Identifying Strategies and Beliefs without Rationality Assumptions

Amos Golan

Department of Economics
Info-Metrics Institute
American University
agolan@american.edu

James W. Bono

Department of Economics
Info-Metrics Institute
American University
bono@american.edu

May 18, 2010

Abstract

In this paper we formulate a solution concept without making assumptions about expected utility maximization, common knowledge or beliefs. Beliefs, strategies and the degree to which players are expected utility maximizers are endogenously determined as part of the solution. To achieve this, rather than solving the game from the players' point of view, we analyze the game as an "observer" who is not engaged in the process of the game. Our approach is an information theoretic one in which the observer utilizes an observation of play and the Maximum Entropy principle. We compare our solution concept with Bayesian Nash equilibrium and offer the entropy ratio test as a method for determining the appropriateness of common modeling assumptions. We also demonstrate that the QRE concept can be significantly generalized when viewed from the observer's perspective. For games of incomplete information we discover that alternative uses of the observer's information lead to alternative interpretations of rationality. These alternative interpretations of rationality may prove useful, especially in the context of *ex post* arbitration, as they indicate who is motivating whom.

Keywords: incomplete information, entropy, information theory, pairwise rationality, QRE, endogenous rationality.

JEL Codes: C70, C79

We would like to thank Jeff Perloff, Steve Kuhn and numerous participants at seminars and conferences. This work builds on an earlier work of Catilina and Golan (2006)

1 Introduction

1.1 Motivation and Objectives

Every model builds on a well-defined set of assumptions. Traditionally, in game theory, one or more of these assumptions concerns what can loosely be regarded as the rationality of the players. Solution concepts such as Nash equilibrium (NE) and Bayesian Nash equilibrium (BNE) require, among other assumptions, that all players are expected utility maximizers [see Nash (1959); Harsanyi (1967)]. The logit quantal response equilibrium (QRE, McKelvey and Palfrey (1995)) assumes that each player is capable of maximizing expected utility only to a certain degree. This is a notion of bounded rationality (and it is the one we adopt here). Other notions of rationality include the number of logical iterations a player can make [see Costa-Gomes and Crawford (2006); Camerer et al. (2006)] and the approximation of equilibrium payoffs [see Radner (1980)]. Alternatively, we develop here a model that allows the researcher to remain agnostic about the degree to which players maximize expected utility. That is, we proceed without any assumptions regarding expected utility maximization or common knowledge and formulate a solution concept where the degree to which players are expected utility maximizers is endogenous; it is a part of our solution. To achieve that objective, rather than solving the game from the players' point of view, we formulate the game as an "observer" who is not engaged in the process of the game.

We observe the actions taken by the players and use this knowledge, together with our understanding of the set of possible games that are played, to construct an "observed solution." This is the only information used here. However, there are many possible solutions that are consistent with the observed actions and our understanding of the set of possible games. Consequently, to identify a particular solution, we have three options: (i) impose additional assumptions until only one of these solutions remains, (ii) employ a decision criterion to select one of these (perhaps infinitely many) solutions or (iii) use a combination of the above two. To proceed with minimal assumptions on the structure of the game we choose option (ii).

In particular, all of the observer's information is constructed in terms of the cross moments of the players' strategy distributions based on the players' observed actions. We then apply the maximum entropy principle to choose a solution. This principle builds on the principle of insufficient reason (Jaynes (1957)) and uses Shannon's entropy (Shannon (1948)) as the criterion to select the solution. Simply stated, we derive our observed solution by maximizing the joint Shannon entropy of players' strategies and beliefs subject to the researcher's information, which includes the researcher's understanding of the set of possible games and the observed actions of the players. In an information-theoretic sense, the solution yields the least informative mixed strategy profile that is consistent with the researcher's information. The resulting solutions can be pure or mixed. For review of information theory within economics and econometrics see the recent work of Golan (2008) or the text by Golan et al. (1996).

Under the approach we take here, which converts our under-determined problem to

a well-posed constrained optimization model, the degree of rationality of each player is determined endogenously. That is, we do not impose player rationality as an identifying assumption. Rather, the degree of rationality is captured via the Lagrange multipliers of the constrained optimization problem. The framework developed here allows one to compare solutions that are based on different rationality, behavioral or informational assumptions, such as solutions representing bounded rationality or common knowledge of expected utility maximization.

1.2 Overview

We first introduce our basic model which is based on the researcher’s observed and known information (Section 2.1). We call this the *observed game*. We then solve the observed game using the Maximum Entropy principle (Section 2.2). This results in the *observed solution*. We show that the Lagrange multipliers resulting from this optimization convey the amount of information in the cross moment constraints. This information is the researcher’s information that includes the observed actions and the assumptions on the basic structure of all potential games played. Like the parameters in the exponents of the logit QRE, these parameters also capture the estimated degree of expected utility maximization done by each player, which indirectly captures the individual’s rationality.

However, unlike logit QRE parameters, which give the degree to which players maximize expected utility given the strategies of all other players, our Lagrange multipliers capture the degree to which players are pairwise rational. For instance, the Lagrange multiplier on the (i, j) cross moment gives the degree to which player i maximizes her utility with regard to player j ’s strategy. A generalization of this model captures the degree to which players are rational with regard to subsets of other players. In this generalization, the Lagrange multiplier on the (i, S) constraint gives the degree to which player i maximizes her utility with regard to the strategies of players in the set S .

Once our solution is introduced, we discuss and compare our solution to other solutions. We start by imposing the expected utilities as the set of constraints rather than the cross moments that are used to obtain the observed solution (Section 2.3). This is similar to the analysis of Wolpert (2003), except that here it is done for games of incomplete information. Again, we study the Lagrange parameters and show that they convey the degree player rationality as captured by the researcher’s observed information.

We then turn to study the relationship between the observed solution and BNE and define the *observed Bayesian solution* as the BNE that is consistent with the cross moments and maximizes entropy (Section 2.4). Simply stated the *observed Bayesian solution* is a special case of the observed solution where BNE requirements are introduced in addition to the cross moments’ conditions within the optimization. There is not always a BNE that is consistent with the observed game. Further, when a BNE is consistent with the observed game, the entropy of the solution is, in general, lower than that of the observed game. This is a reflection of the additional structure that the BNE imposes on the game, implying more information is used. The use of more information translates into lower entropy. The additional structure includes the assumption of a common prior, common

knowledge of the game and common knowledge of expected utility maximization. In more technical terms, we show that the assumptions required to obtain the observed solution are a subset of the assumptions required to obtain a BNE. Pursuing model comparison further, we offer the entropy ratio test as a way for researchers to test the hypothesis that the observed game is consistent with some specific economic theory, e.g. BNE.

All of the preceding assumes that the researcher has no prior knowledge about the players’ strategies, beliefs or level of information. In reality, however, it is quite common for a researcher to have such prior beliefs. These beliefs are the researcher’s own prior information and may come from theory, intuition or from observing an earlier round of a similar game. We discuss this issue in Section 3. To formulate this prior information within our observed game we substitute the Shannon entropy objective with the cross entropy, or Kullback-Liebler divergence measure, and optimize with respect to the same set of constraints [see Kullback and Leibler (1951)]. The result is the *observed solution with priors* which is just a generalization of the *observed solution*.

Generally speaking, under the model with priors, the estimated strategies and beliefs are those that satisfy the observed information (cross moments) and are closest to the observer’s prior beliefs. Under our basic observed solution model, the estimated strategies and beliefs are those that satisfy the observed information (cross moments) and are closest to the uniform distribution. Thus, if the researcher has *no* priors, which is similar to using uniform priors under the information theoretic framework used here, the solution with priors and the observed solution are the same. We also show here how to derive priors that are based on economic theory. For example, we answer the question of “how should the researcher choose prior distributions that reflect the researcher’s belief that the players play a Bayesian game that is fully consistent with the BNE assumptions?” To do this we derive a prior using information theory and call it the *rational prior*. We call the corresponding solution the *observed rational prior solution*.

We demonstrate our basic models and solutions via a simple example that is carried throughout Sections 2 and 3. We use that example to compare the observed solution with the observed Bayesian solution, the observed solution with priors and the observed rational prior solution. This example clarifies the simplicity and generality of our proposed approach.

As mentioned earlier, our model is somewhat related to the logit QRE model. We provide a detailed comparison with that model in Section 4. Building on the work of Wolpert (2003) we construct an information-theoretic derivation of the logit QRE for games of incomplete information (*IIQRE*). Unlike the basic observed game model, in that case, we “fix” the expected utilities, given the strategies of the other players, to obtain an observed solution for each player. Then, since QRE is an equilibrium concept, we solve all of those observed solutions simultaneously. Because this QRE is parameterized by the common prior, we are also able to develop the concept of a *Maximum Entropy IIQRE*. This concept is appropriate when the researcher knows that a QRE is being played but does not know (or is uncertain about) the common prior used by the players.

Finally, we show that under the information theory approach we use, the QRE can be generalized. One way of doing so is by substituting the cross moment constraints with the

expected utility constraints. Just as in the observed solution, the Lagrange multipliers of that model provide a pairwise measure of rationality. This pairwise rationality reflects the degree to which player i maximizes her own utility with regard to the strategy choice of player j . A further generalization of the QRE allows us to inspect the rationality of each player i with regard to any subset of the other players.

