Mechanism Design Based on Beliefs about Responsive Play (Position Paper)

Yevgeniy Vorobeychik and Michael P. Wellman
University of Michigan
Computer Science & Engineering
2260 Hayward
Ann Arbor, MI 48109-2121 USA
{ yvorobey, wellman }@umich.edu

ABSTRACT

In general, identifying a solution concept only incompletely specifies a mechanism design problem. The designer must consider which among a multiplicity of solutions is likely to be played, as well as the possibility that actual play will not correspond to any solution. Given that actual play is the ultimate determiner of a mechanism's success, we advocate that designers embrace the corresponding forecasting problem and evaluate candidate mechanisms with respect to belief distributions over players' response. Solution concepts can play a useful role in delimiting and structuring belief distributions. We propose that membership of prospective strategy profiles in various solution classes be treated as evidence bearing on their likelihood of play. Flexible solution classes, for example based on approximate equilibrium, degree of dominance, or safety level, provide natural measures (e.g., distance from equilibrium) that can be employed in defining belief distributions.

1. MOTIVATION

Mechanism design theory has enjoyed considerable academic acceptance, particularly in the domain of auction design. At the core of the mechanism design framework is application of game-theoretic solution concepts (e.g., Bayes-Nash or dominant strategy equilibrium) to characterize the expected strategy profiles to be played by agents in response to the designer's mechanism choice. The mechanism is then evaluated under the assumption that agents play is consistent with the adopted solution concept.

Conclusions about mechanism quality, then, are only as good as the prediction represented by this solution concept. One problem is that standard equilibrium concepts generally admit multiple solutions. Even if we follow the typical approach within game theory and seek refinements [4], we still generally fail to achieve point predictions. Typical treatment within the mechanism design literature has either made the optimistic assumption that agents will play the most favorable equilibrium (weak implementation), or has approached the problem pessimistically, designing for the worst possible equilibrium outcome (strict implementation) [5, 6]. The former relies on the designer's ability to persuade the agents to play the desirable strategies. The latter allows agents to choose a solution arbitrarily, but as a result may sacrifice good design choices.

A second problem is that actual play may fail to coincide with any prescribed solution. Traditional game-theoretic equilibrium concepts have not always fared well in experimental studies assessing their descriptive validity [11, 2]. The situation might reasonably be expected to be worse in practice (i.e., outside the laboratory), where actual mechanisms and environments may be more complex than can be reflected in design models. Indeed, given realistic bounded agents and complex environments, we should perhaps be suspicious of any sharp prediction about the strategies to be played by agents in response to mechanism designs.

To sum up, few will argue that a mechanism designer is primarily concerned with predicting agent play given a particular mechanism choice, but perfectly accurate point predictions seem elusive both in theory and in practice, while multiplicity of predictions introduces ambiguity into the design process that is not easily resolved.

In this position paper, we suggest that unique prediction of play is generally undesirable. Instead, we advocate that the designer formulate a *belief distribution of agent play* that leverages a flexible solution concept. By using distributions of agent play, we move away from the rigidity of relying on a particular solution concept as *the* prediction, allowing rather that agents may play an arbitrary profile with some probability, albeit perhaps very small. Of course, as a special case we can concentrate all probability mass on a solution concept, while providing flexibility in the distribution over alterative solutions in the class. We can also accommodate relaxations of the traditional weak or strong implementation by adjusting our beliefs in optimistic or pessimistic directions.

The key issue in adopting our approach is how to construct reasonable probability distributions. Our proposal is to define parameterized versions of solution concepts, where the parameters provide "hooks" upon which probabilities of play could hinge. We suggest several such concepts, many of which are relaxations of the traditional ones such as Nash and dominant strategy equilibria.

One important source of distribution models for play may be to capture the uncertainty in our underlying model of strategic interactions between agents, as well as the effect that mechanism choice has upon it. For example, any modeling effort will generally leave open some uncertainty in regards to the agent preferences or preference distributions, insofar as the designer may have made mistaken assumptions or failed to account for important secondary elements that influence agent play. Such uncertainty may be expressed as a distribution of agent payoffs around an estimate expressed by the model. A distribution over agent play may then follow from standard solution concepts applied to the distribution over game models. The difficulty with this approach is that it may require the designer to find a solution correspondence, which is a formidable task in general. Furthermore, in real settings, the designer may have considerable trouble formalizing an explicit distribution over models. To address these problems, albeit imperfectly from the theoretical perspective, we introduce *heuristic beliefs*, that is, beliefs that are derived heuristically from probabilistic indicators of particular game theoretic solutions.

