

A Predictive Theory of Games

David H. Wolpert*

October 3, 2006

Abstract

Conventional noncooperative game theory hypothesizes that the joint (mixed) strategy of a set of reasoning players in a game will necessarily satisfy an “equilibrium concept”. All other joint strategies are considered impossible. Moreover, often the number of joint strategies satisfying that equilibrium concept has measure zero. (Indeed, this is often considered a desirable property of an equilibrium concept.) Under this hypothesis the only issue is what equilibrium concept is “correct”.

This hypothesis violates the first-principles arguments underlying probability theory. Indeed, probability theory renders moot the controversy over what equilibrium concept is correct — while in general there are joint (mixed) strategies with zero probability, in general the set {strategies with non-zero probability} has measure greater than zero. Rather than a first-principles derivation of an equilibrium concept, game theory requires a first-principles derivation of a distribution over joint strategies.

However say one wishes to predict a single joint strategy from that distribution. Then decision theory tells us to first specify a loss function, a function which concerns how we, the analyst/scientist *external* to the game, will use that prediction. We then predict that the game will result in the joint strategy that is Bayes-optimal for that loss function and distribution over joint strategies. Different loss functions — different uses of the prediction — give different such optimal predictions. There is no more role for an “equilibrium concept” that is independent of the distribution and choice of loss function. This application of probability theory *to* games, not just within games, is called Predictive Game Theory (PGT).

This paper shows how information theory provides a first-principles argument for how to set a distribution over joint strategies. The connection of this distribution to the bounded rational Quantal Response Equilibrium (QRE) is elaborated. In particular, taking the QRE to be an approximation to the mode of the distribution, correction terms to the QRE are derived. In addition, some Nash equilibria are not approached by any limiting sequence of increasingly rational QRE joint strategies. However it is shown here that every Nash equilibrium is approached with

*D. Wolpert is with NASA Ames Research Center, Moffett Field, CA, 94035
dhw@ptolemy.arc.nasa.gov

a limiting sequence of joint strategies all of which have non-zero probability. (In general though not all strategies in those sequences are modes of the associated distributions over joint strategies.)

It is also shown that in many games, having a probability distribution with support restricted to Nash equilibria — as stipulated by conventional game theory — is impossible. So the external analyst should never predict a Nash outcome for such games. PGT is also used to derive an information-theoretic (and model-independent) quantification of the degree of rationality inherent in a player’s behavior. This quantification arises from the close formal relationship between game theory and statistical physics. That close relationship is also leveraged to extend game theory to situations with stochastically varying numbers of players. This extension can be viewed as providing corrections to the replicator dynamics of conventional evolutionary game theory.

1 Introduction

Consider any scientific scenario, in which one wishes to predict some characteristic of interest y concerning some physical system. To make the prediction one starts with some information/data/prior knowledge \mathcal{S} concerning the system, together with known scientific laws. One then uses probabilistic inference to transform \mathcal{S} into the desired prediction. In particular, in Bayesian inference we produce a posterior probability distribution $P(y | \mathcal{S})$.

Such a distribution is a far more informative structure than a single “best prediction”. However if we wish to synopsise the distribution, we can distill it into a single prediction. One way to do that is to use the mode of the posterior as the prediction. This is called the Maximum A Posterior (MAP) prediction. Alternatively, say we are given a real-valued loss function, $L(y, y')$ that quantifies the penalty we incur when we predict y' and the true value is y . The Bayes optimal prediction then is the value of y' that minimizes the posterior expected loss, $\int dy L(y, y') P(y | \mathcal{S})$. As an example, say that $y \in \mathbb{R}$ and $L(y, y') = (y - y')^2$. Then the Bayes-optimal prediction is the posterior expected value of y , $\int dy y P(y | \mathcal{S})$. Formally, to predict any other value than the Bayes-optimal prediction violates Cox’s and Savage’s axioms concerning the need to use probability theory when doing science (see [1, 2, 3] and Sec. 2.2 below).

As a technical comment, in practice evaluating the Bayes-optimal prediction may be computationally difficult. In addition doing so often requires that the practitioner specify “prior probabilities”, which when done poorly can lead to bad results. Finally, some non-Bayesian axiomatizations of inference have been offered [4].

Due to these reasons, when \mathcal{S} consists of experimental data, in practice often non-Bayesian techniques (e.g., Fisherian or Waldian minimax) are used instead of pure Bayesian techniques. For example, one might use an unbiased estimate of the data’s mean and an associated confidence interval rather than use a Bayes-optimal predicted mean and associated posterior. More sophisticated

non-Bayesian techniques might use the bootstrap [5] procedure, stacking [6], etc.¹

The general relationship between such non-Bayesian techniques and purely Bayesian techniques is extremely subtle [7]. However in the numeric sciences, even using non-Bayesian techniques, the broad approach to analyze experimental data is to use \mathcal{S} (the experimental data) to generate a probability distribution over the quantity of interest. One can then, if desired, generate an associated single “optimal” prediction. (This is why numeric data is presented with “error bars”.) No one has ever suggested why this broad approach would be appropriate when one is analyzing a physical system without humans and \mathcal{S} is experimental data, but not appropriate when one is analyzing a physical system of a set of human players engaged in a game and \mathcal{S} is the game structure.

Indeed, the Bayesian approach can be motivated in purely game-theoretic terms. Say we have a set of possible games, differing in their utility functions, the players, etc. For each such game γ , let $\Delta_{\mathcal{X}}(\gamma)$ indicate the set of all possible joint mixed strategies in γ . Now consider a two-stage “meta-game” Γ that consists of a scientist (S) playing against Nature (N). In this meta-game N ’s space of possible moves is $\{(\gamma, \Delta_{\mathcal{X}}(\gamma))\}$, i.e., the set of all possible games γ , and for each such game, the set of all possible joint mixed strategies q over γ . The mixed strategy of player N is a distribution over this space, $P(\gamma, q \in \Delta_{\mathcal{X}}(\gamma))$.

At the end of the first stage of the meta-game, N ’s mixed strategy is fairly sampled, producing an outcome (γ', q') . Then S is told both γ' and N ’s mixed strategy, $P(\gamma, q \in \Delta_{\mathcal{X}}(\gamma))$. Together, those give S a posterior over what q' is, $P(q' | \gamma')$.

In the second stage, S makes a move in $\Delta_{\mathcal{X}}(\gamma')$, i.e., picks a mixed strategy over the game γ' . We interpret that move of the scientist S as a prediction of what $q \in \Delta_{\mathcal{X}}(\gamma')$ was produced by the sampling of player Nature, N .

As usual in games against Nature, N has no utility function. However S has a utility function, given by the negative of a loss function that quantifies how accurate her move is as a prediction of N ’s move q' . So to maximize her expected utility the scientist wants to choose her move(s) — her prediction of the joint mixed strategy that governs the game γ' — to minimize her expected loss under the posterior $P(q' | \gamma')$.

Our current situation is exactly this setup, where γ' comprises \mathcal{S} . Intuitively, S must translate her partial information into an inference of q' . How to do this is the crux of PGT. Having done this, just as in conventional non-cooperative game theory, the scientist S choose her move to minimize expected loss (maximize expected utility) in Γ , i.e., she predicts that the joint mixed strategy of the game γ' is the one that, under expectation, is as close to the actual one as possible. That is the scientist’s assessment of the game’s “equilibrium”.

Note that for the same γ' and the same inference by S of the mixed strategy $P(q' | \gamma')$, if we change S ’s loss function, we change her prediction. This is a game-theoretic example of how changing the loss function of the scientist

¹Note though that often such techniques can be cast as approximations to Bayesian techniques [1, 2, 3].

external to the game will in general change the associated equilibrium concept mapping γ' to S 's prediction for the mixed strategy of γ' .

The broad approach of converting \mathcal{S} to a distribution and then — if one has a loss function — converting that distribution to a final prediction is the one that will be adopted in this paper. This paper is about how to use this approach to analyze games. In other words, it about how to infer the mixed strategy $P(q)$ of a Nature whose moves are joint mixed strategies of a game γ' .

As a particular example of the implications of this approach, suppose that our prior information \mathcal{S} concerning the game γ' does not *explicitly* tell us that the players in γ' are all fully rational. Then in general, probabilistic inference will produce a non-delta function distribution over the “rationalities” of the players (however that term is defined). In this way, applying probabilistic inference to games intrinsically results in bounded rationality.

To be more concrete, note that mixed strategies in a non-cooperative game are themselves probability distributions. Therefore probabilistic inference concerning mixed strategies involves probability density functions over probability distributions. Now in Shannon's information theory [8, 9, 1, 10] the fundamental physical objects under consideration are probability distributions, in the form of stochastic communications channels. Accordingly, probabilistic inference in information theory also involves probabilities of probabilities. This makes the mathematical tools for probabilistic analysis in information theory contexts — a topic already well-researched — well-suited to a probabilistic analysis of non-cooperative games.

More precisely, a central concept in information theory is a measure of the amount of information embodied in a probability distribution p , known as the Shannon entropy of that distribution, $S(p)$. Amongst its many other uses, Shannon entropy can be used to formalize Occam's razor based on first-principles arguments. This formalization is known as the minimum information (Maxent) principle, and its Bayesian formulation is embodied in what is known as the entropic prior. This can serve as the foundation of a first-principles formalism for probabilistic inference over probability distributions.

Using entropy to perform probabilistic inference this way has proven extraordinarily successful in an extremely large number of applications, ranging from signal processing to machine learning to statistics. Recently it has also been realized that the mathematics underlying this type of inference can be used to do distributed control and/or optimization. In that context the mathematics is known as Probability Collectives (PC). Preliminary experiments validate PC's power for control of real-world (hardware) systems, especially when the system is large (See collectives.stanford.edu and [11, 12, 13, 14, 15, 16, 17, 18].)

As another example of the successes of entropy-based inference, consider the problem of predicting the probability distribution p over the joint state of a huge number of interacting particles. This is the problem addressed by statistical physics. As first realized by Jaynes, such prediction is an exercise in probabilistic inference of exactly the sort Maxent can be applied to. Accordingly statistical physics can be addressed — and in fact derived in full — using Maxent [19, 1]. In light of all the tests physicists have done of the predictions of

statistical physics, this means that there are (at least) tens of thousands of experimental confirmations of that principle in domains with a very large number of interacting particles.

We can similarly use Shannon’s entropy to do inference of the (the distribution governing the) joint mixed strategy $q(x) = \prod_{i=1}^N q_i(x_i)$ in any game involving N players with pure strategies $\{x_i\}$. Probabilistic inference applied to game theory this way is known as **Predictive Game Theory** (PGT). In PGT, the whole point is to apply probability theory in general and Bayesian analysis in particular *to* games and their outcomes. This contrasts with their use in conventional game theory *within* the structure of individual games (e.g., in correlated equilibria [20]).

1.1 The relation between PGT and conventional game theory: a first look

Before presenting PGT in detail, this section illustrates some of its connection to other work in game theory and statistical physics.

Statistical physics provides a unifying mathematical formalism for the physics of many-particle systems. Any question related to that type of physics can be analyzed, in principle at least, simply by casting it in terms of that formalism, and then performing the associated calculations. There is no need to introduce any new formalism for new questions, new Hamiltonians, etc.

PGT arises from information theory similarly to how statistical physics arises from information theory, only in a different context. Accordingly, PGT can play an unifying role for games analogous to the one statistical physics plays for many-particle systems. All questions related to games can be analyzed, in principle at least, simply by casting them in terms of PGT, and then performing the associated calculations. There is no need to create new formalisms for new game theory issues, new presumptions about the way humans behave in games, etc; one simply casts them in terms of PGT.

Now in PGT the idea of an “equilibrium concept”, so central to conventional game theory, does not directly arise. Let $q(x)$ indicate a joint mixed strategy over joint move x . Generically in PGT, the support of the probability density function over q ’s, $P(q)$, has non-zero measure. In this, one does not allow only a single “equilibrium” q , or even a countable set of q ’s comprising the “equilibrium” joint strategies; the number of allowed q ’s is uncountable.

The fact that the support of P has non-zero measure typically ensures a built-in “bounded rationality” to PGT. This is because typically there will be q that are allowed (i.e., that have non-zero probability) in which one or more of the players is not fully rational.² This aspect of the measure of P has other consequences as well. For example, it means that rather than consider the

²There is other, more traditional work in game theory that shares PGT’s characteristic of allowing so many q with non-zero probability. For example, in Random Utility Models and decision analysis [21, 22] and similar stochastic preference models [23, 24], it may be that only one joint distribution q is possible for some single instantiation of the utilities of the players. However averaged over such utilities, many more joint distributions are allowed.

values of economic quantities of interest at a single (or at most countable set of) equilibrium q , as is conventionally done, one should consider the expected value of such quantities under $P(q)$. This means that attributes of those quantities like how nonlinear they are (which is crucial to approximating the integrals giving their expectation values) have consequences when PGT is used to analyze economics issues, consequences that they do not have when conventional game theory is used.

Are there quantities in PGT that are analogous to conventional game theory’s “equilibrium concept”, even in scenarios where P ’s support has non-zero measure? In other words, are there meaningful ways to identify a small set of special q for a distribution $P(q)$ whose support has non-zero measure? One possible such PGT-based “equilibrium concept” is the Bayes-optimal q associated with $P(q)$. Note though that in general the Bayes-optimal q depends not only on $P(q)$, but also on the loss function of the external scientist making a prediction about the physical system at hand (i.e., about the game). So consider two scenarios. Both concern the exact same game, with the exact same knowledge concerning the game. Therefore we have the exact same distributions over joint mixed strategies. However have the loss function of the external scientist (reflecting how their prediction will be used) differ between the two scenarios. Then the Bayes-optimal prediction will also differ between the two scenarios. So the very choice of “equilibrium concept” is determined (in part) by the external scientist analyzing the game; the “equilibrium” joint mixed strategy is not purely a function of the game itself, but rather also involves the external scientist making predictions about the system.

This dependence on the external scientist exhibited by this choice for PGT’s (analogue of the conventional) notion of a game’s “equilibrium” is not a philosophical preference. The dependence is not something that we have discretion to adopt or not. Rather it is intrinsic to our analyzing games with human players the same we analyze other physical systems in the Bayesian paradigm: by deriving distributions over truths based on partial information, and then (if needed) making single predictions based on that distribution together with an external loss function. Under this interpretation of equilibrium concept, we *have no choice* but to accept the dependence of point predictions on the external scientist making the prediction.

Another possible interpretation of the “equilibrium” of a game is as the posterior

$$\begin{aligned}
 P(x \mid \mathcal{I}) &= \int dq P(x \mid q, \mathcal{I}) P(q \mid \mathcal{I}) \\
 &= \int dq P(x \mid q) P(q \mid \mathcal{I}) \\
 &= \int dq q(x) P(q \mid \mathcal{I}).
 \end{aligned} \tag{1}$$

The Bayes-optimal q reflects two kinds of ignorance. The first is that of us, the external scientists, concerning the game and its players (i.e., the uncertainty

encapsulated in each $P(q)$). The second is the intrinsic noise/randomness in how the players choose their moves (i.e., the uncertainty encapsulated in each q). This is also true for $P(x | \mathcal{S})$. Unlike the Bayes-optimal q though, $P(x | \mathcal{S})$ does not depend on the loss function of the external scientist.

On the other hand, in general $P(x | \mathcal{S})$ will not be a product distribution, i.e., it will not have the moves of the players be independent. This is true even though $P(q | \mathcal{S})$ is restricted to such distributions (a linear combination of product distributions typically is not a product distribution). In addition, say that $P(q | \mathcal{S})$ is restricted to Nash equilibria q . Typically, if there are more than one such equilibria (i.e., the support of P contains more than one point), then under $P(x | \mathcal{S})$ none of the players is playing an optimal response to the mixed strategy over the other players. In other words, even though we might know that all the players are in fact perfectly rational, *our prediction* of their moves has “cross-talk” among the multiple equilibria and does not have perfect rationality.

