Journal of Artificial Intelligence Research 46 (2013) 579-605 Submitted 4/12; published 4/13

Predicting Behavior in Unstructured Bargaining with a Probability Distribution

David H. Wolpert DAVID.H.WOLPERT@GMAIL.COM
Santa Fe Institute,
1399 Hyde Park Road, Santa Fe, NM 87501;
Information Sciences Group,
Los Alamos National Laboratory,
MS B256, Los Alamos, NM 87545

James W. Bono JWBONO@GMAIL.COM

Abstract

In experimental tests of human behavior in unstructured bargaining games, typically many joint utility outcomes are found to occur, not just one. This suggests we predict the outcome of such a game as a probability distribution. This is in contrast to what is conventionally done (e.g, in the Nash bargaining solution), which is predict a single outcome. We show how to translate Nash's bargaining axioms to provide a distribution over outcomes rather than a single outcome. We then prove that a subset of those axioms forces the distribution over utility outcomes to be a power-law distribution. Unlike Nash's original result, our result holds even if the feasible set is finite. When the feasible set is convex and comprehensive, the mode of the power law distribution is the Harsanyi bargaining solution, and if we require symmetry it is the Nash bargaining solution. However, in general these modes of the joint utility distribution are not the experimentalist's Bayes-optimal predictions for the joint utility. Nor are the bargains corresponding to the modes of those joint utility distributions the modes of the distribution over bargains in general, since more than one bargain may result in the same joint utility. After introducing distributional bargaining solution concepts, we show how an external regulator can use them to optimally design an unstructured bargaining scenario. Throughout we demonstrate our analysis in computational experiments involving flight rerouting negotiations in the National Airspace System. We emphasize that while our results are formulated for unstructured bargaining, they can also be used to make predictions for noncooperative games where the modeler knows the utility functions of the players over possible outcomes of the game, but does not know the move spaces the players use to determine those outcomes.

1. Introduction

In game theory, *bargaining* refers to scenarios where two or more people must come to a joint agreement on an outcome. In *structured bargaining*, the scenario is modeled as a noncooperative game, with the players making explicitly delineated alternating moves, e.g., proposals and counter-proposals (Osborne & Rubinstein, 1994; Aumann & Hart, 1992). In contrast, in *unstructured bargaining*, the scenario is modeled without any explicit delineation of the alternating moves. Instead all that is known to the modeler is the *feasible set* of all the joint-utilities that would arise for all the possible bargains the humans might reach. Arguably, in most real-world bargaining scenarios, the interaction between the bargainers is "free form", and so requires an unstructured bargaining analysis. Therefore, to be able to predict and potentially emulate the behavior of interacting humans, we need to have an accurate model of human unstructured bargaining.

In unstructured bargaining there are two spaces; the space of bargains and the space of joint expected utility vectors that the players assign to each such bargain. Note that it may be that more than one bargain results in the same outcome. As shorthand, it is conventional to leave the term "expected" implicit and just refer to "utility". Similar shorthand is to refer to a joint (expected) utility vector as an "outcome".

Most game theoretic modeling to date on unstructured bargaining starts with specification of a "feasible set" of all of the possible outcomes. This is the only knowledge the modeler has concerning the unstructured bargaining scenario. In particular, the modeler is ignorant of the spaces of possible "moves" the players may make to reach a bargain — and is even ignorant of the space of possible bargains.

Traditional unstructured bargaining is concerned with specifying a map that takes any such feasible set S to a single outcome $x \in S$, i.e., a "point-valued solution concept" (Nash, 1950; Harsanyi & Selten, 1972; Kalai & Smorodinsky, 1975; Kalai, 1977). For example, this was the case with Nash's original work, where he showed how for any specified S, his axioms force a unique prediction of an outcome $x \in S$. Some of this work regards the map normatively, as providing fair or reasonable bargaining outcomes. Other work regards it positively, as a prediction for what agreement will be reached by humans bargaining in an unstructured manner (Nydegger & Owen, 1974; Roth & Malouf, 1979; Camerer, 2003). In this paper we are concerned with this positive viewpoint. For an extensive discussion on the interpretation of Nash's solution, see the work of Rubinstein, Safra, and Thomson (1982).

In contradiction to this theoretical positive work, the experimental literature makes clear that in the real world more than one bargain has non-zero probability of being the outcome of any given unstructured bargaining problem (Camerer, 2003; Roth & Malouf, 1979).[1] To accommodate this, in this paper we consider maps that take any unstructured bargaining problem's feasible set S to a probability distribution over S, rather than to a single element of S.[2] Doing this, we derive a parameterized set of possible maps from bargaining problems to distributions over those problems. (Intuitively, the parameters reflect the "bargaining power" of the players.) We call such maps "distributional bargaining concepts", and call their images "utility distributions".

There are many advantages to using distributional bargaining concepts, in addition to their according with experimental data better than point-valued solution concepts. A major one arises if there is an external regulator who can modify some aspects of the bargaining game and has their own utility function over bargaining outcomes. By modifying the game, the regulator changes the associated distribution, and therefore modifies the value of their own expected utility. Accordingly, they can calculate the Bayes-optimal modification to the game.

1. In fact, given the noisy nature of human behavior, it would be stunning if certain physically possible outcomes actually occurred with exactly 0 probability, rather than *some* small, non-zero probability.
2. Our approach allows S to be either finite or infinite; for succinctness we will generically refer to a "probability distribution" even if S is infinite and we properly should refer to a probability density function.

Our approach to deriving a distributional bargaining concept is to "translate" Nash's axioms of unstructured bargaining — conventionally applied to maps that produce a single utility outcome rather than a utility distribution — to apply to distribution-valued maps. To be precise, we use probabilistic versions of Nash's axioms of Scale Invariance (**SI**) and Independence of Irrelevant Alternatives (**IIA**). We also use a probabilistic version of the axiom of Translation-Invariance of utilities (**TI**).

Adopting a Bayesian perspective, we view these axioms as formalizations of the ignorance of the modeler in many scenarios, rather than as assumptions about human behavior: For us, SI means that the modeler does not know anything about how the relative probabilities of outcomes chosen by the players is likely to change if one simply scales the utilities of all possible outcomes. Similarly, for us IIA means the modeler does not know anything about how the relative probabilities of outcomes chosen by the players is likely to change if a subset of of the possible outcomes is removed. And for us, TI means that the modeler does not know anything about how the relative probabilities of outcomes chosen by the players is likely to change if the utilities of all possible bargains are simply translated by the same constant.

To these axioms we add the extra one that all outcomes in which all players do strictly better than the default outcome have Non-Zero probability (**NZ**). (NZ is imposed because it holds in all experiments, due to subject inattention if nothing else.)

There is other work, which, like ours, is also concerned with distributions over S (Peters & Tijs, 1984). However, the work of Peters and Tijs (1984) differs from ours in two important ways. First, while they also use probabilistic versions of SI and IIA, unlike us, they also use the axiom of Pareto optimality (**PAR**). (As discussed in the conclusion, we need not impose the assumption of binding contracts, and therefore need not assume the joint utility is Pareto-optimal.) In addition, unlike us, they do not use NZ, nor do they (explicitly) use TI. Second, Peters and Tijs translate IIA and SI into different probabilistic versions than ours.

As a result of these difference, we arrive at a very different distributional solution concept from that of Peters and Tijs. In particular, our axioms force the distribution to be a power law over the set of joint utility outcomes Pareto superior to the default outcome. In contrast, the axioms used by Peters and Tijs do not result in a tightly characterized solution, power law or otherwise. Accordingly, the approach of Peters and Tijs has not been used to derive the Bayes-optimal modification to a game that might be applied by an external regulator of the game.

Perhaps most importantly, our version of IIA means our solution concept applies to any bargaining game, even finite ones that are non-convex and not comprehensive. This is quite important for a positive bargaining solution concept. In real-world unstructured bargaining, it is common for the bargainers to only consider a finite number of possible outcomes. (Typically, real human bargainers in the field do *not* consider the option of tossing a weighted coin to choose among the possible bargains.) In such scenarios the feasible set is finite.

More generally, since we dispense with PAR, there is nothing that restricts our analysis to scenarios that are traditionally viewed as "bargaining". As we discuss in the conclusion, our results can also be used to predict the outcomes of noncooperative games whenever the modeler only knows the feasible set of that game's joint-utility outcomes, and cannot tractably elaborate the move spaces of the players. In such situations, the best the modeler can do is provide a distribution over the final joint utility outcome. Our distributional bargaining concept provides a way to do this.

The current paper and the work of Peters and Tijs are not the only papers to extend game theory by replacing solution concepts that are point-valued (or set-valued) with solution concepts that are probability distributions. In particular, the work of Wolpert and Bono (2011) introduces a distribution-valued solution concept for noncooperative games when the modeler *does* know the move spaces. That work introduces a map that has an input the specification of an arbitrary non-cooperative game. The output of the map is a distribution over all possible mixed strategy profiles (i.e., all player joint choices) in the input noncooperative game. In contrast, the present paper introduces a map taking the specification of an unstructured bargaining game as its input, and producing a distribution over all player joint utilities as its output.

There are also recent papers in the artificial intelligence literature that, like ours, focus on "non-equilibrium solution concepts" (Brafman & Tennenholtz, 2003; Rezek, Leslie, Reece, Roberts, Rogers, Dash, & Jennings, 2008; Aydoğan & Yolum, 2012; Duan, Doğru, Özen, & Beck, 2012). Rather than merely invoking equilibrium, Brafman and Tennenholtz (2003) employ a reinforcement learning algorithm to efficiently achieve coordination in common interest stochastic games. Rezek et al. (2008) look at game theoretic solution concepts from a machine learning perspective, i.e. they assume players make inferences about their opponents in a Bayesian framework and derive a novel fictitious play algorithm. Recent work in AI that focuses on bargaining, such as that of Aydoğan and Yolum (2012) and Duan et al. (2012), is particularly closely related to the current paper. The main difference is that our paper focuses on an unstructured game, whereas the focus in previous work has been on structured games.

1.1 Contribution of this Paper

Our first contribution in this paper is to derive our distributional bargaining concept, as outlined above. After this we focus on two major advantages of a distributional bargaining concept over point-valued or set-valued bargaining concepts: 1) the ability of a modeler to apply decision theory to predict outcomes, and 2) the ability of an external regulator to employ control theory to optimally regulate the system.