2. PRELIMINARIES

In describing our approach, we restrict our attention to *normal form games*,[1] denoted by $[I, \{R_i\}, \{u_i(r)\}]$, where I refers to the set of players and $m = |I|$ is the number of players. R_i is the set of strategies available to player $i \in I$, with $R = R_1 \times \cdots \times R_m$ representing the set of joint strategies of all players. It is often convenient to refer to a strategy of player i separately from that of the remaining players. To accommodate this, we use r_{-i} to denote the joint strategy of all players other than player i. We define the payoff (utility) function of each player i by $u_i : R_1 \times \cdots \times R_m \to \mathbb{R}$, where $u_i(r_i, r_{-i})$ indicates the payoff to player i to playing strategy r_i when the remaining players play r_{-i}.

We model the strategic interactions between the mechanism designer and participating agents as a two-stage game [13]. The designer moves first by selecting a value, θ, from a set of allowable mechanism settings, Θ. All the participant agents observe the mechanism parameter θ and move simultaneously thereafter. For example, the designer could be deciding between a first-price and second-price sealed-bid auction mechanisms, with the presumption that after the choice has been made, the bidders will participate with full awareness of the auction rules.

Since the participants play with full knowledge of the mechanism parameter, we define a game between them in the second stage as $\Gamma_\theta = [I, \{R_i\}, \{u_i(r, \theta)\}]$. We refer to Γ_θ as a game *induced* by θ. Let $\mathcal{N}(\theta)$ be the set of strategy profiles considered *solutions* of the game Γ_θ. Traditionally, Nash or dominant strategy equilibria have been adopted as the solution concepts, making $\mathcal{N}(\theta)$ the set of equilibria of the appropriate type. Below, we define several alternative solution concepts which may be more appropriate than the traditional concepts for many design scenarios.

Suppose that the goal of the designer is to optimize the value of some welfare function, $W(r, \theta)$, dependent on the mechanism parameter and resulting play, r. Since we will allow for many possible outcomes of agent play, we can evaluate the objective function for a given game abstractly as follows. We define $W_T(\hat{R}, \theta) = T_{\hat{R}} W(r, \theta)$, where T is some functional acting on $W(r, \theta)$. Several examples of T commonly found in the literature are $\inf_{\hat{R}}$ (representing strict implementation) and $\sup_{\hat{R}}$ (representing weak implementation). We have already argued that both of these are somewhat extreme. Instead, we will concentrate on an alternative: we let $T_{F,\hat{R}}$ to be the expectation with respect to some probability distribution F over $\hat{R}$. Then, $W_{T_{F,\hat{R}}}(\hat{R}, \theta) = E_{F,\hat{R}} W(r, \theta)$. Given a description of the solution correspondence $\mathcal{N}(\theta)$ and $W_T(\mathcal{N}(\theta), \theta)$, the designer faces a standard optimization problem.

3. SOLUTION CONCEPTS

In evaluating the welfare function with respect to the distribution of agent play, we relied on a specification of a solution concept which in effect provides the support of this distribution. In this section, we explore a number of ideas about how solution concepts can be defined. We begin with a standard Nash equilibrium concept and its immediate relaxation to approximate Nash equilibria, and go on to relax several other solution concepts in a similar fashion with the hope that we can thereby incorporate all strategy profiles that we may find to be plausible agent play. We then describe a complementary solution concept that may serve as an additional hook to deriving distributions of play by allowing the designer to indirectly model risk aversion of agents. Indeed, all of the solution concepts we propose have relaxation parameters which can serve the designer in defining probabilities of play for various profiles.