Example: To illustrate the foregoing, consider a two player game in which both players have two possible moves, L and R. Indicate any (product distribution) q by two numbers, $q_1(x_1 = L)$ and $q_2(x_2 = L)$. Suppose that

$$P(q | \mathcal{S}) = \frac{\delta(q - (3/4, 3/4)) + \delta(q - (1/4, 1/4))}{2}. \quad (2)$$

Suppose also that we have quadratic loss. For that loss function, as is easy to verify, the Bayes-optimal q is the average q , $\int dq q(x)P(q | \mathcal{S})$. Viewed as a function of x , that particular Bayes-optimal q is the same as $P(x | \mathcal{S})$. Here they equal the distribution $P(L, L) = P(R, R) = 5/16$, $P(R, L) = P(L, R) = 3/16$. Indicate that distribution as p . p is not a product distribution, so $P(p | \mathcal{S}) = 0$. In other words, $P(x | \mathcal{S})$, this game’s “equilibrium”, is a joint mixed strategy that cannot arise.

It is worth noting that a similar phenomenon occurs in simple single-dimensional decision theory. Under quadratic loss, if $P(z)$ is the actual distribution of a random variable, the Bayes-optimal prediction — the prediction that minimizes expected loss under that P — is $y = E_P(z)$. That expectation may even be a point where there is zero probability mass, i.e., it may be that $P(y) = 0$.

Typically for any $P(q | \mathcal{S})$ there is only one (perhaps difficult to evaluate) Bayes-optimal prediction. For example, for quadratic loss functions, that prediction is the (unique) posterior mean, $\int dq qP(q | \mathcal{S})$. More strongly, $\int dq q(x)P(x | \mathcal{S})$ is *always* unique. So under either of this interpretations of “equilibrium concept”, typical games have a unique equilibrium. In this, whichever of these PGT-based interpretations of equilibrium we adopt, all work in conventional game theory that attempts to “fix” the possible multiplicity of conventional concepts of equilibrium (e.g., the many proposed refinements of the Nash equilibrium concept) is rendered moot. The same fate obtains for the different equilibrium concepts that have been proposed in cooperative game theory.

As a practical matter, often calculating the exact Bayes-optimal q can be quite difficult. As a substitute, even if it is not Bayes-optimal, we can calculate the MAP q . When $P(q | \mathcal{S})$ is peaked the MAP q should be a good approximation to the Bayes-optimal q . Indeed, it is common in Bayesian analysis to approximate the Bayes-optimal prediction by expanding the posterior as a Gaussian centered on the MAP prediction.

This MAP q is the minimizer of a Lagrangian functional $\mathcal{L}(q)$. In general this MAP q is a bounded rational equilibrium rather than a Nash equilibrium. As shown below, this MAP bounded rational equilibrium can often be approximated by simultaneously having each player i 's mixed strategy $q_i(x_i)$ be a **Boltzmann distribution** over the values of its expected utility for each of its possible moves:

$$q_i(x_i) \propto e^{\beta_i E_q(u^i | x_i)} \quad \forall i \quad (3)$$

where the joint distribution $q(x) = \prod_i q_i(x_i)$ and $u^i(x)$ is player i 's utility function.

In general there may be more than one solution to the set of coupled equations Eq. 3. (See [25] for examples of closed-form solutions to this set of coupled equations.) In conventional game theory, the set of all such solutions is sometimes called the (logit response) Quantal Response Equilibrium (QRE) [26, 27, 28]. It has been used as a convenient way to encapsulate bounded rationality. Typically approximating the MAP mixed strategy with the QRE should incur less and less error the more players there are in the game. However as discussed below, for small games the QRE may be a poor approximation to the MAP (which itself is an approximation to the Bayes-optimal prediction). Below the correction terms of the QRE (as an approximator of the MAP distribution) are calculated.

Another relation between the QRE and PGT, one that doesn't involve approximations, starts with the fact that at Nash equilibrium each player i sets its strategy q_i to maximize its expected utility $E_{q_i, q_{-i}}(u^i)$ for fixed q_{-i} .³ Consider instead having each player i set q_i to optimize an associated functional, the "maxent Lagrangian":

$$\mathcal{L}_i(q_i) \triangleq E_{q_i, q_{-i}}(u^i) - T_i S(q_i, q_{-i}). \quad (4)$$

For all $T_i \rightarrow 0$ the equilibrium q that simultaneously minimizes $\mathcal{L}_i \forall i$ is a Nash equilibrium [29, 27, 30, 31, 32, 33]. For $T_i > 0$ one gets bounded rationality. Indeed, under the identity $T_i \triangleq \beta_i^{-1} \forall i$ the solution to this modified Nash equilibrium concept turns out to be the QRE.

As discussed in [29], the maxent Lagrangian also arises in statistical physics, where it is called (a mean field approximation to) the "free energy". This formal connection between PGT/QRE and statistical physics can be exploited in several ways. As an example, consider the case where one's prior information consists of the expected energy of a set of interacting particles with joint state r ,

³Throughout this paper the minus sign before a symbol specifying a particular player indicates the set of all of the other players, and similarly for a minus sign before a set of player symbols.

a scenario known as the “canonical ensemble” in statistical physics (CE). In this situation the MAP estimate of the density function $p(r)$ using an entropic prior is the minimizer of $\mathcal{L}(p) = E_p(H) - TS(p)$, where H is the energy of the system of particles. In light of the formula for the maxent Lagrangian, this suggests that bounded rational players in a game can be made formally identical to the particles in the CE. Under this identification, the moves of the players play the roles of the states of the particles, and particle energies are translated into player utilities. Particles are distributed according to a Boltzmann distribution over their energies, and mixed strategies are Boltzmann distributions over expected payoffs.⁴

This connection between PGT and statistical physics raises the potential of transferring some of the powerful mathematical techniques that have been developed in the statistical physics community into game theory. As an example, in the “Grand Canonical Ensemble” (GCE) the number of particles of various types is variable rather than being pre-fixed. One’s prior information is then extended to include the expected numbers of particles of those types. This corresponds to having a variable number of players of various types in a bounded rational game. This suggests how to extend game theory to accommodate games with statistically varying numbers of players.⁵ Among other applications, this provides us with a new framework for analyzing games in evolutionary scenarios, different from evolutionary game theory. (A different type of “GCE game” is analyzed below.)

There are many other aspects of statistical physics that might carry over to PGT. For example, even in the CE, often there are regimes where as some parameter of the system is changed an infinitesimal amount, the character of the system changes drastically. These are known as “phase transitions”. The connection between the math of PGT and that of statistical physics suggests that similar phenomena may arise in games with human players.

PGT has many other advantages in addition to providing a way to exploit techniques from statistical physics in the context of noncooperative games. For example, as illustrated below it provides a natural way to quantify the rationality of experimentally observed behavior of human subjects. One can then, for example, empirically observe the dynamic relationship coupling the rationalities of real players as they play a sequence of games with one another. (Since such correlations are inherently a property of distributions across mixed strategies, they are not readily analyzed using conventional non-distribution-based game theory.)

Another strength of PGT arises if we change the coordinates of the underlying space of joint pure strategies $\{x\}$. After such a change, our mathematics describes a type of bounded rational cooperative game theory in which the moves

⁴Note that having the probability density over mixed strategies follow a Boltzmann distribution does not mean that functionals of that density are Boltzmann-distributed. In particular, the distribution over values of the utility function need not be Boltzmann-distributed.

⁵There are some special cases in conventional game theory that consider, in a full-rationality context, allowing the number of players to vary. See for example [34, 35], which address varying numbers of players in auctions and in bargaining, respectively.

of the players become binding contracts they all offer one another[36, 37]. In this sense, PGT provides a novel relation between cooperative and noncooperative game theory.

1.2 Roadmap

The purpose of this paper, like that of the original work on game theory, is to elucidate a framework for analyzing the reasonably imputed consequences about the behavior of the players when all one knows is the game structure. If possible, this framework should be able to accommodate extra knowledge concerning the game and/or the players if it is available. Loosely speaking, the goal is to provide for game theory the analog of what the canonical and grand canonical ensembles provide for statistical physics: a first-principles mathematical scaffolding into which one inserts one's knowledge concerning the system one is analyzing, to make predictions concerning that system. (See the future work section below for further discussion of this point.)

To do this, the next section starts by cursorily reviewing noncooperative game theory, Bayesian analysis and the entropic prior arising in information theory. In an appendix that prior is illustrating by showing how it can be used to derive statistical physics. In the following section foundational issues of PGT and associated mathematical tools are presented.

The next two sections form the core of the player. The first of them applies the entropic prior to infer mixed strategies of coupled players in a game γ . This application can be viewed as a prescription for how to infer the mixing strategy $P(q | \mathcal{S})$ adopted by a Nature involved in a meta-game with a scientist, where the moves q of Nature are mixed strategies in γ . This section then relates this coupled-players analysis to the QRE. The section after this considers independent players, leveraging the analysis for coupled players.

The following section illustrates some of the breadth of PGT. It is shown there how bounded rationality arises formally as a cost of computation for the independent players scenario. We then present rationality functions. These are a model-independent way to quantify the (bounded) rationality of the mixed strategies followed by real-world players. This section ends by showing how to apply PGT to games with stochastically varying numbers of players.

An appendix discusses the relation between PGT and previous work, and more generally the history of attempts to apply information theory within game theory.

2 Preliminaries

This section first reviews noncooperative game theory. It then reviews information theory and the associated Bayesian analysis. It ends by illustrating that analysis with a review of how it can be used to derive statistical physics. It is recommended that those already familiar with these concepts still read the middle subsection on Bayesian analysis.

2.1 Review of noncooperative game theory

In conventional noncooperative normal form game [38, 35, 39, 40, 41] theory one has a set of N independent **players**, indicated by the natural numbers $\{1, 2, \dots, N\}$. Each player i has its own finite set of allowed **pure strategies**, each such pure strategy written as $x_i \in X_i$. We indicate the the size of that space of possible pure strategies by player i as $|X_i|$. The set of all possible joint strategies is $X \triangleq X_1 \times X_2 \times \dots \times X_N$ with cardinality $|X| \triangleq \prod_{i=1}^N |X_i|$, a generic element of X being written as x .

A **mixed strategy** is a distribution $q_i(x_i)$ over player i 's possible pure strategies, $\{x_i\}$. In other words, it is a vector on the $|X_i|$ -dimensional unit simplex, Δ_{X_i} . Each player i also has a **utility function** (sometimes called a “payoff function”) u^i that maps the joint pure strategy of all N of the players into a real number.

As a point of notation, we will use curly braces to indicate an entire set, e.g., $\{\beta_i\}$ is the set of all values of β_i for all i . We will also write $\Delta_{\mathcal{X}}$ to refer to the Cartesian product of the simplices Δ_{X_i} , so that mixed joint strategies (i.e., product distributions) are elements of $\Delta_{\mathcal{X}}$. We will sometimes refer to u^i as player i 's “payoff function”, and to player i 's pure strategy x_i as its “move”. x is the joint move of all N players. As mentioned above, we will use the subscript $-i$ to indicate all moves / distributions / utility functions, etc., other than i 's. We will use the integral symbol with the measure implicit, so that it can refer to sums, Lebesgue integrals, etc., as appropriate. In particular, given mixed strategies of all the other players, we will write the expected utility of player i as $E(u^i) = \int dx \prod_j q_j(x_j) u^i(x)$. As a final point of notation, we will write \vec{a} to mean a finite indexed set all of whose components are either real numbers are infinite (greater than any real number). We will then write $\vec{a} \succeq \vec{b}$ to indicate the generalized inequality that $\forall i$, either a_i and b_i are real numbers and $a_i \geq b_i$, both a_i and b_i are infinite, or b_i is a real number and a_i is infinite. Also, in the interests or expository succinctness, we will be somewhat sloppy in differentiating between probability distributions, probability density functions, etc.; generically, “ $P(\dots)$ ” will be one or the other as appropriate.

Much of noncooperative game theory is concerned with **equilibrium concepts** specifying what joint-strategy one should expect to result from a particular game. In particular, in a **Nash equilibrium** every player adopts the mixed strategy that maximizes its expected utility, given the mixed strategies of the other players. More formally, $\forall i, q_i = \operatorname{argmax}_{q'_i} \int dx q'_i \prod_{j \neq i} q_j(x_j) u^i(x)$ [42, 38, 39].

One problem with the Nash equilibrium concept is its assumption of **full rationality**. This is the assumption that every player i can both calculate what the strategies $q_{j \neq i}$ will be and then calculate its associated optimal distribution.⁶ This requires in particular that each player calculate the entire joint distribution $q(x) = \prod_j q_j(x_j)$. If for no other reasons than computational limitations of real

⁶Here we use the term “bounded rationality” in the broad sense, to indicate any mixed strategy that does not maximize expected utility, regardless of how it arises.

humans, this assumption is essentially untenable. This problem is just as severe if one allows statistical coupling among the players [20, 38].

For simplicity, throughout each analysis presented in this paper we will treat N , the pure strategy spaces, the associated utility functions, and the statistical independence of the pure strategies chosen by the players, as fixed parts of the problem definition rather than random variables. Further we will impose no *a priori* restrictions about whether the players have encountered one another before, what information they have about one another and the game they’re playing, whether they have engaged in the game before, what their information sets are, whether there are any social norms at work on them, etc. We do not even require, *a priori*, that the players be prone to human psychological idiosyncracies. To incorporate any information of this sort into the analysis would mean modifying the priors and/or likelihoods considered below in a (mostly) straightforward way, but is beyond the scope of this paper.

2.2 Review of Bayesian analysis and decision theory

Consider any scenario in which we must reason about attributes of a physical system without knowing in full all salient aspects of that system. This is the basic problem of inductive inference. How should we do this reasoning? Many different desiderata, arising from work by De Finetti, Cox, Zellner and many others, lead to the same conclusion: if our goal is to assign real-valued numbers to the different hypotheses concerning the system at hand, we should use the rules of Probability theory[43, 44, 45, 4, 46, 47, 1, 2, 3]. In particular, this implies that we should use Bayes’ theorem to calculate what we want to know from what we are told/assume/observe/know:

$$P(\text{truth } z \mid \text{data } d) \propto P(d \mid z)P(z) \tag{5}$$

where the proportionality constant is set by the requirement that $P(\text{truth } z \mid \text{data } d)$ be normalized, and “data” means everything we are told/assume/observe/know concerning the system. $P(\text{truth } z \mid \text{data } d)$ is called the **posterior** probability, $P(\text{data } d \mid \text{truth } z)$ is called the **likelihood**, and $P(z)$ is called the **prior**.

Say that rather than a full posterior distribution, for some reason we must predict a single one of the candidate hypotheses z . According to Savage’s axioms, to do this we must be provided with a **loss function** $L(y, z)$ that maps any pair of a truth z and a prediction y to a real-valued loss (see [43] and various chapters in [35]). Then the associated **Bayes optimal** prediction is $\text{argmin}_y E_P(L(y, z))$ where the expectation is over the posterior distribution $P(\text{truth } z \mid \text{data } d)$.⁷

⁷There is controversy about the precise details of Savage’s axioms and their implications, the precise way priors should be chosen, and even the precise physical meaning of “probability”[7]. Such details are not important for current purposes. Other choices can be made, based on other desiderata. However typically the broad outlines of any approach based on such alternatives is the same: to do inference one constructs a probability distribution over possible truths and then, if needed, distills that distribution into a single prediction.

Note that the loss function is determined by the scientist external to the system who is making the prediction; it is not specified in the definition of the system under consideration.

According to the foregoing, to do statistical inference for a particular physical scenario our first task is to translate the “particular physical scenario” we’re considering into a mathematical formulation of possible truths z , data d , etc. Having done that, we can employ mathematical tools like Bayes’ theorem, approximation techniques for finding Bayes-optimal predictions, etc. to analyze our mathematical formulation. After doing this we use our translation to convert all this back into the physical scenario. This translational machinery is how we couple the abstract mathematical structure of probability theory to our particular physical inference problem.