To elaborate on the first advantage, consider that the modeler of the outcome of the bargaining will often have a loss function, measuring the quality to them of predicting the joint expected utility outcome is x when the actual outcome is x'.[3] Given such a loss function, and given a distribution over outcomes, there is a well-defined Bayes-optimal prediction for a single joint utility outcome. This Bayes-optimal prediction will vary with the loss function. Accordingly, whatever distributional bargaining concept is used, in general the associated Bayes-optimal prediction is neither the most likely joint expected joint utility outcome (the Harsanyi solution) nor the most likely bargain.

Let s be the vector-valued function taking bargains to the associated joint utility outcomes. In general, s need not be invertible. (E.g., if the feasible set is a set of K possible

3. As shorthand, we will often abbreviate "joint expected utility outcome" to "joint utility outcome", or even "outcome".

joint choices among the N players, where $K \geq N + 1$.) In such cases we cannot go from a prediction for the joint utility (whether made with a conventional bargaining concept or a distributional one) to a prediction for a joint bargain. One would need a likelihood function $P(b \mid x)$ giving the relative probabilities of all bargains b given the joint utility x to invert a prediction for joint utility into a prediction for the bargain reached.

On the other hand, if s *is* invertible, then the distribution over joint utilities does fix the distribution over bargains. However, even in this case, if the Jacobian of s is non-uniform, then the most likely joint utility is not the most likely joint bargain. Intuitively, s may "concentrate" probability density on some regions of possible joint utility outcomes (those corresponding to relatively many possible bargains), and diminish it in others.

Note though, that when the feasible set is countable and s is invertible, the issue of a non-uniform Jacobian of s disappears. This provides yet another benefit to using a distribution bargaining concept — like ours — that is applicable to countable (and even finite) feasible sets, not just convex and comprehensive ones.

As an illustration of the foregoing, we consider a scenario naturally modeled as unstructured bargaining. Often the path of an aircraft through the National Airspace System is renegotiated inflight, e.g., due to unforeseen weather. Such negotiations do not follow any particular protocol — they are unstructured. Accordingly, we use them to illustrate the Nash distributional bargaining model. We emphasize the fact that even in this simple scenario, the map s from bargains to joint utility outcomes is non-invertible. As a result, while it is straightforward to evaluate the probability density function over possible joint utility outcomes, the same cannot be said for evaluating relative probabilities of various bargains.

To address the second major advantage of our distribution-valued solution concept, i.e. the ability of an external regulator to employ control theory to optimally regulate the system, we consider an external regulator who has a real-valued welfare function defined over the bargains of the players, and who can modify some parameters of the bargaining game. We show how a distributional solution concept allows such an external regulator to perform Bayes-optimal configuration of the bargaining game, i.e., how the regulator can set the parameters of the game under their control to optimize the expected welfare of the resultant bargaining outcome. In the context of the Nash distributional bargaining concept, we call this approach to setting parameters "Nash distributional bargaining management".

As an example, suppose that the regulator can modify the set of allowed bargains within some range. In this case, the Nash distributional bargaining concept can be used to determine the regulator's optimal modification. Similar external interventions would be changing the default bargain, or even modifying the relative bargaining power of the players.

A more complicated type of external intervention is where the regulator presents a non-binding suggested bargain to the players. We view such a suggested bargain as a Schelling-like focal point for the negotiations, and introduce a model mapping any such suggested bargain to an associated distortion over the distribution over bargaining outcomes. This model provides the regulator with a well-defined algorithm for choosing the suggested bargain that optimizes the expected welfare of the outcome of the negotiations. We use the application of flight rerouting negotiations to demonstrate Nash distributional bargaining management using such suggested bargains.

The type of optimal control of strategic agents captured by the Nash distributional bargaining management concept is a recent theme in the artificial intelligence literature. For example, there is a growing literature focused on optimally managing negotiations between strategic, self-interested agents (Brafman & Tennenholtz, 1996; Chalamish & Kraus, 2012; Lopez-Carmona, Marsa-Maestre, Klein, & Ito, 2012). The primary difference between our paper and the current literature is that our bargaining game is unstructured, while the others use a structured game. The management of unstructured flight rerouting negotiations can also be considered a type of constrained automated mechanism design, like the framework studied by Vorobeychik, Reeves, and Wellman (2012). That is, the mechanism in our paper is the bargaining scenario as modified by the (automated) external regulator for each unique flight rerouting game. Similarly, in the domain of air traffic management, the work of Agogino and Tumer (2012) studies a multi-agent approach to managing air traffic flows in which agents are reinforcement learners. Once again, the main difference between our paper and these others is the difference between controlling structured and unstructured games.

Despite the recent interest in optimal control of strategic agents, the foundation for these ideas can be found in more conventional concepts such as correlated equilibrium. As Ashlagi, Monderer, and Tennenholtz (2008) point out, correlated equilibrium is often conceptualized as arising from some external party that provides recommended actions to game participants but cannot enforce its recommendations. Ashlagi et al. look at the value of correlation, i.e. the welfare improvement arising from an external parties recommendations, which is directly analogous to the welfare improvement arising from Nash distributional bargaining management.

1.2 Roadmap of this Paper

In section 1.3 we start with an overview of our notation. Then in section 2, we provide a general definition of distributional bargaining concepts. After this we focus on the "Nash distributional bargaining concept", which is defined by NZ and probabilistic versions of TI, SI and IIA. We next prove that Nash distributional bargaining concepts take the form of a power law distribution. This result holds for any bargaining games, even non-convex, non-comprehensive ones, and even ones with a finite set of possible bargains.

In section 3, we discuss the Nash distributional bargaining concept. We focus on its mathematical structure, its physical meaning, and its normative implications for an external regulator who can change the feasible set. We show that the mode of the Nash distributional bargaining concept is the Harsanyi bargaining solution. We then point out that if we also impose Nash's symmetry axiom, we instead get the Nash bargaining solution as the mode of the distribution over outcomes.

In section 4 we introduce our model of flight rerouting negotiations. We show that this setting yields a feasible set that is not convex, comprehensive, or even connected. Here we also demonstrate the NDB model and its use in prediction. Here we also explore the issues of invertibility of the bargaining set mentioned above.

In section 5 we develop the concept of Nash distributional bargaining management, where an external regulator can make a recommendation to the bargainers or modify some aspect of the unstructured bargaining game. We demonstrate these concepts using the flight rerouting scenario and show that the external regulator can make welfare improvements.

We conclude with a discussion of the results and directions for future research.

1.3 Notation

We consider conventional N-dimensional unstructured bargaining problems. Any such problem is the pair of a bounded feasible set, $S \subset \mathbb{R}^N$, and a special disagreement point, $d \in S$. We assume S contains at least two elements, and that there exists an $x \in S$ such that $x \succ d$ (i.e, such that the generalized strict inequality $\{x_i > d_i \; \forall i\}$ holds). We refer to such a problem as **simple** if $d = 0$ and no point $x \in S$ (i.e., no set of N utility values for that bargaining problem) has a component less than 0. Given a bargaining problem (S, d), we use the term "outcome" to refer to any element of S. We will refer to (S, d) as **standard** if S contains an open set and is convex and comprehensive[4].

We define $A + B$ for two sets $A, B \subset \mathbb{R}^N$ to mean the set of all x that can be written as $a + b$ for some $a \in A, b \in B$, and similarly for $A - B$. (So in particular, $A - B$ is not the same as $A \setminus B$.) We will also assume the usual topology over $\mathbb{R}^N$, etc. Given any $A \subset \mathbb{R}^N$ and $k \in \mathbb{R}^N$ where $k \succ 0$, we define kA to be the subset of of $\mathbb{R}^N$ given by replacing each $x \in A$ by the Hadamard (i.e., component by component) product $kx \triangleq (k_1x_1, k_2x_2, \ldots)$. We define AB for any two sets $A, B \subset \mathbb{R}^N$ similarly.

As shorthand, throughout we use "$\int$" with the measure implicit. In particular, when feasible sets are finite, expressions like "$\int dx \ldots$" implicitly use the point mass measure, i.e., are equivalent to sums.

2. Nash Distributional Bargaining Concepts

From a Bayesian perspective, we are interested in the posterior probability density function, $P(x \mid S, d)$. Here we do not proceed by decomposing that distribution into a prior and a likelihood. Instead, in analogy to Nash's approach, we model $P(x \mid S, d)$ directly, by using Nash's axioms to restrict its form.

In doing this it will be convenient to simplify notation, so that (again following Nash) we talk in terms of maps from (S, d) to distributions over x rather than in terms of $P(x \mid S, d)$:

Definition 1. *An N-dimensional **(distributional) bargaining concept** is a map from any N-dimensional bargaining problem (S, d) to a probability density function with support restricted to S. (When S is countable, the image is a probability distribution rather than a probability density function, and we will generically refer to a "distribution" with the implicit cardinality of S making clear whether we mean a probability distribution or probability density function.)*

We will generically indicate a distributional bargaining concept with the symbol μ, and indicate its value for a problem (S, d) as $\mu_{S,d}$. So $\mu_{S,d}(x)$ is the non-negative real number produced by applying μ to (S, d) and then evaluating the resultant density function at x.

To use any unstructured bargaining approach to make a prediction about human bargaining behavior in a particular physical setting (laboratory or field), in addition to explicitly specifying the bargaining problem (S, d), there must also be a specification of the physical setting. Often this specification is only implicit, but here we make ours explicitly. We restrict attention to bargaining problems (S, d) and associated physical settings such that every $x \in S$ is physically possible. We then go further and also restrict attention to

4. Recall that S is comprehensive if $\forall x \in S, x \succeq y \succeq d$, $y \in S$.

physical settings such that all those physically possible $x \in S$ where $x \succ d$ occur with some non-zero probability. We justify this by noting that in most laboratory and field settings, even if the bargaining were structured with a Pareto efficient outcome that is a dominant Nash equilibrium of the underlying game, there would still be a non-zero probability that other outcomes arise, due to bounded rationality of the players. (A stronger version of this reasoning — which we do not invoke here — would require that even outcomes where some player does worse than at the disagreement point occur with non-zero probability.)