3.1 Nash and Approximate Nash Equilibria

Perhaps the most common solution concept for games is *Nash equilibrium*, defined as follows:

DEFINITION 1. *A strategy profile* $r = (r_1, \ldots, r_m)$ *constitutes a* Nash equilibrium *of game* $[I, \{R_i\}, \{u_i(r)\}]$ *if for every* $i \in I$, $r_i' \in R_i$,

$$u_i(r_i, r_{-i}) \geq u_i(r_i', r_{-i}).$$

When R is the set of all pure strategies, the above defines a *pure strategy Nash equilibrium*; alternatively, if R is the set of all mixed strategies, the definition describes a *mixed strategy Nash equilibrium*.

We often appeal to the concept of an *approximate*, or *ϵ-Nash equilibrium*, where ϵ is the maximum benefit to any agent for deviating from the prescribed strategy.

DEFINITION 2. *A strategy profile* $r = (r_1, \ldots, r_m)$ *constitutes an* ϵ-Nash equilibrium *of game* $[I, \{R_i\}, \{u_i(r)\}]$ *if for every* $i \in I$, $r_i' \in R_i$,

$$u_i(r_i, r_{-i}) + \epsilon \geq u_i(r_i', r_{-i}),$$

where $\epsilon \geq 0$.

Naturally, for any ϵ, the set of ϵ-Nash equilibria contains the set of Nash equilibria. In this sense, it is a weaker concept. The philosophical difficulty, of course, is that while Nash equilibrium is a mutual best response, at least one agent has an incentive to deviate from an ϵ-Nash equilibrium. A typical justification of the concept is that agents may be indifferent to small improvements in payoff, but have a great desire for coordination (so, for example, if the designer "offers" them an ϵ-Nash equilibrium to play, such that ϵ is very small, they'll happily agree and no one would deviate). Alternatively, we may ascribe to agents a cost to finding an actual best response which is greater than ϵ. If this is the case, no agent will be willing to deviate once an ϵ-Nash strategy profile is common knowledge and no better alternative is available.

Yet another interpretation might be that agents exert some (bounded) effort to find a better response, and their likelihood of succeeding is inversely related to ϵ. In this and other interpretations, the key property is that ϵ provides a parametric hook for describing probability of play. The greater the incentive agents have to deviate from a particular profile (all else equal), the less likely we deem such a profile to be played.

3.2 Alternatives to (ϵ)-Nash Equilibria

While the Nash equilibrium has long been the primary solution concept for games, many have expressed dissatisfaction with it in the context of real mechanism design problems and agent strategic considerations. For example, Erev and Roth [2] provide experimental evidence that a reinforcement learning algorithm tends to be a better predictor of actual play in games with a unique equilibrium. In a similar vein, Selten [11] presents a series of arguments

[1] By employing the normal form, we model agents as playing a single action, with decisions taken simultaneously. The general approach and arguments presented here could also apply to games in extensive form, or indeed to any game form and its associated solution concepts.

against Bayesian rationality as a reasonable predictor and, even, as an effective normative system. Rosenthal has proposed restricting strategies considered to *rules of thumb* [10, 9], that is, simple patterns of play, conditioned on context, that have proven effective over time. Many studies in multiagent systems research effectively take this approach, experimentally estimating what Walsh et al. [14] call a *heuristic strategy payoff matrix*. Our own group has further developed this methodology under the heading *empirical game-theoretic analysis* [8, 13, 15].

In the heuristic spirit of these restricted-strategy approaches, we can also develop rules for further narrowing a space of profile candidates based on strategic analysis. In this section, we describe several concepts that may prove useful for this purpose.

3.2.1 Nearly Dominant Profiles

Dominant strategies are often regarded as especially compelling solutions, and dominant strategy equilibrium is a commonly used solution concept in mechanism design. Unfortunately, aside from the common textbook examples, such as the famed Prisoner's Dilemma, dominant strategies rarely exist in real strategic scenarios. In comparison, a Nash equilibrium is guaranteed to exist in all finite games.

An intuitive property of a rule of thumb is that it *usually* works *reasonably* well. The way we can translate this idea into a solution concept is by introducing ϵ-dominant strategies.