To assist us in making this translation, we imagine an infinite set of instances of our physical scenario. All of those instances share the physical characteristics of our scenario that fix our statistical inference problem. *Every* other physical characteristic is allowed to vary across those instances. In this way the set of all of those instances define our statistical inference problem [7]. Formally, we define the **invariant** of our inference problem as the set of exactly those characteristics of the physical scenario, and no others, that would necessarily be the same if we were presented with a novel instance of the exact same inference problem. Equivalently, we can define the invariant as the set of all the physical instances consistent with those characteristics, and no instance that is inconsistent with those characteristics. By explicitly delineating an inference problem’s invariant, we can mathematically formalize that problem.

As discussed in the introduction, this Bayesian perspective is inherent in much of conventional game theory. Most obviously, all the work on Bayesian games, correlated equilibria, etc. adopts elements of the Bayesian perspective. In addition, as described above, we can define a meta-game Γ of a Scientist S playing against Nature N , in which the possible states of N are the possible games γ we wish to analyze and the associated mixed strategies q of that game. The move of S in Γ is interpreted as a prediction of N ’s move, i.e., of the mixed strategy of γ . The crucial issue in this meta-game is how, based on knowing that the game is γ , the external scientist should infer N ’s mixed strategy, $P(q)$. This is a problem of how to infer a distribution over distributions.

This inference, done using invariants, is the topic of PGT. In this paper we will only explore such use of invariants in conjunction with the entropic prior, since that prior is directly concerned with inferring distributions over distributions. However other priors also merit investigation.

2.3 Review of the entropic prior

Shannon was the first person to realize that based on any of several separate sets of very simple desiderata, there is a unique real-valued quantification of the amount of syntactic information in a distribution $P(y)$. He showed that this amount of information is (the negative of) the Shannon entropy of that

distribution, $S(P) = - \int dy P(y) \ln[\frac{P(y)}{\mu(y)}]$.⁸ Note that for a product distribution $P(y) = \prod_i P_i(y_i)$, entropy is additive: $S(P) = \sum_i S(P_i)$.

So for example, the distribution with minimal information is the one that doesn't distinguish at all between the various y , i.e., the uniform distribution. Conversely, the most informative distribution is the one that specifies a single possible y .

Say that the possible values of the underlying variable y in some particular probabilistic inference problem have no known *a priori* stochastic relationship with one another. For example, y may not be numeric, but rather consist of the three symbolic values, {red, dog, Republican}. Then simple desiderata-based counting arguments can be used to conclude the prior probability of any distribution $p(y)$ is proportional to the **entropic prior**, $\exp(-\alpha S(p))$, for some associated non-negative constant α .⁹

Intuitively, absent any other information concerning a particular distribution p , the larger its entropy the more *a priori* likely it is.¹⁰

If the possible y have a more overt mathematical relationship with one another, the situation is often not so clear-cut. For example, symmetry group arguments are often invoked in such situations, and can give more refined predictions. Despite this, for the most important scenarios it considers, scenarios where it has had such great successes, statistical physics simply uses the entropic prior, as described below. In accord with this, in this paper attention will be restricted to the entropic prior.

Say we have some information \mathcal{I} concerning p . Then by Bayes' theorem, the posterior probability of distribution p is

$$P(p | \mathcal{I}) \propto \exp(-\alpha S(p)) P(\mathcal{I} | p). \quad (6)$$

The associated MAP prediction of p based on \mathcal{I} is $\operatorname{argmax}_p P(p | \mathcal{I})$.

Intuitively, Eq. 6 pushes us to be conservative in our inference. Of all hypotheses p equally consistent (probabilistically) with our provided information, we are led to prefer those that contain minimal extra information beyond that which is contained in the provided information. This is a formalization of Occam's razor.

Physically, \mathcal{I} is all characteristics of the system that would not change if we were presented with a novel instance of the exact same inference problem.

⁸ μ is an *a priori* measure over y , often interpreted as a prior probability distribution. Unless explicitly stated otherwise, here we will always assume it is uniform, and not write it explicitly. See [19, 1, 8].

⁹The issue of how to choose α — or better yet how to integrate over it — is quite subtle, with a long history. See in particular work on ML-II [44] and the “evidence procedure” [10].

¹⁰Note that this is different from saying that the larger s is, the more *a priori* likely it is that the entropy of p is larger:

$$\begin{aligned} P_S(s) &= \int dp \delta(S(p) - s) P(p) \\ &= \frac{\int dp \delta(S(p) - s) \exp(-\alpha S(p))}{\int dp \exp(-\alpha S(p))}. \end{aligned}$$

From a frequentist perspective, it is an invariant across a set of experiments: \mathcal{S} delineates what characteristics of the system are fixed in those experiments, while all characteristics not in \mathcal{S} are allowed to vary. In essence, \mathcal{S} is the invariant that defines the inference problem.

In particular, \mathcal{S} includes any functions of p , $F(p)$, such that we know (by the specification of the precise inference problem at hand) that $F(p)$ would not change if we confronted a novel instance of the same inference problem. In general we may not know the actual value of $F(p)$ that is shared among the instances specified by \mathcal{S} ; we may only know that that value is the same in all of those instances.¹¹

Note that \mathcal{S} cannot specify p , the precise state of the system — there must be some salient characteristics of the system that are not fixed by \mathcal{S} . If this were not the case the likelihood $P(\mathcal{S} | p)$ would be a delta function, and therefore the prior would be irrelevant. In such a case, statistical inference would reduce to the truism “whatever happens happens”. Accordingly, we never have \mathcal{S} contain a set of functions $\{F_i(p)\}$ whose values jointly fix p exactly.

An important example of the foregoing occurs in statistical physics, where \mathcal{S} is the observed temperature T of a physical system. T is taken to fix a function $F(p)$, namely the expected energy $H(x)$ of the system under distribution $p(x)$: $F(p) = \int dx H(x)p(x)$ is fixed by \mathcal{S} . It is the application of the entropic prior to this situation that results in the canonical ensemble mentioned above. The number of experiments validating this application is extraordinary; including experiments in high school labs, it is probably on the order of 10^8 (at least). In this paper that application serves as a touchstone for how to translate \mathcal{S} into a distribution over distributions, and therefore as the primary analogy mentioned in the derivation of PGT. However in the interests of expediting the flow of this paper, that application is relegated to an appendix.

3 Predictive Game Theory - general considerations

3.1 The two types of game theory

Say we are presented with a noncooperative normal form game for N players other than us, and a set of N subjects who will fill the roles of the players in a fixed manner. We wish to make predictions concerning the outcome of that game when played by that set of players other than us. As discussed above, for us to obey Cox’s axioms, when making those predictions we must use probability theory. If we wish to distill the probability distribution over outcomes into a

¹¹Relating this back to the mathematics of probability theory, in such a case that value of $F(p)$ is known as a **hyperparameter**. Formally, hyperparameters have their own priors. To get a final posterior over what we wish to infer — p — we must marginalize over possible values of all hyperparameters. Implicitly, the reason that here we simply choose one value of a hyperparameter and discard all others is that we expect the posterior distribution of the hyperparameter to be highly peaked, so that we do not need to carry out such marginalization. See the discussion of ML-II in [44, 43, 10], and also [10].

single prediction, then to obey Savage’s axioms we must have a loss function and use decision theory.

Note that these normative axioms differ fundamentally from those that can be used to derive various equilibrium concepts. The axioms underlying equilibrium concepts concern external physical reality, namely the players of the game. They concern something outside of our control. In other words, such axioms are hypothesized physical laws. Like any other such laws, they can be contradicted or affirmed by physical experiment, at least in theory. In fact, behavioral game theory [48] essentially does just that, experimentally determining what such hypothesized physical laws are valid. (See also all the work on behavior economics [49, 50].)

In contrast, Cox’s and Savage’s axioms are normative. They tell us how best to make predictions about the physical world. They make no falsifiable claims about the real world; it is under our control whether they will be followed, not the external world’s control. Violating them in our analysis is akin to performing an analysis in which we violate the axioms defining the integers.

Before showing some ways to arrive at a distribution over outcomes for this situation, we must clarify what the space of “outcomes” of the game is. There are two broad types of such spaces to consider, with associated types of game theory.

In type I game theory, what a player chooses in any particular instance of the game is its move in that instance. (This is the analog of the variable y in App.9.) In general, in the real world a particular human’s choice of move will vary depending on their mood, how distracted they are, etc. Physically, this variability arises from variability in the dynamics of neurotransmitter levels in the synapses in their brain during their decision-making, associated dynamical variability in the firing potentials of their neurons during that process, etc.

Due to this variability the choice of each player is governed by a probability distribution. (This is the analog of the variable q in App.9). So the joint choice of the players is also described by a distribution. We write that joint distribution as

$$\begin{aligned}
 q(x) &= P(x | q) \\
 &= P(x_N | q) \prod_{i=1}^{N-1} P(x_i | q, x_{i+1}, x_{i+2}, \dots) \\
 &= P(x_N | q) \prod_{i=1}^{N-1} P(x_i | q) \\
 &= \prod_{i=1}^N P(x_i | q) \\
 &= \prod_{i=1}^N q_i(x_i), \tag{7}
 \end{aligned}$$

where the third equality follows from the statistical independence of the players’

choices. So the probability of a particular joint move x is given by a product distribution $q(x) \triangleq \prod_i q_i(x_i)$.¹² If the interaction between the humans is not, physically, a conventional normal form noncooperative game, then the space of allowed q must be expanded to allow q that statistically couple the moves. Such extensions are beyond the scope of this paper.

q incorporates the subconscious biases of the players, the day-to-day distribution of their moods, and more generally the full physical stochastic nature of their separate decision-making algorithms. It also reflects what the players know about each other, whether they have directly interacted before, what they know about the game structure, their information sets, etc. In short, q is the physical nature of the game setup, *in toto*. Note that we cannot examine the precise states of all the neurons and neurotransmitters in the brains of the players, and even if we could we cannot precisely evaluate the associated stochastic dynamics. Accordingly, while it is physically real, in practice it is impossible for us to ascertain a distribution q_i exactly.

In contrast, x is the quantity that the players consciously determine, and it is observable. q is instead the physical process that specifies how the players select that observable. Both of these quantities differ from our (limited) knowledge about q . That knowledge is embodied in probability distributions $P(q)$. P reflects *us* as much as the players.

We imagine an infinite set of instances of our setup, i.e., an infinite set of q 's. The invariant \mathcal{S} specifies all characteristics of the system — and only those aspects — such that if they had been different in some particular instance, we would know it. In this paper those aspects include the players' utility functions, their identities, etc. In going from one instance to the next, we assume the problem is “reset”, i.e., there is no information conveyed from one instance to the next. (In particular the players' minds are “wiped clean” between instances.) Our inference problem is to predict q based on such an invariant, i.e., formulate the posterior $P(q | \mathcal{S})$.¹³

In type II game theory, what player i chooses in any particular instance of a game is a mixed strategy, $q_i(x_i)$. Each player i 's mixed strategy is separately randomly sampled, “by Nature”, to get the player's move x_i . So as in type I games, $q(x) \triangleq \prod_i q_i(x_i)$.

¹²Loosely speaking, when used as an approximation in statistical physics, such product distributions are called “mean field theory”. See [51].

¹³In the real world we don't know a player i 's (risk-neutral) utility function, but must infer it based on what actually is known/observed about player i 's relative preferences, d_i . Formally, this means $\{d_i\} \in \mathcal{S}$, not u^i , and we must integrate over the “hidden variable” u^i to get the quantity we are ultimately interested in:

$$\begin{aligned} P(q | \mathcal{S}) &= P(q | \{d_i\}, \dots) \\ &= \int d\{u^i\} P(q | \{d_i\}, \{u^i\}, \dots) P(\{u^i\} | \{d_i\}, \dots) \\ &= \int d\{u^i\} P(q | \{u^i\}, \dots) P(\{u^i\} | \{d_i\}, \dots) \end{aligned}$$

where “...” indicates the contents of \mathcal{S} other than $\{d_i\}$. For simplicity, in this paper we bypass this integral by assuming $d_i = u_i \forall i$, i.e., we assume \mathcal{S} includes the $\{u_i\}$ directly.

In general, in real world type II games, a particular human’s choice of mixed strategy will vary depending on their mood, how distracted they are, etc. Accordingly the joint choice of the players is described by a distribution $\pi(q)$. π reflects the stochastic nature of the players, what they know about each other, what they know about the game structure, etc. In short, π is the physical nature of the setup. In contrast, our (limited) knowledge about π is embodied in a distribution $P(\pi)$.

As usual, we imagine an infinite set of instances of our setup, i.e., an infinite set of π ’s, with no information conveyed between instances. The invariant \mathcal{S} specifies all characteristics of the system — and only those aspects — such that we would know if they had been different in some particular instance. Our inference problem is to predict π , i.e., formulate the posterior $P(\pi | \mathcal{S})$. So an “entropic prior” concerns the probability of any such joint distribution π , our likelihood must concern π , etc. Note that the invariant setting the likelihood, $P(\mathcal{S} | \pi)$, is a characteristic of an entire distribution over joint mixed strategies (namely π), not (directly) of the joint mixed strategies themselves.

In both game types, since q is a product distribution, *if one is given q* then knowing one player’s move provides no extra information about another player’s move. (Formally, $P(x_i | q, x_j) = P(x_i | q)$.) However in the absence of knowing q in full, knowing just the q_i of some player i may provide information about other the q_j of other players j . For example, this could be the case if those mixed strategies q_i and q_j are determined in part based on previous interactions between the players. Similarly, even if the players have never previously interacted, if there is overlap in what they each know about the game (e.g., they each know the utility functions of all players), that might couple members of the set $\{q_i\}$. Accordingly, in type I game theory $P(q | \mathcal{S})$ need not be a product distribution (over the q_i) in general, and in type II game theory π need not be a product distribution.

In general, which game type one uses to cast the problem is set by the problem at hand. If the players all consciously choose mixed strategies — if that’s how their thought processes work — then we have a type II game. If the players choose moves, we have a type I game. One can even have mixed game types, in which some players choose moves, and some choose mixed strategies. Our lacking knowledge of what scenario we face is analogous to lack of knowledge concerning the payoff structure: our inference problem is not fully specified.¹⁴

For the reasons elaborated above, we will adopt the entropic prior for both game types. Note that for either game type, the entropic prior evaluated for a product distribution is itself a product, i.e., if $q(x) = \prod_i q_i(x_i)$, then $e^{\alpha S(q)} =$

¹⁴Of course, the lack of knowledge underlying both game types can in principle be addressed by setting a prior probability distribution over the underlying unknown and defining an associated likelihood function. Here that would mean distributions over whether each player chooses moves or chooses mixed strategies. No such analysis which would essentially mix the two game types is considered in this paper.

$\prod_i e^{\alpha S(q_i)}$. As a result, by symmetry the associated marginal over x ,

$$\int dx q(x)P(x) \propto \int dx \prod_i q_i(x_i) e^{\alpha S(q_i)}, \quad (8)$$

must be uniform over x .

In some situations $P(q | \mathcal{S})$ will not be of interest, but rather the associated posterior

$$P(x | \mathcal{S}) = \int dq P(q | \mathcal{S}) \prod_i q_i(x_i) \quad (9)$$

will be. Now for the entropic prior, we know what the associated prior $P(x)$ is (it's uniform). This suggests one formulate a likelihood $P(\mathcal{S} | x)$. One could then use Bayes' theorem with the uniform $P(x)$ to arrive at the posterior $P(x | \mathcal{S})$ directly, rather than arrive at it via the intermediate variable q . This would constitute a third type of game, in addition to the other two presented above, in which instances would be x 's rather than q 's or π 's.

