Selecting Strategies Using Empirical Game Models: An Experimental Analysis of Meta-Strategies

Christopher Kiekintveld
University of Michigan
Computer Science and Engineering
Ann Arbor, MI 48109-2121 USA
ckiekint@umich.edu

Michael P. Wellman
University of Michigan
Computer Science and Engineering
Ann Arbor, MI 48109-2121 USA
wellman@umich.edu

ABSTRACT

In many complex multi-agent domains it is impractical to compute exact analytic solutions. An alternate means of analysis applies computational tools to derive and analyze empirical game models. These models are noisy approximations, which raises questions about how to account for uncertainty when analyzing the model. We develop a novel experimental framework and apply it to benchmark *meta-strategies* – general algorithms for selecting strategies based on empirical game models.

We demonstrate that modeling noise is important; a naïve approach that disregards noise and plays according to Nash equilibrium yields poor choices. We introduce three parameterized algorithms that factor noise into the analysis by predicting *distributions* of opponent play. As observation noise increases, rational players generally make less specific outcome predictions. Our comparison of the algorithms identifies logit equilibrium as the best method for making these predictions. Logit equilibrium incorporates a form of noisy decision-making by players. Our evidence shows that this is a robust method for approximating the effects of uncertainty in many contexts. This result has practical relevance for guiding analysis of empirical game models. It also offers an intriguing rationale for behavioral findings that logit equilibrium is a better predictor of human behavior than Nash equilibrium.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Artificial Intelligence—*Distributed Artificial Intelligence, Multi-Agent Systems*

General Terms

Experimentation, Economics

Keywords

Computational game theory, meta-strategy analysis, empirical game modeling, multi-agent systems, uncertainty, strategy selection, experimentation, simulation, methodology

1. INTRODUCTION

Game theory provides a rich set of tools for analyzing multi-agent decision problems. These methods require a complete and precise specification of the players, their strategy sets, and the payoffs for the possible outcomes. For complex domains with large strategy spaces or stochastic outcomes, it may be implausible to specify such a model fully. In these cases we can facilitate analysis by applying computational methods to derive an *empirical* model of the game. Versions of this approach have been used to study domains including TAC Travel [17], TAC SCM [23, 20], continuous double auctions [16], and simultaneous ascending auctions [17]. The success of these studies raises additional questions about which methods yield the best results. We develop a novel framework for evaluating alternative approaches and apply it to study ways selecting strategies using empirical models.

Players using empirical game models to make strategy choices are playing a game of *incomplete information* [9], similar to the *global games* introduced by Carlsson and van Damme [6, 14]. The common solution concept for games of incomplete information is Bayes-Nash Equilibrium (BNE), which enforces consistent beliefs and simultaneous payoff maximization for all players. Unfortunately, computing exact Bayes-Nash equilibria is generally infeasible for the games we are interested in. Finding BNE is known to be intractable in the general case [7]. In this specific case, the information available to players is very complex, implying intricate beliefs. Carlsson and van Damme [6] note that global games are mathematically challenging beyond the the 2 player, 2 action model they study. Moreover, the details of how players can learn about the game are likely to vary across applications, limiting the value of any exact solution that depends on a specific observation model.

Our approach focuses on evaluating general solution algorithms using empirical evidence, seeking to identify the state of the art. Each algorithm takes as input an empirical game model and returns a strategy to play. We evaluate performance using *regret* – how much a player could gain by using a different algorithm. This measure is related to how well an algorithm approximates a BNE of the full game of incomplete information. One of the important advantages of this approach is that we can easily vary the class of game and the observation model. This yields insights into the robustness of the algorithms to different forms of uncertainty.

We begin by introducing *meta-games* to model the empirical observation process. Our solution algorithms are *meta-strategies* in this model. We show that ignoring the uncertainty in empirical observations and naïvely applying Nash equilibrium yields poor strategy choices. We then introduce three parameterized algorithms that generalize equilibrium notions to predict distributions over the possible outcomes. These predictions depend on the noise in an intuitive way; as noise increases, broader predictions are preferable. Our evaluation shows that an algorithm based on *logit equilibrium* [13] consistently outperforms the others.

Cite as: Selecting Strategies Using Empirical Game Models: An Experimental Analysis of Meta-Strategies, Christoper Kiekintveld and Michael P. Wellman, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.),May,12-16.,2008,Estoril,Portugal,pp.1095-1102.

This result has immediate implications for applications of empirical game-theoretic analysis, such as the previously-cited studies of auctions and market games. Identifying point Nash equilibria of empirical games is not very informative; instead, analysis should focus on distributional outcome predictions. We identify logit equilibrium as the best known method for making such predictions. This work also has an interesting connection to behavioral game theory. Logit equilibrium was originally used to explain differences between human behavior and Nash equilibrium predictions in laboratory experiments [13, 1, 8]. Our results provide a stronger rationale for this behavior as a robust method to deal with many forms of uncertainty.

2. THE META-GAME MODEL

Figure 1 depicts an abstract model of the interaction between players using empirical game models to make decisions. This model is intended to capture the salient features of typical applications of empirical game-theoretic reasoning, including payoff uncertainty and the absence of common knowledge. The essence of the model is a standard one-shot normal-form game embedded within an observation process.[1] A *base game* representing the true underlying payoffs is selected at random from a specified game class. Each player independently observes a different *empirical* game, representing the player's noisy estimate of the base game. The players select strategies based on the empirical games and receive payoffs defined by their joint choices and the base game payoffs.