Given the foregoing, it is straightforward to translate Nash's unstructured bargaining axioms into the context of distributional bargaining concepts. In particular, SI, TI and IIA get translated as follows:

Definition 2. *An N-dimensional distributional bargaining concept μ is a **Nash** distributional bargaining (NDB) concept if for any $S \subset \mathbb{R}^N$,*

pSI $\forall x, y, d \in S$ *such that* $x, y \succ d$, $\forall k \succeq 1$,

$$\frac{\mu_{S,d}(x)}{\mu_{S,d}(y)} = \frac{\mu_{(kS,kd)}(kx)}{\mu_{(kS,kd)}(ky)},$$

pTI $\forall x, d \in S$ *such that* $x \succ d$, $\forall a \in \mathbb{R}^N$, $\mu_{S,d}(x) = \mu_{S-\{a\},d-a}(x-a)$,

pIIA $\forall T \subseteq S$, $\forall x, y, d \in T$ *such that* $x, y \succ d$,

$$\frac{\mu_{S,d}(x)}{\mu_{S,d}(y)} = \frac{\mu_{T,d}(x)}{\mu_{T,d}(y)}$$

NZ $\forall x, d \in S$ *such that* $x \succ d$, $\mu_{S,d}(x) \neq 0$

pSI means that if we rescale all coordinates of a problem to get a new problem, the relative probability of any two points x and y in the original problem is the same as the relative probability of the rescaled versions of those points in the new (rescaled) problem. This is the "translation" of Nash's scaling invariance axiom, SI, into the context of distributional bargaining concepts. For the results in this paper, we could have weakened pSI to the following:

pSI* *For any simple, standard, closed problem* $(S, 0)$, $\forall x, y \in S$ *such that* $x, y \succ 0$, $\forall k \succeq 1$,

$$\frac{\mu_{S,0}(x)}{\mu_{S,0}(y)} = \frac{\mu_{(kS,0)}(kx)}{\mu_{(kS,0)}(ky)} \tag{1}$$

For expository simplicity though, we use the stronger version given in Def. 2.

pTI implements the assumption of translation-invariance of utility functions. (Recall that in our notation, the subtraction operator on sets is not the same as set subtraction.) pIIA means that if we remove potential solutions other than x and y from a bargaining problem, the relative probabilities of x and y don't change. This is the "translation" of Nash's IIA axiom into the context of distributional bargaining concepts. Finally, as mentioned earlier, the NZ axiom says that all outcomes in which all players do strictly better than they do under the default outcome receive some positive probability of occurring. This axiom requires the solution concept to conform to the physical reality of human behavior.

We emphasize that pIIA and pSI are *not* assumptions about human behavior. Rather they are reflections of ignorance by the modeler, just as prior probabilities in Bayesian modeling are reflections of such ignorance. If the modeler does not know of a preferred scale with which to predict player behavior, then they want to use a distribution that is invariant under rescalings (Jaynes & Bretthorst, 2003). That is what pIIA formalizes. Similarly, if the modeler does not know how removing a set of alternatives would affect the relative probabilities of the remaining ones, then they want to use a distribution that does not affect those relative probabilities. That is what pSI formalizes.[5]

In the remainder of this section we work through a derivation of our main mathematical result, presented in Thm. 1. Loosely speaking, this result says that a bargaining concept μ must be a product (over the players) of power law distributions, if it is to obey the axioms presented above.

Given any two N-dimensional problems (S, d) and (T, d), define $W \equiv S \cup T$. So $S \subseteq W$. Therefore by pIIA, for any N-dimensional Nash bargaining concept μ, $\forall x, y \in S$ such that $x, y \succ d$, $\mu_{S,d}(x)/\mu_{S,d}(y) = \mu_{W,d}(x)/\mu_{W,d}(y)$. Similarly, $T \subseteq W$, and therefore $\forall x, y \in T$ such that $x, y \succ d$, $\mu_{T,d}(x)/\mu_{T,d}(y) = \mu_{W,d}(x)/\mu_{W,d}(y)$. This establishes the following:

Lemma 1. *Fix any two N-dimensional problems (S, d) and (T, d). $\forall x, y \in S \cap T$ such that $x, y \succ d$, for any N-dimensional NDB concept μ,*

$$\frac{\mu_{S,d}(x)}{\mu_{S,d}(y)} = \frac{\mu_{T,d}(x)}{\mu_{T,d}(y)} \tag{2}$$

The kinds of geometric distortions underlying existence proofs of the Nash, Kalai-Smorodinsky and egalitarian solutions cannot be applied when (S, d) is non-standard, and especially when S contains a finite number of elements. This inability to handle finite S has been one of the major obstacles of the Nash bargaining theory. In real-world unstructured bargaining, it is quite common for people to bargain over a finite number of possible outcomes, without ever considering the possibility of using a randomization device to decide the final bargain.[6]

In contrast, Lemma 1 means that if we can establish the form of $\mu_{T,d}$ for a particular class of problems (T, d) with infinite T, then we have established its form for any pair $(S \subset T, d)$. This is true even if S is neither convex nor comprehensive. Indeed, S can even be finite.[7] The following result provides such a form of $\mu_{T,d}$ for a particular class of problems (T, d).

5. Of course, if the modeler *does* have knowledge about preferred scales and / or about how the human players would change their distribution if some alternatives were removed, then that knowledge should replace pSI and /or pIIA, respectively.
6. Indeed, one could argue that real bargaining scenarios cannot involve an uncountably infinite number of bargaining outcomes. After all, real human beings cannot specify all digits of any real number in a given interval. When communicating with one another, real human beings are only able to specify numbers having a finite precision, together with a finite set of infinite-precision real numbers, like π, $\sqrt{2}$, etc.
7. Other authors have examined bargaining solutions on non-standard domains. For example, many papers extend the Nash solution to non-convex S (Kaneko, 1980; Herrero, 1989; Conley & Wilkie, 1996; Zhou,

Theorem 1. *Let μ be an N-dimensional NDB concept such that for any closed, simple, standard problem $(T, 0)$, $\mu_{T,0}(x)$ is differentiable throughout the interior of T. Then there are constants $a_i, i = 1, \ldots n$, independent of S and d, such that for all problems (S, d) and all $x \in S$ such that $x \succ d$, $\mu_{S,d}(x) \propto \prod_i (x_i - d_i)^{a_i}$.*

Proof. Define $S'' \equiv S - \{d\}$ and $x' \equiv x - d$. Then use pTI to write $\mu_{S,d}(x) = \mu_{S'',0}(x')$. So it suffices to prove the theorem for all $x' \succ 0$ for the bargaining problem $(S'', 0)$.

Let $S' \equiv \{S'' \setminus \{x'' \in S'' : \exists i \text{ where } x''_i \leq 0\} \cup \{0\}\}$. So S' is S'' without any x'' for which at least one player is not doing better than the disagreement point, unioned with the disagreement point. $(S', 0)$ is a simple bargaining problem. In addition, by pIIA, for all $x', y' \succ 0$

$$\frac{\mu_{S',0}(x')}{\mu_{S',0}(y')} = \frac{\mu_{S'',0}(x')}{\mu_{S'',0}(y')}$$

Therefore it suffices to prove the theorem for problem $(S', 0)$.

Let T be the convex, comprehensive closure of S'. So $(T, 0)$ is a simple, standard bargaining problem whose interior contains $(S', 0)$. Therefore by pIIA, it suffices to prove the proposition for $(T, 0)$.

Let Ω_T be the interior of T. Note that by hypothesis, $\mu_{T,0}(.)$ is a differentiable function over Ω_T.

For any N-dimensional $z \succ 0$, define the N-dimensional vector $\ln[z] \triangleq (\ln(z_1), \ln(z_2), \ldots)$. Make similar definitions for e^z and for the sets $\ln[A]$ and e^A where A is any subset of $\mathbb{R}^N$. Note that with the exception of points with zero-valued components, any point in T can be written as the vector e^z for some $z \in \mathbb{R}^N$, since $(T, 0)$ is a simple bargaining problem. So in particular, every point in Ω_T can be written e^z for some $z \in \mathbb{R}^N$.

Say that z is a point such that $\mu_{T,0}(e^z)$ is non-zero (so that e^z is in T). Define the associated scalar $\mathcal{P}(z) \triangleq \ln[\mu_{T,0}(e^z)]$. Note that $\mu_{T,0}(.)$ is non-zero over Ω_T by Def. 2(i). In addition, it has a continuous derivative throughout Ω_T by our hypothesis. Therefore $\mathcal{P}(.)$ has continuous derivative throughout $\mathcal{N} \equiv \ln[\Omega_T]$.

Choose a vector k all of whose components equal 1 except for component i, which is in the interval (0, 1]. Define the set of all such k as K_i. Note that since T is simple and standard, $K_i \Omega_T \subseteq \Omega_T$. Therefore $\ln[k] + \mathcal{N} \subseteq \mathcal{N}$.

Combining pSI and pIIA, $\forall x', y \in \Omega_T$

$$\frac{\mu_{T,0}(kx')}{\mu_{T,0}(ky)} = \frac{\mu_{T,0}(x')}{\mu_{T,0}(y)}$$

Taking logarithms we get

$$\mathcal{P}(\ln[k] + \ln[x']) - \mathcal{P}(\ln[k] + \ln[y]) = \mathcal{P}(\ln[x']) - \mathcal{P}(\ln[y])$$

$\forall x', y \in \Omega_T, k \in K_i$.

1997; Mariotti, 1998a). They generally do so by replacing certain of Nash's axioms. More recent work furthers the analysis of non-standard problems to include finite domains (Mariotti, 1998b; Xu & Yoshihara, 2006; Kıbrıs & Sertel, 2007; Peters & Vermeulen, 2010). However, a general conclusion of this work is that its goal is out of reach; this work typically fails to find a convincing solution concept that closely resemble Nash's solution and is single-valued for all finite problems.

Given that $\mathcal{P}$ must be differentiable across $\mathcal{N}$, we can take the partial derivative of both sides of this equation with respect to $\ln[k_i]$. Evaluate those derivatives in the limit as $k_i \to 1$. Doing this for all allowed x', y establishes that $\frac{\partial \mathcal{P}(u)}{\partial u_i} - \frac{\partial \mathcal{P}(v)}{\partial v_i} = 0 \; \forall u, v \in \mathcal{N}$, i.e., it establishes that $\frac{\partial \mathcal{P}(u)}{\partial u_i}$ is constant across $\mathcal{N}$. Therefore $\mathcal{P}$ is linear in coordinate i across $\mathcal{N}$.