Unfortunately, it is hard to see how to formulate the likelihood $P(\mathcal{S} | x)$ without employing q or π . Recall that an invariant \mathcal{S} is the set of all physical instances that can occur in our inference problem, and no other instances. However in formulating $P(\mathcal{S} | x)$ instances are specified by values of x , and for almost any inference problem, all x 's may occur. So in any such inference problem, the associated $P(\mathcal{S} | x)$ does not exclude any x at all, i.e., it is vacuous, as far as inference of x is concerned. It is also hard to see what might be gained by using such an alternative game type. Indeed, since it conflates the distribution q concerning physical reality with the distribution P concerning our (lack of) knowledge about that reality, one would expect substantial losses of insight if one used such an alternative game type for one's analysis.

To distinguish it from distributions like q or π , we will refer to a distribution P that describes our inference as a **predictive distribution**.¹⁵ So for example, both $P(q | \mathcal{S})$ and $P(\pi | \mathcal{S})$ are predictive distributions. As other examples, $P(x | \mathcal{S}) = \int dq P(q | \mathcal{S})q(x)$, and $P(q | \mathcal{S}, x_j)$ are also predictive distributions. Predictive distributions reflect *our* knowledge/insight/ignorance concerning the game. This contrasts with distributions like q and π , which reflect the "physical" distributions of the players in the game.

As a final notational comment, we will use the following shorthand for each i 's "effective utility", sometimes called i 's **environment**:

$$U^i(x_i) \triangleq E(u^i | x_i). \quad (10)$$

In type I game theory this reduces to

$$\begin{aligned} U^i(x_i) &= E_{q_{-i}}(u^i | x_i) = \int dx_{-i} q_{-i}(x_{-i} | x_i) u^i(x_i, x_{-i}) \\ &= \int dx_{-i} u^i(x) \prod_{j \neq i} q_j(x_j) \quad (\text{since } q \text{ is a product distribution}) \end{aligned} \quad (11)$$

¹⁵This use of the term "predictive distribution" should not be confused with the one arising in Bayesian statistics. We will not have need to use the latter term in this paper.

We will write $E_{q_{-i}}(u^i | x_i)$ as $U_{q_{-i}}^i(x_i)$ when the q_{-i} defining U^i needs to be made explicit. We will also write

$$E_q(u^i) = q_i \cdot U^i \tag{12}$$

when working with type I games. (The expansion of U^i for type II games proceeds analogously.)

3.2 Needed Mathematical Tools

This section presents some preliminary mathematical tools from statistical physics that are useful for performing Bayesian analysis of games using entropic priors. In essence, these tools amount to a suite of relationships involving the Boltzmann distribution, entropy, and optimization. Though it is a bit laborious to work through these tools, they are crucial for understanding bounded rational players in general, and for understanding the QRE in particular. We will focus on type I games; as usual, similar considerations apply for type II games.

(1) Start by noting that if we take its logarithm, any distribution $q_i(x_i)$ can be expressed as an exponential of some function over x_i . So in particular we can write any MAP q_i that way:

$$\operatorname{argmax}_{q_i} P(q_i | \mathcal{J}) \propto e^{\beta_i f_i(x_i)} \tag{13}$$

for some appropriate function f_i and constant $\beta_i \geq 0$.¹⁶

(2) We now relate the formulation of an MAP q_i in exponential form (as in (1)) to a particular choice for our game theory problem’s invariant.

Say we associate with each player i a “guess” she makes (potentially explicitly, potentially not) for her environment function, $U_{q_{-i}}^i$. Write that guessed function as $f_i(x_i)$. We presume that we can view the player’s behavior as though she were trying to perform well for that (guessed) environment. Formally, we presume that in each instance of our inference problem, the mixed strategy of player i results in the same (invariant) value K_i for what $E(U^i)$ would be if player i ’s guess for her environment U^i were correct, i.e., if U^i equalled f_i . So the invariant of the game for player i is

$$q_i \cdot f_i = K_i. \tag{14}$$

Intuitively, with this invariant, as one goes from one instance of the inference problem to the next, we presume that player i is always just as smart, as measured with the (potentially counterfactual) environment function f_i .

For this invariant, the likelihood $P(\mathcal{J} | q_i)$ restricts q_i to lie on the hyperplane of distributions obeying Eq. 14:

$$P(\mathcal{J} | q_i) = \delta(q_i \cdot f_i - K_i). \tag{15}$$

¹⁶Of course, there is always freedom to absorb some portion of any β_i into the associated f_i , but that is irrelevant for current purposes.

So given our use of the entropic prior, the posterior for player i is

$$P(q_i | \mathcal{S}) \propto e^{\alpha S(q_i)} \delta(q_i \cdot f_i - K_i). \quad (16)$$

Accordingly the MAP q_i is given by the q maximizing the so-called **maxent Lagrangian**,

$$\mathcal{L}(q_i) \triangleq S(q_i) + \beta_i [q_i \cdot f_i - K_i] \quad (17)$$

where in the usual way the β_i are the Lagrange parameters, here divided by α .¹⁷

Solving for the q_i minimizing this Lagrangian, $q_i^{\beta_i}$, we get the distribution of Eq. 13, with β_i a function of K_i . Equivalently, we can take K_i to be a function of β_i . This is what we will do below.

Now parallel the conventionally nomenclature of statistical physics, and define the **partition function**

$$Z_{f_i}(\beta_i) \triangleq \int dx_i e^{\beta_i f_i(x_i)}. \quad (18)$$

(Note that the partition function is the normalization constant of Eq. 13.) Then using Eq. 13 to express the Boltzmann distribution $q_i^{\beta_i}$, our constraint Eq. 15 means that

$$K_i(\beta_i) = f_i \cdot q_i = \frac{d \ln(Z_{f_i}(\beta_i))}{d \beta_i} \quad (19)$$

as is readily verified by evaluating the derivative. As shorthand, sometimes we will absorb β_i into f_i , and simply write $Z(V) \triangleq Z_V(1)$.

So say we are given some distribution q_i . Take its logarithm to get a function f_i and exponent β_i . Then use Eq. 19 to translate that f_i and β_i into a value of K_i . Using these choices of K_i , f_i and β_i , formulate the associated invariant \mathcal{S} given by Eq. 14. As shown by Eq. 13, the MAP distribution for that \mathcal{S} is our starting distribution q_i . In this way we can view that \mathcal{S} as the “effective” invariant for this (arbitrary) starting q_i . We can translate *any* q_i into an MAP distribution by choosing an appropriate \mathcal{S} this way.

(3) We will refer to the function $K_i(\cdot)$ arising in Eq. 19, which maps β_i to the expected value of f_i under $q_i^{\beta_i}(\cdot)$, as the **Boltzmann utility** for player i , where f_i is implicit. We now present some general characteristics of the Boltzmann utility, characteristics that are particularly important for understanding the QRE.

First, with slight abuse of terminology, we will sometimes write the Boltzmann utility with f_i explicitly listed as the first argument and the subscript i

¹⁷As an aside, say that we replaced Eq. 14 with the inequality constraint $q_i \cdot f_i > K_i$. The entropy function is concave, and so is this inequality constraint. Accordingly, by Slater’s theorem, there is zero duality gap [52] and we can apply the KKT conditions to get a solution. In other words, for this modified invariant the maxent Lagrangian still applies, and therefore so does the solution of Eq. 13.

dropped, i.e., as $K(f, \beta \in \mathbb{R}^+)$. In this case it is the domain of the first argument of the Boltzmann utility $K(f, \beta)$ that (implicitly) sets the space X_i to be integrated over to evaluate $K(f, \beta)$.¹⁸

As is readily verified, the variance (over f_i values) of the Boltzmann distribution of Eq. 13 is given by the derivative of $K_i(\beta_i)$ with respect to β_i . Since variances are non-negative, this means that $K_i(\beta_i)$ is a non-decreasing function. In fact, for fixed f_i , so long as f_i is not a constant-valued function (i.e., not independent of its argument), the associated Boltzmann utility $K_i(\cdot)$ is a monotonically increasing bijection with domain $\beta_i \in [0, \infty)$ and associated range $[\frac{\int dx_i f_i(x_i)}{|X_i|}, \max_{x_i} f_i(x_i)]$.¹⁹

We extend the domain of definition of K_i by adding to it the special value “ ∞^* ”, and defining $K_i(\infty^*) = \max_{x_i} f_i(x_i)$. This makes K_i a bijection whose domain is $\beta_i \in [0, \infty) \cup \infty^*$ when $f_i(x_i)$ is not a constant, and is the singleton $\{\infty^*\}$ otherwise. In both cases the range is $[\frac{\int dx_i f_i(x_i)}{|X_i|}, \max_{x_i} f_i(x_i)]$.

With some abuse of notation from now on we will extend the meaning of the linear ordering “ \geq ” to have $\infty^* \geq k \forall k \in \mathbb{R}$. We will also drop the asterisk superscript from “ ∞^* ”, relying on the context to make the meaning of “ ∞ ” clear. We will engage in more abuse by writing “ \vec{b} ” even if some component $b_i = \infty$ (so that \vec{b} is not a Euclidean vector, properly speaking).

Just as expected f_i cannot decrease as the Boltzmann exponent rises, the entropy of a Boltzmann distribution $e^{\beta_i f_i(x)} / Z_{f_i}(\beta_i)$ cannot increase as its Boltzmann exponent β_i rises.²⁰ So the picture that emerges is that as β_i increases, the Boltzmann distribution gets more peaked, with lower entropy. At the same time, it also gets higher associated expected value of f_i .

(4) We now extend the discussion to allow q_i not to be $q_i^{\beta_i}$, the Boltzmann distribution over values of f_i with exponent β_i (the distribution in Eq. 13). This extension will prove important in quantifying the rationality of a player based solely on their mixed strategy and environment (Sec. 6.2 below).

First expand $S(q_i)$ for the case where in fact q_i does equal the Boltzmann distribution $q_i^{\beta_i}$. Then using Eq. 14, we see that for this Boltzmann distribution

¹⁸Note that despite the terminology, the Boltzmann utility is not a “utility function” in the sense of a mapping from x to \mathbb{R} . Rather it’s what expected utility would be for a particular type of mixed strategy, in a particular environment, as a function of parameters of that mixed strategy.

¹⁹To see this, note that the variance is non-zero for all $\beta_i < \infty$, so long as $f_i(x_i)$ is not a constant. Accordingly, under such circumstances $K_i(\beta_i)$ is invertible.

²⁰To see this say we replace the invariant $q_i \cdot f_i = K_i(\beta_i)$ with $q_i \cdot f_i \geq K_i(\beta_i)$. Then for fixed q_{-i} , the MAP q_i is the q_i that maximizes $S(q_i)$ subject to that inequality constraint that $q_i \cdot f_i \geq K_i(\beta_i)$. The entropy is a concave function of its argument, as is this inequality constraint, so our problem is concave. Therefore the critical point of the associated Lagrangian is the MAP q_i . Now if we increase β_i , and therefore increase K_i , the feasible region for our new invariant decreases. This means that when we do that the maximal feasible value of S cannot increase. So the entropy of the critical point of the Lagrangian for our new invariant cannot increase as β_i does. However that critical point is just the Boltzmann distribution $q_i(x_i) \propto \exp(-\beta_i f_i(x_i))$, i.e., it is the MAP q_i for original equality invariant, $q_i \cdot f_i = K_i$. So the property that increasing β_i cannot increase the entropy must also hold for the original equality invariant.

q_i ,

$$q_i \cdot f_i + \frac{S(q_i)}{\beta_i} = \frac{\ln(Z_{f_i}(\beta_i))}{\beta_i}. \quad (20)$$

Comparing with Eq. 73 in App. 9, we see that the quantity on the left-hand side is essentially identical to the free energy of statistical physics.²¹ Accordingly, we call it the **free utility** of the player.

Note that the free utility is a function of q_i and β_i , and is defined even when q_i is *not* a Boltzmann distribution over values of f_i . In contrast, the quantity on the right-hand side of Eq. 20 is only a function of β_i . For fixed β_i , that quantity on the right-hand side of Eq. 20 is an upper bound on the free utility of the player.²² For that fixed β_i , the **free utility gap** of player i is defined as the difference between its actual free utility and the maximum possible at that β_i , $\frac{\ln(Z_{f_i}(\beta_i))}{\beta_i}$. That gap is zero — player i 's free utility is maximized — at player i 's associated equilibrium (Boltzmann) mixed strategy. Intuitively, player i “tries to” maximize free utility rather than expected utility, insofar as it “tries to” achieve its MAP mixed strategy.

(5) Finally, we present a restriction that simplifies our discussion below of how the support of $P(q | \mathcal{S})$ covers $\Delta_{\mathcal{X}}$ (Sec. 4.7 below).

We say that a particular q is **benign** for utilities $\{u^i\}$ if for all players i , we can write the associated expected utility $q_i \cdot U_{q_{-i}}^i$ as $K(U_{q_{-i}}^i, \beta_i)$ for a $\beta_i > 0$ with $U_{q_{-i}}^i$ defined in terms of u^i and q_{-i} in the usual way. In this paper, for simplicity we will only consider benign q 's. This means in particular that we assume that for all players, their expected utility is not worse than the one they would get for a uniform mixed strategy (which corresponds to infinite β_i). While the analysis can be extended to allow negative β_i (where player i adopts a worse-than-uniform mixed strategy), there is no need for such considerations here.

3.3 Effective invariants and the QRE

As discussed above, every product distribution q can be specified by saying that each of its marginalizations q_i is an MAP prediction for some associated “guessed environment” f_i and β_i . But not every q can be expressed this way if we demand that the guessed environments f_i for each player are their actual environments. In other words, only for a subset of all q 's will it be the case that $f_i(x_i) \triangleq U_{q_{-i}}^i(x_i) \forall i$.

Demanding such self-consistency in q results in a coupled set of nonlinear equations for q . This is the set of equations that specifies the QRE, Eq. 3.²³ It

²¹Free energy has a different sign on the entropy. This just reflects the fact that players work to raise utility whereas physical systems work to minimize energy.

²²This follows from the fact that the q_i that maximizes the free utility for our β_i is just the associated Boltzmann distribution.

²³In [27], $U_{q_{-i}}^i$ is called “a statistical reaction function”, and the set of coupled equations giving that solution is called the “logit equilibrium correspondence”.

was first derived in this manner, as the self-consistent solution to a set of MAP inferences, in [11, 29, 25].

Note that there is no particular decision-theoretic significance to the QRE derived in this manner. In particular, it is not the Bayes-optimal solution to any inference problem. Nor is it derived as the MAP solution to any (single) problem. Rather it is given by a set of MAP solutions, each for a separate inference problem. There is one such problem for each separate player. We then posthoc “tie together” those separate problems, by requiring that our solutions to them are consistent with one another.

Unfortunately, such a two-stage process has no clear justification in terms of Cox’s and Savage’s axioms. More generally, it is hard to formally justify the approach of enforcing consistency among a set of separate inference problems rather than considering a single aggregate inference problem. (Recall that the inference is being done by the external scientist, and that that scientist is external to the system.) To address that single inference problem, we must use a single invariant that concerns the entire joint system.

In such an alternative to the QRE’s posthoc “tying together” one analyzes all players’ distributions simultaneously, from the very beginning. This means that we analyze the distribution over full joint strategies that involve all the players. In this approach, to get any particular player i ’s distribution we would marginalize the distribution over joint strategies, rather than (as in the QRE) start with those marginal distributions and try to tie them together.

The natural invariant for this aggregate inference problem is the “aggregate invariant” that $q_i \cdot f_i = K_i \forall i$. However as shown below, the QRE is not even the MAP of the posterior over the space of joint strategies under this invariant (never mind being Bayes-optimal.) An analysis of this aggregate problem is the subject of the next few sections. A discussion of the historical context of the QRE can be found in an appendix.