2 The Observed Game and Solution

2.1 The Observed Game

Consider the following structure. An external observer, called Sun (super-natural), has identified a situation where players may be interacting strategically. Sun may be an economist, a researcher, a moderator or anyone that is interested in understanding the strategic interaction of some set of individuals. Sun knows the following basic information about this situation, called Γ :

- A finite set of players $I = \{1, \dots, I\}$.
- For each player i , a set of κ_i pure strategies X_i .
- For each player i , a set of τ_i possible types Θ_i , where $\sum_i \tau_i = \tau$.
- A set of possible games $\Theta = \Theta_1 \times \dots \times \Theta_I$.
- For each player i , a type-contingent payoff function $u_i(\cdot | \theta_i) : X \mapsto \mathbb{R}$.

Note that the description of the strategic situation, Γ , does not contain any description of what the players know about the others, whether they are expected utility maximizers, or whether they share a common prior, etc. Instead, Γ is simply the strategic situation as understood by Sun.

In addition to knowing Γ , Sun observes a single instance of play and records his observations in Y . Y represents the observed actions taken by all players. The data are arranged so that if player i played pure strategy $x_i \in X_i$, then $Y_i(x_i) = 1$, and $Y_i(x'_i) = 0$ for all $x'_i \neq x_i$. Together with Γ , we list the *observed game* as $\langle \Gamma, Y \rangle$. From the observed game, Sun wishes to infer two things:

1. for each player i a type-contingent mixed strategy $P_{\theta_i} \in \Delta(X_i)$, and
2. for each player i a probability distribution $\alpha_i \in \Delta(\Theta)$.

To estimate the beliefs and type-contingent strategies for each player, Sun makes two assumptions.

O1 Each player i knows the set of actions available to him, X_i .

O2 Each player i may have prior beliefs that can be summarized as a probability distribution, α_i , over the games in Θ .

We emphasize again that the analysis we are doing here is from the perspective of the external observer, Sun, who possesses only the above information. Sun does not know, for example, whether the players maximize expected utility. In this regard, estimating a type-contingent mixed strategy for each player does not imply that players are expected utility maximizers or even that players necessarily use type-contingent strategies. It simply allows for the possibility that a player's strategy depends on her payoffs. Alternatively, if Sun did not allow type-contingent strategies, then this would rule out the possibility that players are expected utility maximizers.

Sun also does not know the prior beliefs of the players. Further, since Sun does not know that all players share common knowledge of the prior probabilities over Θ , he does not assume this. Doing so would rule out the possibility that players have different beliefs. Instead, Sun estimates beliefs for each player. Doing so does not, in principle, rule out the possibility that there is a common prior.

Sun knows that Y_i , the observed action of i , contains some information about P_i and α_i . Sun does not know the exact form of the relationship between Y_i , P_i and α_i . However, he knows that this relationship may depend on i 's utility function. That is, player i may be adjusting her strategy to increase (or decrease) her own expected utility. The relationship between Y_i , P_i and α_i may also depend on the identities of the other players. So player i may be increasing (or decreasing) her expected utility with regard to the strategy choice of player j , effectively ignoring the strategy choices of the other players. Similarly, player i may be choosing her strategy to increase (or decrease) player j 's expected utility.

In light of these considerations and since the exact parametric functional form is unknown to Sun, he chooses to incorporate these relationships via the cross-moments for each ordered-pair of players (i, j) , $i \neq j$. This means that Sun imposes the least possible structure. In particular, *Sun requires that i 's expected utility, given i and j 's strategies and that all other players randomize uniformly, is equal to i 's expected utility, given i and j 's observed actions and that all other players randomize uniformly.* This restriction is given by:

$$\sum_{\theta} \sum_x P_{\theta_i}(x_i) P_{\theta_j}(x_j) \alpha_i(\theta) u_i(x|\theta_i) = \sum_{\theta} \sum_x Y_i(x_i) Y_j(x_j) \alpha_i(\theta) u_i(x|\theta_i). \quad (2.1)$$

where $x_{-i,j}$ is the profile of all players' actions except i and j . The cross moment given by (i, j) is i 's expected utility from playing against j where every other players' strategy is uniform. With equation 2.1 we allow for a wide range of interaction effects between the players' beliefs, strategies and payoffs.

Note that there are many alternative ways of incorporating the information from Y . For instance, Sun might have chosen to simply require that each player's expected utility is equal to $\sum_{\theta} \sum_x Y(x) \alpha_i(\theta) u_i(x|\theta_i)$, the observed expected utility. However, this would prevent Sun from capturing the many pairwise interactions that may (or may not) be

underlying the behavior of the players. This alternative is discussed in section 2.4.

2.2 The Observed Solution

Taking the cross moments, equation 2.1, together with the requirements that all probability distributions are proper,

$$\sum_{x_i} P_{\theta_i}(x_i) = 1 \quad \text{and} \quad \sum_{\theta} \alpha_i(\theta) = 1, \quad (2.2)$$

Sun has $I \times I + \tau$ equations and $(\sum_i \tau_i \times \kappa_i) + I \times |\Theta|$ unknowns, where $|\Theta|$ stands for the number of elements in Θ . Since there are more equations than unknowns, there are many choices of (P, α) that are consistent with the observed game as captured by the cross moments. Thus, Sun is facing an under-determined problem. A logical way of choosing among the different potential solution is to apply the classical Maximum Entropy formalism (Jaynes (1957)) which chooses the set of distributions (P, α) that maximize the Shannon entropy of P and α with respect to the observed information (constraints), equation 2.1.

Specifically, Sun's problem is:

$$\begin{aligned} \max_{P, \alpha} H(P, \alpha) = & - \left[\sum_i \sum_{\theta_i} \sum_{x_i} P_{\theta_i}(x_i) \ln(P_{\theta_i}(x_i)) + \sum_i \sum_{\theta} \alpha_i(\theta) \ln(\alpha_i(\theta)) \right] \quad (2.3) \\ \text{s.t. equation 2.1,} \\ & \sum_{x_i} P_{\theta_i}(x_i) = 1 \quad \forall i, \theta_i \quad \text{and} \\ & \sum_{\theta} \alpha_i(\theta) = 1 \quad \forall i \end{aligned}$$

The Lagrangian for problem 2.3 is:

$$\begin{aligned} \mathcal{L} = & H(P, \alpha) + \sum_i \sum_{\theta_i} \lambda_{\theta_i} \left[\sum_{x_i} P_{\theta_i}(x_i) - 1 \right] + \sum_i \lambda_{\alpha_i} \left[\sum_{\theta} \alpha_i(\theta) - 1 \right] \dots \\ & + \sum_{ij} \lambda_{ij} \left[\sum_{\theta} \sum_x P_{\theta_i}(x_i) P_{\theta_j}(x_j) \alpha_i(\theta) u_i(x|\theta_i) - Y_i(x_i) Y_j(x_j) \alpha_i(\theta) u_i(x|\theta_i) \right] \end{aligned}$$

The solutions to problem 2.3 are then

$$P_{\theta_i}^*(x_i) \propto \exp \left[\sum_{j \neq i} \sum_{\theta_{-i}} \sum_{x_{-i}} P_{\theta_j}^*(x_j) [\lambda_{ij}^* \alpha_i^*(\theta) u_i(x|\theta_i) + \lambda_{ji}^* \alpha_j^*(\theta) u_j(x|\theta_j)] \right] \quad (2.4)$$

and

$$\alpha_i^*(\theta) \propto \exp \left[\sum_{j \neq i} \sum_x \lambda_{ij}^* \left[P_{\theta_i}^*(x_i) P_{\theta_j}^*(x_j) - Y_i(x_i) Y_j(x_j) \right] u_i(x|\theta_i) \right]. \quad (2.5)$$

The normalization constant for $P_{\theta_i}^*$ is

$$\Omega(\theta_i) = \sum_{x_i} \exp \left[\sum_{j \neq i} \sum_{\theta_{-i}} \sum_{x_{-i}} P_{\theta_j}^*(x_j) \left[\lambda_{ij}^* \alpha_i^*(\theta) u_i(x|\theta_i) + \lambda_{ji}^* \alpha_j^*(\theta) u_j(x|\theta_j) \right] \right],$$

and the normalization constant for α_i^* is

$$\Omega_i = \exp \left[\sum_{j \neq i} \sum_x \lambda_{ij}^* \left[P_{\theta_i}^*(x_i) P_{\theta_j}^*(x_j) - Y_i(x_i) Y_j(x_j) \right] u_i(x|\theta_i) \right].$$

Solutions 2.4 and 2.5 are referred to as the *observed solutions*.

2.3 Discussion, Interpretation and Rationality

Within the setup of our *observed solutions*, there is no estimate of $\alpha^*(\theta)$, but rather we have $\alpha_i^*(\theta)$. That is, following the logic of imposing minimal structure, Sun chooses to estimate the beliefs of each agent, allowing beliefs to differ across players, though Sun's model does not force players' beliefs to differ. This represents a significant departure from the equilibrium concept approach where a common prior is generally required for logical consistency. This is because equilibrium solutions require that player i has beliefs about player j 's beliefs about player i 's beliefs and so on. Traditionally, such an infinite regress is avoided with the assumption that there is a common prior, which reduces the game of incomplete information to one of imperfect information. Then, the players are endowed with behaviorally-based response correspondences, and a fixed point (or points) provides the solution.