DEFINITION 3. *A profile, r is ϵ-dominant if* $\forall i \in I, \forall t \in R_{-i}$,

$$u_i(r_i, t) + \epsilon \geq u_i(r', t),\ \forall r' \in R_i.$$

While dominant strategies are rare, strategies that are nearly dominant may be more common, and may indeed provide a solid basis for certain rules of thumb. Of course, there will always be an ϵ-dominant strategy if we set ϵ to be high enough (as long as the payoff functions are bounded). However, once ϵ is large, such strategies are no longer nearly dominant in any meaningful way. Still, this solution concept may be a useful and reasonably compelling way to model agent play without appealing to the hyperrational Nash equilibrium. From the players' viewpoint, nearly dominant pure strategies are easy to find with respect to a given game, although the algorithmic question of finding the entire set of (pure and mixed) nearly dominant profiles may be a bit more involved and will be subject of future work. From a designer's viewpoint, they are reasonable things to expect agents to play.

As in the case of ϵ-Nash equilibria, an important advantage of this relaxation of the dominant strategy profiles is that we can derive distributions of play based on "how dominant" a particular profile is. That is, we can assess a relatively low likelihood of play to profiles in which at least one agent may be significantly better off by playing another strategy for some deviation of other agents. Alternatively, we may fix ϵ and assess zero probability of play to profiles which are not ϵ-dominant (assuming, of course, that there is at least one strategy profile that is).

3.2.2 Nearly Undominated Profiles

In his seminal work, Pearce [7] describes the notion of rationalizable strategies. While the set of all rationalizable strategies is not always identical to the set of strictly undominated strategies, the two concepts are closely related, and are indeed appealing on similar grounds. The argument of Pearce was that the Nash equilibrium concept was too strong to describe actual behavior. By weakening it to a set of *plausible* strategy profiles that may be observed, actual behavior may be explained, although no longer modeled precisely.

While the idea that players are unlikely to play a profile that is strictly dominated (or not rationalizable) is very intuitive, there is experimental evidence to suggest that dominated strategies (for example, cooperative play in Prisoner's Dilemma) may indeed be played in practice [1]. As a consequence, we introduce here an even weaker concept of nearly undominated or ϵ-undominated strategies, which include strategies that, while dominated, are very close to being optimal for some strategy profile that other agents may play.

DEFINITION 4. *A profile, r, is ϵ-undominated if* $\forall i \in I, \exists t \in R_{-i}$, *such that*

$$u_i(r_i, t) + \epsilon \geq u_i(r', t),\ \forall r' \in R_i.$$

Since this solution concept is very weak, it allows the designer to retain most strategy profiles in a game as plausible rules of thumb, eliminating only those that are clearly extremely harmful to at least one agent. The assumption that a very poorly performing strategy would never be played is quite plausible in real situations. For example, a strategy that consistently loses the agent one million dollars, as compared to any other strategy that incurs no loss, carries considerable appeal of being assessed probability zero of play.

Given a set of nearly undominated strategies, we can also imagine that a likelihood of nearly dominant or nearly Nash profiles would be greater based on their nearness to the corresponding solution concept. Thus, we can actually combine all three of the relaxed solution concepts we have so far discussed to obtain a distribution of agent play.

3.2.3 Safety of Pure and Mixed Profiles

Risk aversion is a common feature of preferences, and given a precise understanding of agent preferences can be accounted for in payoff functions. However, such precise models of risk attitude may not be easy to come by. As an alternative, some game-theoretic studies appeal to worst-case criteria such as safety level. To capture this approach in a flexible solution concept, we define the notion of δ-safety.

DEFINITION 5. *Let R_{-i} be the joint space of deviations of agents other than i. A profile r is δ_i-safe for agent i if*

$$\delta_i(r) \geq \max_{t \in R_{-i}} (u_i(r) - u_i(r_i, t)).$$

A profile r is then δ-safe if it is δ_i-safe for all agents, that is, if

$$\delta(r) \geq \max_{i \in I} \max_{t \in R_{-i}} (u_i(r) - u_i(r_i, t)).$$

Alternatively, r is δ-safe if, for every player $i \in I$,

$$u_i(r) \leq u_i(r_i, t) + \delta,\ \forall t \in R_{-i}$$

We do not see the notion of δ-safety as having much independent value. Instead, we view it as a useful way to distinguish particular types of rule-of-thumb strategies that players may consider. For example, we can imagine that in a set of approximate Nash equilibria, there may be profiles that would be extremely sensitive to deviations by players, and, therefore, have a higher bound on δ-safety. The notion of δ-safety thus provides a mechanism designer with an additional assessment of likelihood of play by indirectly accounting for risk aversion of agents without having to quantify it.