4 Coupled players

Recall that the posterior is given by the prior and the likelihood. Since (for both game types) we’ve chosen the prior, our next task is to set the (\mathcal{I} -based) likelihood. We want that likelihood to have the same form as the likelihood underpinning the CE of statistical physics: a Heaviside theta function that restricts attention to a subset of all possible systems, with the distribution across that subset then set by our prior (see the appendix). However as elaborated below, the likelihood theta function appropriate for games is more complicated than the one that arises in statistical physics. This is because there are multiple payoff functions in games, each with its own effect on the system’s theta function, whereas there is only one Hamiltonian in a statistical physics system.

In this section we illustrate how to set the likelihood in a scenario where the players may have knowingly interacted with each other before the current game. (In the next section we use these results to address the case where the players have not previously knowingly interacted.) In general those previous interac-

tions are allowed to vary from one instance to another; the invariant restricting our instances will also be what restricts the possible previous interactions.

4.1 Invariants of human players

In general of course, \mathcal{I} does not specify everything about our inference problem, and in particular it does not specify the value of that which we want to infer. Here what we wish to infer is the actual joint strategy of the players. (See Sec. 2.3.) So the joint strategy is not specified by \mathcal{I} . Therefore a player’s payoff is also not specified for any particular one of its moves, since that payoff will depend on the moves of the other players in general; that payoff may vary between instances.

Instead, here we stipulate that any player will try to maximize her expected utility, to the best of her computational abilities, the best of her insights into the other players and the game structure, etc.²⁴ Intuitively, this means we assume a “pressure” embodied in the distribution over q ’s biasing the distribution to have q_i that achieve high values of $U_q^i \cdot q_i$. This pressure is matched by counterpressures from the other players affecting U_q^i .

What we consider invariant is that from one instance to the next player i does not change, and therefore how insightful player i is into the other players (based on her previous interactions with them), how computationally powerful i is, etc., does not change. In other words, how well player i performs, *in light of her (varying) environment of possible payoffs* (i.e., in light of U^i) is the same in all instances. In short, “how smart” every player i is does not change from one instance to the next. As in the case of statistical physics though, here our invariant need not specify precisely how smart each player is *a priori*, only that how smart each of them is doesn’t vary from one instance to the next.

As an example, consider the situation where the players knowingly are repeatedly playing the game with each other, forming a sequence of games. Say we are considering the distribution over joint mixed strategies q at some fixed (invariant) sequence index t . In this scenario it is an entire sequence of games leading up to game t that constitutes “an instance of our inference problem”. We must determine what is invariant from one instance of that problem to another.

Note that in game t of any instance the players’ actual moves are independent, tautologically. (This is reflected in having q be a product distribution.) However in general the q at t will change from one instance to the next. Indeed, consider any time $t' \leq t$. At that time, in every sequence of games, each player modifies its mixed strategy based on the history of move-payoff pairs in that sequence for times previous to t' , i.e., each player tries to learn what strategy is best based on its history and adapts its strategy accordingly. Since the move-payoff pairs of that history are formed by statistical sampling (of the joint mixed strategy), they will not be the same in all sequences. Accordingly, in general the modification i makes to its mixed strategy at t' will not be the

²⁴This is not the case in situations like Allais’ paradox; see below.

same in all sequences. Therefore the final joint mixed strategy q will vary from one sequence to the next.

As a result of this sampling, across the set of all instances (i.e., all sequences) there will be some statistical coupling between the time t mixed strategies of the separate players. This means in particular that in general the time t MAP q , $\operatorname{argmax}_{q^t} P(q^t \mid \text{invariant})$, is not a product of the individual time t MAP q_i ,

$$\begin{aligned} \prod_i \operatorname{argmax}_{q_i^t} P(q_i^t \mid \text{invariant}) &= \prod_i \operatorname{argmax}_{q_i^t} \int dq_{-i}^t P(q_i^t, q_{-i}^t \mid \text{invariant}) \\ &\neq \operatorname{argmax}_q P(q^t \mid \text{invariant}) \end{aligned} \quad (21)$$

(This is in contrast to the case with independent players considered in Sec. 5).

Since the final q varies across the instances, in general we can't expect that for each player i , its environment U^i will be the same at the end of each sequence. Indeed, even consider the case where play evolves to a Nash equilibrium at time t . If the game has multiple such Nash equilibria, then in general which one holds for a particular sequence of games will depend on the history of moves and payoffs in that sequence. Accordingly, U^i will depend on that sequence.

Formalizing all this means formalizing our invariant \mathcal{S} of “how smart” a player is. Here we consider how to do this for type I game theory, where inference is of q . The discussion for type II game theory proceeds *mutatis mutandi*.

Consider just those instances of our inference problem in which player i is confronted with some particular vector of move-conditioned expected utility values, U^i . We say that that i is “as smart” in any one of those instances as another if in each of them separately, on average, the move i chooses has the same payoff. In other words, how smart i 's is is the same in all of those instances if i 's expected utility, $q_i \cdot U^i$, has the same (potentially unknown) value in all of them. We write that value as $\epsilon_i(U^i)$. As an example, at a Nash equilibrium $\epsilon_i(U^i) = \max_{x_i} U^i(x_i) \forall i$.

Note how conservative this restriction on q_i is. In particular, so long as $\epsilon_i(U^i) < \max_{x_i} U^i(x_i) \forall i$, then we are guaranteed that multiple q_i satisfy this restriction. This is true even if the game has only a single Nash equilibrium.

Our invariant is simply that the functions $\{\epsilon_i\}$ are the same in all instances. This invariant does not concern the joint choices (moves) of the players across the instances (which is given by the x 's). Rather it concerns q , which is the physical nature of the process driving the players to make those choices. However the invariant does not specify that process. In particular it does not stipulate how the players reason concerning each other. For example, it does not stipulate how many levels of analysis of the sort “I know that you know that I know that you prefer ...” any of the players go through (if any levels at all). All that \mathcal{S} stipulates is that certain high-level encapsulations of that decision-making, given by the $\{\epsilon_i\}$, are the same in all instances.

As a result of this invariance, even though the moves $\{x_i\}$ of the players are independent in any particular instance (since q is a product distribution), our (!) lack of knowledge concerning the set of all the instances might result in

a posterior $P(q \mid \mathcal{S})$ in which the distributions $\{q_i\}$ are statistically coupled. (Recall that q reflects the players, and P reflects our inference concerning them.)

Note that for the entropic prior $P(q)$ there is no statistical coupling between x_i and x_j in the prior distribution $P(x)$. (Recall that for that prior, $P(x) = \int dq P(x \mid q)P(q)$ must be uniform, by symmetry.) However the potential coupling between the $\{q_i\}$ means that in the posterior distribution, the moves are *not* statistically independent (assuming one doesn't condition on q). Formally,

$$\begin{aligned} P(x_i \mid \mathcal{S}) &= \int dq_i P(x_i \mid q_i)P(q_i \mid \mathcal{S}) \\ &= \int dq_i q_i(x_i)P(q_i \mid \mathcal{S}) \end{aligned} \tag{22}$$

so

$$\prod_i P(x_i \mid \mathcal{S}) = \int dq \prod_i P(q_i \mid \mathcal{S})q_i(x_i). \tag{23}$$

On the other hand, recall that

$$P(x \mid \mathcal{S}) = \int dq P(q \mid \mathcal{S}) \prod_i q_i(x_i). \tag{24}$$

So if $P(q \mid \mathcal{S})$ is not a product distribution, then in general $P(x \mid \mathcal{S}) \neq \prod_i P(x_i \mid \mathcal{S})$, i.e., in this situation $P(x \mid \mathcal{S})$ — which is the distribution over joint moves reflecting our understanding of the system — is not a product distribution either. In such a situation, *to us*, x_i and x_j are statistically coupled.

Such coupling also typically arises in the Bayes-optimal prediction for the distribution over joint strategies. Indeed, say we adopt a quadratic loss function, so that if we guess the joint distribution is q'' , when in fact it is q' , our loss is $(q'' - q')^2$. Then given the posterior $P(q \mid \mathcal{S})$, the associated Bayes-optimal prediction for q — the prediction that minimizes our posterior expected quadratic loss — is

$$p_{quad} \triangleq \int dq q P(q \mid \mathcal{S}). \tag{25}$$

This is the same as the joint mixed strategy given by Eq. 24. (This is not the case for other loss functions.) Accordingly, our conclusion about coupling of the $\{x_i\}$ holds for this Bayes-optimal joint mixed strategy.²⁵

Consider changing the cognitive process of some player $j \neq i$ in way that does not change ϵ_j . Also do not change anything concerning all the other players. Do all this in such a way that how that player j chooses moves at time t changes, but nothing else changes about j 's behavior. So in particular, for any fixed vector U^i , the q_j that governs player j at time t and is consistent with that

²⁵Note the slight abuse of terminology; the moves of the players are statistically coupled in this “joint mixed strategy”, which is why we do not write that Bayes-optimal distribution over x as q but a p .

U^i will change.²⁶ Now the distribution over possible q_i at time t is based on behavior of player i and of other players for times $t' \leq t$. Since those factors are unchanged by our change to q_j at time t , so is the distribution over possible q_i then. Accordingly, the change in q_j will in general change the expected utility of player i at time t . In other words, changing player $j \neq i$ will in general change ϵ_i . This illustrates that our invariance is implicitly determined by the set of players as a whole. This is in addition to its reflecting how the players have interacted, the structure of the game, etc.²⁷

4.2 Specifying the function ϵ_i

Now in general for any player i , our invariant doesn't force all instances to have the same vector U^i . So to complete the quantification of how smart a player is we need to specify the function ϵ_i . To do this we use a Gedanken experiment; we consider how player i would behave in a counterfactual "game against Nature" inference problem. In that new problem we focus on just one player i , fixing the others. Formally, our invariant is expanded from that of the original problem, to a new invariant \mathcal{S}' that also include U^i . Since the invariant still stipulates that $E(u^i) = \epsilon_i(U^i)$, having U^i also invariant means that the expected utility $E(u^i)$ does not change between instances of this new problem.

Write the (potentially unknown) value of that invariant expected utility as v_i . Since we use the entropic prior over q_i this new inference problem has the usual entropic posterior. Also as usual, the MAP q_i is given by a Boltzmann distribution:

$$q_i^*(x_i) \propto e^{b_i U^i(x_i)} \quad (26)$$

where the Lagrange parameter going into b_i enforces the constraint, namely that $q_i^* \cdot U^i = v_i$. (See Sec. 3.2 for the more general way that this constraint arises and some useful equalities relating b_i , v_i , etc.)

We must now consider how b_i changes as U^i changes. The lowest order case is where b_i is a constant, independent of U^i . This means that for real-valued b_i , ϵ_i is identical to the Boltzmann utility K_i discussed in Sec. 3.2, with U^i playing the role that f_i does in the definition of Boltzmann utility, and b_i playing the role of β_i . Just as we extend the domain of definition of $K_i(\cdot)$ to include ∞ , we

²⁶We mean "consistent" in the sense that even though $q_{j \neq i}$ has changed, it is still true that

$$\begin{aligned} U^i(x_i) &= \int dx'_{-i} u^i(x_i, x'_{-i}) q_{-i}(x'_{-i}) \\ &= \int dx'_j dx'_{-\{i,j\}} u^i(x_i, x'_j, x'_{-\{i,j\}}) q_j(x'_j) q_{-\{i,j\}}(x'_{-\{i,j\}}). \end{aligned}$$

²⁷Note how problematic it would be to try to encapsulate our invariance in a traditional, non-Jaynesian "bath"-based approach to statistical physics. In such an approach, the invariants are the sums, across both the system under consideration and an external "bath", of certain physical quantities. In other words, the aggregate amount of those quantities across the system and the external bath is taken to be conserved. For example, the CE arises if one takes the aggregate energy to be conserved. It is not at all clear how one could express our invariants as the values of such conserved quantities.

do the same for ϵ_i and for q_i^* : For $b_i = \infty$, q_i^* is the distribution that is uniform over the set $\operatorname{argmax}_{x_i} U^i(x_i)$, and zero elsewhere.

Below we will use the shorthand $q^*(x) \triangleq \prod_i q_i^*(x_i)$ where for each i the U^i arising in the definition of q_i^* is understood to be based on the q_{-i}^* (i.e., each U^i means $U_{q^*}^i$). So the definition of q^* reflects coupling between the player's mixed strategies (though not necessarily between their moves): a change to q_j^* for some particular j in general will modify the strategies $q_{k \neq j}^*$. q^* is the Quantal Response Equilibrium (QRE) solution, discussed in Sec. 3.2. In general, for any particular game and \vec{b} , there is at least one, and may be more than one associated q^* . This follows from Brouwer's fixed point theorem [27, 29].

As a point of notation, the expression $\mathcal{S}_{\vec{b}}$ is defined to be the invariant that $\forall i, q_i \cdot U_{q_{-i}}^i = K(U_{q_{-i}}^i, b_i)$, where it is implicitly assumed that $\vec{b} \succeq \vec{0}$. For any such \vec{b} there is always at least one q that satisfies $\mathcal{S}_{\vec{b}}$ (e.g., the QRE).

4.3 The impossibility of a Nash equilibrium

We now consider the second interpretation of “equilibrium joint mixed strategy”, not as the Bayes-optimal q given \mathcal{S} , but rather as $P(x | \mathcal{S})$. We do this for the coupled players invariant, \mathcal{S} .

As discussed above, the advantage of this second interpretation of “equilibrium” is that it does not depend on the loss function of the external scientist. Rather it only depends on properties of the game and the players. Unfortunately, as elaborated here, in some situations the joint mixed strategy $P(x | \mathcal{S})$ cannot equal a Nash equilibrium of the underlying game. So under this interpretation, a Nash equilibrium is impossible.

To illustrate this, say that (the support of) $P(q | \mathcal{S})$ is restricted to the Nash equilibria of the underlying game, so that the players are perfectly rational. (See Sec. 4.7 below.) Say that there are multiple such equilibria, written q^1, q^2, \dots , with $P(q | \mathcal{S}) \triangleq \sum_j a^j \delta(q - q^j)$. So the a^i form a probability distribution over the equilibria. Since the entropic prior extends over all $q \in \Delta_{\mathcal{X}}$, in general none of the a^i will equal zero exactly.

Since the equilibria are all product distributions, using Eq. 24 we can write

$$P(x | \mathcal{S}) = \sum_j a^j \prod_k q_k^j(x_k) \quad (27)$$

so that

$$\begin{aligned} P(x_i | \mathcal{S}) &= \int dx_{-i} P(x | \mathcal{S}) \\ &= \sum_j a^j q_i^j(x_i). \end{aligned} \quad (28)$$

Consider the case where the Nash equilibria are not interchangeable, so $P(q | \mathcal{S})$ is not a product distribution. This means that $P(x | \mathcal{S})$ is not a product distribution in general, so that the players appear to be coupled, to us. (See the discussion just below Eq. 24.)

At an intuitive level, such coupling is analogous to the consistency-among-players coupling that underlies the concept of a Nash equilibrium. However because it mixes the Nash equilibria with each other, in general the sum in Eq. 28 is not a best response mixed strategy for the product distribution $\prod_{j \neq i} P(x_j | \mathcal{I})$. Formally, $p^i(x_i) \triangleq P(x_i | \mathcal{I})$ does not maximize the dot product

$$\int dx_i p^i(x_i) \left[\int dx_{-i} u(x_i, x_{-i}) \prod_j P(x_j | \mathcal{I}) \right]. \quad (29)$$

Similarly $P(x_i | \mathcal{I})$ is not a best-response mixed strategy for $P(x_{-i} | \mathcal{I})$. So when the underlying game has multiple non-interchangeable equilibria, then even if the players are perfectly rational, in general we will not predict distributions governing the moves of the players that are best-response mixed strategies to each other.

Example: Say we have a team (common payoff) symmetric game involving two players, each having a move space $\{A, B\}$. Let the shared utility function be $u(A, A) = 2, u(A, B) = u(B, A) = 0, u(B, B) = 1$. This game has three Nash equilibria: $(A, A), (B, B)$, and the mixed strategy where each player makes move A with probability $1/3$. The first two of those q have entropy 0 (they are delta functions). The associated value of the entropic prior, $\exp(\alpha S(q))/Z(\alpha)$, is just $[Z(\alpha)]^{-1}$. The last Nash equilibrium has entropy $\ln[3] - 2/3\ln[2]$.