As opposed to the equilibrium concept approach, in our approach Sun does not explicitly model the way in which players respond to one another. Instead he merely states that, for each pair (i, j) , P_i and P_j have the property that the (i, j) and (j, i) cross moments are equal to the corresponding observed cross moments. So the observed solutions are not interpreted as a fixed point of a system of behaviorally-derived response correspondences of the players (i.e. best response, quantal response). The observed solutions are simply Sun's best estimate of player behavior and beliefs given the observed game. As such, Sun does not encounter the same inconsistencies one encounters when modeling an equilibrium in which players have different prior beliefs. It is, of course, possible to formulate Sun's problem so that $\alpha_i^* = \alpha_j^*$ for all $i, j \in I$. However, this assumption is not necessary, and given Sun's agnostic approach to rationality and common knowledge, it

is not in the spirit of the observed solution concept. Nevertheless, if the observed game is consistent with a common prior and that common prior maximizes the joint entropy of (P, α) , then the resulting solutions will have a common prior. That is, whether or not there is a common prior will be determined from the observed game rather than assumed *a priori*.

To further distinguish between our approach and equilibrium concepts, we note that the observed solution does not provide Sun's prediction for the probability of x , $P(x)$. According to Bayesian decision theory, for Sun to form such a prediction, he would first need to stipulate his *own* beliefs about the relative likelihood of states in Θ . However, Sun does not do so in our model. In our model, he simply recovers (based on the available information) the *players'* beliefs. To be more concrete, recall that equilibrium concepts generally assume independent moves, so that $P(x)$ is the product distribution, $P(x) = \prod_i P_i(x)$. However, here we have that $P_i^*(x_i) = \sum_{\theta} \alpha_i^*(\theta) P_{\theta_i}^*(x_i)$. As detailed above, α_i is potentially different for each i . This means that, in our model, $P^*(x) = \prod_i \sum_{\theta} \alpha_i^*(\theta) P_{\theta_i}^*(x_i)$ does not represent Sun's prediction for the probability of observing x .

The above argument helps to clarify the proper interpretation of the observed solution. It is not a prediction of player behavior because Sun does not have an estimate of $\alpha(\theta)$. Rather, the purpose of the observed solution is to facilitate an understanding of a few things: the players's beliefs and degree of expected utility maximizing behavior, as well as the players' strategies. As is already evident, the beliefs of the players are understood through the α_i . The degree of expected utility maximizing behavior is understood through examination of the Lagrange parameters λ_{ij}^* . In addition, our framework is natural for combining observed data and theory and allows us to test whether the players are engaged in a specific game, i.e. whether the assumptions of an equilibrium solution concept are consistent with the observed information at hand. We discuss this issue at the end of Section 2.

The Lagrange parameters λ_{ij}^* capture the marginal impacts of the constraints 2.1 on the objective. In the model we present here, the Lagrange multipliers reflect the amount of information in the observed moments, equations 2.1. To see this, recall that the (i, j) constraint represents player i 's expected utility of playing against j when everyone else randomizes uniformly. Since the objective is the Shannon entropy - an information criterion, each λ_{ij}^* is the marginal amount by which the (i, j) cross moment forces the observed solution, (P_i^*, P_j^*) , to diverge from the uniform distribution - a state where Sun is completely uncertain about the game. The larger the divergence from uniformity, as measured by $|\lambda_{ij}^*|$, the more the information contained in that constraint. Intuitively, when λ_{ij}^* is very large in magnitude, the component distributions of $P_i^* = \{P_{\theta_i}^*\}_{\theta_i}$ and $P_j^* = \{P_{\theta_j}^*\}_{\theta_j}$ are far from uniform. Let the vertices of $\Delta(X_i)$ be given by the Dirac delta functions $\{\delta_i(x_i)\}_{x_i}$, where $\delta_i(x_i)$ is zero everywhere except at x_i where it is one. For large $|\lambda_{i,j}^*|$, the (i, j) cross moment is forcing $P_{\theta_i}^*$ and $P_{\theta_j}^*$ toward the vertices of $\Delta(X_i)$ and $\Delta(X_j)$, where entropy is minimized and the gradient is infinite. Conversely, when $|\lambda_{i,j}^*|$ is close to zero, the constraint does not force the solution to diverge much from the uniform solution. When all $\lambda_{i,j}^*$'s equal zero, the observed solution is simply the uniform

solution, meaning Sun has no insight into the behavior or beliefs of the players.

Further inspection of equation 2.4 reveals that λ_{ij}^* is related to the degree to which i maximizes expected utility given j 's strategy. When λ_{ij}^* is large and positive, $P_i(x_i)$ will be greater for x_i that result in greater expected utility for i given j 's strategy and that all other players randomize uniformly. A large negative value of λ_{ij}^* indicates that i intentionally minimizes her expected payoff with regard to j 's strategy. That is, a large negative value of λ_{ij}^* results in more weight being given to strategies x_i that do comparatively worse against P_j^* for each θ_i . Note that $P_{\theta_i}^*(x_i)$ also depends on λ_{ji}^* . It enters via the term $P_{\theta_j}^*(x_j)\alpha_j^*(\theta)u_j(x|\theta_j)$. If λ_{ji}^* is positive, then the greater j 's expected utility, given x_i and that all other players randomize uniformly, the greater is $P_{\theta_i}^*(x_i)$.

Furthermore, it is clear by looking at equations 2.4 and 2.5 that when the Lagrange parameters are all equal to zero, the observed solutions are all uniform. This is because the uniform distribution has the maximal entropy. Here, the interpretation is that none of the players are maximizing (or minimizing) utility with respect to the strategy choices of any other player.

This interpretation is similar to the interpretation of Wolpert (2003) for the QRE Lagrange parameter, β_i . The QRE Lagrange parameter β_i represents the degree to which i maximizes expected utility with regard to the strategies of the other players. However, Wolpert's information theoretic derivation of the QRE comes from fixing expected utilities for each player (rather than cross moments) and the additional constraint that the players are in equilibrium. Simply stated, the QRE equilibrium solution is based on more information than the model proposed here uses.

This formulation allows us to discuss a new view of rationality, the degree to which one player maximizes utility with respect to *one* other player's strategy choice. A further generalization of this view is that in our formulation we can capture the degree to which player i maximizes utility with respect to the strategy choices of any *subset* of the other players. This is a more realistic view of player behavior. For example, player 1 may pay great attention and take great care to maximize her expected utility with respect to the actions of players 2 and 3 but may not pay attention to the actions of players 4 and 5. This formulation might be useful in situations where a mediator is interested in determining which of the players are motivating the actions of a given player. Suppose $S_{-i} = 2^{I \setminus i}$ is equal to the set of all subsets of players not containing i . Then for each i and each element $s \in S_{-i}$, Sun will impose a constraint of the form:

$$\sum_{\theta} \sum_x P_{\theta_i}(x_i) P_{\theta_s}(x_s) u_i(x|\theta_i) = \sum_{\theta} \sum_x Y_i(x_i) Y_s(x_s) u_i(x|\theta_i).$$

Maximizing entropy $H(P, \alpha)$ subject to these constraints will produce a solution with Lagrange parameters $\lambda_{i,s}$. A large positive value of $\lambda_{i,s}$ indicates that i 's strategy does well in maximizing i 's expected utility given P_s and that all other players are randomizing uniformly. In other words, a large positive value of $\lambda_{i,s}$ indicates that i chooses P_i to intentionally maximize expected utility with regard to the strategy choices of players $s \subset (I \setminus i)$. A large negative value of $\lambda_{i,s}$, however, indicates that i intentionally minimizes

expected utility with regard to the strategy choices of players $s \subset (I \setminus i)$. Thus, the larger the magnitude of the Lagrange parameters, $|\lambda_{i,s}|$, the more rational player i is toward the group s . The following two remarks further generalize the observed solution.

Remark 1: It should be straightforward to introduce a noise component to problem 2.3 so that the observed solution converges to the true (P, α) as the number of independent random samples of pure actions goes to infinity. To retain the focus of the current paper, we do not pursue this noisy formulation here. Rather we will revisit the topic of consistency in future work. A more interesting issue is how to incorporate learning within a repeated games structure under the formulation proposed here. This is also beyond the scope of this work and will be developed in future work.

Remark 2: Suppose Sun wants to know if a particular theory (e.g. NE, BNE, QRE) is consistent with the observed game. For the applied researcher this translates into a testable hypothesis on the theory. Fortunately, our framework offers a straightforward method for testing such a hypothesis. First find the observed solution, (P^*, α^*) . Then, impose the structure of the hypothesized theory. One does this by adding constraints to problem 2.3 that appropriately characterize such a theory. For example, if the hypothesized theory is BNE, the additional constraints are the equality constraints that force a common prior as well as the inequalities that characterize expected utility maximization. If the two solutions are practically the same (in terms of their maximal entropy values - the values of their objective functions) then the observed game is consistent with the hypothesized theory. If the entropy values of the solutions are very different, then the hypothesized theory is inconsistent with the observed game. Statistically, this testable hypothesis can be formulated as an entropy ratio test (two times the absolute difference between the two values) which has a Chi Square limiting distribution (Golan (2008)).