We refer to the full model as a *meta-game*, and strategies in the meta-game as *meta-strategies*. The general strategy selection algorithms we are interested in map directly to meta-strategies in this model. We can view players as selecting meta-strategies (solution algorithms) that act on their behalf once game observations are made. Modeling solution algorithms as strategies in a larger game motivates evaluation of the algorithms with game-theoretic concepts, such as regret and equilibrium.

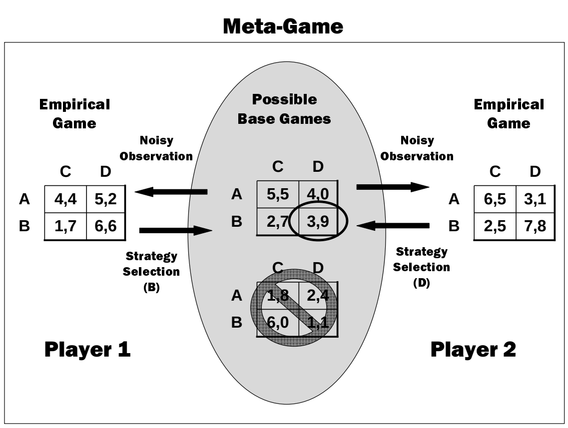

Figure 1: A graphical depiction of a meta-game showing selection of the base game, generation of empirical games, and strategy selections based on the empirical games. Here, player 1 chooses strategy B and receives payoff 3, while player 2 chooses strategy D and receives payoff 9.

[1]We study the one-shot case in part to make a clean distinction between empirical learning of game parameters and opponent modeling, which are both possible in repeated interactions.

Meta-games are games of incomplete information [9] in which the signals are empirical game models. There are also many interesting issues in directing simulations to derive an empirical model [10]. We consider these beyond the scope of this paper and take the empirical model as given. Global games [6] also embed an underlying game within a noisy observation model. Analytic results for global games show that uncertainty can motivate the selection of risk-dominant over payoff-dominant equilibria. In ensemble games [3], players also play collections of game instances. Cognitive limitations can lead agents to sub-optimally choose similar strategies across game instances. Our approach of using game-theoretic criteria to evaluate solution algorithms is similar in spirit to efficient learning equilibrium [4], which proposes an equilibrium between learning algorithms in repeated games.

2.1 Normal-form Games

A normal-form game is defined by a tuple $\{I, \{S_i\}, \{u_i(s)\}\}$. I is the set of players, which has size m. S_i is the set of pure strategies available to player i. The payoff function u_i gives the payoffs player i receives for each outcome profile, which specifies a pure strategy for each player, $s \in \{S_1 \times \cdots \times S_m\}$. The notation s_{-i} denotes a profile of strategies for all players except i. A homogeneous profile is one in which all players play the same strategy. A mixed strategy σ_i assigns a probability of playing each of player i's pure strategies; pure strategies are special case of mixed strategies. Mixed strategy outcome profiles are denoted by $\Sigma = \{\sigma_1 \times \cdots \times \sigma_m\}$.

A Nash equilibrium is a mixed-strategy profile Σ with the property that no player would benefit by deviating to a different strategy. That is, for all players i and all strategies $s_i \in S_i$, $u_i(\Sigma_i, \Sigma_{-i}) \geq u_i(s_i, \Sigma_{-i})$. Nash equilibria are stable in the sense that no player regrets their strategy choice once the choices of all players are revealed. We define the regret, ϵ, for each strategy profile to be the maximum benefit any player could gain by deviating (which may be negative). Any profile with $\epsilon \leq 0$ is a Nash equilibrium. An ϵ-Nash equilibrium is an approximate equilibrium with maximum regret ϵ.

2.2 Meta-Games

A meta-game, Γ, encompasses a base game, $\mathcal{G}$, and the empirical games observed by the players, Ω_i. The base game and empirical games are normal-form games with identical sets of players and strategies but potentially different payoff functions. A meta-game is defined by a tuple $\{I, \{S_i\}, U, f, \mathcal{O}, \Theta\}$. The sets of players and strategies, I and S_i, mirror those in a normal form game. The set of possible payoff functions for $\mathcal{G}$ is U. Payoffs for $\mathcal{G}$ are drawn randomly from U according to the density f. The set of possible empirical games is $\mathcal{O}$. The conditional probability function $\Theta = \Pr\{\Omega_i|\mathcal{G}\}$ gives the relationship between the base game and the empirical observations.[2] Typically, we define $\mathcal{O}$ and Θ implicitly by defining a noise process used to generate empirical games from the base game. The meta-game has three phases:

1. **Game Selection:** Nature draws the base game payoffs $\mathcal{G}$ from U according to f.

2. **Observation:** Nature generates the empirical games Ω_i, conditioned on $\mathcal{G}$.

3. **Strategy Selection:** Players select mixed strategies σ_i, conditioned on Ω_i. Payoffs are determined by the outcome Σ and base game $\mathcal{G}$.

[2]We assume that the observations are the same for all players for notational convenience.

To distinguish strategies in the base game $\mathcal{G}$ from strategies in the meta-game Γ, we refer to the latter as meta-strategies. Meta-strategies map from empirical games to strategies in the base game. This is precisely the function of a typical solution algorithm applied to an empirical game. *Strategy selection algorithms are meta-strategies.* For example, an algorithm that finds a Nash equilibrium and plays according to this equilibrium maps every empirical game into a strategy choice, so it is a meta-strategy. Another simple meta-strategy is to play the uniform mixed strategy regardless of the observation.

In principle, players may have knowledge of the underlying game class and observation model which could be used to derive explicit Bayesian beliefs. The meta-strategies we consider in this work do not make use of any specific information about the underlying model. This has an important advantage. These meta-strategies are generic, and can be applied without modification regardless of the specifics of the model. For many real domains there are difficult questions about how players would acquire such knowledge or what form it would take. Nevertheless, we would expect more targeted meta-strategies that make use of specific distributional information to yield better choices. It would be interesting to investigate this type of meta-strategy in future work.