Since this is true for all coordinates i, $\mathcal{P}(x') = \ln[\mu_{S',0}(e^{x'})]$ is a hyperplane across $\mathcal{N}$. Accordingly, $\mu_{S',o}(x')$ must be a product of monomials across Ω_T, i.e., $\mu_{T,0}(x') \propto \prod_i (x'_i)^{a_i}$ across Ω_T. (The coefficients of the hyperplane give the powers $\{a_i\}$ of the monomials.) Converting from x' back to x gives the claimed result. □

We refer to a distributional bargaining concept that meets the conditions in Thm. 1 as a **power law** distributional bargaining concept.

The axioms giving a power law distribution do not always hold in the real world. As a simple example, the default bargain d may serve as a focal point, in which case one might presume that $\mu_{S,d}(d) > 0$. In such cases, either the differentiability assumption of Theorem 1 must be relaxed, or one of the axioms defining NDB's must be. More generally, even when differentiability is assumed and we do not model d as a focal point, it may be possible to motivate other distributional bargaining concepts besides the power law bargaining concept. In particular, there are several axiomatic arguments that motivate predicting the behavior of a single decision maker according to a logit distribution in their utilities (Train, 2003). Logit distributions over utilities can also be motivated for modeling behavior of multiple interacting players in a noncooperative game, e.g., with the arguments originally used to motivate the logit Quantal Response Equilibrium (McKelvey & Palfrey, 1995), or with arguments based on maximum entropy inference (Wolpert, Harre, Olbrich, Bertschinger, & Jost, 2012). These arguments suggest (but do not formally derive) the idea of predicting behavior as a product of logit distributions, $\mu_{S,d}(x) \propto \prod_i exp(x_i)$, rather than with a product of monomials. (One advantage of such a distributional bargaining concept is that it does not require specification of d, since replacing x_i with $x_i - d_i$ in the exponents does not change the value of the probability distribution.)

For the rest of this paper, unless specified otherwise, we will restrict attention to power law distributional bargaining concepts. Note that the computational complexity of evaluating such a power law distributional concept is not an issue: one simply needs to evaluate the product of monomials to get the relative probabilities. If the normalization constant is also desired, at worst, it can be gleaned via standard Monte Carlo methods.

For (a power law bargaining concept) $\mu_{S,d}$ to be normalized when S is a standard problem, each a_i must exceed -1. More generally though, S could be countably infinite, with the utilities of player i given by $x_i \in \{1/k + d_i : k = 1, 2, \ldots\}$. For $\mu_{S,d}$ to be normalized in this case, a_i must exceed 0. Therefore by Lemma 1, a_i must always exceed 0.

3. Discussion of Power Law Distributional Bargaining Concepts

In this section we discuss assorted mathematical characteristics of power law distributional bargaining concepts, together with some of their physical implications.

3.1 The Relation between Power Law Distributional Bargaining Concepts and Point-Valued Bargaining Concepts

It should be emphasized that to derive the functional form of Thm. 1 there was no need for analogs of the last two of Nash's unstructured bargaining axioms, PAR and SYM. However it is straightforward to incorporate them if desired. One natural way to translate PAR into the context of distributional unstructured bargaining would be to require that improving every player's utility (i.e., changing any $x \succ d$ in a way that shrinks no component x_i) cannot decrease probability density. Imposing this on an NDB concept would rule out having any a_i less than 0.[8]

Similarly SYM can be translated to mean that all the a_i must have the same value. We will call a power law bargaining concept that meets this additional requirement a **fully Nash** distributional bargaining concept. For any fully NDB concept μ and any (S, d), $\mu_{S,D}$ over its support is the distribution $\prod_i (x_i - d_i)^\alpha$ for some scalar α. Now for fixed d, the level curves in x of $\prod_i (x_i - d_i)^\alpha$ are independent of α (assuming $\alpha \neq 0$). So in particular the maximum of $\prod_i (x_i - d_i)^\alpha$ is independent of α. Accordingly, to find that maximum, we can take $\alpha = 1$. Accordingly, the mode of $\mu_{S,d}(x)$ is the maximum over $x \in S$ of $\prod_i (x_i - d_i)$. In other words, the Nash bargaining solution of a bargaining problem (S, d) is the mode of any fully NDB concept applied to (S, d).

In addition, for any fixed $\delta > 0$, as we send $\alpha \to \infty$, the probability of the event $\{x$ lies more than δ away from the Nash bargaining solution$\} \to 0$. In this sense, all solutions other than the Nash solution become impossible in that limit.

However, in general if the feasible set of possible bargains is uncountable, the function taking joint bargains to joint utility outcomes will have a non-uniform Jacobian. In such cases, the most likely bargain is not the bargain corresponding to the Nash solution.

Furthermore, for finite α, in general the Bayes-optimal guess for x is *not* the Nash bargaining solution. For example, for quadratic loss functions, the Bayes-optimal guess for x is $\int dx\ \mu_{S,d}(x)x$. This differs from the Nash solution, $\text{argmax}_x \mu_{S,d}(x)$. Furthermore, consider expanding the i'th coordinate of the border of S, while leaving its border along axis j unchanged. Then for quadratic loss, this change to S will in general change the Bayes-optimal guess of x_j, even if it does not change the Nash bargaining solution point.

As a final comment, recall that is we do not impose the SYM requirement on our power law bagaining concept, the α_i of the players may differ, i.e., the players are allowed to be heterogenous. In this case the mode of the solution concept is not the Nash bargaining solution but the "weighted Nash solution" of (Harsanyi & Selten, 1972), evaluated for the very constants α_i that are the power law exponents of the NDB concept.

3.2 The Mathematical Structure of Power Law Bargaining Concepts

Nowhere does the definition of power law bargaining concept explicitly refer to differences $x_i - d_i$. So why do those differences arise in Thm. 1? Ultimately, the reason is that the definition of a bargaining problem requires specifying both the feasible set and a special

8. A more extreme way to impose PAR would be to make predictions using $\mu_{PAR(S),d}$, where $PAR(S)$ is the Pareto frontier of S. By Lemma 1, this is equivalent to "masking" $\mu_{S,d}$ to the Pareto frontier of S, i.e., to replacing it with a new distribution $\mu'_{S,d}$ whose support is restricted to the Pareto frontier of S, where for all x, x' on that frontier, $\mu'_{S,d}(x)/\mu'_{S,d}(x') = \mu_{S,d}(x)/\mu_{S,d}(x')$.

disagreement point d within the feasible set. Due to this, the translation invariance condition translates the point d along with the point x where μ gets evaluated. This causes the difference between those two points to be an invariant of μ, as reflected in Thm. 1.

Note also that $\mu_{S,d}$ is a product distribution over its support. In this sense, independence of the players is automatic. Such independence is not an explicit part of the definition of a power law bargaining concept. Ultimately, this independence arises from the fact that the SI axiom of Def. 2 involves scaling each player's utility independently.[9]

One must be careful in interpreting this player independence. It does *not* mean that under a distribution $\mu_{S,d}$, the utility values of the players are statistically independent. Unless S is a box, in general the border of S will serve to couple the players. For example, when $d = 0$, we have

$$\begin{aligned}
\prod_i \mu_{S,d}(x_i) &= \prod_i [\int_S dx_{-i} \; \mu_{S,d}(x_i, x_{-i})] \\
&\propto \prod_i [\int_S dx_{-i} \prod_j (x_j)^{a_j}] \\
&\not\propto \prod_i [x_i^{a_i}] \\
&= \mu_{S,d}(x),
\end{aligned} \tag{3}$$

i.e., the product of the marginal distributions of the player utilities does not equal the joint distribution of the player utilities.

As an example of the implications of this, say the border of S is changed by having the range of possible utility values for player i grow for some particular range of utilities of the other players, but stay the same elsewhere.[10] Then the statistical coupling between the players will generally change. Intuitively, after this modification to S, what I can infer about the likely value of $x_{j \neq i}$ given a particular value of x_i will have changed. (This is despite the independence of irrelevant alternatives axiom.)

3.3 Physical Meaning of Power Law Distributional Bargaining Concepts

Thm. 1 does not specify the values of $\mu_{S,d}(x)$ if $x \not\succ d$. However, it is often the case that for any $x \in S$ where $x \not\succ d$, there is an x' such that $x' \succ x, d$ where $\prod_i (x'_i - d_i)^{a_i}$ is arbitrarily close to zero.[11] Moreover, given that $x' \succ x$, it would be quite peculiar if in experiments it were the case that $\mu_{S,d}(x) > \mu_{S,d}(x')$. This strongly suggests that if μ meets the conditions in Thm. 1, then we should stipulate that $\mu_{S,d}(x) = 0$ for any $x \not\succ d$. (Formally though, we do not need to make that requirement in the analysis of this paper.)

9. For example, if we modified that axiom by first rotating the space $\mathbb{R}^N$, then applying the scaling operators, and then rotating back, we would no longer have a product distribution in the individual x_i, but rather in linear combinations of the x_i.
10. Formally, we specify some box over the values x_{-i}: $M \equiv \{x_j \in [b_j, t_j] : j \neq i, t_j > b_j \forall j \neq i\}$. Then over M alone, we extend the associated border of S along coordinate i, i.e., $\forall x_{-i} \in M$, we modify S by expanding the set of x_i such that $(x_i, x_{-i}) \in S$.
11. As an example, let S be a sphere centered on d, and consider any x such that $x_1 < 0$ while at the same time $x_2, x_3, \ldots > 0$. Then the distribution of the point $(\epsilon, x_2, x_3, \ldots)$ can be arbitrarily close to zero by taking ϵ small enough.

How the components a_i in Thm. 1 are physically interpreted is determined by how the distribution $\mu_{S,d}$ is physically interpreted. Ultimately, that has to do with what we physically mean by an "unstructured bargaining scenario".

One possible physical meaning is that $\mu_{S,d}$ is a population average over all humans and over all unstructured games that have feasible set S and disagreement point d. In this interpretation, if one adopts a power law bargaining concept, it makes sense to have the constants a_i be identical. Another possibility is that the distribution is not interpreted as a population average, but rather to refer to a set of N particular individuals involved in the bargaining problem at hand. In this case, if one adopts a power law bargaining concept, the constants a_i will differ in general. This is for several reasons. One is that different people have different bargaining styles, and different powers of persuasion. More generally, in some bargaining scenarios certain players will only have weak power to affect the outcome, or at least not be able to affect all aspects of the outcome.

In general these two interpretations are not mutually consistent. That's because averaging (a population of) different distributions all proportional to a product of monomials will not give a product of monomials, in general. Which interpretation one adopts ultimately depends on which interpretation one feels best characterizes the bargaining scenario under consideration.