If we define $w(\alpha) \triangleq \exp(\alpha \{\ln[3] - 2/3\ln[2]\})$, then the prior probability of the first two (pure strategy) Nash equilibria are $1/[2 + w(\alpha)]$, and the prior probability of the last (mixed strategy) Nash equilibrium is $w(\alpha)/[2 + w(\alpha)]$. Since all three equilibria have the same likelihood (namely, 1), these prior probabilities are also their respective posterior probabilities, $P(q | \mathcal{I})$, i.e., they are the three values of a^j . Accordingly,

$$\begin{aligned} P(x = (A, A) | \mathcal{I}) &= \frac{1}{2 + w(\alpha)} + \frac{w}{(2 + w)} \left(\frac{1}{3}\right)^2 \\ &= \frac{9 + w(\alpha)}{9(2 + w(\alpha))} \\ P(x = (B, B) | \mathcal{I}) &= \frac{1}{2 + w(\alpha)} + \frac{w}{(2 + w)} \left(\frac{2}{3}\right)^2 \\ &= \frac{9 + 4w(\alpha)}{9(2 + w(\alpha))} \end{aligned} \quad (30)$$

and

$$\begin{aligned} P(x = (A, B) | \mathcal{I}) &= P(x = (B, A) | \mathcal{I}) \\ &= \frac{2w}{9(2 + w(\alpha))} \end{aligned} \quad (31)$$

Not only is this distribution $P(x | \mathcal{I})$ not a Nash equilibrium; neither player i plays best-response to $P(x_{-i} | \mathcal{I})$. In fact, $P(x | \mathcal{I})$ is not even a product

distribution in general. (To see this, evaluate the marginals of $P(x | \mathcal{S})$ for the two players, and compare the product of those two marginals with $P(x | \mathcal{S})$.)

Note that these conclusions do not depend critically on our choice of ϵ_i , or even on our choice of encapsulating \mathcal{S} in terms of such functions ϵ_i . (After all, we're explicitly allowing the case where $P(q | \mathcal{S})$ is restricted to Nash equilibria.) Rather it comes from the fact that our prior allows non-zero probability for all of the Nash equilibria.

4.4 The QRE and ϵ_i

Unlike the usual motivation of the QRE, the motivation for our choice of ϵ_i does *not* say that q_i must be a Boltzmann distribution. It does not say that the probability distribution over possible q_i is a delta function about a Boltzmann distribution q_i . Rather it says that q_i^* , the most likely q_i for the single-player inference problem, is a Boltzmann distribution. It then uses that fact to motivate a functional form for ϵ_i in the multi-player scenario. Here we only assume that the relation between $E_q(u^i)$ and U^i given by that functional form is consistent with $q_i = q_i^*$. In general the invariant $E_q(u^i) = \epsilon_i(U^i; b_i)$ holds for many q_i in addition to the Boltzmann distribution.

Indeed, fix q , and consider any i and the associated $U_{q_{-i}}^i$. Recall that we are restricting attention to benign q 's (cf. the discussion at the end of Sec. 3.2). So no matter what it is, our q_i is consistent with our invariant for that $U_{q_{-i}}^i$, for some b_i . Since this is true for all i , any q is consistent with our full invariant for some \vec{b} . Furthermore, for any finite α , the support of the entropic prior is all $\Delta_{\mathcal{X}}$. This means that every q has non-zero posterior probability $P(q | \mathcal{S}_{\vec{b}})$ for some \vec{b} .

In contrast not every q_i is a Boltzmann distribution, i.e., not every q_i is part of a QRE. In other words, to assume a system is in a QRE is to make a restrictive assumption about the physical system q , an assumption that may or may not be correct. This is not the case with our invariant.

Finally, it turns out that the QRE can be viewed as an approximation to the MAP prediction for our \mathcal{S} . A detailed discussion of this is presented in Sec. 4.6 below.

4.5 The MAP q

Given our invariant, our likelihood is

$$\begin{aligned} P(\mathcal{S} | q) &= \prod_i \delta(E_q(u^i) - \epsilon_i(U_q^i)) \\ &= \prod_i \delta(q_i \cdot U^i - \epsilon_i(U_q^i)). \end{aligned} \quad (32)$$

Recall that with the canonical ensemble the likelihood stipulates a linear constraint on the underlying probability distribution. In contrast, due to the non-linearity of ϵ_i , here the likelihood stipulates a non-linear constraint on q .

As usual, if we wish we can distill the associated posterior into a single prediction for q , e.g., into the MAP estimate. Naively, one might presume that q^* is that MAP estimate. After all, q^* respects our constraints that $E_q(u^i) = \epsilon_i(U_q^i) \forall i$, and it maximizes the entropy of each player's strategy considered in isolation of the others. However in general q^* will not maximize the entropy of the joint mixed strategy subject to our constraints. In other words, while MAP for each individual player's strategy, in general it is not MAP for the joint strategy of all the players. The reason is that setting each separate q_i to maximize the associated entropy (subject to having q obey our invariant), in a sequence, one after the other, will not in general result in a q that maximizes the sum of those entropies. So it will not in general result in a q that maximizes the entropy of the joint system.

Proceeding more carefully, the MAP estimate of the mixed strategy q is given by the critical point of the Lagrangian

$$\mathcal{L}(q, \{\lambda_i\}) = S(q) + \sum_i \lambda_i (q_i \cdot U^i - \epsilon_i(U^i)) \quad (33)$$

where the λ_i are the Lagrange parameters enforcing the constraints provided by the likelihood function of Eq. 32. The critical point of this Lagrangian must satisfy

$$\begin{aligned} 0 &= \frac{\partial \mathcal{L}}{\partial q_i(x_i)} \\ &= -1 - \ln[q_i(x_i)] + \lambda_i E(u^i | x_i) + \sum_{j \neq i} \lambda_j [E(u^j | x_i) - \frac{\partial \epsilon_j(U^j)}{\partial q_i(x_i)}] \\ &= -1 - \ln[q_i(x_i)] + \lambda_i E(u^i | x_i) + \\ &\quad \sum_{j \neq i} \lambda_j [E(u^j | x_i) - \int dx_j \frac{\partial \epsilon_j(U^j)}{\partial U^j(x_j)} \frac{\partial U^j(x_j)}{\partial q_i(x_i)}] \\ &= -1 - \ln[q_i(x_i)] + \lambda_i E(u^i | x_i) + \\ &\quad \sum_{j \neq i} \lambda_j [E(u^j | x_i) - \int dx_j \frac{\partial \epsilon_j(U^j)}{\partial U^j(x_j)} E(u^j | x_i, x_j)]. \end{aligned} \quad (34)$$

Accordingly, at the MAP solution, for all players i ,

$$q_i(x_i) \propto e^{\lambda_i E(u^i | x_i) + \sum_{j \neq i} \lambda_j [E(u^j | x_i) - \int dx_j \frac{\partial \epsilon_j(U^j)}{\partial U^j(x_j)} E(u^j | x_i, x_j)]} \quad (35)$$

This is a set of coupled nonlinear equations. The solution will depend on the functional form of each ϵ_j . The form being investigated here is Boltzmann utility functions, so we must plug that into Eq. 34 to evaluate $\frac{\partial \epsilon_j(U^j)}{\partial U^j(x_j)}$. After doing that, interchange the order of the two differentiations, to differentiate with respect to $U^j(x_j)$ before differentiating with respect to b_j . Carrying through

the algebra one gets

$$\begin{aligned} \frac{\partial \epsilon_j(U^j)}{\partial U^j(x_j)} &= q_j^*(x_j)[1 + b_j\{U^j(x_j) - E_{q_j^*}(U^j)\}] \\ &= q_j^*(x_j)[1 + b_j\{E_{q_{-j}}(u^j | x_j) - E_{q_j^*, q_{-j}}(u^j)\}] \end{aligned} \quad (36)$$

We must now plug this into the integrals occurring in Eq.'s 34 and 35. Each such integral becomes

$$\begin{aligned} \int dx_j \frac{\partial \epsilon_j(U^j)}{\partial U^j(x_j)} E_q(u^j | x_i, x_j) \\ = \int dx_j q_j^*(x_j) E_q(u^j | x_i, x_j) [1 + b_j\{E_{q_{-j}}(u^j | x_j) - E_{q_j^*, q_{-j}}(u^j)\}] \end{aligned} \quad (37)$$

Together with the constraints $\{E_q(u^j) = \epsilon_j(U^j)\}$, Eq. 35 now gives us a set of coupled nonlinear equations for the parameters $\{\lambda_j\}$ and the q_j . The solution to this set of equations gives our MAP q .

4.6 The relation between the MAP q and the QRE

Ultimately the only free parameters in our solution for the MAP q are \vec{b} . The QRE solution q^* is also a set of coupled nonlinear equations parameterized by \vec{b} . In general there is a very complicated relation between the MAP $q(x)$ and $q^*(x)$, one that varies with \vec{b} (as well as with the $\{u^j\}$, of course). In particular, in general the two solutions differ.

Intuitively, the reason for the difference between the two solutions is that each player i does not operate in a fixed environment, but rather in one containing intelligent players trying to adapt their moves to take into account i 's moves. This is embodied in the likelihood of Eq. 32. In contrast to that likelihood, the likelihoods of the QRE each implicitly assume that the associated player i operates in a fixed environment.

Formally, if we make a change to q_i , then the likelihood of Eq. 32 will induce a change to q_{-i} , to have the invariant for the players other than i still be satisfied. This change to q_{-i} will then induce a “second order” follow-on change to q_i , to satisfy the invariant for player i . This second-order effect will not arise in the likelihood associated with the QRE q_i^* , which treats the other players as fixed.

Note that with the likelihood of Eq. 32 the second-order effect will induce a further change to q_{-i} , to ensure the invariant is still satisfied, which will then cause a third order change to q_i , and so on. This back-and-forth is a direct mathematical manifestation of the “I know that you know that I know that you prefer ...” feature at the core of game theory. This is the phenomenon that distinguishes game theory as a subject from decision theory. The difference between the QRE and the MAP q is an encapsulation of this distinguishing feature.

There are other ways to view the intuitive nature of the relationship between the QRE and the MAP q . For example, in deriving the MAP q one follows standard probability theory and multiplies likelihoods concerning the separate players to get a likelihood concerning the full joint system. The mode of (the product of the prior and) that joint likelihood gives the single most likely solution to our inference problem. In contrast, the QRE q starts by separately finding the most likely solutions to each of many different inference problems (one problem for each player). It then multiplies those solutions concerning different problems together. It is not apparent what justifying formal argument (i.e., one based on Savage's axioms) there is for taking that product of solutions of different problems as one's guess for the solution to a single joint problem.

The mathematical relationship between the QRE and the MAP q is a complicated one. Here we consider the simplifying approximation that *under the integral of Eq. 37*, we can equate $q(x)$ and $q^*(x)$. In other words, assume we can use the mean-field approximation within integrands. Exploiting this, we can evaluate the integral in Eq. 37:

$$\begin{aligned}
& \int dx_j \frac{\partial \epsilon_j(U^j)}{\partial U^j(x_j)} E_q(u^j | x_i, x_j) \\
& \cong \int dx_j [q_j(x_j) E_q(u^j | x_i, x_j) + \\
& \quad q^*_{-j}(x_j) E_{q^*}(u^j | x_i, x_j) b_j \{E_{q^*}(u^j | x_j) - E_{q^*}(u^j)\}] \\
& = E_q(u^j | x_i) - \\
& \quad b_j [E_{q^*}(u^j) E_q(u^j | x_i) - \int dx_j q_j(x_j) E_{q^*}(u^j | x_j) E_{q^*}(u^j | x_j, x_i)]
\end{aligned} \tag{38}$$

where we have used the fact that q is a product distribution. Plugging this into Eq. 34 gives

$$\begin{aligned}
0 & = -1 - \ln[q_i(x_i)] + \lambda_i E_q(u^i | x_i) + \\
& \quad \sum_{j \neq i} \lambda_j b_j [E_{q^*}(u^j) E_q(u^j | x_i) - \\
& \quad \int dx_j q^*_{-j}(x_j) E_{q^*}(u^j | x_j) E_{q^*}(u^j | x_j, x_i)] \tag{39}
\end{aligned}$$

as our equation for q_i in terms of q_{-i} and q^* .

So consider the situation where for all j ,

$$E_{q^*}(u^j) E_{q^*}(u^j | x_i) = \int dx_j q^*_{-j}(x_j) E_{q^*}(u^j | x_j) E_{q^*}(u^j | x_j, x_i). \tag{40}$$

In this situation, in light of Eq. 35, we recover for the MAP q the very QRE solution that we assumed when we made the mean-field approximation, where $b_i = \lambda_i \forall i$. Accordingly, if the QRE solution obeys Eq. 40, it is an MAP solution.

If the QRE only approximately obeys Eq. 40, then the exact MAP solution can be found by expanding around the QRE via Eq. 39.

The difference between the two sides of Eq. 40 is a covariance, evaluated according to q_j^* , between the random variables $E_{q^*}(u^j | x_j, x_i)$ and $E_{q^*}(u^j | x_j)$.²⁸ Comparing Eq.'s 37 and 35, this provides the following result concerning our mean-field approximation:

Theorem 1: The QRE q^* is the MAP of $P(q | \mathcal{I})$ with the vector equality $\lambda = b$ iff $\forall i$,

$$\sum_{j \neq i} (b_j)^2 \text{Cov}_{q_j^*} [E_{q^*}(u^j | x_j, x_i), E_{q^*}(u^j | x_j)]$$

is independent of x_i , where Cov is the covariance operator:

$$\text{Cov}_p[a(y), b(y)] \triangleq \int dy p(y) a(y) b(y) - \int dy p(y) a(y) \int dy p(y) b(y).$$

Particularly for very large systems (e.g., a human economy), it may be that $E_{q^*}(u^j | x_j, x_i) = E_{q^*}(u^j | x_j)$ for almost any i, j and associated moves x_i, x_j . In this situation the move of almost any player i has no effect on how the expected payoff to player j depends on j 's move. If this is in fact the case for player i and all other players j , then the covariance for each j, x_j, x_i reduces to the variance of $E_{q^*}(u^j | x_j)$ as one varies x_j according to q_j^* .

This variance is given by the partition function:

$$\begin{aligned} \text{Var}_{q_j^*}(E_{q^*}(u^j | x_j)) &= \text{Var}_{q_j^*}(U_{q^*}^j) \\ &= \frac{\partial^2 \ln(Z_{U_{q^*}^j}(b_j))}{\partial (b_j')^2} \Big|_{b_j' = b_j}. \end{aligned} \quad (41)$$

In particular, for $b_j \rightarrow \infty$ — perfectly rational behavior on the part of agent j — the variance goes to 0. So if every i is “decoupled” from all other agents, then in the limit that all such agents become perfectly rational, the expression in Thm. 1 generically goes to 0. (The b_j -dependence in the covariance occurs in an exponent, and therefore generically overpowers the $(b_j)^2$ multiplicative factor.) So the QRE approaches the MAP solution in that situation.

On the other hand, if the players have bounded rationality, their variances are nonzero. In this case the expression in Thm. 1 is nonzero for each i, x_i . Typically for fixed i the precise nonzero value of that variance will vary with x_i . In this case, by Thm. 1, we know that the QRE differs from the MAP solution.

There are many special game structures (e.g., zero-sum games) in which one can make some arguments about the likely form of the sum in Thm. 1. An elaboration of those arguments is the subject of ongoing research.

²⁸Note that that second random variable is just the average (according to q_i^*) of the first one. So we can rewrite the covariance another way, as a covariance evaluated according to $q_j^*(x'_j)q_i^*(x'_i)$, between the random variables $E_{q^*}(u^j | x'_j, x_i)$ and $E_{q^*}(u^j | x'_j, x'_i)$.