3. EXPERIMENTAL FRAMEWORK

Our goal is to identify the best methods for selecting strategies based on noisy empirical game models. Analytic approaches are limited for two reasons. Generally, solving realistic meta-games for exact Bayes-Nash equilibria is intractable, both conceptually and analytically. A more subtle issues is that applications present a wide variety of game distributions and observation mechanisms. In practice, these are often very difficult to estimate accurately. Even if we were able to solve a particular meta-game model analytically, there remain important questions about how robust the solution is to variations in these parameters.

We develop an experimental framework for benchmarking candidate meta-strategies. This approach produces useful results even when the meta-game is too complex to solve exactly. It can be used to identify best practices, setting a baseline for further advances. An important benefit of this approach is that it is relatively easy to vary the game class and observation model. This provides important evidence about the robustness of the solutions that is difficult to obtain analytically. Similar benchmarking approaches are common in other decision problems where optimal algorithms are unknown, including many machine learning problems. Recently, experimental benchmarking has also been applied to learning in repeated games [11].

We begin by selecting a set of candidate meta-strategies (solution algorithms) to analyze.[3] These may be solutions to simplified problems, algorithms that appear in the literature, or simple heuristic approaches. It is often useful to define spaces of parameterized meta-strategies to explore specific hypotheses. To instantiate the meta-game model we must also specify a class of games and an observation model. In our experiments we will vary these as parameters of the meta-game model.

Given a set of implemented meta-strategies and an instantiated meta-game model, we can use simulation to estimate the payoffs in the meta-game. The steps follow those outlined in Section 2.2. To determine the payoffs for a single combination of meta-strategies (one for each player), we start by sampling from the distribution of games and generate observations for each player. We run the meta-strategies to select strategies based on the observations. These strategies are played against one another to determine the payoffs for each player and the associated meta-strategy. To achieve an accurate estimate, this process is repeated many times.

We simulate payoffs in this manner for each possible combination of meta-strategies. The end result is an estimate of the expected meta-game payoffs for this restricted set of meta-strategies. These payoffs can be represented in a symmetric, normal-form payoff matrix. We apply game-theoretic solution concepts to analyze this payoff matrix. For instance, we can evaluate the regret a player has for using a meta-strategy given the opponents' meta-strategies. We can also identify constrained meta-strategy equilibria and dominance relationships.

3.1 Game Classes and Observation Models

We present experimental results for combinations of three game classes and two observation models. All games are 2-player, 4-strategy games with asymmetric payoffs normalized to the interval $[0, 1]$. The three classes of games are *uniform random*, *common interest*, and *constant sum*. Instances of these game classes are generated using the Gamut toolkit [15]. Uniform random games are generated using the random game generator, while the common interest and constant sum game classes are instances of covariance games, with covariance of 1 and -1, respectively.[4] We generated 2500 instances of each class; the same instances are used in each experiment.

We consider two observation models, *stochastic observation* and *incomplete observation*. For stochastic observations we generate the empirical game by adding mean-zero Gaussian noise to each payoff in the base game. We vary the standard deviation of the distribution to vary the noise level. For the incomplete observations the empirical game reveals the exact payoffs for a randomly-selected subset of the pure-strategy profiles. The payoffs for the remaining profiles are interpreted as "unobserved," which in our model means that they contain only the mean payoff for the game class as an estimate (0.5 for all classes we consider). The noise level in this case is determined by the number of profiles observed. In the stochastic model, each outcome has an identical level of associated uncertainty. The incomplete model is characterized by very uneven levels of uncertainty for the individual outcomes. While these game classes and observation models are clearly not exhaustive, we intend them to present a diverse set of plausible conditions.

4. NAÏVE EQUILIBRIUM ANALYSIS

We begin with a simple experiment to demonstrate the significance of first-order payoff uncertainty to strategic decisions. Consider two candidate meta-strategies. *Naïve Pure Strategy Nash Equilibrium* (NPSNE) plays according to the most stable pure-strategy Nash equilibrium of the empirical game.[5] The "most stable" criteria generically selects a unique equilibrium for the class of games we test. This meta-strategy is naïve in that it ignores uncertainty and applies the perfect-information Nash equilibrium solution to the empirical game. *Best-Response to Uniform* (BRU) plays a best-response to uniform random opponent play, making no attempt to

[3] Analyzing a restricted set of strategies is sometimes called *heuristic strategy analysis* [22]. One method for approximating Bayes-Nash equilibria uses "constrained equilibrium" of restricted strategy spaces [2].

[4] Payoffs for each outcome of a covariance game are drawn from a multivariate normal distribution, fixing the correlation between each pair of player's payoffs.

[5] The results in this section exclude 612 of the 2500 instance that have no PSNE, ensuring that NSPSE will find an equilibrium. The same qualitative result is observed without this screening, and for the other game classes and observation models.

predict opponent behavior using payoff information.

We tested the meta-strategies on random games with the stochastic observation model, varying the noise level. Each noise level corresponds to a different meta-game, represented by a set of symmetric 2x2 payoff matrices. In Figure 2 we plot the payoffs in these meta-games against the level of observation noise. Each line corresponds to the payoffs for a meta-strategy playing against either itself or the other meta-strategy.

