joint probability of success. Although our approach allows asymmetric copulas, they are not meant as a tool to capture modelling issues concerning the native problem or its relation with the derived problem.

5.5 Negotiation analysis

This last section touches upon the relationships of our approach with the practice of negotiations. When Roth (1979, Section I.B) writes that “a solution can be interpreted as a model of the bargaining process”, he argues for considering its descriptive implications that can be tested empirically. However, he points out that there is an alternative prescriptive approach which views the solution as “a rule which tells an arbitrator what outcome to select”; e.g., Border and Segal (1997) aim to this second goal.

Our model is concerned with a mediator rather than an arbitrator: the latter can impose a solution, while the former can only recommend it. This puts the mediator in a hybrid position. Maximising the joint probability of success appears very natural, but its prescriptive value for the mediator may not be shared by the agents who presumably worry about more than just striking an agreement. This is not the place for a full discussion, but we wish to add a few remarks to highlight the potential of our approach.

Subramanian (2010, pp. 109–110) finds “that the implications for negotiation strategy change dramatically when we move away from the assumption that dealmakers will accept deals that are just better than their BATNA [Best Alternative to a Negotiated Agreement, equivalent to the disagreement point] to the more realistic and nuanced assumption that the likelihood the other side will say yes increases with the incentives to do so.” The target-based interpretation makes the role of individual acceptance probabilities explicit. The mediator should acquaint herself with the agents well enough to understand how each of them views the native bargaining problem and code her understanding in terms of individual acceptance probability.

A second step requires the mediator to make a conscious choice about the possible connections between the individual acceptance probabilities and the joint probability of success. While the assumption of stochastic independence seems by far the simple and most natural one, it is not the only possible one. Different copulas may rationalise alternative recommendations, and thus this element should be given attention by the mediator.

References

- [1] S. Alon and E. Lehrer (2014), “Competitive equilibrium as an ordinal bargaining solution: An axiomatic approach”, December, *mimeo*.

- [2] E. Alvoni and P.L. Papini (2007), “Quasi-concave copulas, asymmetry and transformations”, *Commentationes Mathematicae Universitatis Carolinae* **48**, 311–319.
- [3] K.J. Arrow (1963), *Social Choice and Individual Values*, 2nd edition, New York: Wiley.
- [4] K. Binmore, A. Rubinstein, and A. Wolinsky (1986), “The Nash bargaining solution in economic modelling”, *The RAND Journal of Economics* **17**, 176–188.
- [5] C. Blackorby, W. Bossert, and D. Donaldson (1994), “Generalized Ginis and cooperative bargaining solutions”, *Econometrica* **62**, 1161–1178.
- [6] K. Border and U. Segal (1997), “Preferences over solutions to the bargaining problem”, *Econometrica* **65**, 1–18.
- [7] R.F. Bordley and M. LiCalzi (2000), “Decision analysis using targets instead of utility functions”, *Decisions in Economics and Finance* **23**, 53–74.
- [8] E. Castagnoli and M. LiCalzi (1996), “Expected utility without utility”, *Theory and Decision* **41**, 281–301.
- [9] X. Cao (1982), “Preference functions and bargaining solutions”, *Proceedings of the 21st IEEE Conference on Decision and Control* **1**, 164–171.
- [10] A. Dhillon and J.-F. Mertens (1999), “Relative utilitarianism”, *Econometrica* **67**, 471–498.
- [11] P. Fishburn (1970), *Utility Theory for Decision Making*, New York: Wiley.
- [12] A. Francetich (2013), “Notes on supermodularity and increasing differences in expected utility”, *Economics Letters* **121**, 206–209.
- [13] D. Glycopantis and A. Muir (1994), “On the Nash bargaining solution: An alternative approach”, *British Review of Economic Issues* **16**, 55–71.
- [14] J.-M. Grandmont (1972), “Continuity properties of a von Neumann-Morgenstern utility”, *Journal of Economic Theory* **4**, 45–57.
- [15] J.C. Harsanyi and R. Selten (1972), “A generalized Nash solution for two-person bargaining games with incomplete information”, *Management Science* **18**, P80–P106.
- [16] I.N. Herstein and J. Milnor (1953), “An axiomatic approach to measurable utility”, *Econometrica* **21**, 291–297.
- [17] E. Kalai (1977), “Proportional solutions to bargaining situations: Interpersonal utility comparisons”, *Econometrica* **45**, 1623–1630.
- [18] E. Kalai (1977a), “Nonsymmetric Nash solutions and replications of 2-Person bargaining”, *International Journal of Game Theory* **6**, 129–133.
- [19] E. Kalai and M. Smorodinsky, “Other solutions to Nash’s bargaining problem”, *Econometrica* **43**, 513–518.
- [20] M. Kaneko (1980), “An extension of the Nash bargaining problem and the Nash social welfare function”, *Theory and Decision* **12**, 135–148.
- [21] M. Kaneko and K. Nakamura (1979), “The Nash social welfare function”, *Econometrica* **47**, 423–435.
- [22] R. Kihlstrom, A.E. Roth, and D. Schmeidler (1981), “Risk aversion and solutions to Nash’s bargaining problem,” in: O. Moeschlin and D. Pallaschke (eds.), *Game Theory and Mathematical Economics*, Amsterdam: North-Holland Publishing Company, 65–71.

- [23] M. LiCalzi (1999), “A language for the construction of preferences under uncertainty”, *Revista de la Real Academia de Ciencias Exactas, Físicas y Naturales* **93**, 439–450.
- [24] E. Liebscher, “Construction of asymmetric multivariate copulas”, *Journal of Multivariate Analysis* **99**, 2234–2250.
- [25] E. Miyagawa (2002), “Subgame-perfect implementation of bargaining solutions”, *Games and Economic Behavior* **41**, 292–308.
- [26] J.F. Nash (1950), “The bargaining problem”, *Econometrica* **18**, 155–162.
- [27] R.B. Nelsen (2006), *An Introduction to Copulas*, 2nd edition, New York: Springer.
- [28] R.B. Nelsen (2007), “Extremes of nonexchangeability”, *Statistical Papers* **48**, 329–336.
- [29] A.E. Roth (1979), *Axiomatic Models of Bargaining*, Berlin: Springer, 1979.
- [30] A. Rubinstein, Z. Safra, and W. Thomson (1992), “On the interpretation of the Nash bargaining solution and its extension to non-expected utility preferences”, *Econometrica* **60**, 1171–1186.
- [31] A. Sklar (1959), “Fonctions de répartition à n dimensions et leurs marges”, *Publications de l’Institut de Statistique de l’Université de Paris* **8**, 229–231.
- [32] G. Subramanian (2010), *Negotiauctions*, New York: Norton.
- [33] W. Thomson (1994), “Cooperative models of bargaining,” in: R.J. Aumann and S. Hart (eds.), *Handbook of Game Theory*, vol. 2, Amsterdam: North-Holland Publishing Company, 1237–1284.
- [34] W. Trockel (2008), “The Nash product is a utility representation of the Pareto ordering”, *Economics Letters* **99**, 220–222.