There is also a third possible interpretation, in which one averages not only over the bargainers, but also over a set of *structured* bargaining scenarios. In this interpretation, the distribution of "unstructured bargaining" is interpreted as an average over distributions of structured bargaining scenarios. Nash's bargaining axioms would then be interpreted as reflecting the ignorance of the external modeler concerning the structure of the game the players are engaged in.

3.4 Knitting Together 2-Player Distributions to Get Multi-Player Distributions

There are problematic aspects to requiring that IIA applies to the full joint distribution of all N players' utilities when $n > 2$. However it seems less objectionable to stipulate that IIA applies to the distribution of any two players' utilities, *conditioned on the utilities of the other player(s)*. For example, we can require that $P(x_1, x_2 \mid x_3, \ldots x_N, S, d)$ obeys IIA.[12] We can also impose the other conditions defining power law bargaining concepts to any such conditional distribution. The result is that any such conditional distribution is a product of monomials, i.e., for all $i, j \neq i$,

$$P(x_i, x_j \mid x_{-i,-j}, S, d) \quad \propto \quad (x_i - d_i)^{a_i(x_{-i,-j})}(x_j - d_j)^{a_j(x_{-i,-j})} \tag{4}$$

where for full generality the exponents can vary with the value of $x_{-i,-j}$.

This is a set of $n(n-1)$ equations involving the full joint distribution $\mu_{S,d} = P(x \mid S, d)$, parameterized by the matrices $\{a_i(x_{-i,-j})\}$. There are an infinite number of joint distributions $P(x \mid S, d)$ that obey all of these equations simultaneously for *some* set of matrices $\{a_i(x_{-i,-j})\}$: by inspection, any distribution which is a power law bargaining

12. Note that this condition does *not* concern a of scenario where players 3 through N somehow fix their utilities to the values $x_3, \ldots x_N$, and after that the players 1 and 2 bargain. Rather it concerns the full N-player bargaining scenario where all N players bargain together.

concept (i.e., of the form $\mu_{S,d}(x) \propto \prod_i (x_i - d_i)^{a_i}$) obeys those equations for the degenerate case where the matrices are all constants.

More generally, requiring that all Eq.'s (4) simultaneously hold for some given set of matrices $\{a_i(x_{-i,-j})\}$ provides constraints on how the conditioned two-player bargaining concepts $P(x_i, x_j \mid x_{-i,-j}, S, d)$ "knit together" to give a full N-player bargaining concept. To be precise, given matrices $\{a_i(x_{-i,-j})\}$, the full joint distribution $\mu_{S,d}(x) = P(x \mid S, d)$ must obey the following equations:

$$\begin{aligned} \int dx \; \mu_{S,d}(x) &= 1, \\ \forall i, j \neq i, \;\; \mu_{S,d}(x) &\propto (x_i - d_i)^{a_i(x_{-i,-j})}(x_j - d_j)^{a_j(x_{-i,-j})} \int dx'_{-i,-j} \; \mu_{S,d}(x_i, x_j, x'_{-i,-j}) \end{aligned} \tag{5}$$

Future work involves investigating the properties of such "knitting together" conditional bargaining concepts.

4. NDB Flight Path Rerouting

In this section, we demonstrate NDB using the example of flight rerouting negotiations in the National Airspace System. Flight rerouting negotiations take place between the humans in an aircraft's cockpit and those manning air traffic control (ATC) when severe weather or air traffic result in the need for changes to the scheduled flight path while the aircraft is already airborne. Such en route rerouting is often referred to as "tactical rerouting", to distinguish it from "strategic rerouting", which takes place between the airlines operation center and ATC when a flight must be rerouted before it is airborne.

Tactical rerouting negotiations can be initiated by either the cockpit or ATC. Though they generally have a back-and-forth, offer/counter type of feel, these negotiations do not follow any set protocol. Therefore, an unstructured bargaining approach is appropriate for studying tactical rerouting negotiations. (Strategic rerouting negotiations might also be well suited to an unstructured bargaining approach.) In addition, it is worth emphasizing that although there is a natural tendency for a pilot to defer to ATC, in light of the greater knowledge of the latter, legally all responsibility ultimately lies with the pilot.

Tactical rerouting negotiations generally result in a distance/heading pair that determines a way-point through which the pilot agrees to fly. For example, a negotiation might look like this:

Cockpit: Denver Center, United 1492. Request 20 degrees right for weather.

Controller: United 1492, how long will you need that heading?

Cockpit: Looks like about 40 miles or so.

At this point the controller might grant permission, or he might make a counter proposal, such as "can you accept 20 [degrees] left?". In the latter case, negotiations often continue in the manner above. The distance/heading pair can be summarized by (l, θ), where $\theta \in$

$[-90, 90]$ is an angle in degrees, and $l \in [0, N]$ is a distance in miles.[13] This constitutes the set of bargains B, i.e., $B = [0, E] \times [-90, 90]$.

The cockpit and ATC have preferences over the bargains they can make, which are summarized by utility functions over the set of available bargains, B. For example, the cockpit doesn't want to fly too far off course, but it also doesn't want to fly too close to a storm center. In many cases, ATC might want what is best for the cockpit. Yet in other cases, ATC might be concerned about air traffic or other things (e.g., the impact of the rerouting on other traffic) that the cockpit does not know or care about.

To evaluate the utilities of a bargain (l, θ) for the cockpit and ATC, we need to evaluate certain features of the flight path that results from it. To determine the flight path that results from (l, θ) we first translate that bargain into a way-point in a Cartesian coordinate system. In this Cartesian coordinate system, we say that the flight's current position is $w_1 = (0, 0)$, and that the flight is pointing in the direction of the positive horizontal axis. Let θ_r be the radians equivalent of θ. This means that the agreed way-point (l, θ), is located at $w_2 = (l \cos\theta_r, l \sin\theta_r)$ in Cartesian coordinates.

After meeting its way-point, w_2, the flight will return to a "fix" further along its original flight path. In some rerouting negotiations, this fix is also part of the negotiated bargain. However, such bargains are relatively uncommon. To simplify our model, we assume that, whatever the bargain, the flight will return to the fix located along the horizontal axis at point $w_3 = (E, 0)$ in Cartesian coordinates. Using linear interpolation to connect w_1, w_2 and w_3, we create a constant-speed, constant-altitude 3D flight path, $\Lambda_{l,\theta}(t) = (e_{l,\theta}(t), n_{l,\theta}(t))$. In order to simplify notation, we will refer to the components of $\Lambda_{l,\theta}(t)$ as $e(t)$ and $n(t)$ with the dependence on (l, θ) implicit.

We now use $\Lambda_{l,\theta}(t) = (e(t), n(t))$ to calculate utility-relevant features of the bargain (l, θ). These features include the total length of $\Lambda_{l,\theta}(t)$, given by

$$\mathcal{L}_{l,\theta} = \int dt \sqrt{\frac{\partial e(t)}{\partial t}^2 + \frac{\partial n(t)}{\partial t}^2},$$

and whether or not $\Lambda_{l,\theta}(t)$ maintains a safe distance from weather with center C and radius R,

$$\mathcal{D}_{l,\theta} = \begin{cases} 1 & \text{if } \min_t ||\Lambda_{l,\theta}(t) - C|| > R \\ 0 & \text{otherwise.} \end{cases}$$

We assume the cockpit's utility is a linear combination of total length and maintenance of safe distance, given by:

$$s_c(l, \theta) = \alpha_{\mathcal{L}} \mathcal{L}_{l,\theta} + \alpha_{\mathcal{D}} \mathcal{D}_{l,\theta}$$

for real numbers $\alpha_{\mathcal{L}} \leq 0$ and $\alpha_{\mathcal{D}} \geq 0$.

We assume ATC's utility is a linear combination of maintenance of safe distance and the new flight path's impact on existing air traffic. We express this through the traffic penalty function $H(e, n)$, defined over Cartesian coordinates. The total traffic penalty for flight path $\Lambda_{l,\theta}(t)$ is

$$\mathcal{H}_{l,\theta} = \int dt \; H(e(t), n(t)).$$

13. We limit the angle to $[-90, 90]$ because heading changes of more than 90 degrees are extremely uncommon.

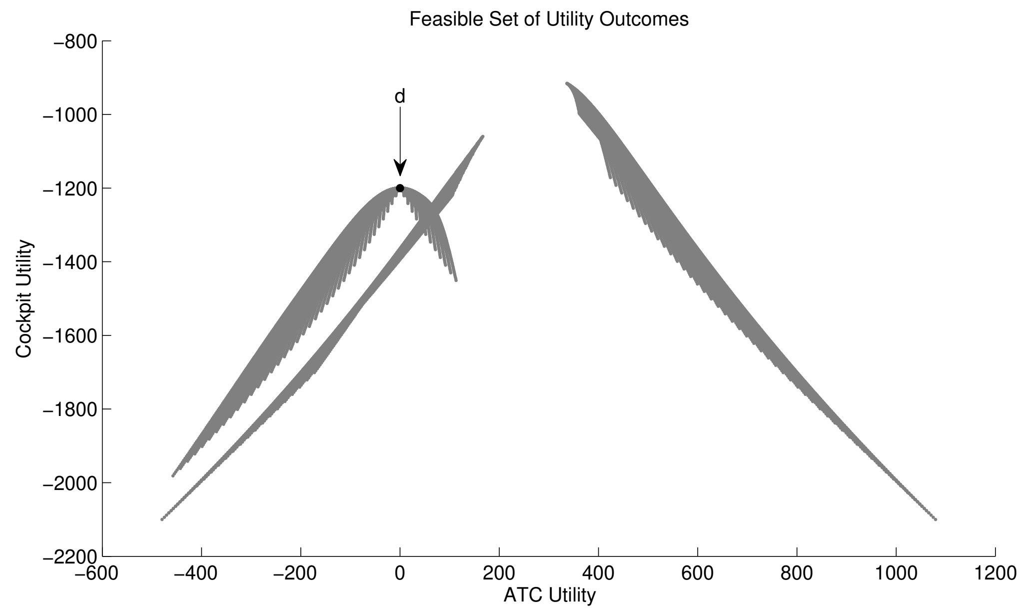

Figure 1: Feasible set of joint utility outcomes for the rerouting unstructured bargaining problem. Not only is the feasible set non-convex and not comprehensive, it is not even connected. Parameters are $a_c = a_{atc} = 1$, $\alpha_{\mathcal{L}} = -4$, $\alpha_{\mathcal{D}} = 300$, $\beta_{\mathcal{H}} = -0.01$ and $\beta_{\mathcal{D}} = 300$. The weather is a large circle with center $C = (150, 20)$ and radius $R = 40$. The traffic penalty is given by $H(e, n) = n$.