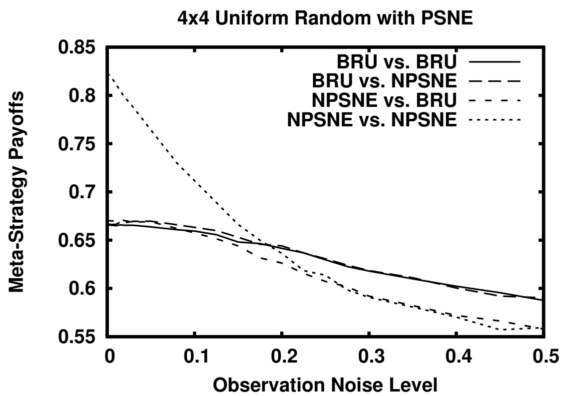

Figure 2: Payoffs to NPSNE and BRU with varying noise levels, for random games and stochastic observations.

In the limiting case where observations contain no noise, the NPSNE strategy will always select and play the same Nash equilibrium. Therefore, we know that in this case the profile where both players select strategies using NPSNE must be an equilibrium of the meta-game. We can see this in Figure 2. For the noise level of 0, the NPSNE achieves a very high payoff of approximately 0.83 when the other player is also using NPSNE. Changing to play BRU against NPSNE results in a payoff of roughly 0.67 – substantially worse. The equilibrium analysis is still beneficial as small amounts of noise are introduced. As the noise level increases, the performance of NPSNE relative to BRU steadily decreases. For noise levels higher than roughly 0.2, BRU is a *dominant* meta-strategy; regardless of whether the opponent is using NPSNE or BRU, the player achieves a higher payoff using BRU. This is a striking demonstration of the fact that ignoring payoff uncertainty in equilibrium analysis can lead to poor strategy choices. When there is sufficient uncertainty, even the extreme method of predicting completely random opponent play results in better choices than making strategic predictions based on faulty assumptions.

5. CANDIDATE META-STRATEGIES

It is useful to think of meta-strategies as having two stages: *prediction* and *best-response*. A meta-strategy first predicts how opponents will play, and then computes a best response to this prediction. Both NPSNE and BRU take this form, but the predictions they make lie at opposite ends of a spectrum. At one extreme lie completely uninformed predictions, at the other, predictions based on idealized strategic reasoning. In between we envision methods that factor noise into the strategic analysis in some intermediate way to make *distributional* predictions about the possible outcome. Vorobeychik and Wellman [21] advocate using similar distributional solution concepts in the context of mechanism design.

We introduce three classes of meta-strategies that interpolate between uninformed and equilibrium predictions using a single parameter. These meta-strategies allow us to explore an inherent tension in the prediction task. Specific predictions allow a player to compute a targeted best-response. However, if the prediction is inaccurate this may be a very poor choice. Noisy payoff observations increase the risk of inaccurate predictions, because players share less knowledge of the game. We explore this tradeoff by analyzing the relationship between observation noise and the most stable parameter settings for these meta-strategies. In general, we expect that as noise increases, the most stable meta-strategies will make broader predictions about the possible outcomes.

5.1 Approximate Pure-Strategy Nash Equilibrium

The ϵ-*Nash Solver* (ENS) generalizes NPSNE by best-responding to a distributional prediction that places greater weight on profiles with lower regret. The motivation is that low-regret profiles may be likely outcomes, even if they are not exact equilibria of the empirical game. Approximate equilibria may in fact be equilibria of the base game, or appear to be equilibria in the empirical games observed by other players. We construct a distributional prediction from the regret measure for each profile, denoted $\epsilon(s)$, using a Boltzmann distribution:

$$\Pr(s) = \frac{e^{-\epsilon(s)/\tau}}{\sum_{s' \in (S \setminus s)} e^{-\epsilon(s')/\tau}} \quad (1)$$

The "temperature" parameter τ interpolates between the uniform distribution as $\tau \to \infty$ and the most stable outcome distribution as $\tau \to 0$. This method may produce distributions that require correlated strategies. For example, the two outcomes where both players choose strategy 1 or both players choose strategy 2 could each occur with probability 0.5. This cannot occur unless the players have access to a common randomization device. Our model does not explicitly allow correlation, so we interpret this as an approximation of an uncorrelated outcome distribution.

5.2 Approximate Nash Equilibrium

The *Replicator Dynamics Solver* (RDS) searches for approximate Nash equilibria using replicator dynamics [18, 17]. This method can find mixed-strategy equilibria, so it finds better approximations on average than ENS. The replicator dynamics model an evolving population where better strategies have greater representation in the subsequent population. We use the update rule:

$$p_g(s_i) \propto p_{g-1}(s_i) \cdot (EP(s_i) - W), \quad (2)$$

where $p_g(s_i)$ is the fraction of the population for player i playing pure strategy s_i in generation g, $EP(s_i)$ is the expected payoff to s_i in generation g, and W is the lowest payoff. RDS runs for 1000 generations, applying random restarts if the population is not close to equilibrium after every 5 generations.[6] The lowest-ϵ profile found is the equilibrium estimate. We use a generic parameterization to interpolate between this prediction and the uniform distribution. We form a weighted combination, placing weight δ on the uniform distribution and $1 - \delta$ on the equilibrium distribution.

[6]These search parameters empirically identified good equilibrium approximations quickly; the average ϵ is roughly 0.3% of the maximum payoff.

5.3 Logit Equilibrium

The *Logit Equilibrium Solver* (LES) plays a best response to predictions made using logit equilibrium [13]. Logit equilibrium is a fixed point where all players use stochastic choice functions that place higher probability on playing strategies with higher payoffs, but no strategy is selected with certainty. This can be interpreted as a Bayes-Nash equilibrium where each player privately observes a vector of payoff perturbations for each pure strategy and maximizes their perturbed payoffs. Players know that their opponents have perturbed payoffs, but do not observe the values. The magnitude of the perturbations is defined by the parameter λ. As $\lambda \to 0$, noise is infinite and play is uniform random. As $\lambda \to \infty$, noise is 0 and the logit equilibrium is a Nash equilibrium. We use the Gambit solver to compute logit equilibria [19, 12]. This solver generically identifies a unique Nash equilibrium for large λ.