4. CONSTRUCTING DISTRIBUTIONS OF AGENT PLAY

Our goal in this section is to offer several ideas about constructing belief distributions of agent play that rely on game-theoretic solution concepts (thus taking the players' incentives seriously), but admit varying degrees of commitment to their defining criteria.

The first approach is for the designer to decide exactly which solution concept is the best model for the strategic behavior of the agent pool. The designer may choose a very weak concept and make as few assumptions as possible about agent rationality and common knowledge, and we provided some guidance about such choices in the preceeding section. Once the solution concept is chosen, the designer will need to assess the relative likelihood of solutions, for example, modeling each solution as equally likely to be played. Alternatively, the designer may wish combine a solution concept with some notion of δ-safety as we previously defined, and thereafter devise a distribution that puts higher probability on solutions with low δ. Note that inherent in this approach is the assumption that any profile that is not a solution will be played with probability zero.

Yet another approach is to develop a distribution of play based on the relaxation parameter within a solution concept. Previously, we defined several ϵ- and δ-concepts (ϵ-Nash, ϵ-dominant, ϵ-undominated, δ-safe). Each pure strategy profile in the game will most certainly be any such ϵ-concept for some value of ϵ (the same is true of δ-safety). We could then assess the probability of play for a particular profile to be inversely proportional to its value of ϵ or δ for the selected solution concept. Typically, our assessment of probability of play will be positive for every pure strategy profile in the game. We could also develop similar probabilistic models of agent play based on combinations of these solution concepts, thus allowing the designer to hedge its bets among competing criteria.

The approaches above presume the designer is certain about the game model itself. This, of course, is suspect, and indeed some effort has been made within the game theory community to assess the quality of particular solution concepts based on how well they survive such modeling noise [3]. Here, we take another approach, and assume that payoffs specified by the designer are unbiased samples from a Normal distribution with some variance. Variance here would need to be specified by the designer, based on confidence in the model. This judgment too entails modeling assumptions, but we expect that adding the extra degree of freedom will generally be helpful.

The distributions of play that we described above are all *conditional* on a particular game. Thus, in order to find the distribution of play in the face of uncertainty about the actual game, the designer would need to take the expectation of this conditional distribution with respect to the distribution of games. Let us try to say this in a somewhat more formal language. Suppose we fix the game, Γ, that the agents will play, and choose a solution concept, $\mathcal{C}$. We designate the distribution of agent play conditional on Γ and the solution concept by

$$P_{\Gamma,\mathcal{C}}(r) = \Pr\{r \text{ is played}|\Gamma, \mathcal{C}\},$$

where r is a pure strategy profile. This may be specified by the designer as suggested above. Now, in order to derive the posterior probability of agents playing a pure strategy profile r given a particular solution concept, we would simply take the expectation with respect to the distribution of games:

$$P_{\mathcal{C}}(r) = \Pr\{r \text{ is played}|\mathcal{C}\} = E_{\Gamma}[P_{\Gamma,\mathcal{C}}].$$

This can be done for every pure strategy profile to obtain a distribution of agent play.

There are several important shortcomings in the approach just suggested. The first is simply that it requires numerical integration, as we do not have a closed-form expression of this expectation for any solution concept that we have discussed. This happens to be a relatively significant problem, since numerical techniques here would require computing the entire set of solutions for each of some finite set of games, which may quickly become impractical when the game is relatively large. Another limitation is that the designer is still required to specify a model of beliefs given a game, even though he may be uncertain about his model of the players' payoffs. Often, the designer may find himself incapable of doing even that very sensibly, but would instead like to have a more heuristic, though systematic, approach to modeling agent play. To this end, we propose a *heuristic* distribution of play, which does not require any modeling on designer's part, except his choice of a solution concept.