We express ATC's utility as $s_{atc}(l, \theta) = \beta_{\mathcal{D}}\mathcal{D} + \beta_{\mathcal{H}}\mathcal{H}_{l,\theta}$ for some real numbers $\beta_{\mathcal{D}} \geq 0$ and $\beta_{\mathcal{H}} \leq 0$.

The only thing left to specify is the outcome when negotiations break down, i.e., the default bargain, d_b. It might be that, in the event negotiations break down, the cockpit will choose the path that maximizes its own objectives without getting clearance from ATC. In this scenario, the punishment the cockpit receives as a result of changing its course without ATC approval should also enter the cockpit's objective. However, there are likely many complicated and variable factors to consider when modeling the effect of punishment, such as pilot attitude, idiosyncratic airline rules, etc. Hence, for simplicity, we simply take the default bargain to be the original flight path, i.e., the straight line between $(0, 0)$ and $(E, 0)$. Therefore, the bargaining game is given by

$$S = \{(x_c, x_{atc}) \in \mathbb{R}^2 : x_c = s_c(l, \theta) \text{ and } x_{atc} = s_{atc}(l, \theta) \text{ for some } (l, \theta) \in B\}$$

and $d = (s_c(E, 0), s_{atc}(E, 0))$.

Figure 1 shows the feasible set of utility outcomes associated with the flight rerouting model. Note that this set is neither convex nor comprehensive. Furthermore, there are breaks, so that the feasible set consists of disconnected subsets. Despite these irregularities of the feasible set, the NDB distribution can be applied trivially, giving

$$\mu_{S,d}(x) \propto (x_c - d_c)^{a_c} (x_{atc} - d_{atc})^{a_{atc}}$$

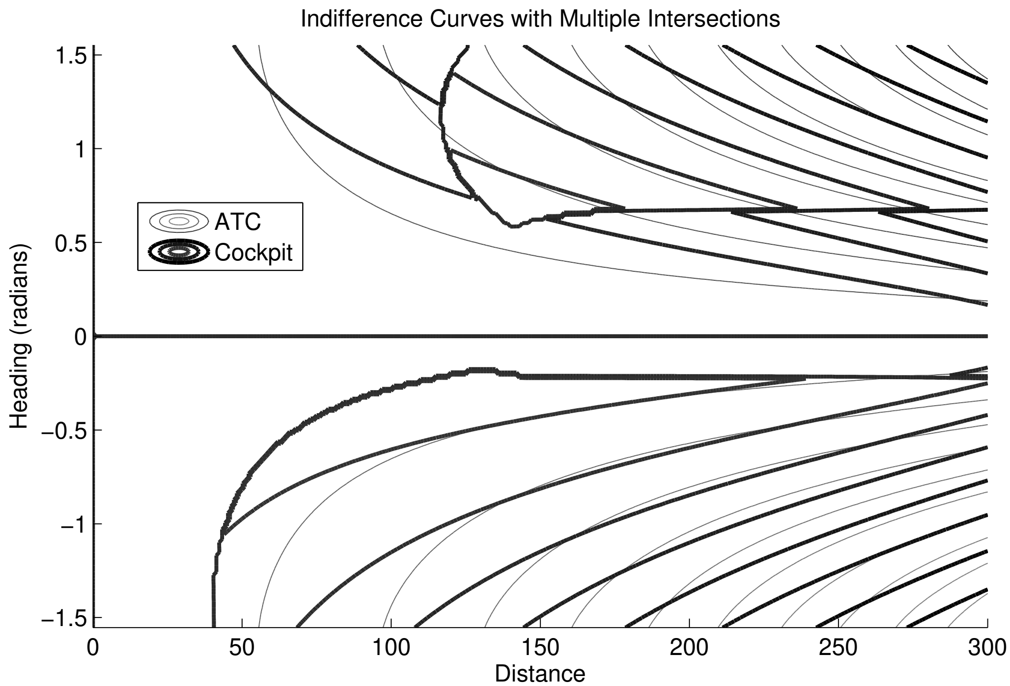

Figure 2: Indifference curves of s_c and s_{atc} plotted in the space of bargains for the flight path rerouting problem. $s(.,.)$ is not invertible because the indifference curves have multiple crossings. Parameters are $a_c = a_{atc} = 1$, $\alpha_{\mathcal{L}} = -4$, $\alpha_{\mathcal{D}} = 300$, $\beta_{\mathcal{H}} = -0.01$ and $\beta_{\mathcal{D}} = 300$. The weather is a large circle with center $C = (150, 20)$ and radius $R = 40$. The traffic penalty is given by $H(e, n) = n$.

for all $x \in S$ such that $x \succ d$, and $\mu_{S,d}(x) = 0$ otherwise.

As discussed in the introduction, if we want to use the NDB distribution over joint utilities to make predictions about the bargains in $b \in B$, but don't have a likelihood function $P(b \mid x)$, then we must be able to invert the mapping $s = (s_c, s_{atc})$. Unfortunately, s is not invertible, as can be seen by noting from figure 2 that the indifference curves of s_c and s_{atc} have multiple crossings.

Recall that if the map s were invertible, then to translate from the distribution over joint utilities to the distribution over bargains we would have to use the Jacobian of s. To be precise, if (l, θ) is a point in B, then the NDB distribution over bargains, $\mu^B_{S,d}$, evaluated at (l, θ) is given by

$$\mu^B_{S,d}(l, \theta) = \mu_{S,d}\left(s_c(l, \theta), s_{atc}(l, \theta)\right) |J(l, \theta)|,$$

where $J(l, \theta)$ is the determinant of the Jacobian of s,

$$\begin{vmatrix} \frac{\partial s_c}{\partial l} & \frac{\partial s_c}{\partial \theta} \\ \frac{\partial s_{atc}}{\partial l} & \frac{\partial s_{atc}}{\partial \theta} \end{vmatrix}.$$

Note though that even if s were invertible, the Jacobian might not be well-defined. Indeed, our flight-rerouting s contains discontinuities where flight paths cross from outside to inside

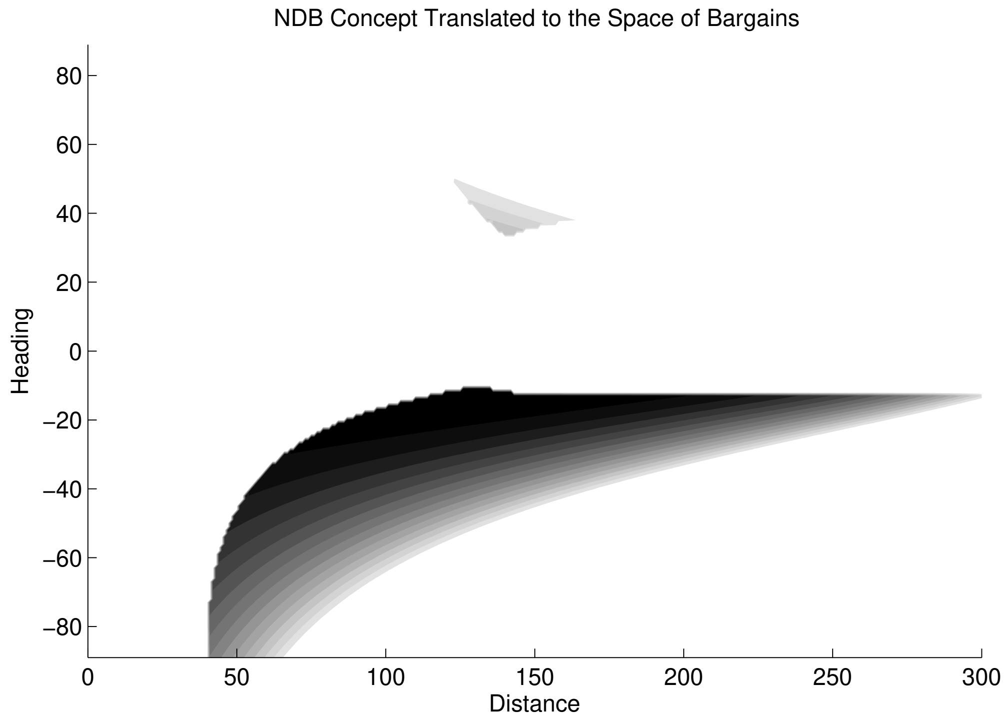

Figure 3: NDB concept mapped back into the space of bargains $B' = [0, 1, ..., N] \times [-90, -89, ..., 89, 90]$, for $a_c = a_{atc} = 1$, $\alpha_{\mathcal{L}} = -4$, $\alpha_{\mathcal{D}} = 300$, $\beta_{\mathcal{H}} = -0.01$ and $\beta_{\mathcal{D}} = 300$. The weather is a large circle with center $C = (150, 20)$ and radius $R = 40$. The traffic penalty is given by $H(e, n) = n$.

the weather radius, and the variable $\mathcal{D}_{l,\theta}$ jumps from 1 to zero. Partial derivatives of s_c and s_{atc} do not exist at such points in B.

Say that the set of bargains were countable, e.g., a grid imposed by the technological constraints of aircraft or the cognitive constraints of pilots and air traffic controllers. In this case there would be no need to worry about discontinuities or how to calculate the Jacobian. s might still be non-invertible however, so that we cannot transform the distribution over joint utilities to a distribution over bargains. In fact this happens with the flight-rerouting s when the set of bargains is modified to be a grid of integer distances and degrees, i.e., $B' = [0, 1, ..., N] \times [-90, -89, ..., 89, 90]$.

Regardless of the cardinality of B, one *can* translate a distribution over joint utilities to a distribution over bargains even when s is non-invertible, so long as one knows the likelihood of bargains given joint utility outcomes. As an example, say that B is countable, and all bargains that give rise to any particular joint utility x are equally likely. In this case, if there are $|B_x|$ bargains that give rise to x, then the probability of any such bargain is simply $\mu_{S,d}(x)/|B_x|$. Using this assumption and the grid $B' = [0, 1, ..., N] \times [-90, -89, ..., 89, 90]$, figure 3 shows the NDB distribution for the flight-rerouting problem in the space of bargains.