Logit equilibrium is a particularly interesting meta-strategy for two reasons. It can be interpreted as an exact solution to a simplified version of the meta-games we study. A significant restriction of logit equilibrium is that noise terms are associated only with individual pure strategies, rather than strategy profiles. Logit equilibrium is also of broader interest because it explains many qualitative aspects of human behavior that are not predicted by Nash equilibrium [13, 1, 5].

5.4 Noise and Distributional Predictions

The parameters of these meta-strategies all control how specific the predicted distribution of opponent play is, as shown for LES in Figure 3. This captures a key tradeoff: a more specific prediction allows a fine-tuned best-response, but risks tuning for the wrong outcome. Intuitively, observation noise makes prediction more difficult and the distribution should reflect this uncertainty. We test this hypothesis using our meta-strategy parameter spaces.

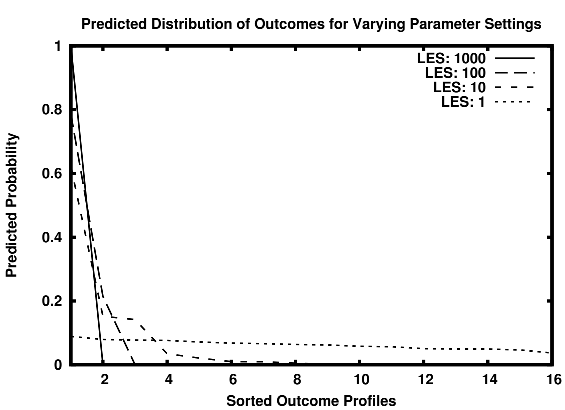

Figure 3: As the LES parameter λ is decreased, the predicted outcome distributes weight more evenly across the outcomes. Each line shows a distribution where outcomes are first sorted by probability, so the probability associated with the most likely outcome is on the far left.

We ran experiments to determine the most stable parameter setting across different noise levels for each class of meta-strategies in self-play. The stability measure we use is the regret for the homogeneous profile where all players use the same parameter setting.[7] We tested roughly 15 parameter settings for each meta-strategy class on random games with both observation models. The meta-games for each noise level are averages of 12,500 samples, with 5 observations each for 2500 base game instances.

Figure 4 shows the most stable parameter settings for each noise level. There is a clear pattern that holds for all meta-strategies, game classes, and observation models. As noise increases, the most stable parameter settings make broader (less specific) predictions about the outcome. This supports the intuition that our meta-strategy parameterizations are capturing an important aspect of strategic reasoning under uncertainty. It also suggests a general preference for distributional solution concepts over point predictions when uncertainty is a factor.

6. META-STRATEGY COMPARISON

Our final experiment directly compares the quality of the strategy choices made by the candidate meta-strategies. This comparison has immediate practical relevance for applications and sets a baseline for future algorithms to improve upon. We present results for the three game classes paired with both observation models, again varying the noise level. We take 12,500 samples for each meta-game, using 5 observations for each of 2500 base game instances. The meta-strategy candidates include one instance each of the ENS, RDS, LES, and BRU meta-strategies. The parameters used for each noise level correspond to the most stable settings given in Figure 4. There are two reasons we use a single parameter setting for each noise level. One is to keep the computational costs for simulating and analyzing the meta-game manageable. The other is that the measures we present are based on the benefit to deviating to an alternate strategy. In this section we wish to focus analysis of deviations between classes of meta-strategies, rather than deviations to different parameter settings.

The results of the meta-strategy analysis are presented in Figure 5. To visualize performance differences across different noise levels, we need a measure that summarizes the results of the meta-game for each noise level. We use the regret for the homogeneous profile where both players use the same meta-strategy as the performance metric. This measure roughly summarizes the stability of each meta-strategy, and allows us to identify meta-strategies for which the homogeneous profile is a Nash equilibrium of the constrained meta-game.[8] Equilibria of the constrained game may not be equilibria of the full game, since there may be beneficial deviations to meta-strategies that are not in the set of candidates. However, any profile that is not an equilibrium of the constrained game can be definitively rejected as an equilibrium of the full game. Equilibria of the constrained game should be interpreted as candidate solutions that have survived all known challengers.

Typically, ENS, RDS, and LES are all homogeneous meta-strategy equilibria for low noise levels. This is not surprising, since all three identify and play approximate equilibria of the base game in this case. The fact that deviating to a different meta-strategy results in a net loss is evidence of a coordination effect, due to the tendency of the different classes of meta-strategies to select different equilibria. We generally observe larger regret magnitudes at lower noise levels, with regret converging to zero in the high noise conditions. This is also intuitive; under very noisy conditions, the outcome of any analysis depends largely on the noise and the final selections are close to random for all meta-strategies.

[7]Typically, these homogeneous profiles are equilibria and among the most stable profiles overall. Alternative stability measures yield qualitatively identical results.

[8]No one-dimensional measure can fully capture the strategic analysis of a game. We evaluated several other measures, each of which revealed similar performance patterns.