4.7 The posterior q covers all Nash equilibria

Not all q can be cast as a QRE for some appropriate \vec{b} . So in particular, a q that occurs in the real world will in general differ, even if only slightly, from all possible QRE's. This can be viewed as a shortcoming of the QRE (a shortcoming that applies to all equilibrium concepts with a sufficiently small number of parameters).

Now as $\vec{b} \rightarrow \infty$, the QRE reduces to some mixed strategy Nash equilibrium. Different sequences of the \vec{b} going to the infinity vector can lead to different Nash equilibria. However in general starting from the point where all $b_j = 0$ and continuously increasing the components of \vec{b} can only lead to one particular equilibrium, and other Nash equilibria are not the limit of such a sequence [27]. This too can be viewed as a short-coming of the QRE.

However from the perspective of PGT, there is far more to the posterior distribution specified by a particular vector \vec{b} than some single q chosen using that posterior, be that q the associated Bayes-optimal q , the MAP q , or an approximation to the MAP q like the QRE. In this, the potential impossibility of one particular sequence of such q 's approaching some particular one of the game's Nash equilibria is not necessarily a reason for concern.

To formalize this we start with the following result:

Proposition 1: Define $\mathcal{Q}(\vec{b}) \triangleq \{q \in \Delta_{\mathcal{X}} : \forall i, P(q_i | \mathcal{S}_{\vec{b}}) > 0 \text{ for some } \vec{b}' \succeq \vec{b}\}$. Let B be some sequence of \vec{b} values that converges to ∞ , i.e., such that for all $\vec{b} \in B$ having no infinite components, $\exists \vec{b}' \in B$ where $\vec{b}' \succ \vec{b}$. Then all members of $\cap_{\vec{b} \in B} \mathcal{Q}(\vec{b})$ are Nash equilibria of the game.

Proof: Hypothesize $\exists \tilde{q} \in \cap_{\vec{b} \in B} \mathcal{Q}(\vec{b})$ which is not a Nash equilibrium. Then $\exists i$ such that $U_{\tilde{q}-i}^i$ is not constant valued. In addition, we know that $\tilde{q}_i \cdot U_{\tilde{q}-i}^i \equiv v_i < \max_{x_i} U_{\tilde{q}-i}^i(x_i)$. However recall from Sec. 3.2 that if $U_{\tilde{q}-i}^i$ is not constant-valued, the Boltzmann utility $K(U_{\tilde{q}-i}^i, \cdot)$ is a monotonically increasing bijection with domain $[0, \infty)$ and range $[\frac{\int dx_i U_{\tilde{q}-i}^i(x_i)}{\int dx_i 1}, \max_{x_i} U_{\tilde{q}-i}^i(x_i)]$. Since v_i falls within that range, this means that we can invert $K(\cdot, U_{\tilde{q}-i}^i)$ to get a unique finite value \tilde{b}_i that is consistent with \tilde{q} . Accordingly, $P(\mathcal{S}_{\vec{b}} | \tilde{q})$ is non-zero only if $b_i = \tilde{b}_i$, and therefore so is $P(\tilde{q} | \mathcal{S}_{\vec{b}})$.

However by definition \tilde{q} must be a member of $\mathcal{Q}(\vec{b})$ for all \vec{b} in the limiting sequence. That means in particular that it must be a member of $\mathcal{Q}(\vec{b}')$ for some \vec{b}' where $b'_i > \tilde{b}_i$. But by definition, all members q of $\mathcal{Q}(\vec{b}')$ have $P(q | \mathcal{S}_{\vec{b}'}) > 0$ for some b such that $b_i \geq b'_i > \tilde{b}_i$. Since we know $P(\tilde{q} | \mathcal{S}_{\vec{b}'})$ is non-zero only if $b_i = \tilde{b}_i$, this means that $\tilde{q} \notin \mathcal{Q}(\vec{b}')$, contrary to hypothesis. **QED.**

Conversely, every q has a non-infinitesimal posterior probability (density), for some (potentially infinite) \vec{b} that specifies that posterior. More formally,

Proposition 2: For any benign $q \in \Delta_{\mathcal{X}}$ there is a unique \vec{b} and associated invariant $\mathcal{S}_{\vec{b}}$ such that $P(q | \mathcal{S}_{\vec{b}}) \neq 0$. For that \vec{b} , for all $q' \in \Delta_{\mathcal{X}}$,

$$\frac{P(q | \mathcal{S}_{\vec{b}})}{P(q' | \mathcal{S}_{\vec{b}})} \geq |X|^{-\alpha} \quad (42)$$

where α is the exponent of the entropic prior.

Proof: First recall that any q has non-zero posterior probability $P(q | \mathcal{S}_{\vec{b}})$ for some \vec{b} , assuming finite entropic prior constant α . (See Sec. 4.3.) So to prove the first part of the proposition we must establish the uniqueness of that \vec{b} .

Consider any i and the given q . Say $q_i \cdot U_{q_{-i}}^i \equiv v_i \neq \max_{x_i} U_{q_{-i}}^i(x_i)$. This means that $U_{q_{-i}}^i$ is not the constant function that is independent of its argument. Now recall from Sec. 3.2 that for any such v_i and fixed non-constant U^i , there is always a unique $b_i \in [0, \infty)$ such that $K_i(b_i)$ equals v_i . On the other hand, as explained in the discussion in that subsection, if $v_i = \max_{x_i} U_{q_{-i}}^i(x_i)$, then regardless of whether $U_{q_{-i}}^i$ varies with its argument, $b_i = \infty$. So there is a unique b_i consistent with q , which we write as b_i^* . Since this holds simultaneously for all i , the entire vector \vec{b}^* with components $\{b_i^*\}$ is unique.

This means that the likelihood $P(\mathcal{S}_{\vec{b}^*} | q) = 1$. On the other hand, $P(\mathcal{S}_{\vec{b}^*} | q') \leq 1$ for any q' . Accordingly, the ratio in the proposition is bounded above by the ratio of the exponential prior at q to that at q' . However the ratio of $e^{\alpha S(q')}$ between any two points q'' is bounded below by $\frac{\exp(\alpha \cdot 0)}{\exp(\alpha \ln(|X|))}$. **QED**

In particular, this result holds for Nash equilibrium q ; such equilibria arise for $\vec{b} = \vec{\infty}$. The relative probabilities of those Nash q are given by the ratios of the associated prior probabilities, i.e., by (the exponential of) the associated entropies, $S(q)$. This reflects our presumption that it is *a priori* more likely that the adaptation/learning processes that couple the players results in a Nash equilibrium with broad q than that it results (for example) in a “golf hold” pure strategy q . (Generically, such golf hole solutions are more difficult to find for any broadly applicable learning process.)

Prop. 2 also holds for any particular q infinitesimally close to one of the Nash equilibria. In this sense, the posterior probability is arbitrarily tightly restricted to any one of the Nash equilibria for some appropriate \vec{b} .

The picture that emerges then is that $\forall \vec{b}, \exists$ proper submanifold of $\Delta_{\mathcal{X}}$ that is the support of the posterior. There is no overlap between those submanifolds (one for each \vec{b}), and their union is all of $\Delta_{\mathcal{X}}$, including the Nash equilibria q 's (for which $\vec{b} = \vec{\infty}$). Those Nash equilibria are the limit points of those submanifolds (in a sequence of increasing \vec{b}).

Within any single one of the submanifolds no q has too small a posterior (cf. Prop. 2). This is because all q within a single submanifold have the same value (namely 1) of their likelihoods. Accordingly, the ratios of the posteriors of the q 's within the submanifold is given by the ratios of (the exponentials of) the entropies of those q 's.

Finally, consider the case that the submanifolds get a unique maximum as $\vec{b} \rightarrow \infty$. This means that the mode of the posterior — the MAP q — necessarily goes to a single one of the Nash equilibria in that limit. In this sense, “only one Nash equilibrium is picked out by that limit”. In particular, this limiting behavior holds for the QRE approximation to the MAP q . As mentioned, this has been seen as a problematic aspect of the QRE equilibrium concept. However from the prospect of PGT there is nothing untoward about this behavior. After all, all of the Nash equilibria have non-zero posterior in that limit (cf. Prop. 2); it just so happens that the QRE ends up at a single one of those equilibria.

4.8 Alternative choices of \mathcal{I}

Of course, one can always design “learning” algorithms for players to follow in such a way that our assumed invariants don’t hold. After all, in the extreme case you can design “learning” algorithms that are intentionally stupid, giving higher probability to moves with lower expected utility. Less trivially, there are many algorithms that are of interest in the game theory community even though they would never be considered by anyone in the machine learning community applying learning algorithms to real-world problems (e.g., fictitious play). It may well be that such algorithms don’t obey the assumed invariants exactly for some $\{U^i\}$.

However this issue also obtains, at least as strongly, for alternative encapsulations of rationality like Nash equilibrium, trembling hand, quantal response, etc. It is trivial to design “learning algorithms” that guarantee that those equilibria cannot arise. More generally, in *all* statistical inference — in other words, in all of science — any formalization of invariants may well have some error. This is even true in statistical physics, and is an intrinsic feature to any predictive science.

All of this notwithstanding, there are a number of alternative choices of ϵ_i to the one considered here that should be investigated in depth. To give a simple example, the \mathcal{I} considered here is only a “lowest order” choice for an invariant. In particular, as mentioned above, our choice of \mathcal{I} assumes b_i is independent of U^i . A more sophisticated analysis than can be fully developed here would consider possible couplings between b_i ’s and U^i ’s; in this paper b_i ’s and U^i ’s are independent.

However one does not need to couple b_i ’s and U^i ’s to get reasonable alternative choices of ϵ_i . For example, due to the U^i -independence of b_i with our \mathcal{I} , whatever b_i is, for some q the associated likelihood $P(\mathcal{I}_b | q) = 0$ (just like whatever the temperature of a physical system, some phase space distributions are incompatible with that temperature). To avoid this, in many scenarios we might want to allow how smart a player i is to vary from one instance to the next, even without considering detailed mathematical structures relating variations in b_i with those in U^i .

One way to do this would mean allowing b_i to vary in an U^i -independent

manner, with only its average value fixed.²⁹ This simple generalization of \mathcal{S} can be accommodated by switching the analysis to involve type II games. Although the details of that analysis (like all other details of type II game theory) is beyond the scope of this paper, it is worth making some broad comments on it.

That type II analysis starts by extending the definition of an environment vector to type II games in the obvious way: indexed by q'_i , the type II environment is defined by

$$U_{\pi_{-i}}^i(q'_i) \triangleq \int dx' dq'_{-i} \pi_{-i}(q'_{-i}) q'_{-i}(x'_{-i}) q'_i(x'_i) u^i(x'_i, x'_{-i}) \quad (43)$$

so that the expected value of u^i is given by

$$\begin{aligned} \pi_i \cdot U_{\pi_{-i}}^i &= \int dq'_i \pi_i(q'_i) U_{\pi_{-i}}^i(q'_i) \\ &= E_{\pi_i, \pi_{-i}}(u^i) \end{aligned} \quad (44)$$

The analysis also extends the definition of $K(., .)$ to type II games in the obvious way: $K(U_{\pi_{-i}}^i, B_i)$ is what $\pi_i \cdot U_{\pi_{-i}}^i$ would be if π_i were the associated Boltzmann distribution, $\pi_i(q'_i) \propto \exp(B_i U_{\pi_{-i}}^i(q'_i))$.

The new invariant would then be that for all i ,

$$\pi_i \cdot U_{\pi_{-i}}^i = K(U_{\pi_{-i}}^i, B_i) \quad (45)$$

This invariant allows any q_i to occur, thereby allaying the potential shortcoming with the invariant this paper focuses on. It is now certain π_i that are excluded rather than certain q_i .

To motivate our next alternative for the invariant \mathcal{S} , consider the distribution induced by $q_{-i}(x_{-i})$ over player i 's move-specified utilities $u^i(x_i, .)$ (one such distribution for each x_i),

$$P_{q_{-i}}(u^i(x_i, .) = u) = \int dx'_{-i} q_{-i}(x'_{-i}) \delta(u - u^i(x_i, x'_{-i})). \quad (46)$$

\mathcal{S} implicitly assumes that those aspects of i 's behavior that it is safe for us to presume are only those that involve the first moments of these distributions,

$$\begin{aligned} U_{q_{-i}}^i(x_i) &= \int du P_{q_{-i}}(u^i(x_i, .) = u) u \\ &= \int dx'_{-i} q_{-i}(x'_{-i}) u^i(x_i, x'_{-i}). \end{aligned} \quad (47)$$

²⁹Indeed, in practice each b_i is at best loosely known. So formally speaking it is a random variable with its own distribution, and so even within a type I game it must be marginalized out to get our posterior $P(q | \mathcal{S})$. This type of random variable is known as a ‘‘hyperparameter’’ [44, 43]. (A more common example of a hyperparameter is the typically unknown width of a Gaussian noise process that corrupts some data.) In particular, in almost all of this paper we are implicitly assuming that the posterior over each b_i is quite peaked, so that in our analysis we can simply set b_i to a constant, albeit an unknown one.

In this it simply emulates conventional game theory.

However in many real-world coupled-players scenarios the higher moments, reflecting the breadth and overlaps of the distributions over $u^i(x_i, \cdot)$, will have a major impact on our inference of q_i . Intuitively, if those distributions — each a function purely of q_{-i} — maintain the same mean but get broader with more overlap between them, that will increase the variability of what inferences i makes concerning those means and their linear ordering. (For example, that is the case if i makes its inference of those means based on empirical samples of the distributions.) This will make our associated distribution over q_i broader — there are more q_i that we can conceive of i arriving at. Similarly, such broadening of the distributions over the $u^i(x_i, \cdot)$ would often be evident to i . That might make i realize it can have less confidence in its inference of the ordering of the means of those distributions. In such a situation, many real-world players i would become more conservative in formulating their mixed strategy, q_i . So not only might the distribution over q_i get broader, but it may also shift, if q_{-i} changes to cause this kind of broadening of the distributions.

To be more quantitative, say the variances of the $U^i(x_i, \cdot)$,

$$V_{q_{-i}}^i(x_i) \triangleq \left\{ \int dx_{-i} q_{-i}(x_{-i}) [u^i(x_i, x_{-i})]^2 \right\} - [U_{q_{-i}}^i(x_i)]^2, \quad (48)$$

are increased, and that the overlap between the distributions over each $u_i(x_i, \cdot)$ (measured for example via Kullback-Leibler distance between those distributions) also are increased. Then there is often increased uncertainty on our part about the relationships between i 's sample-driven preferences among the x_i . This often means we are less sure in our inference of what i 's current mixed strategy is, which means our posterior over q_i should get broader.

In addition, under such broadening in the $u^i(x_i, \cdot)$ there is increased uncertainty about what i 's best move would be for the actual move x_{-i} that will be formed by sampling $q_{-i}(x_{-i})$. Typically this means that the information that i has gleaned via its previous interactions with the other players is not as helpful to i for determining its best move for the current game. Intuitively, when these distributions are broader i faces worse signal-to-noise in discerning the relation between the $U^i(x_i)$ based on limited data. This will often manifest itself by changes to what mixed strategy i is most likely to adopt.