2.4 Using Expected Utility Instead of Cross Moments

So far we have incorporated the observed information (actions) with minimal structure, or minimal *a priori* information. Nonetheless, we can extend the observed solution to examine the implications of alternative methods of incorporating the observations, Y , when more information, or structure, is known or assumed. Wolpert (2003) gives the general solution for a game of complete information where the researcher knows the expected utilities of the players. Here we extend Wolpert's solution to the context of the observed game, $\langle \Gamma, Y \rangle$. The Lagrangian from the observed solution is now modified to be

$$\begin{aligned} \mathcal{L} = & H(P, \alpha) + \sum_i \sum_{\theta_i} \lambda_{\theta_i} \left[\sum_{x_i} P_{\theta_i}(x_i) - 1 \right] + \sum_i \lambda_{\alpha_i} \left[\sum_{\theta} \alpha_i(\theta) - 1 \right] \dots \\ & + \sum_i \lambda_i \left[\sum_{\theta} \sum_x P_{\theta_i}(x_i) P_{\theta_{-i}}(x_{-i}) \alpha_i(\theta) u_i(x|\theta_i) - \sum_{\theta} \alpha_i(\theta) u_i(Y|\theta_i) \right], \end{aligned}$$

where the constraint captures the researcher's information about expected utilities. The solution is

$$\tilde{P}_{\theta_i}(x_i) \propto \exp \left[\sum_{k=1}^I \tilde{\lambda}_k \sum_{\theta_{-i}} \sum_{x_{-i}} \tilde{P}_{\theta_{-i}}(x_{-i}) \tilde{\alpha}_k(\theta) u_k(x_i, x_{-i} | \theta_k) \right] \quad (2.6)$$

$$\tilde{\alpha}_i(\theta) \propto \exp \left[\tilde{\lambda}_i \left(\left[\sum_x \tilde{P}_\theta(x) u_i(x | \theta_i) \right] - u_i(Y | \theta_i) \right) \right]. \quad (2.7)$$

$\tilde{P}_{\theta_i}(x_i)$ depends on the sum of expected utilities for all players, given that i is type θ_i and chooses x_i . This parallels the analysis in Wolpert (2003). However, in our formulation, each player has her own beliefs about which game is being played, given by α_i . So, when calculating $\tilde{P}_{\theta_i}(x_i)$, the sum of expected utilities for all players is calculated according to i 's beliefs, $\tilde{\alpha}_i$. Intuitively, this is because i 's likelihood of choosing x_i given θ_i can depend on i 's beliefs about the relative likelihoods of each θ_{-i} .

Looking at 2.6, it is straightforward to see that as the Lagrange parameter $\tilde{\lambda}_i$ gets larger, i 's strategy, $\tilde{P}_i$, gets closer to a best response to $\tilde{P}_{-i}$. In other words, if i 's expected utility given x_i and θ_i is greater than i 's expected utility given x'_i and θ_i , then as $\tilde{\lambda}_i$ goes to infinity, the ratio $\tilde{P}_{\theta_i}(x'_i)/\tilde{P}_{\theta_i}(x_i)$ goes to zero. Intuitively, this arises because as i 's strategy gets closer to a best response, it gets further from the uniform distribution. So if the expected utility constraint says that player i is choosing a strategy close to a best response, then the marginal impact on the objective will be very large.

More precisely, the gradient of the Shannon entropy is zero (entropy is maximized) at the uniform distribution, where P is such that $P_{\theta_i}(x_i) = \frac{1}{|\kappa_i|}$ for all i and θ_i . The gradient of the Shannon entropy is infinite (where entropy is minimized) for P such that P_{θ_i} is a delta function for each i and θ_i . It is smooth over the interior of the set of P 's and globally concave over the set of P 's. So as a constraint forces P_{θ_i} closer to a vertex of $\Delta(X_i)$, the marginal impact of that constraint grows. If a constraint requires P_{θ_i} to approach a best response to P_{-i} , then, generally speaking, P_{θ_i} approaches a vertex or edge of the simplex $\Delta(X_i)$. This is because i 's best response will generally be a subset of her pure strategy for each P_{-i} and θ_i . In this case, the marginal impact of the constraint will be very large, and the entropy of P_{θ_i} will approach zero.

The parameter $\tilde{\lambda}_j$ also has an effect on $\tilde{P}_i$. That is, $\tilde{\lambda}_j$ determines the way j 's observed action impacts Sun's estimate of i 's strategy. If j 's expected utility given x_i and θ_i is greater than j 's expected utility given x'_i and θ_i , then as $\tilde{\lambda}_j$ goes to infinity, the ratio $\tilde{P}_{\theta_i}(x'_i)/\tilde{P}_{\theta_i}(x_i)$ goes to zero. However, $\tilde{\lambda}_j$ shows up in $\tilde{P}_{\theta_i}$ for all i and θ_i . So as $\tilde{\lambda}_j$ goes to infinity, $\tilde{P}_{-j} = \cup_{i \neq j} \{\tilde{P}_{\theta_i}\}_{\theta_i}$ converges to a set of delta functions $\tilde{\delta}_{-j} = \cup_{i \neq j} \{\tilde{\delta}_{\theta_i}\}_{\theta_i}$ that maximize j 's expected utility given $\tilde{P}_{-i}$ and α_i (for each $i \neq j$). Thus, as $\tilde{\lambda}_j$ goes to infinity, every other player selects the strategy that maximizes j 's expected payoff given what everyone else is doing. So a large value of $\tilde{\lambda}_j$ not only has the effect of making $\tilde{P}_j$ a best response to $\tilde{P}_{-j}$, but it also has the effect of making each $\tilde{P}_i$ ($i \neq j$) a best response for j to $\tilde{P}_{-i}$. The opposite holds as $\tilde{\lambda}_i$ goes to negative infinity.

As expected, Sun's estimate of player i 's beliefs, $\tilde{\alpha}_i(\theta)$, also depends on i 's expected utility. More precisely, it depends on the difference between i 's expected utility given θ_i and i 's observed utility given θ_i . If $\tilde{\lambda}_i$ is positive, then Sun estimates that i gives more weight to games in which his expected utility is greater than his observed payoff. The greater this difference, the greater the prior weight assigned to θ . As $\tilde{\lambda}_i$ goes to infinity, the ratio $\tilde{\alpha}_i(\theta')/\tilde{\alpha}_i(\theta)$ converges to zero for θ and θ' such that i 's expected utility given θ is greater than i 's expected utility given θ' . Intuitively, this has the same interpretation as before. The constraint on i 's utility forces $\tilde{\alpha}_i$ toward a vertex of the simplex, $\Delta(\Theta)$, where the gradient is infinite. Suppose θ_i^* maximizes the difference

$$\left[\sum_x \tilde{P}_\theta(x) u_i(x|\theta_i) \right] - u_i(Y|\theta_i),$$

then as $\tilde{\lambda}_i$ goes to infinity, $\tilde{\alpha}_i(\theta)$ converges to zero for all $\theta \in \Theta$ such that the i 'th component is $\theta_i \neq \theta_i^*$. Conversely, all $\theta \in \Theta$ such that the i 'th component is θ_i^* receive equal weight, i.e. $\tilde{\alpha}_i(\theta_i^*, \theta_{-i}) = \frac{1}{|\Theta_{-i}|}$. We also note that the beliefs of player i depend only indirectly on the utilities of the other players, through $\tilde{P}_\theta$.

2.5 Example 1

This simple example demonstrates the observed game and solution. The matrix below (Table 1) summarizes what Sun knows about the strategic setting. This is Γ . There are two players, $\{Mom, Dad\}$, each with two pure strategies, $\{opera, soccer\}$ (or $\{O, S\}$). Each player may be one of two possible types, $\{nice, cruel\}$, and has corresponding type-contingent payoffs. Altogether, there are four possible games. Note that Sun does not know the player's beliefs about the relative likelihood of each game. Sun also does not know (and will not assume) that the players are expected utility maximizers.

		Dad				
		nice		cruel		
		O	S	O	S	
Mom	nice	O	2,1	0,0	2,0	0,2
		S	0,0	1,2	0,1	1,0
	cruel	O	0,1	2,0	0,0	2,2
		S	1,0	0,2	1,1	0,0

Table 1: Normal form representation of Γ .

Mom's overall strategy for this game of imperfect information P_m consists of a strategy for each of her possible types, nice and cruel. When she is nice, we denote by $P_{m|n}(O)$ the probability that she chooses *opera* (O), and by $P_{m|n}(S)$ the probability that she chooses

soccer (S). When she is cruel, these probabilities are $P_{m|c}(O)$ and $P_{m|c}(S)$. For Dad, the probabilities are $P_{d|n}(O)$, $P_{d|n}(S)$, $P_{d|c}(O)$, and $P_{d|c}(S)$ respectively.

Mom believes that she has a probability $\alpha_m(\theta_m = \text{nice})$ of being nice and probability $1 - \alpha_m(\theta_m = \text{nice})$ of being cruel. She believes that Dad has probability $\alpha_m(\theta_d = \text{nice})$ of being nice and $1 - \alpha_m(\theta_d = \text{nice})$ of being cruel. Dad's beliefs are given by α_d . Because Sun does not assume that players are rational, he does not need to assume that they share the same prior. Therefore, we have that $\alpha_m(\theta_i)$ is not necessarily equal to $\alpha_d(\theta_i)$. Finally, both players believe that players' types are independent. Therefore, Mom's belief that both players are nice is $\alpha_m(\theta_m = \text{nice}, \theta_d = \text{nice}) = \alpha_m(\theta_m = \text{nice})\alpha_m(\theta_d = \text{nice})$.

Sun has one additional piece of information: the observed actions Y . Using that observation Y and what he knows about the game (or potential set of games) Γ , Sun wishes to know (P, α) . To do so, Sun builds on the principle of insufficient reason and the maximum entropy formulation 2.3. He finds the set of distributions (P, α) that maximize total entropy subject to the cross moments captured by Equations 2.1.

Table 2 gives the observed solution for each possible observation as well as the entropy of that solution. Note that since there are only two players, the cross moments of equation 2.1 are equivalent to the constraints on expected utility from subsection 2.4. Therefore, we have that the observed solution is the same for both methods of incorporating the information of the observed game.