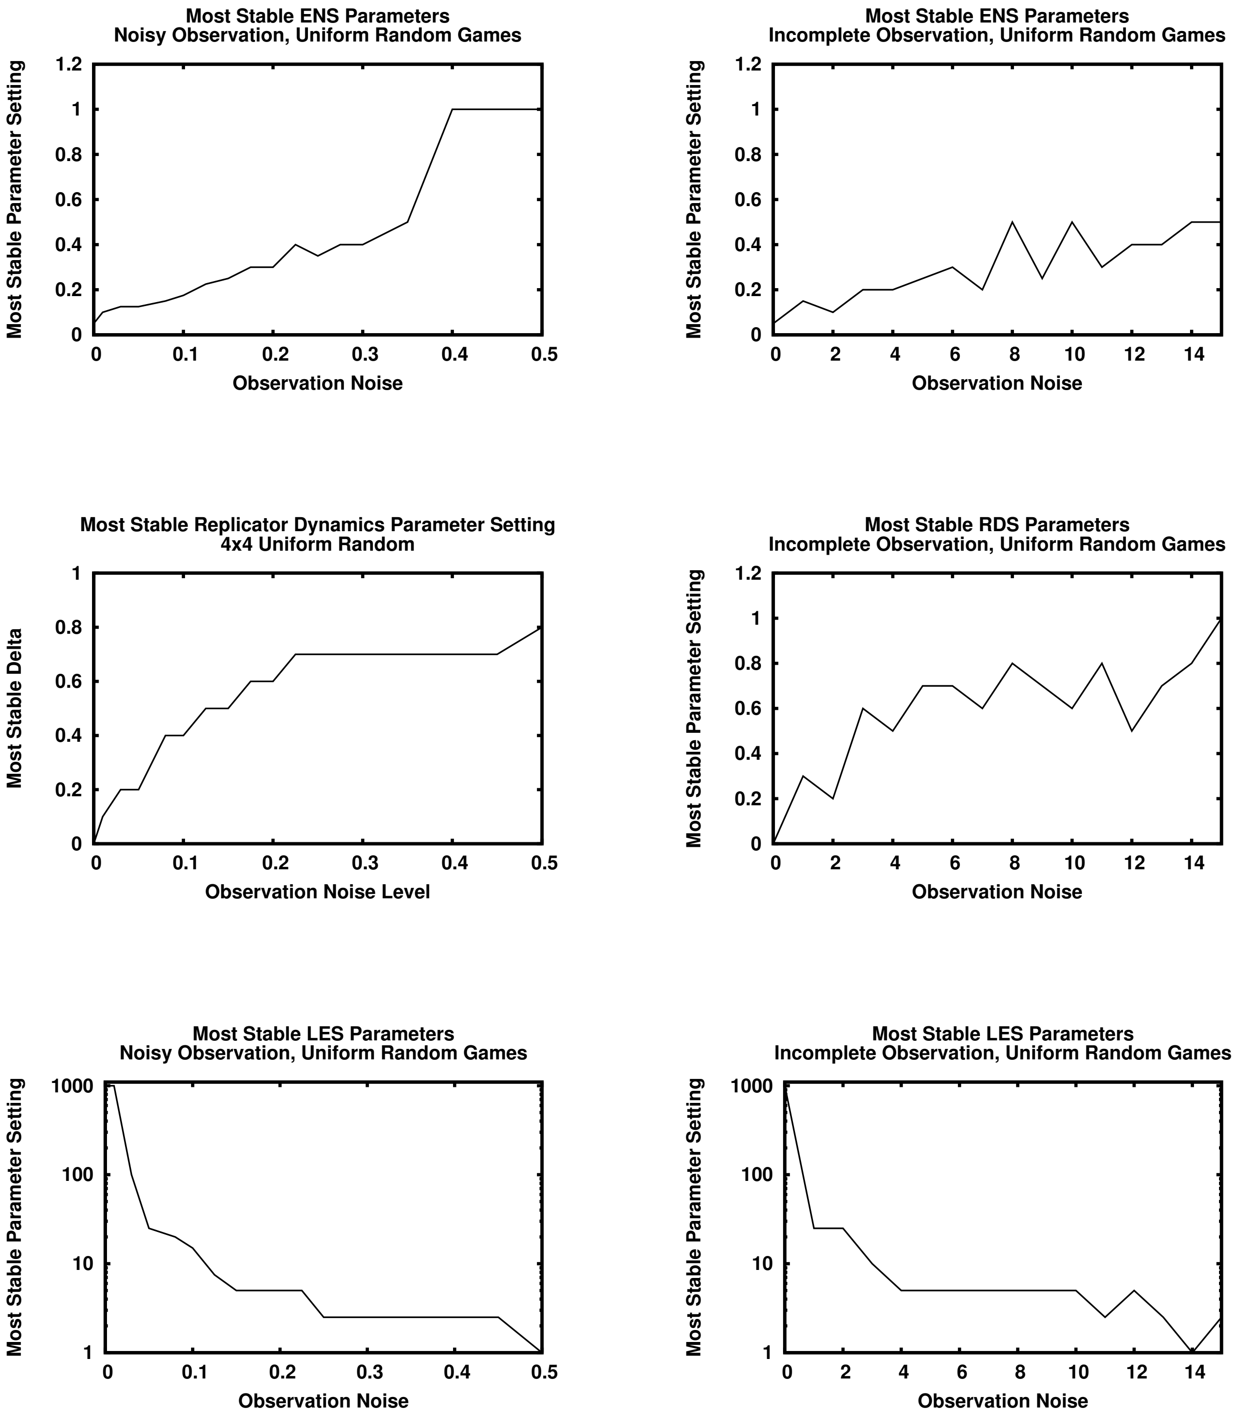

Figure 4: The parameter settings for ENS, RDS, and LES with the most stable homogeneous profiles for different noise levels. Results are shown for both the stochastic and incomplete observation models on random games. For ENS and RDS, higher parameter settings correspond to *less* specific predictions, while higher settings for LES correspond to *more* specific predictions.