5. NDB Management

Suppose a regulator, external to an unstructured bargaining game, has preferences over the joint utility of the bargain reached, formalized as a real-valued social welfare function $U(x \in S)$. Suppose further that they can change some characteristics ρ of (S, d). For example, they might be able to change the joint utility d of the default bargain in certain ways, or replace S with one of several $S' \subset S$. In such cases, assuming they model the game with an NDB $\mu_{S,d}$, the regulator should choose

$$\begin{aligned}\rho &= \text{argmax}_\rho \mathbb{E}_\rho(U(x)) \\ &= \text{argmax}_\rho \int dx\; U(x)\mu_{S_\rho,d_\rho}(x)\end{aligned} \tag{6}$$

More generally, ρ may affect $s(.), d_B$, and/or B, not just their image (S, d). In this more general setting the regulator should choose

$$\begin{aligned}\rho &= \text{argmax}_\rho \mathbb{E}_\rho(U) \\ &= \text{argmax}_\rho \int dx\; U(x)\mu_{s_\rho(B^\rho),s_\rho(d_B^\rho)}(x)\end{aligned} \tag{7}$$

More generally still, the social welfare function U might be replaced with a function W defined over the space of bargains B. So long as s_ρ is invertible for all ρ, Eq. 7 still holds for this variant. (Just define $U(x \in s_\rho(B_\rho)) \equiv W(s_\rho^{-1}(x))$.) However if s_ρ is non-invertible for certain ρ, this equation must be modifed. In this situation, assuming they have the associated likelihood function, the regulator should choose the action

$$\begin{aligned}\rho &= \text{argmax}_\rho \mathbb{E}_\rho(W(b)) \\ &= \text{argmax}_\rho \int dxdb\; P(b \mid x)W(b)\mu_{s_\rho(B^\rho),s_\rho(d_B^\rho)}(x)\end{aligned} \tag{8}$$

In addition to these types of actions by the regulator, there are others that are not considered in typical axiomatic analyses of unstructured bargaining — including the analysis of this paper up to this point. An example is where the regulator's intervention is simply to suggest a bargain to the players before they start to bargain, i.e., to provide them a "focal bargain" for their bargaining. To advise the regulator in such cases, we need to model the effects of these types of action by the regulator, and use that model in Eq. 7 .

We refer to any of these types of optimal choice by the external regulator as "NDB management". We do not mean to claim that it is always possible to implement NDB management. In some situations an external regulator may not be able to do anything that will have a substantial effect on the distribution over possible bargains by the players. In the rest of this section we present some preliminary analysis of NDB management, concentrating on an example where the regulator's possible action is to provide a focal bargain.

5.1 NDB Management in Rerouting Negotiations

Here we demonstrate a simple form of NDB management using the example of rerouting negotiations from Section 4, related to the focal point concept introduced by Schelling (1960). Specifically, we assume that the regulator can make a suggestion for a flight rerouting to both the cockpit and ATC before they begin rerouting negotiations. The idea is that this suggestion will serve as a focal bargain for negotiations, having the effect of raising the relative probability of bargains similar to the suggested one. This suggestion by an "external regulator" could be implemented by some automated software that operates both in the cockpit and in ATC.[14]

Use $\mu^B_{S,d}(l,\theta)$ to indicate the distribution over $(l,\theta) \in B$ induced by the NBD for game (S,d). (So if needed, $\mu^B_{S,d}(l,\theta)$ implicitly depends on the Jacobian of s and a likelihood function $P(b \mid x) = P(l,\theta \mid x)$.) Then we assume the regulator models the effects of a suggested focal bargain (l_0,θ_0) on the distribution over B as

$$P_{l_0,\theta_0}(l,\theta) \propto \left(\frac{\psi_{l_0,\theta_0}(l,\theta)}{Z(B)}\right)^{\gamma} \mu^B_{S,d}(l,\theta)^{1-\gamma}$$

where ψ is a bivariate Gaussian distribution with mean (l_0,θ_0) and covariance matrix Σ, $\gamma \in [0,1]$ measures the strength of the impact of the manager's suggestion on the players' behavior, and $Z(B)$ is the normalization constant for ψ. We call this the "NDB management distribution". It is clearly an *ad hoc* model of focal bargains; we use it here only as a simple way to demonstrate NDB management.

We assume that to minimize her computational requirements, the regulator requires that the space of bargains be a grid of integer distances and degrees, i.e., $B' = [0, 1, ..., E] \times [-90, -89, ..., 89, 90]$.[15] As mentioned in Section 4, even in this finite set of possible bargains there are instances where multiple bargains map to the same joint utility outcomes. As usual, to address this we must specify $P(b \mid x)$; here we make the simple assumption that this likelihood is uniform over all b that map to x, and zero for all others. Accordingly,

$$\mu^B_{S,d}(l,\theta) = \frac{\mu_{S,d}(s(l,\theta))}{|B_x|},$$

where $|B_x|$ is the number of bargains in B that give rise to outcome x.

As usual, the regulator will choose her action — which here means suggest a heading and distance — to maximize the resultant expected social welfare. Writing the social welfare function as $W(l,\theta)$, this optimal action is

$$\text{argmax}_{(l_0,\theta_0)\in B}\mathbb{E}_{l_0,\theta_0,S,d}\,(W) = \int_B dl\ d\theta\ W(l,\theta)P_{l_0,\theta_0}(l,\theta).$$

In general, W might depend on the impact of the rerouting on traffic flow over the entire National Airspace System, in addition to depending on the utility of the cockpit and of ATC. For simplicity though, here we take $W(l,\theta) = \tau s_c(l,\theta) + (1-\tau)s_{atc}(l,\theta)$ where $\tau \in [0,1]$.

Figure 4 depicts the effect of a regulator's recommendation on the predicted flight path. The left panel is the NDB distribution over bargains, $\mu^B_{S,d}$. It represents the prediction before a regulator makes a suggestion. The NDB modal bargain produces the detoured flight path denoted in the center panel by the sequence of two lines joined by an open

14. Experimental studies find empirical support for a theory of focal points and bargaining (Binmore, Swierzbinski, Hsu, & Proulx, 1993).
15. In practice, this may not need to be an explicit requirement, since it is unlikely that non-integer values would be considered by the players.

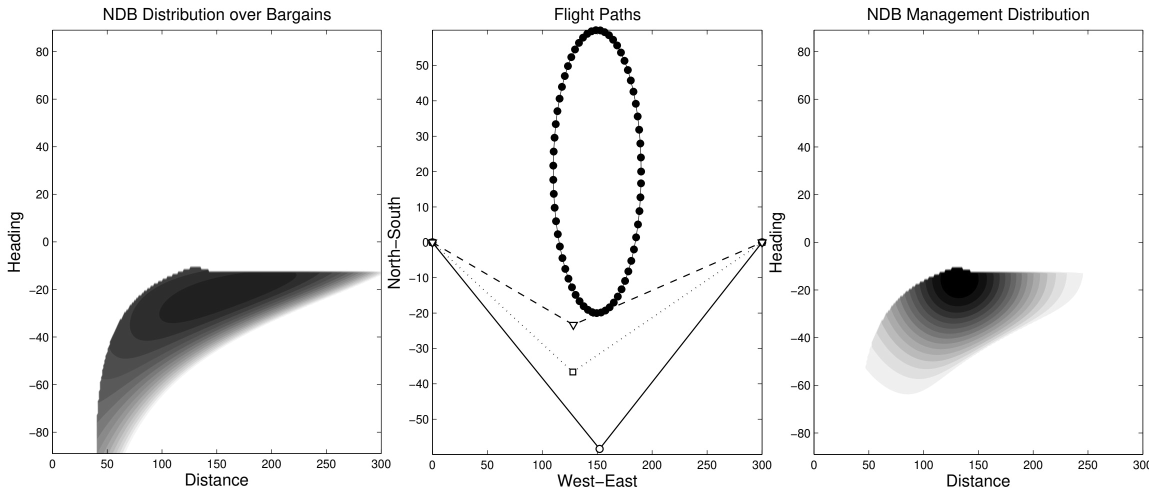

Figure 4: Left Panel: NDB distribution over $B' = [0, 1, ..., E] \times [-90, -89, ..., 89, 90]$ with $a_c = 1$, $a_{atc} = 3$, $\alpha_{\mathcal{L}} = -4$, $\alpha_{\mathcal{D}} = 300$, $\beta_{\mathcal{H}} = 0.01$ and $\beta_{\mathcal{D}} = 300$. Center Panel: Weather is the large circle with center $C = (150, 20)$ and radius $R = 40$, NDB modal flight path (circle line), NDB management modal flight path (square line), NDB management recommended waypoint (triangle line). Right Panel: NDB management distribution with $\tau = .5$, $\gamma = .5$, $\sigma_\theta = 360$ and $\sigma_l = 1,200$. The traffic penalty is given by $H(e, n) = n$.

circle; the flight begins its detour at the point $(0, 0)$ and returns to its original flight path at $(300, 0)$. The reason the modal flight path does not pass closer to the weather is that the traffic penalty is given by $H(e, n) = n$, i.e., traffic is denser in the north. This means ATC is rewarded for flight paths that pass further to the south. And since $a_{atc} = 3 > 1 = a_c$, ATC is relatively successful at convincing the cockpit to accept bargains that result in longer, more southerly flight paths.

The NDB management distribution is shown in the right panel. It represents the prediction after the regulator has made an optimal suggestion. In this case, the optimal suggestion is $(130, -10)$. The path associated with the optimal suggestion is denoted in the center panel by the pair of lines joined by a triangle. The modal bargain of the NDB management distribution induced by that suggestion is denoted by the pair of lines joined by a square. Note that the regulator's suggestion has the effect of skewing the distribution toward bargains with shorter distances l and smaller heading changes. This is because the regulator assigns equal weight to ATC and cockpit utility, whereas the NDB distribution is skewed toward outcomes that are better for ATC. As a result, the modal bargain shifts significantly to the north.

6. Conclusion and Future Work

In both laboratory and field unstructured bargaining experiments, it is typically found that more than one bargain can arise. To accommodate this, in this paper we consider maps that take any unstructured bargaining problem S to a probability distribution over S, rather than to a single element of S. Our approach is to "translate" Nash's axioms of unstructured bargaining to apply to this distribution-valued map. Doing this, we derive the "Nash distributional bargaining concept", which maps any feasible set of joint utility outcomes to a power law over that set. This power law nature is intriguing due to the ubiquity of empirical power laws in the real economy, e.g., in wage distributions, city sizes, etc.