	(O, O)	(O, S)	(S, O)	(S, S)
$H(P^*, \alpha^*)$	5.402	5.483	5.331	5.402
$P_{m n}^*(O)$	0.528	0.514	0.444	0.436
$P_{m c}^*(O)$	0.527	0.508	0.437	0.467
$P_{d n}^*(O)$	0.564	0.486	0.556	0.472
$P_{d c}^*(O)$	0.533	0.492	0.563	0.474
$\alpha_m^*(\theta_m = n)$	0.378	0.623	0.288	0.719
$\alpha_m^*(\theta_d = n)$	0.5	0.5	0.5	0.5
$\alpha_d^*(\theta_m = n)$	0.5	0.5	0.5	0.5
$\alpha_d^*(\theta_d = n)$	0.719	0.623	0.288	0.378

Table 2: Observed solutions and associated entropy values for each profile of strategies.

2.6 Comparing the Observed Solution with BNE

We next motivate an alternative to the observed solution by comparing the observed solution with the traditional equilibrium concept of Bayesian Nash equilibrium (BNE). The BNE arises as a result of additional information that is unavailable to Sun. This additional information is incorporated as assumptions on common knowledge, expected utility maximization, and a common prior. The game of imperfect information, Γ^α , is given by

- A finite set of players $I = \{1, \dots, I\}$.

- For each player i , a set of κ_i pure strategies X_i .
- For each player i , a set of τ_i possible types Θ_i , where $\sum_i \tau_i = \tau$.
- A set of possible games $\Theta = \Theta_1 \times \dots \times \Theta_I$.
- A commonly known prior distribution over games α .
- For each player i , a type-contingent payoff function $u_i(\cdot|\theta_i) : X \mapsto \mathbb{R}$.

As a reminder, we use the notation $\delta_i(x_i)$ to denote the Dirac delta function for player i centered about x_i .

Definition 2.1. A *Bayesian Nash equilibrium (BNE)* of the game of imperfect information, Γ^α , is a profile of “type-contingent” strategies P^b such that for all i , θ_i and x'_i

$$\sum_{\theta_{-i}} \sum_x P_{\theta_i}^b(x_i) P_{\theta_{-i}}^b(x_{-i}) \alpha^b(\theta_{-i}) u_i(x|\theta) \geq \sum_{\theta_{-i}} \sum_x \delta_i(x'_i) P_{\theta_{-i}}^b(x_{-i}) \alpha^b(\theta_{-i}) u_i(x|\theta). \quad (2.8)$$

To obtain the BNE from Γ^α , the following assumptions are required:

B1 common knowledge of Γ ,

B2 common knowledge of expected utility maximization.

However, Γ and Γ^α differ only in that Γ^α includes the common prior. Therefore, to obtain the BNE from Γ , the following additional assumption is required:

B3 common priors

This is in contrast to the assumptions required to obtain the observed solution, $O1$ and $O2$ together with the observed cross moments information that Sun uses to capture the relevant information from the observed game, $\langle \Gamma, Y \rangle$. Due to the strong informational asymmetry of the observed and BNE solutions, it would be surprising if the two ever coincide. For this to happen, we need $\alpha_i^* = \alpha_j^* = \alpha^b$ for all $i, j \in I$ and $P_{\theta_i}^* = P_{\theta_i}^b$ for all i and θ_i . However, Sun may want to consider (and test) the hypothesis that the observed game satisfies assumptions $B1$ to $B3$. To do so, Sun may write these as additional

constraints for problem 2.3. Sun's problem is then:

$$\begin{aligned}
\max_{P, \alpha} H(P, \alpha) = & - \left[\sum_i \sum_{\theta_i} \sum_{x_i} P_{\theta_i}(x_i) \ln(P_{\theta_i}(x_i)) + \sum_i \sum_{\theta} \alpha(\theta) \ln(\alpha(\theta)) \right] \\
\text{s.t.} \quad & \sum_{\theta} \sum_x Y_i(x_i) Y_j(x_j) \alpha(\theta) u_i(x|\theta) = \sum_{\theta} \sum_x P_{\theta_i}(x) P_{\theta_j}(x) \alpha(\theta) u_i(x|\theta) \\
& \sum_{x_i} P_{\theta_i}(x_i) = 1 \quad \forall i, \theta_i \\
& \sum_{\theta} \alpha(\theta) = 1 \quad \text{and} \\
& \sum_{\theta_{-i}} \sum_x P_{\theta_i}(x_i) P_{\theta_{-i}}(x_{-i}) \alpha(\theta_{-i}) u_i(x|\theta) \geq \sum_{\theta_{-i}} \sum_x \delta_i(x'_i) P_{\theta_{-i}}(x_{-i}) \alpha(\theta_{-i}) u_i(x|\theta) \dots \\
& \dots \forall i, \theta_i \quad \text{and} \quad x'_i.
\end{aligned} \tag{2.9}$$

This says that out of all distributions that are consistent with the observed moments, Sun searches for the one that maximizes entropy *and* is consistent with the BNE requirements. This yields the BNE solution (P^b, α^b) . Note that it is entirely possible (perhaps likely) that no such BNE exists because the first and last constraints may be incompatible. That is, there may not exist a BNE that has the same cross moments as the observations. In any event, the Lagrangian is

$$\begin{aligned}
\mathcal{L} = & H(P, \alpha) + \sum_{\theta_i} \lambda_{\theta_i} \left[\sum_{x_i} P_{\theta_i}(x_i) - 1 \right] + \lambda_{\alpha} \left[\sum_{\theta} \alpha(\theta) - 1 \right] \dots \\
& + \sum_{ij} \lambda_{ij} \left[\sum_{\theta} \sum_x P_{\theta_i}(x_i) P_{\theta_j}(x_j) \alpha(\theta) u_i(x|\theta_i) - Y_i(x_i) Y_j(x_j) \alpha(\theta) u_i(Y|\theta_i) \right] \dots \\
& + \sum_i \sum_{\theta_i} \sum_{x'_i} \lambda_{\theta_i, x'_i} \left[\sum_{\theta_{-i}} \sum_x P_{\theta_i}(x_i) P_{\theta_{-i}}(x_{-i}) \alpha(\theta_{-i}) u_i(x|\theta_i) - \delta_i(x'_i) P_{\theta_{-i}}(x_{-i}) \alpha(\theta_{-i}) u_i(x|\theta_i) \right]
\end{aligned}$$

where λ_{θ_i, x'_i} is the Lagrange multiplier for the BNE constraint that requires that P_{θ_i} does better than the Dirac delta function around x'_i . Note that the common prior assumption, B3, is directly incorporated into the Lagrangian, i.e. $\alpha(\theta)$ is not subscripted by i . From

this we can write the solution to problem 2.9 as

$$P_{\theta_i}^{b*}(x_i) \propto \exp \left[\sum_{j \neq i} \sum_{\theta_{-i}} \sum_{x_{-i}} P_{\theta_j}^{b*}(x_j) \alpha^{b*}(\theta) [\lambda_{ij}^{b*} u_i(x|\theta_i) + \lambda_{ji}^{b*} u_j(x|\theta_j)] \right] \dots \quad (2.10)$$

$$\begin{aligned} & \exp \left[\sum_{x'_i} \lambda_{\theta_i, x'_i}^{b*} \left[\sum_{\theta_{-i}} \sum_{x_{-i}} P_{\theta_{-i}}^{b*}(x_{-i}) [1 - \delta_i(x'_i)] \alpha^{b*}(\theta_{-i}) u_i(x|\theta_i) \right] \right] \dots \\ & \exp \left[\sum_{j \neq i} \sum_{\theta_j} \sum_{x'_j} \lambda_{\theta_j, x'_j}^{b*} \left[\sum_{\theta_{-i,j}} \sum_{x_{-i}} P_{\theta_{-i,j}}^{b*}(x_{-i,j}) [P_{\theta_j}^{b*}(x_j) - \delta_j(x'_j)] \alpha^{b*}(\theta_{-j}) u_j(x|\theta_j) \right] \right] \end{aligned}$$

$$\alpha^{b*}(\theta) \propto \exp \left[\sum_{ij} \lambda_{ij}^{b*} \sum_x \left[P_{\theta_i}^{b*}(x_i) P_{\theta_j}^{b*}(x_j) - Y_i(x_i) Y_j(x_j) \right] u_i(x|\theta_i) \right] \dots \quad (2.11)$$

$$\exp \left[\sum_i \sum_{x'_i} \lambda_{\theta_i, x'_i}^{b*} \sum_x P_{\theta_{-i}}^{b*}(x_{-i}) [P_{\theta_i}^{b*}(x_i) - \delta_i(x'_i)] u_i(x|\theta_i) \right].$$

If a feasible solution exists, then (P^{b*}, α^{b*}) is the BNE that is consistent with the moment conditions and assumptions *B1* to *B3*. We call this the *observed Bayesian solution*. We can now make a comparison between the observed solution, (P^*, α^*) , and the observed Bayesian solution, (P^{b*}, α^{b*}) . The two solutions differ in the amount of information used to derive them. Less information is used for solving the observed solution. The difference between $H(P^*, \alpha^*)$ and $H(P^{b*}, \alpha^{b*})$ is the value of the informational content in assumptions *B1* to *B3*. In fact, under the framework we develop here, assumptions *B1* to *B3* can be tested to see whether they are consistent with what the players know and play. This can be done by means of the Entropy Ratio test discussed in Remark 2.

The following result sheds further light on the relationship between the observed and observed Bayesian solutions.