Observe that the distribution over agent payoff functions induces a probability that each pure strategy profile r is a solution $\mathcal{C}$. Elsewhere, we derived closed-form expressions for these probabilities [12], which we can now use to obtain a very simple heuristic distribution over play by normalizing as follows:

$$\Pr\{r \text{ is played}|\mathcal{C}\} = \frac{\Pr\{r \text{ is } \mathcal{C}\}}{\sum_{r' \in R} \Pr\{r' \text{ is } \mathcal{C}\}}.$$

Since we have a closed-form expression for $\Pr\{r' \text{ is } \mathcal{C}\}$, computation is greatly simplified.

5. CONCLUSION

In this work, we focused on an important practical shortcoming of mechanism design theory: lack of effective methods for point-predictions of actual play. Indeed, we believe that making very precise predictions of agent play is undesirable, as complexity of real strategic settings will generally make models imperfect enough that even the most appropriate solution concept will not necessarily make a good predictor of play. Instead, we argued for the need to evaluate mechanism choices with respect to belief distributions of play which are based on flexible solution concepts. Thus, game theoretic notions may eagerly enter the distributions of play, but need not define them entirely.

It seems most useful in devising distributions of agent play to have solution concepts with relaxation parameters which allow relative assessment of probabilities on different strategy profiles. We proposed the degree of approximation of a solution concept as an example of such a parameter. Additionally, we suggested a complementary solution concept which may be used to incorporate the designer's beliefs about the risk aversion of agents in the distribution of play.

Finally, we proposed a number of examples of how solution concepts that we suggested may be used in deriving distributions of agent play. These approaches can be easily extended to incorporate uncertainty about the designer's model of the strategic scenario, although incurring considerable computational effort. To alleviate this difficulty, as well as to consider yet another alternative sensible way to develop distributions of play, we introduced heuristic distributions, which are based on closed-form expressions of probabilities that each profile is a solution.

We believe that there is still a considerable gap between theoretical mechanism design and its practical applications. In this work, we suggested that this gap may be narrowed if the designer has effective ways to determine the distribution of agent play based on game-theoretic notions. While we proposed a number of methods to this end, much work needs to be done to verify whether these are truly effective in practical settings, or whether others need to be developed in their place. We hope is that our criticisms and approach will stimulate further research in this direction.

6. REFERENCES

[1] Russell Cooper and Thomas W. Ross. Cooperation without

reputation: Experimental evidence from prisoner's dilemma games. *Games and Economic Behavior*, 12:187–218, 1996.

[2] Ido Erev and Alvin E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed-strategy equilibria. *American Economic Review*, 88:848–881, 1998.

[3] Drew Fudenberg, David M. Kreps, and David K. Levine. On the robustness of equilibrium refinements. *Journal of Economic Theory*, 44:354–380, 1988.

[4] David M. Kreps. *Game Theory and Economic Modelling*. Oxford University Press, 1990.

[5] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[6] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[7] David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52:1029–1050, 1984.

[8] Daniel M. Reeves, Michael P. Wellman, Jeffrey K. MacKie-Mason, and Anna Osepayshvili. Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, 39:67–85, 2005.

[9] Robert W. Rosenthal. Barbaining rules of thumb. *Journal of Economic Behavior and Organization*, 22:15–24, 1993.

[10] Robert W. Rosenthal. Rules of thumb in games. *Journal of Economic Behavior and Organization*, 22:1–13, 1993.

[11] Reinhard Selten. Evolution, learning, and economic behavior. *Games and Economic Behavior*, 3:3–24, 1991.

[12] Yevgeniy Vorobeychik. Mechanism design and analysis using empirical game models. Technical report, University of Michigan, 2006.

[13] Yevgeniy Vorobeychik, Christopher Kiekintveld, and Michael P. Wellman. Empirical mechanism design: Methods, with application to a supply chain scenario. In *Seventh ACM Conference on Electronic Commerce*, Ann Arbor, MI, 2006.

[14] William E. Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, Edmonton, 2002.

[15] Michael P. Wellman. Methods for empirical game-theoretic analysis (extended abstract). In *Twenty-First National Conference on Artificial Intelligence*, Boston, 2006.