The most striking result is the strong performance of the LES meta-strategy across all conditions tested. The homogeneous profile of this meta-strategy is almost always an equilibrium of the meta-game. Often, it is the only equilibrium. Where there are multiple equilibria, this is often the one with the greatest loss for deviating. The performance of LES relative to the other meta-strategies is especially strong when there is moderate noise. In these cases, there is typically a large incentive to deviate from the other meta-strategies to LES. This constitutes strong evidence that logit equilibrium is an effective method for factoring uncertainty into strategic reasoning. Our results also provide important evidence that this is a robust approximation for a wide range of games and observation mechanisms.

7. DISCUSSION

We define a meta-game model where players make decisions based on empirical observations of an underlying game. Though we are not able to solve this model directly, it motivates an experimental framework for benchmarking general solution algorithms applied to empirical game models. We demonstrate that the payoff noise captured by our model is consequential for strategy choice. While naïve equilibrium predictions perform well for low levels of noise, they are dominated by uninformed predictions when there is sufficient observation noise.

We describe three candidate meta-strategies that make distributional predictions. These methods are parameterized to interpolate between the extremes of uninformed predictions and complete-information equilibrium predictions. There is a systematic relationship between the parameter settings that minimize regret in the meta-game and the level of observation noise. As noise increases, the most stable parameters correspond to less specific outcome predictions. This supports the hypothesis that these parameter spaces capture a key tradeoff for strategic reasoning under uncertainty.

Our final experiment provides a direct comparison of the candidate meta-strategies. The performance of a meta-strategy based on logit equilibrium is very compelling. It is especially encouraging that this method is robust to substantial variations in the parameters that define the underlying uncertainty. In addition to offering guidance for analysis of empirical models, this result provides a new perspective on the success of logit equilibrium in explaining behavioral data in humans.

8. ACKNOWLEDGMENTS

This work was supported in part by grant IIS-0414710 from the National Science Foundation.

9. REFERENCES

[1] S. P. Anderson, J. K. Goeree, and C. A. Holt. Minimum-effort coordination games: stochastic potential and logit equilibrium. *Games and Economic Behavior*, 34:177–199, 2001.

[2] O. Armantier, J.-P. Florens, and J.-F. Richard. Approximation of Bayesian Nash equilibrium. *Journal of Applied Economics*, 2006. Submitted.

[3] J. Bednar and S. Page. Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. *Rationality and Society*, 19(1):65–97, 2007.

[4] R. I. Brafman and M. Tennenholtz. Efficient learning equilibrium. *Artificial Intelligence*, 159(1–2):27–47, 2004.

[5] C. M. Capra, J. K. Goeree, R. Gomez, and C. A. Holt. Anomalous behavior in a traveler's dilemma? *American Economic Review*, 89(3):678–690, 1999.

[6] H. Carlsson and E. van Damme. Global games and equilibrium selection. *Econometrica*, 61(5):989–1018, 1993.

[7] V. Conitzer and T. Sandholm. Complexity results about Nash equilibrium. In *18th International Joint Conference on Artificial Intelligence*, pages 765–771, 2003.

[8] J. K. Goeree, C. A. Holt, and T. R. Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 8(4):347–367, 2005.

[9] J. Harsanyi. Games with incomplete information played by Bayesian players, Parts I, II, and III. *Management Science*, 14:159–182, 320–334, 486–502, 1967–1968.

[10] P. R. Jordan, Y. Vorobeychik, and M. P. Wellman. Searching for approximate equilibria in empirical games. In *AAMAS-08*, 2008. To appear.

[11] A. Lipson. An empirical evaluation of multiagent learning algorithms. Master's thesis, University of British Columbia, 2005.

[12] R. D. McKelvey, A. M. McLennan, and T. L. Turocy. Gambit: Software tools for game theory, version 0.2006.01.20, 2006. http://econweb.tamu.edu/gambit.

[13] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38, 1995.

[14] S. Morris and H. S. Shin. *Global Games: Theory and Applications*, volume Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society). Cambridge University Press, 2003.

[15] E. Nudelman, J. Wortman, K. Leyton-Brown, and Y. Shoham. Run the GAMUT: A comprehensive approach to evaluating game-theoretic algorithms. In *AAMAS-04*, 2004.

[16] S. Phelps, M. Marcinkiewicz, S. Parsons, and P. McBurney. A novel method for automatic strategy acquisition in n-player non-zero-sum games. In *AAMAS-06*, pages 705–712, Hakodate, 2006.

[17] D. M. Reeves. *Generating Trading Agent Strategies: Analytic and Empirical Methods for Infinite and Large Games*. PhD thesis, University of Michigan, 2005.

[18] P. Taylor and L. Jonker. Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 16:76–83, 1978.

[19] T. L. Turocy. A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51:243–263, 2005.

[20] Y. Vorobeychik, C. Kiekintveld, and M. P. Wellman. Empirical mechanism design: Methods, with application to a supply-chain scenario. In *Seventh ACM Conference on Electronic Commerce*, 2006.

[21] Y. Vorobeychik and M. P. Wellman. Mechanism design based on beliefs about responsive play (position paper). In *ACM EC-06 Workshop on Alternative Solution Concepts for Mechanism Design*, 2006.

[22] W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2002.

[23] M. P. Wellman, J. Estelle, S. Singh, Y. Vorobeychik, C. Kiekintveld, and V. Soni. Strategic interactions in a supply chain game. *Computational Intelligence*, 21:1–26, 2005.