Theorem 1. *Assumptions O1 and O2 are implied by assumptions B1 and B3.*

Proof. O1 follows from B1, because B1 says there is common knowledge of Γ , hence each player knows Γ . Furthermore, Γ contains X_i for each i . So each player i knows X_i . O2 follows from B3 because the common prior is a probability distribution over Θ . $\square$

In other words, the difference between the entropies of the observed solution and the observed Bayesian solution is the information content of *B2* (EU maximization) together with the information in *B1* that is not in *O1* as well as the information in *B3* that is not in *O2*. The above theorem reinforces the motivation for the observed solution developed here. It is intended for situations in which certain strong assumptions regarding common knowledge of perfect rationality, common knowledge of Γ and common priors are not necessarily satisfied or are not known to the researcher. If the weaker assumptions *O1*

and $O2$ are satisfied, the observed solution provides a starting point for analyzing the observed game.

2.7 Example 2

We may wish to compare the observed solutions from the example of Section 2.5 with the BNE of Γ given a common prior, α^b . The first thing to notice is that, under the observed solution, Mom and Dad do not have the same beliefs. That is, out of all (P, α) pairs that are consistent with the moment conditions, the pair that maximizes entropy is such that the players' beliefs are different. In contrast, a BNE requires a common prior. Therefore, there is no BNE such that the beliefs, α^b , coincide with the beliefs from the observed solution, α^* .

We also note that, under the observed solution, Mom and Dad have fully mixed strategies for every observed pair of actions. This is in contrast to what can occur in a BNE. To see this, note that in order for Mom to be willing to play a mixed strategy when she is nice, she must be indifferent between her expected payoffs from playing O and S ; $U_m(O, P_d|nice) = U_m(S, P_d|nice)$. This equation implies

$$\alpha(\theta_d = nice) = \frac{1 - 3P_{d|c}(O)}{3(P_{d|n}(S) - P_{d|c}(O))}. \quad (2.12)$$

We can write the analogous equation for when Mom is cruel; $U_m(O, P_d|cruel) = U_m(S, P_d|cruel)$. This equation implies

$$\alpha(\theta_d = nice) = \frac{2 - 3P_{d|c}(O)}{3(P_{d|n}(S) - P_{d|c}(O))}. \quad (2.13)$$

Equations 2.12 and 2.13 are incompatible. Therefore, there is no strategy of Dad, P_d , such that Mom has a best response to randomize both when she is nice and when she is cruel. In other words, in a BNE, if Mom randomizes when she is nice, then she does not randomize when she is cruel. The same holds true for Dad. Hence there are no BNE, P^b , that coincide with the observed solutions, P^* , for any observed action pair, Y in that example.

Despite this fact, we can examine what happens when we force the observed solution to be a BNE. That is, we require not only that our observed solution satisfies the moment conditions from equation 2.1, but also that our observed solution is a BNE. This is the observed Bayesian solution described in subsection 2.6 above, and we denote such solutions as (P^{b*}, α^{b*}) . We present the observed Bayesian solutions to our simple example in Table 3 for each pair of observed actions.

Because (P^{b*}, α^{b*}) resulted from an optimization with more constraints than (P^*, α^*) , it contains at least as much information as the information in (P^*, α^*) and probably more. In our example, this additional information is reflected by the fact that the maximized

	(O, O)	(O, S)	(S, O)	(S, S)
$H(P^{b*}, \alpha^{b*})$	3.474	4.045	2.903	3.474
$P_{m n}^{b*}(O)$	1	0.333	1	0.167
$P_{m c}^{b*}(O)$	0.333	1	0.167	1
$P_{d n}^{b*}(O)$	0.833	0.666	0	0
$P_{d c}^{b*}(O)$	0	0	0.833	0.666
$\alpha^{b*}(\theta_m = n)$	0.5	0.5	0.2	0.8
$\alpha^{b*}(\theta_d = n)$	0.8	0.5	0.2	0.5

Table 3: Observed Bayesian solutions and associated entropy values for each profile of strategies.

entropy with the BNE constraints, $H(P^{b*}, \alpha^{b*})$, is less than without the BNE constraints, $H(P^*, \alpha^*)$. Once again, the difference between these numbers is interpreted as the information content of assumptions $B1$ to $B3$ over and above the information content of assumptions $O1$ and $O2$, or as the additional structure imposed on the solution.

3 Observed Solution with Priors

3.1 The Priors

Consider now the case where Sun has some prior beliefs about the distributions (P, α) . We represent these prior distributions as (P^0, α^0) . These priors are Sun's belief about the strategies and beliefs of the players *before* Sun observes Y . They may come from a previous round of the game, from intuition or theory or from any other source of information. These prior beliefs are *different* from the prior beliefs of the players, which are represented by α , a distribution over Θ .

The priors are the only additional information Sun has regarding 2.3. So, Sun's objective is now to incorporate these priors within the observed game. However, Sun wants to do so in such a way that if the observed data are inconsistent with these priors, the data will determine the solution. Stated differently, Sun uses the observed data, represented as the cross moments 2.1, to update her prior beliefs. This means the priors must be incorporated in the objective function (and not as constraints in the optimization model). Thus, out of all distributions that are consistent with 2.1 and 2.2, Sun is searching for the one that is closest to the prior beliefs where "closeness" is defined in terms of an entropy divergence measure. This entropy divergence is given by

$$K(P, \alpha || P^0, \alpha^0) = \left[\sum_i \sum_{\theta_i} \sum_{x_i} P_{\theta_i}(x_i) \ln \left(\frac{P_{\theta_i}(x_i)}{P_{\theta_i}^0(x_i)} \right) + \sum_i \sum_{\theta} \alpha_i(\theta) \ln \left(\frac{\alpha_i(\theta)}{\alpha_i^0(\theta)} \right) \right]. \quad (3.1)$$

This measure captures the entropy distance between the prior (P^0, α^0) and (P, α) . Sun chooses (P, α) to minimize equation 3.1 subject to the constraints 2.1 and 2.2. The solution is (P^{0*}, α^{0*}) , called the *observed solution with priors*.

Intuitively, before observing Y , Sun begins with some beliefs (P^0, α^0) . After observing Y and combining it with Γ through the cross moments, Sun realizes that his beliefs may not be consistent with the newly observed information. Therefore, he searches for distributions (P^{0*}, α^{0*}) that *are* consistent with the newly observed information *and* are closest to his prior beliefs. In the event that Sun's prior beliefs are fully consistent with the moment conditions, we get that $(P^{0*}, \alpha^{0*}) = (P^0, \alpha^0)$. In such a case, the KL distance is minimized at zero because the newly observed data have no new information relative to Sun's prior beliefs: the priors are reconfirmed by the newly observed actions of all players.

The observed solution with priors is a simple generalization of the observed solution from problem 2.3. In the latter, we choose the (P, α) that is closest to the uniform distribution and satisfies 2.1. For the observed solution with priors, we choose the pair (P, α) that is closest to the priors and satisfies 2.1. Hence, using uniform priors within equation 3.1 reduces the model to the original problem 2.3. The difference between the two models is that when prior beliefs are not available for Sun, or if Sun refuses to use such priors, Sun searches for the most uniform solution. By maximizing the entropy, Sun maximizes her missing information about the game and the players.

Prior beliefs also make it possible to revisit a comparison of BNE and the observed solution. If Sun's prior beliefs are that players will choose a particular BNE, these beliefs can easily be introduced as the priors (P^0, α^0) . Then, the observed solution (P^{0*}, α^{0*}) will be the distribution that is consistent with the moment conditions and minimizes KL divergence from that BNE. In this case, the observed solution with priors is "tilted" in the direction of expected utility maximization and common priors. This is in contrast to being "tilted" in the direction of a uniform distribution, as is the case with the observed solution. However, unlike the observed Bayesian solution, which is constrained to be a BNE, the observed solution with a BNE prior is not generally a BNE. Rather, the observed solution with a BNE prior is the pair (P, α) that is closest to the BNE and satisfies the observed moments 2.1.

A basic question that Sun is facing is how to choose her priors. Naturally the answer depends on what Sun knows. We provide here a simple example of one case. We can consider what happens when, instead of beginning with prior distributions (P^0, α^0) , Sun begins with just the belief that assumptions $B1$ - $B3$ hold. Note that this belief is not in the form of probability distributions, (P^0, α^0) . Rather, there are many probability distributions that satisfy $B1$ - $B3$, the set of BNE. In order to select a single BNE Sun invokes again the maximum entropy principle. That is, out of the set of probability distributions that satisfy assumptions $B1$ - $B3$, Sun chooses as her prior the probability distributions that maximize entropy. We call this the *rational prior* and write it as (P^r, α^r) . Put differently, Sun chooses as her prior the BNE that has minimal information content (maximum entropy).

Having determined probability distributions to represent her prior beliefs, Sun then

chooses (P^{r*}, α^{r*}) to minimize the KL divergence from the rational prior, (P^r, α^r) . We refer to (P^{r*}, α^{r*}) as the *observed rational prior solution*.

We can now compare the informational contributions of the different priors on the observed solutions. To do so, we compare the entropies of the different solutions, or the entropy distance (the KL measure) between each prior and the resulting solution. Note that for the basic observed solution, the comparison prior is the uniform distribution.

3.2 Example 3

Here we extend Example 1 (Section 2.5) to demonstrate the impact of Sun’s prior beliefs on the observed solutions. Suppose Sun’s prior beliefs, (P^0, α^0) , are summarized by the rightmost column of Table 4. These priors are not meant to be economically significant. In other words, for the purpose of the example, we offer no explanation for why Sun’s prior beliefs are given by the numbers in the rightmost column of Table 4. These priors are simply numbers between zero and one such that P_{θ_i} and α_i are proper probability distributions. For example, these priors say that Sun believes that when Mom is nice she will choose *opera* with probability $P_{m|n}^{0*}(O) = 0.2$ and *soccer* with probability 0.8.