Future work involves trying to translate variants of Nash's axioms into a distributional bargaining concept. For example, it may prove possible to translate the weak monotonicity axiom of the Kalai-Smorodinsky solution concept (Kalai & Smorodinsky, 1975) into probabilistic terms. If so, then by combining it with our probabilistic versions of the remaining Kalai-Smorodinsky axioms (which are shared with the axioms of the Nash bargaining concept), we may be able to produce a "Kalai-Smorodinsky distributional bargaining concept".

There are many advantages to using distributional bargaining concepts, in addition to their according with the experimental fact that multiple bargains can arise for any given game. One advantage is that such concepts seamlessly accommodate feasible sets that are not convex and comprehensive, and even finite feasible sets. Another advantage arises if there is an external regulator who can modify some aspects of the bargaining game and has their own utility function over bargaining outcomes. By modifying the game, the regulator changes the associated distribution, and therefore modifies the value of their own expected utility. Accordingly, they can calculate the Bayes-optimal modification to the game.

We emphasize that there is nothing in the formal definition of a bargaining game or in our translation of Nash's axioms that restricts our analysis to scenarios that are traditionally viewed as "bargaining". In particular, we do not assume binding contracts. Therefore there is no reason to assume that the outcome of the game has to be Pareto optimal. (Without a binding contract on a player's behavior, that player will typically have both an incentive and the ability to change her behavior in a way that helps her but hurts her opponent, and thereby moves the final outcome off of the Pareto frontier.) For this reason, distributional bargaining concepts can be applied to model any noncooperative game where the modeler only knows the feasible set of the game's joint utility outcomes, and does not know (or cannot tractably elaborate) the full underlying extensive form of the game. In modeling noncooperative game behavior with such limited information, the best the modeler can do is provide a distribution over the final joint utility outcome. Intuitively speaking, such a distribution amounts to an "average" over all extensive form games that have the associated feasible set of possible joint utilities. The Nash distributional bargaining concept provides a way to form such a distribution over outcomes, using axioms that seem applicable in many scenarios not traditionally viewed as unstructured bargaining.[16]

16. The only subtlety with using the Nash distributional bargaining concept this way is that there needs to be an outcome of the noncooperative game that can reasonably be expected to behave as the "default point" d. However even this restriction can be dispensed with using some distributional bargaining concepts, e.g., using the logit one mentioned at the end of Sec. 2.

As an illustration, in the work of Smith, Suchanek, and Williams (1988) an artificial speculative market was run for T rounds. This created a bubble which burst, leaving some test subjects poorer and some richer. Predicting the moves in the full extensive form game the subjects engaged in over the T rounds is laborious, at best. As an alternative, one could imagine predicting the experiment's outcome based only on knowing its feasible set, i.e., knowing all ways the money could end up divided among the subjects at the end of the experiment. Future work involves analyzing the Nash distributional bargaining concept for such situations.

Such market experiments could prove to be a useful source of data for testing alternative distributional bargaining theories, in addition to experiments explicitly viewed in terms of unstructured bargaining. In particular, if it is possible to produce a Kalai-Smorodinsky distributional bargaining concept, as discussed above, it may be possible to compare the predictions of such an alternative theory against the NDB concept using data from market experiments. Such a comparative study would be analogous to the types of meta-studies conducted to compare non-cooperative solution concepts using experimental data (Wright & Leyton-Brown, 2010).

As discussed above, the function s plays a prominent role in making predictions about bargains rather than about joint utilities. Arguably, knowledge of s should even affect the predictions we make over joint utility outcomes. As an example, say the set of possible bargains is so large that the players cannot be expected to examine it completely, and that s concentrates the vast majority of bargains onto a single joint utility x, e.g., one on the Pareto frontier. In this case, it would seem reasonable to ascribe higher probability to that x than we would if only a single bargain mapped to it. However traditional approaches to unstructured bargaining ignore s when predicting joint utilities, relying on axioms solely concerned with the space of joint utilities. For the most part, that is the approach we adopted here. (The section on optimal recommendations diverged from this in modeling the effect of the recommendation.) Future work involves incorporating knowledge of s directly into the predictions of joint utility outcomes.

Acknowledgments

We would like to thank Sylvia Thoron for helpful discussion.

References

Agogino, A., & Tumer, K. (2012). A multiagent approach to managing air traffic flow. *Autonomous Agents and Multi-Agent Systems*, *24*, 1–25. 10.1007/s10458-010-9142-5.

Ashlagi, I., Monderer, D., & Tennenholtz, M. (2008). On the value of correlation. *Journal of Artificial Intelligence Research*, *33*, 575–613.

Aumann, R., & Hart, S. (1992). *Handbook of Game Theory with Economic Applications*. North-Holland Press.

Aydoğan, R., & Yolum, P. (2012). Learning opponent's preferences for effective negotiation: an approach based on concept learning. *Autonomous Agents and Multi-Agent Systems*, *24*, 104–140. 10.1007/s10458-010-9147-0.

Binmore, K., Swierzbinski, J., Hsu, S., & Proulx, C. (1993). Focal points and bargaining. *International Journal of Game Theory*, *22*, 381–409. 10.1007/BF01240133.

Brafman, R. I., & Tennenholtz, M. (1996). On partially controlled multi-agent systems. *Journal of Artificial Intelligence Research*, *4*, 477–507.

Brafman, R. I., & Tennenholtz, M. (2003). Learning to coordinate efficiently: A model-based approach. *Journal of Artificial Intelligence Research*, *19*, 11–23.

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.

Chalamish, M., & Kraus, S. (2012). Automed: an automated mediator for multi-issue bilateral negotiations. *Autonomous Agents and Multi-Agent Systems*, *24*, 536–564. 10.1007/s10458-010-9165-y.

Conley, J. P., & Wilkie, S. (1996). An extension of the nash bargaining solution to nonconvex problems. *Games and Economic Behavior*, *13*, 26–38.

Duan, L., Doğru, M., Özen, U., & Beck, J. (2012). A negotiation framework for linked combinatorial optimization problems. *Autonomous Agents and Multi-Agent Systems*, *25*, 158–182. 10.1007/s10458-011-9172-7.

Harsanyi, J., & Selten, R. (1972). A generalized nash solution for two-person bargaining games with incomplete information. *Management Science*, *18*, 80–106.

Herrero, M. J. (1989). The nash program: Non-convex bargaining problems. *Journal of Economic Theory*, *49*(2), 266 – 277.

Jaynes, E. T., & Bretthorst, G. L. (2003). *Probability Theory : The Logic of Science*. Cambridge University Press.

Kalai, E. (1977). Proportional solutions to bargaining situations: Interpersonal utility comparisons. *Econometrica*, *45*(7), 1623–1630.

Kalai, E., & Smorodinsky, M. (1975). Other solutions to nashs bargaining problem. *Econometrica*, *43*(3), 513518.

Kaneko, M. (1980). An extension of the nash bargaining problem and the nash social welfare function. *Theory and Decision*, *12*, 135–148. 10.1007/BF00154358.

Kıbrıs, O., & Sertel, M. (2007). Bargaining over a finite set of alternatives. *Social Choice and Welfare*, *28*, 421–437. 10.1007/s00355-006-0178-z.

Lopez-Carmona, M., Marsa-Maestre, I., Klein, M., & Ito, T. (2012). Addressing stability issues in mediated complex contract negotiations for constraint-based, non-monotonic utility spaces. *Autonomous Agents and Multi-Agent Systems*, *24*, 485–535. 10.1007/s10458-010-9159-9.

Mariotti, M. (1998a). Extending nash's axioms to nonconvex problems. *Games and Economic Behavior*, *22*(2), 377 – 383.

Mariotti, M. (1998b). Nash bargaining theory when the number of alternatives can be finite. *Social Choice and Welfare*, *15*, 413–421. 10.1007/s003550050114.

McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, *10*, 6–38.

Nash, J. (1950). The bargaining problem. *Econometrica*, *18*(2), 155162.

Nydegger, R. V., & Owen, G. (1974). Two-person bargaining: An experimental test of the nash axioms. *International Journal of Game Theory*, *3*, 239–249. 10.1007/BF01766877.

Osborne, M., & Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press, Cambridge, MA.

Peters, H., & Tijs, S. (1984). Probabilistic bargaining solutions. In *Operations Research Proceedings*. Springer-Verlag.

Peters, H., & Vermeulen, D. (2010). WPO, COR, and IIA bargaining solutions. accepted by International Journal of Game Theory.

Rezek, I., Leslie, D. S., Reece, S., Roberts, S. J., Rogers, A., Dash, R. K., & Jennings, N. R. (2008). On similarities between inference in game theory and machine learning. *Journal of Artificial Intelligence Research*, *33*, 259–283.

Roth, A. E., & Malouf, M. W. K. (1979). Game-theoretic models and the role of information in bargaining. *Psychological Review*, *86*(6), 574–594.

Rubinstein, A., Safra, Z., & Thomson, W. (1992). On the interpretation of the nash bargaining solution and its extension to non-expected utility preferences. *Econometrica*, *60*(5), 1171–1186.

Schelling, T. (1960). *The strategy of conflict*. Harvard university press.

Smith, V. L., Suchanek, G. L., & Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, *56*(5), pp. 1119–1151.

Train, K. E. (2003). *Discrete Choice Methods with Simulation*. Cambridge University Press.

Vorobeychik, Y., Reeves, D. M., & Wellman, M. P. (2012). Constrained automated mechanism design for infinite games of incomplete information. accepted by Journal of Autonomous Agents and Multiagent Systems.

Wolpert, D. H., Harre, M., Olbrich, E., Bertschinger, N., & Jost, J. (2012). Hysteresis effects of changing parameters in noncooperative games. *Physical Review E*, *85*, 036102. DOI: 10.1103/PhysRevE.85.036102.

Wolpert, D. H., & Bono, J. W. (2011). Distribution-valued solution concepts. working paper.

Wright, J. R., & Leyton-Brown, K. (2010). Beyond equilibrium: Predicting human behavior in normal form games. In *Twenty-Fourth Conference on Artificial Intelligence (AAAI-10)*. forthcoming.

Xu, Y., & Yoshihara, N. (2006). Alternative characterizations of three bargaining solutions for nonconvex problems. *Games and Economic Behavior*, *57*(1), 86 – 92.

Zhou, L. (1997). The nash bargaining theory with non-convex problems. *Econometrica*, *65*(3), 681–685.