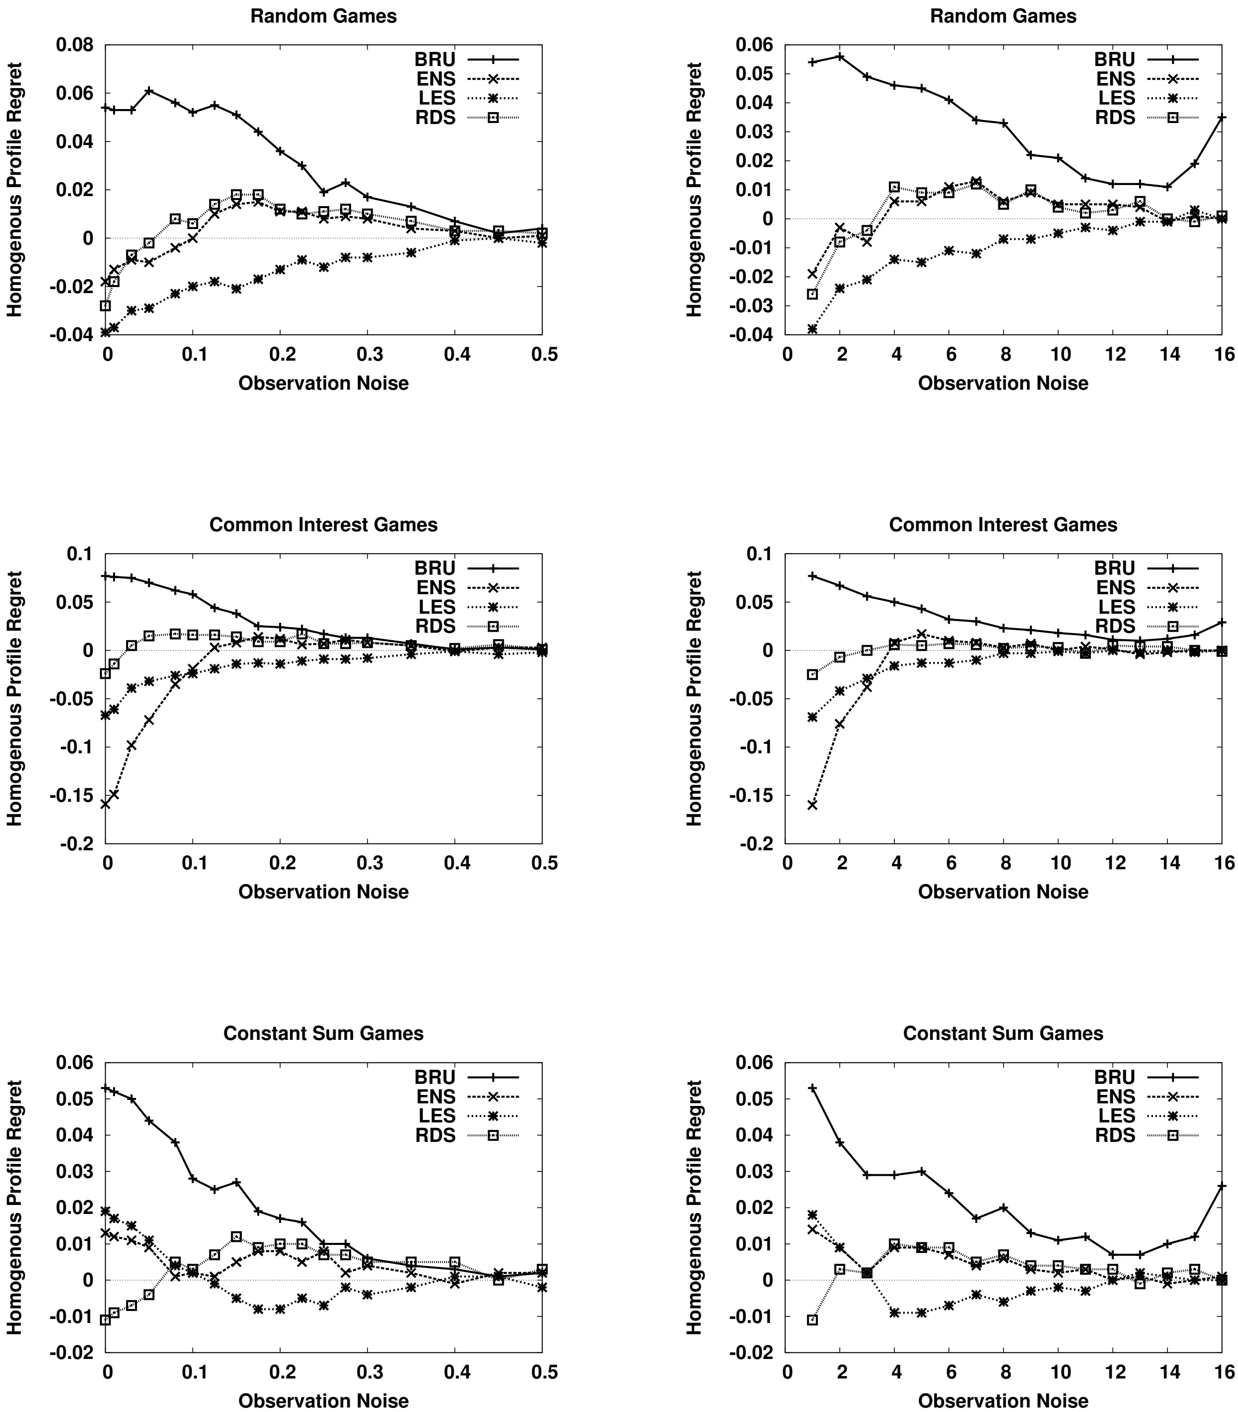

Figure 5: Comparison of four classes of meta-strategies across different noise levels for six combinations of game class and observation model. The stochastic observation model results are on the left, and the incomplete observation model on the right.