	(O, O)	(O, S)	(S, O)	(S, S)	(P^0, α^0)
$K(P^{0*}, \alpha^{0*} P^0, \alpha^0)$	0.162	0.326	0.879	0.596	
$H(P^{0*}, \alpha^{0*})$	4.611	4.907	4.227	4.111	4.587
$P_{m n}^{0*}(O)$	0.228	0.215	0.052	0.059	0.2
$P_{m c}^{0*}(O)$	0.355	0.306	0.271	0.224	0.3
$P_{d n}^{0*}(O)$	0.446	0.545	0.866	0.104	0.4
$P_{d c}^{0*}(O)$	0.493	0.475	0.384	0.618	0.5
$\alpha_m^{0*}(\theta_m = n)$	0.331	0.692	0.493	0.658	0.6
$\alpha_m^{0*}(\theta_d = n)$	0.701	0.699	0.721	0.685	0.7
$\alpha_d^{0*}(\theta_m = n)$	0.798	0.800	0.843	0.841	0.8
$\alpha_d^{0*}(\theta_d = n)$	0.913	0.628	0.630	0.642	0.9

Table 4: Observed solutions given priors and minimized KL divergence and entropy values for each profile of strategies.

For each of the four possible observed strategy profiles, Sun finds the distribution that minimizes the KL distance to (P^0, α^0) . This is the observed solution with priors. These solutions are displayed in the columns of Table 4 along with the associated KL values $K(P^{0*}, \alpha^{0*} || P^0, \alpha^0)$ and the entropy values $H(P^*, \alpha^*)$. Note that the observed solutions with priors (P^{0*}, α^{0*}) are “tilted” in the direction of (P^0, α^0) . Also note that the entropy values $H(P^{0*}, \alpha^{0*})$ can be greater or less than $H(P^0, \alpha^0)$. However, they are always less than the entropy of the observed solution, $H(P^*, \alpha^*)$.

Similarly, we investigate the rational prior solution for the setup of Example 1. The first step is to determine the rational prior, (P^r, α^r) . Recall this is done by maximizing

the entropy subject to assumptions $B1$ to $B3$. The rational prior for Γ is shown in the rightmost column of Table 5. Using the rational prior, we find the observed rational prior solution, (P^{r*}, α^{r*}) , for each Y . These solutions are displayed in the columns of Table 5. As in Table 4, the entropy values $H(P^{r*}, \alpha^{r*})$ can be greater or less than $H(P^r, \alpha^r)$. However, also as in Table 4, the entropy values are always less than the entropy of the observed solution, $H(P^*, \alpha^*)$, because the rational prior represents additional information.

	(O, O)	(O, S)	(S, O)	(S, S)	(P^r, α^r)
$K(P^{r*}, \alpha^{r*} P^r, \alpha^r)$	0.136	0.009	0.609	0.136	
$H(P^{r*}, \alpha^{r*})$	3.901	4.082	3.261	3.901	4.094
$P_{m n}^{r*}(O)$	0.373	0.383	0.157	0.396	0.399
$P_{m c}^{r*}(O)$	1	1	1	1	1
$P_{d n}^{r*}(O)$	0.604	0.617	0.843	0.627	0.601
$P_{d c}^{r*}(O)$	0	0	0	0	0
$\alpha_m^{r*}(\theta_m = n)$	0.502	0.513	0.385	0.795	0.555
$\alpha_m^{r*}(\theta_d = n)$	0.544	0.572	0.778	0.584	0.555
$\alpha_d^{r*}(\theta_m = n)$	0.584	0.572	0.778	0.544	0.555
$\alpha_d^{r*}(\theta_d = n)$	0.795	0.513	0.385	0.502	0.555

Table 5: Rational prior and observed rational prior solutions for each profile of strategies.

4 Observed Solution and QRE

Building on the analysis of Wolpert (2003) for games of complete information, we use the approach of this paper to derive a version of QRE for games of incomplete information. To do so, consider the situation where the researcher knows the common prior α and the type-contingent strategy distributions of all players except i , P_{-i} . This information can be used together with the fact that player i achieves expected utility equal to $\sum_{\theta} \alpha(\theta) u_i(Y|\theta_i)$ (observed utility). As before, to find P_i , we employ the maximum entropy formulation. The Lagrangian is

$$\begin{aligned} \mathcal{L}_i = & H(P_i) + \beta_i \left[\sum_{\theta} \sum_x P_{\theta}(x) \alpha(\theta) u_i(x|\theta_i) - \sum_{\theta} \alpha(\theta) u_i(Y|\theta_i) \right] \\ & + \sum_{\theta_i} \beta_{\theta_i} \left[\sum_{x_i} P_{\theta_i}(x_i) - 1 \right]. \end{aligned}$$

Maximizing the Lagrangian, $\mathcal{L}_i$, with respect to P_{θ_i} yields the solution

$$P_{\theta_i}^\alpha(x_i) \propto \exp \left[\beta_i^\alpha \sum_{\theta_{-i}} \sum_{x_{-i}} P_{\theta_{-i}}(x_{-i}) \alpha(\theta) u_i(x|\theta_i) \right]. \quad (4.1)$$

We can formulate a complete system of equations by writing the solution for each player. Then, by Brouwer's fixed point theorem, a solution, P^α , for the system of equations exists. In keeping with tradition, we refer to this solution as the *Imperfect Information QRE* (or IIQRE).¹ This solution is similar to the HQRE of Rogers et al. (2009). The difference is with regard to what the players are uncertain about. In particular, the HQRE assumes that players are uncertain about each other's rationality, whereas the IIQRE assumes players are uncertain about each other's preferences.

Note that the IIQRE is a set of solutions parameterized by α , the common prior. So, if Sun truly does not know the common prior, then he can once again apply the Maximum Entropy principle to choose it. That is, Sun chooses among all possible IIQRE by finding the distribution α that maximizes the entropy of the solutions $H(P^\alpha, \alpha)$. Doing so yields the solutions $(\tilde{P}^q, \tilde{\alpha}^q)$. Here $\tilde{\alpha}^q$ is the common prior that maximizes $H(P^\alpha, \alpha)$ and $\tilde{P}^q$ is the solution P^α evaluated at the maximal entropy value of α , i.e. $\tilde{P}^q = P^{\tilde{\alpha}^q}$. We refer to $(\tilde{P}^q, \tilde{\alpha}^q)$ as the *Maximum Entropy IIQRE*. Note that the Maximum Entropy IIQRE solution concept selects a single BQRE from the set of all games of incomplete information given by Γ .

The parameter $\tilde{\beta}_i^q$, which corresponds to $\beta_i^{\tilde{\alpha}^q}$ (i.e. β_i^α at the maximal entropy value of α), gives the degree to which $\tilde{P}_i^q$ is a best response to $\tilde{P}_{-i}^q$, just as in McKelvey and Palfrey (1995). However, in our information-theoretic derivation, $\tilde{\beta}_i^q$ results from the constraint on i 's expected utility, rather than from assumptions about the log-Weibull distribution that governs the errors i makes in estimating the expected utility of each of her pure strategies. Hence, in the same way that Wolpert provides a purely information-theoretic derivation of logit QRE, we provide a purely information theoretic derivation of logit IIQRE.

Note that simply by writing the system of equations corresponding to equation 4.1, we are already assuming that the players are in equilibrium. This assumption is in addition to the assumption that i 's expected utility is equal to i 's observed utility. Hence, more structure is imposed on the game so the Maximum Entropy IIQRE contains more information than the observed solution from Equations 2.6. In particular, $H(\tilde{P}^q, \tilde{\alpha}^q) \leq H(\tilde{P}, \tilde{\alpha})$. However, to make sure the two solutions are compared with the same dimensionality, we evaluate I copies of the entropy of the common prior $\tilde{\alpha}^q$. That is, we compare $H(\tilde{P}^q, \tilde{\alpha}^q) = H(\tilde{P}^q) + I(H(\tilde{\alpha}^q))$ with $H(\tilde{P}, \tilde{\alpha})$.

¹The term Bayesian QRE (or BQRE) is already defined, in Esaray et al. (2008), to be a model of an extensive bargaining game in which players choose quantal responses (instead of best responses) in each subgame (given that the other players are also choosing quantal responses) and update beliefs according to Bayes rule. Perhaps Perfect Bayesian QRE (PBQRE) would have been a more appropriate name for that model since it seems to be the QRE analog of perfect Bayesian equilibrium (PBE), whereas BQRE indicates the QRE analog of Bayesian Nash equilibrium (BNE) as is presented here.

Since the equilibrium assumption carries the requirement that players share a common prior, Sun has an estimate of $\tilde{\alpha}^q(\theta)$ for each θ in Θ . This implies that the value $\tilde{P}^q(x)$ is well-defined for all $x \in X$. It is simply

$$\tilde{P}^q(x) = \prod_i \sum_{\theta} \tilde{\alpha}^q(\theta) \tilde{P}_{\theta_i}^q(x_i).$$

The existence of a coherent prediction $\tilde{P}^q$ is in contrast to the observed solutions, P^* and $\tilde{P}$, where the lack of an equilibrium assumption means that we allow each player to have her own beliefs. The lack of a common prior then implies that $P(x)$ is not well defined in terms of $\{P_i(x)\}_i$. Once again, this is because the lack of a common prior means that players are effectively playing different games, and the observed solution framework does not formulate an estimate of Sun's *own* beliefs about the relative likelihoods of each game, i.e. α_{Sun} .

There are other interesting differences between the observed solution of equation 2.6 and the Maximum Entropy IIQRE. For example, unlike $\tilde{P}_{\theta_i}$, $\tilde{P}_{\theta_i}^q$ does not depend on the expected utility of any other player j . This arises because Sun takes the strategies P_{-i} as given when maximizing the entropy of P_i . In other words, to obtain $\tilde{P}$, Sun maximizes the entropy of (P, α) , the sum of the entropies of all P_i and α_i . To obtain $\tilde{P}^q$, Sun maximizes the entropy of each P_i subject to P_{-i} . The difference is analogous to the difference between global and distributed optimization.

Now suppose that Sun observes more than just i 's expected utility. Assume that Sun knows, for each j , the cross moments (i, j) given by Equations 2.1, where α_i is replaced by α , the common prior. Proceeding exactly as above, we can write the Lagrangian as

$$\begin{aligned} \mathcal{L}_i = & H(P_i) + \sum_j \beta_{ij} \left[\sum_{\theta} \sum_x P_{\theta_i}(x_i) P_{\theta_j}(x_j) \alpha(\theta) u_i(x|\theta_i) - Y_i(x_i) Y_j(x_j) \alpha(\theta) u_i(Y|\theta_i) \right] \\ & + \sum_{\theta_i} \beta_{\theta_i} \left[\sum_{x_i} P_{\theta_i}(x_i) - 1 \right]. \end{aligned}$$

Maximizing the Lagrangian, $\mathcal{L}_i$, with respect to P_{θ_i} yields the solution

$$P_{\theta_i}^{\alpha'}(x_i) \propto \exp \left[\sum_j \beta_{ij}^{\alpha'} \sum_{\theta_{-i}} \sum_{x_{-i}} P_{\theta_j}(x_j) \alpha(\theta) u_i(x|\theta_i) \right]. \quad (4.2)$$

Just as before, Sun recognizes that $P^{\alpha'}$ is a system of equilibrium behavior parameterized by α . We refer to this as the *IIQRE with cross moments*.

However, as before, Sun may not have explicit information about α . So, Sun once again maximizes the entropy of $H(P^{\alpha'}, \alpha)$ with respect to the common prior α . Doing so yields the solutions (P^{*q}, α^{*q}) , called the *Maximum Entropy IIQRE with cross moments*. Here α^{*q} is the common prior that maximizes $H(P^{\alpha'}, \alpha)$ and P^{*q} is the solution $P^{\alpha'}$

evaluated at the maximal entropy value of α , i.e. $P^{*q} = P\alpha^{*q}$.

In contrast to the Maximum Entropy IIQRE, the Maximum Entropy IQRE with cross moments incorporates the pairwise notion of rationality that arises in the observed solution, equations 2.4 and 2.5. That is, when we formulate the QRE via information theory using the cross moments as constraints, as opposed to imposing the expected utility, we get the degree to which i maximizes expected utility with regard to each one of the other players $j \neq i$. So, for $P_{\theta_i}^{*q}(x_i)$, the Lagrange parameter, β_{ij} , gives the degree to which x_i maximizes i 's expected utility against P_j when all other players are using uniform strategies.

This pairwise rationality, β_{ij} , is a far more detailed view of i 's utility maximizing behavior than the rationality β_i , which, in a sense, summarizes i 's utility maximizing behavior with regard to all $j \neq i$. This increased detail *appears* to come at an information cost. In particular, we are using all cross moments for each player i instead of all expected utilities. There are $I - 1$ cross moments for each i . So there are $I \times (I - 1)$ cross moment constraints for the system represented by equation 4.2. However, there are only I expected utility constraints for the IIQRE, the system represented by equation 4.1. So it appears that the use of cross moments carries an additional information cost. In fact, a more accurate view is that this information exists in the observed game (Γ, Y) and is being wasted when unused.

Finally, recall that the observed solution, $P_{\theta_i}^*(x_i)$, depends on the expected utility of j given i 's strategy and that every other player randomizes uniformly. Therefore, if λ_{ji}^* goes to infinity, then i 's strategy P_i^* maximizes j 's expected utility given j 's strategy P_j^* and that every other player randomizes uniformly. In contrast, the analogous term from the IIQRE with cross moments, $P_{\theta_i}^{*q}(x_i)$, does not depend on β_{ji}^{*q} . This is precisely because the IIQRE with cross moments includes the assumption that players are in equilibrium. The observed solution, on the other hand, does not include the equilibrium assumption. The difference is between maximizing total entropy (as is the case with the observed solution) and maximizing the entropy of each player's strategy given the strategies of the other players (as is the case with the IIQRE). This difference, in addition to the fact that i 's strategy does not depend (directly) on β_{ji}^{*q} and j 's expected payoff, means that the equilibrium assumption of the IIQRE with cross moments contains additional information/structure that may be inconsistent with the game played. Put differently, $H(P^{*q}, \alpha^{*q}) \leq H(P^*, \alpha^*)$. As before, the common prior of the IIQRE with cross moments forces us to compare $H(P^*, \alpha^*)$ with $H(P^{*q}, \alpha^{*q})$ defined as $H(P^{*q}) + I(H(\alpha^{*q}))$ so that the dimension involved in calculating entropy is the same on both sides of the inequality.

5 Conclusion

We find that, when studying games, it may be unnecessary to make assumptions regarding expected utility maximization, common knowledge and common priors. In particular, if a researcher observes an instance of play, the researcher can use that observation along with the maximum entropy principle to determine a unique set of strategies, beliefs and

rationality called the observed solution. Such an information-theoretic approach also facilitates an in-depth analysis of common assumptions regarding rationality, common priors, common knowledge and equilibrium through the use of the entropy ratio test.

As discussed above, the observed solution is not intended as a prediction technology. In this way it differs from most solution concepts. Rather it is intended as an inference technology used by an external observer to make judgments with minimal assumptions. Therefore, the observed solution may be most appropriately used in the context of a courtroom or other arbitration scenario. In such a scenario, an *ex post* analysis of parties' perceptions (beliefs), intentions (strategies) and efficacy (rationality) is desired, and standard equilibrium assumptions cannot usually be confirmed.

Our analysis strengthens the connection between information theory and game theory. In particular, we demonstrate that the QRE concept can be significantly generalized. For games of incomplete information we discover that alternative moment conditions lead to alternative interpretations of rationality. These alternative interpretations of rationality are economically intuitive and may prove useful, especially in the context of *ex post* arbitration, as they indicate who is motivating whom.

Future empirical work will apply the observed solution and other ideas explored here to real-world data sets. In such empirical work, the goal is to statistically test common modeling assumptions. This perspective offers an alternative to the extensive work in structural modeling that generally adopts an equilibrium model (and its assumptions) for identification purposes.

Finally, one challenge in making the observed game approach applicable to a broad class of games and, hence, data sets, is that the current formulation is only appropriate for one-shot games. This is because repeated games generally involve learning, which must be formally accounted for within the observed game framework. Future theoretical work will address this challenge.

References

- Camerer, Colin, Teck-Hua Ho, and Juin-Kuan Chong**, “A Cognitive Hierarchy Model of Games,” *Quarterly Journal of Economics*, 2006, *119* (3), 861–898.
- Catilina, E. and A. Golan**, “What is the Game,” 2006. Department of Economics, American University.
- Costa-Gomes, Miguel and Vincent Crawford**, “Cognition and Behavior in Two-Person Guessing Games: An Experimental Study,” *American Economic Review*, December 2006, *96* (5), 1737–1768.
- Esaray, Justin, Bumba Mukherjee, and Will H. Moore**, “Strategic Interaction and Interstate Crises: A Bayesian Quantal Response Estimator for Incomplete Information Games,” *Political Analysis*, 2008, *16*, 250273.
- Golan, A., G. Judge, and D. Miller**, *Maximum Entropy Econometrics: Robust Estimation with Limited Data*, John Wiley & Sons, 1996.
- Golan, Amos**, “Information and Entropy Econometrics A Review and Synthesis,” *Foundations and Trends in Econometrics*, 2008, *2* (1-2), 1–145.
- Harsanyi, John**, “Games with Incomplete Information Played by Bayesian Players, I-III. Part I. The Basic Model,” *Management Science*, 1967, *14* (3). Theory Series.
- Jaynes, E.T.**, “Information theory and statistical mechanics,” *Phys. Rev.*, 1957, *106*, 620.
- Kullback, S. and R.A. Leibler**, “On Information and Sufficiency,” *The Annals of Mathematical Statistics*, 1951, *22* (1), 7986.
- McKelvey, Richard D. and Thomas R. Palfrey**, “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, July 1995, *10* (1), 6–38.
- Nash, John F.**, “Equilibrium Points in N-person Games,” *Proceedings of the National Academy of Sciences*, 1959, *36*, 48–49.
- Radner, Roy**, “Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives,” *Journal of Economic Theory*, April 1980, *22* (2), 136–154.
- Rogers, B.W., T.R. Palfrey, and C.F. Camerer**, “Heterogenous quantal response equilibrium and cognitive hierarchies,” *Journal of Economic Theory*, 2009.
- Shannon, Claude E.**, “A Mathematical Theory of Communication,” *Bell System Technical Journal*, July, October 1948, *27*, 379–423, 623–656.
- Wolpert, D. H.**, “Theory of Collective Intelligence,” in K. Tumer and D. H. Wolpert, eds., *Collectives and the Design of Complex Systems*, Springer New York 2003.