A standard illustration of both of these effects arises if one compares two extreme scenarios. The first is the “US economy game” that any particular US citizen i repeatedly engages in with the 300 million other human players in the US. The second is a simple game against Nature that i repeatedly engages in where there is no variance in Nature’s choice of move. Our inference of q_i is far easier in the second scenario. Similarly, typically i will have an easier time discerning its best move in the second scenario.³⁰

One approach to incorporate such effects would be to have the set of all $\{V^i(x_i)\}$ (running over i as well as the associated x_i) and overlaps between

³⁰See [11, 53, 29, 15, 54, 55], and references therein to “Collective Intelligence” for a discussion of how this second type of effect can be addressed for mechanism design and distributed control.

the distributions over the $\{u^i(x_i, \cdot)\}$ help specify \vec{b} . Such an approach could obviously address the second of the effects we’re concerned with, involving how much information i has managed to glean concerning the other players. It is not a fully satisfactory approach to addressing the first effect however. This is because once \vec{b} is set — however that is done — some q are excluded, i.e., some q have posterior probability equal to 0. Typically to change \vec{b} to allow those previously excluded q — and thereby broaden the distribution over q_i — the Bayes-optimal (or MAP) q also changes. Moreover, such a modification invariably excludes some q that were previously allowed (see Sec. 4.7). Instead what we want is our increase of the breadth of the posterior over q to allow previously excluded q , while still allowing all q we did earlier.

The exclusionary character of the posterior over q that is causing this difficulty can be removed by casting the analysis in terms of type II games rather than (as in the exposition above) type I games. After all, in general the π that obeys the conventional type II game invariant has support extending over all q (cf. Eq. 45). A detailed exploration of how to use type II games to incorporate the effects of the $\{V^i(x_i)\}$ and overlaps between the $\{u^i(x_i, \cdot)\}$ into our posterior is beyond the scope of this paper however.

The discussion so far has focused on variants of \mathcal{S} that are at most loosely based on empirical data. Those variants incorporate none of the insight of behavior economics, prospect theory, or behavioral game theory [49, 50, 56, 48]. Crucially important future investigations involves incorporating that work, and more generally the entire field of user-modeling and knowledge-engineering, into our choice of \mathcal{S} .

It is worth noting that there are alternative \mathcal{S} ’s that don’t involve the numeric values of the u^i ’s, but rather only require that each u^i provides an ordering over the q . The idea here is to consider what is invariant if i stays “just as smart”, while U^i undergoes a non-affine monotonically increasing transformation, and q^i changes accordingly. For example, one might argue that q_i would be “just as smart” after such a transformation if the fraction of alternative q'_i such that $q'_i \cdot U^i > q_i \cdot U^i$ is the same before and after the transformation. Formally, this would mean that $\int dq'_i \Theta(U^i \cdot [q_i - q'_i])$ is a constant, rather than (as in the choice considered in this paper) $U^i \cdot q_i$. Intuitively, under this choice, “how smart” i is reflects how good she is at ruling out some of the candidate q'_i as inferior to the final q_i she uses.³¹ This formalization of how smart i is is essentially identical to what is called **intelligence** in work on Collective Intelligence [57, 54, 58].

Ultimately, to use first-principle reasoning to derive a predictive game theory of real-world humans (i.e., to settle on a particular choice of the form of \mathcal{S}), we should start with evolutionary game theory and related disciplines. After all, it is in our evolutionary history that the biases in human behavior were generated. So it would be first-principles reasoning in that domain that (hopefully!) could

³¹One obvious variation of this measure of how smart i is to replace the uniform measure in the integral $\int dq'_i \Theta(U^i \cdot [q_i - q'_i])$ with a non-uniform one, for example emphasizing those q'_i having larger dot product with U^i . A related variation would replace the Heaviside function in the integrand with some smooth increasing function, e.g., a logistic function.

be used to derive those behavioral biases. Once one has derived those biases, they would then be incorporated into the likelihood function of PGT. This would be the chain of first-principles reasoning one would ultimately like to follow to derive a fully-justified first-principles PGT. As an alternative, one could short-circuit this process, and use behavior game theory, user-modeling, and the like to set that likelihood empirically. Both approaches are the subject of future work; here we are content to use “introspective” arguments like those of conventional game theory to drive our reasoning.

5 Independent players

5.1 Basic formulation

When the players have never previously knowingly interacted, there is no statistical coupling between the associated mixed strategies, $\{q_i\}$. In this case the setup for coupled players (Sec. 4) does not apply. Instead we must separately specify likelihoods for each of the players. The joint likelihood is then the product of those separate likelihoods.

Here for simplicity we consider a game of complete information, so that every player knows the move spaces and utility functions for all players. Intuitively, those players are not just unaware physical particles, without any “goals” that they are “trying” to achieve. Rather each is a reasoning entity, trying to maximize its own utility, and it knows the same holds for the other players. This results in the “I know that you know that I know that you prefer ...” common knowledge feature that lies at the core of many views of game theory.³²

Intuitively, this overlap in knowledge among the players acts as a “virtual coupling” between the players. However it is not a formal statistical coupling. After all, as mentioned above, $P(\mathcal{S} | q) = \prod_i P(\mathcal{S} | q_i)$ for our independent players invariant \mathcal{S} . Therefore (for an entropic prior) the posterior distributions over mixed strategies are statistically independent:

$$P(q | \mathcal{S}) = \prod_i P(q_i | \mathcal{S}). \quad (49)$$

Given this independence, how do we capture the “virtual coupling”, so crucial to noncooperative game theory, in the independent-players invariant \mathcal{S} ?

To answer this, concentrate on some particular player i . As a surrogate for virtual coupling, say we had a game of actual coupling, as in Sec. 4. That would

³²See [59] for a fine-grained distinction between such “common knowledge” and “mutual knowledge”; such distinctions are not important for current purposes. Also see [60] for related, qualitative discussion.

set up a distribution over the joint moves of the players other than i ,

$$\begin{aligned}
P(x'_{-i} | \mathcal{I}_c) &= \int dx'_i P(x'_i, x'_{-i} | \mathcal{I}_c) \\
&= \int dx'_i dq q(x'_i, x'_{-i}) P(q | \mathcal{I}_c) \\
&= \int dq \left[\int dx'_i q(x'_i, x'_{-i}) \right] P(q | \mathcal{I}_c) \\
&\propto \int dq q_{-i}(x'_{-i}) \prod_j e^{a^i S(q_j)} \delta(q_j \cdot U_q^j - \epsilon_j^i(U_q^j)) \quad (50)
\end{aligned}$$

where the subscript c on the invariant indicates it's the invariant for a counterfactual coupled players scenario, a^i is an associated entropic prior constant for player i , and each ϵ_j^i is an associated Boltzmann utility function, with (implicit) Boltzmann constant b_j^i .

Now if player i makes move x_i , and the remaining players make move x'_{-i} , then the utility for player i is $u^i(x_i, x'_{-i})$. Accordingly, if the distribution over x'_{-i} were actually given by Eq. 50, then the expected utility for player i for making move x_i would be

$$\begin{aligned}
U_c^i(x_i) &\triangleq \int dx'_{-i} u^i(x_i, x'_{-i}) P(x'_{-i} | \mathcal{I}_c) \\
&= \frac{\int dx'_{-i} dq u^i(x_i, x'_{-i}) q_{-i}(x'_{-i}) \prod_j e^{a^i S(q_j)} \delta(q_j \cdot U_q^j - \epsilon_j^i(U_q^j))}{\int dq \prod_j e^{a^i S(q_j)} \delta(q_j \cdot U_q^j - \epsilon_j^i(U_q^j))} \quad (51)
\end{aligned}$$

Note that U_c^i implicitly depends on an associated value a^i , as well as on the values $\{b_j^i\}$ parameterizing the set of functions $\{\epsilon_j^i\}$.

Say that in choosing its move player i assumes that its actual utility U^i is well-approximated by U_c^i for some appropriate a^i and $\{b_j^i\}$. This means that the reasoning of player i reflects the ‘‘I know that you know ...’’ common knowledge feature of game theory; it makes its move under the presumption that the counterfactual coupled players scenario gives a good approximation to its actual environment. With this approach, there is no infinite regress difficulty like that underlying other approaches to the issue of common knowledge. (This reliance on counterfactual coupling to formalize that common knowledge feature can be viewed as an alternative to approaches like Aumann’s epistemic knowledge [61].)

Note (as discussed just below Eq. 24) that $p_{quad} \triangleq \int dq q P(q | \mathcal{I}_c)$ is the Bayes-optimal distribution over joint moves under quadratic loss and the invariant \mathcal{I}_c . So the distribution $P(x_{-i} | \mathcal{I}_c)$ underlying U_c^i is the same as the distribution induced by sampling that single Bayes-optimal distribution. Also recall that $p_{quad}(x)$ is not a product distribution; under it the moves of the players are not statistically independent. So we are modeling every player i as though she achieves a certain performance level for a counterfactual game in which all the players (herself included) make their moves according to the

(coupled) distribution p_{quad} — but in reality she is free to make moves according to a different distribution.

Say that player i makes the perfectly rational move for the counterfactual game. In this situation, player i chooses her moves on the presumption that the other players all behave according to that counterfactual game. The coupling in that counterfactual game can be viewed as how player i 's implements the common knowledge reasoning underlying much of conventional game theory. Our presumption, formalized below, is that while the behavior of player i will not necessarily be perfectly rational for the counterfactual game, that behavior can be approximated as though she is trying to behave that way.

Say that all players i go through the kind of counterfactual reasoning outlined above for associated values of a^i and $\{b_j^i\}$ that do not vary much between them. Then they will all have used very similar distributions $P(x | \mathcal{I}_c)$ to choose their moves. This commonality in their reasoning will not statistically couple their moves; Eq. 49 will still hold. However it will generate the virtual coupling inherent in the “I know that you know ...” feature. Intuitively, it is because they all model the “I know that you know ...” phenomenon in terms of similar statistical coupling scenarios that they are virtually coupled.

Now in practice, no player i will exactly evaluate such a counterfactual coupling scenario to get a guess for U^i (and indeed may not even be able to, for example due to computational limitations). But we can presume that each such player will go through reasoning not too different from such an evaluation, for some particular a^i and $\{b_j^i\}$. Accordingly, as a surrogate for each player i 's actual reasoning, and the associated virtual coupling among all the players, we can stipulate that each player's reasoning results in a mixed strategy q_i that is highly consistent with a counterfactual statistical coupling scenario given by Eq. 51.

To formalize this we must define what it means to have q_i be “highly consistent” with U_c^i . One natural way to do that is by stipulating that $q_i \cdot U_c^i = K_i$ for some parameter K_i , exactly as in the discussion of effective invariants in Sec. 3.2. In other words, we stipulate that $E_{p_{quad}}(u^i)$ be an i -dependent constant. Plugging it in, this definition of “highly consistent” gives us our invariant for player i , i.e., it gives us the likelihood over q for each player i .

In Sec. 5.3 we will replace each K_i with an equivalent parameter β_i that is easier to work with. This parameter will just be the parameter saying how smart q_i is for utility U_c^i , as in Eq. 14 and the associated discussion of effective invariants in Sec. 3.2. To have the notation reflect this alternative parameterization we will sometimes write $K(U_c^i, \beta_i)$ (again, just as in Sec. 3.2). One of the major advantages of parameterizing the i 'th likelihood with β_i rather than K_i is that β_i always ranges from 0 to $+\infty$, for any game, for any player i , and independent of what q_{-i} is. This is not the case for K_i ; its range of values will depend on q_{-i} in general. Intuitively, β_i is simply K_i normalized to account for this.

As mentioned above, since the players are independent, the joint likelihood is the product of the separate individual likelihoods. Using our notation for K_i ,

we can write this likelihood as

$$\begin{aligned}
P(\mathcal{I} | q) &= \prod_i P(\mathcal{I} | q_i) \\
&= \prod_i \delta(q_i \cdot U_c^i - K(U_c^i, \beta_i)), \tag{52}
\end{aligned}$$

with each U_c^i given by Eq. 51.

Comparing this with Eq. 32, and recalling that $\epsilon_i(U_q^i) = K(U_q^i, b_i)$, we arrive at an alternative motivation for the choice of Eq. 52 for the independent-players likelihood. Our presumption for the independent players scenario is that each player is coupled to an environment in the exact same way as in the coupled players scenario, via the function ϵ_i for some appropriate Boltzmann exponent (labeled β_i for the independent players scenario, and b_i for the coupled players scenario). However in the coupled players scenario the environment of each player i is set by the actual q_{-i} . In contrast, in the independent players scenario, each player i 's environment is set by a counterfactual q_{-i} . Intuitively, we are presuming that each player i acts just as *we* do, when we make predictions for a coupled players scenario.

Plugging in, the posterior for the independent players scenario is given by

$$\begin{aligned}
P(q | \mathcal{I}) &\propto e^{\alpha S(q)} P(\mathcal{I} | q) \\
&= \prod_i e^{\alpha S(q_i)} \delta(q_i \cdot U_c^i - K(U_c^i, \beta_i)). \tag{53}
\end{aligned}$$

Plugging Eq. 51 into this result, we get the posterior probability over q for independent players:

$$\begin{aligned}
P(q | \mathcal{I}) &\propto \prod_i e^{\alpha S(q_i)} \delta[K(U_c^i, \beta_i) - \\
& q_i \cdot \frac{\int dx_{-i} dq' u^i(\cdot, x_{-i}) q'_{-i}(x_{-i}) \prod_j e^{\alpha^i S(q'_j)} \delta(q'_j \cdot U_{q'}^j - \epsilon_j^i(U_{q'}^j))}{\int dq' \prod_j e^{\alpha^i S(q'_j)} \delta(q'_j \cdot U_{q'}^j - \epsilon_j^i(U_{q'}^j))}]. \tag{54}
\end{aligned}$$

Next we plug in the usual coupled players ϵ_j^i :

$$\begin{aligned}
P(q | \mathcal{I}) &\propto \prod_i e^{\alpha S(q_i)} \delta[K(U_c^i, \beta_i) - \\
& q_i \cdot \frac{\int dx_{-i} dq' u^i(\cdot, x_{-i}) q'_{-i}(x_{-i}) \prod_j e^{\alpha^i S(q'_j)} \delta(q'_j \cdot U_{q'}^j - K(U_{q'}^j, b_j^i))}{\int dq' \prod_j e^{\alpha^i S(q'_j)} \delta(q'_j \cdot U_{q'}^j - K(U_{q'}^j, b_j^i))}]. \tag{55}
\end{aligned}$$

where the K function is as defined in Sec. 3.2 with U_c^i given by Eq. 51 and parameterized by a^i and the set of values $\{b_j^i\}$. As usual, the posterior over x is given by $\int dq q(x) P(q | \mathcal{I})$ and is identical to the Bayes-optimal q under a quadratic loss function.

Intuitively, for fixed i , the $\{b_j^i\}$ are how smart player i imputes the other players in the counterfactual game to be, which she uses to encapsulate the common knowledge aspect of the game. So it encapsulates how she thinks the other players will choose their moves. In particular, she presumes that in formulating their mixed strategies, the other players will consider how smart she is to be b_i^i . Player i then uses a^i to set the relative probabilities of the q 's that are all consistent with those $\{b_j^i\}$. More carefully, a^i and the $\{b_j^i\}$ serve as our presumptions of the values of these quantities inherent to player i . Properly speaking, we do not really presume that she explicitly has such quantities and uses them to calculate a counterfactual game. Rather we presume that her behavior can be well-approximated by such a common-knowledge type of reasoning by her.

In contrast, β_i reflects our assessment of how well player i carries out such reasoning. It measures how smart we believe she is in evaluating the counterfactual game, and even the degree to which that game really guides her choice of move. α then controls the relative probabilities of the q 's that are all consistent with our assessment of a^i and the $\{b_j^i\}$ for all players i .

Note that since $P(q | \mathcal{I})$ is a product distribution, if $P(q_i | \mathcal{I})$ changes, there is no effect on $P(q_{j \neq i} | \mathcal{I})$. This is true even if the players are all fully rational, both in the actual game and the counterfactual game, so that with probability 1 the system is at a Nash equilibrium of the original game. Accordingly, issues like whether such a Nash equilibrium is “stable” do not occur in the independent players scenario. Any changes to (the distribution governing) player i 's mixed strategy has no effect on the (distribution governing) the mixed strategy of any other player j . This is because such a player j is playing best-response to the counter-factual game, not to the actual game.

5.2 Independent players and the impossibility of a Nash equilibrium

Since $P(q | \mathcal{I})$ for independent players is a product,

$$\begin{aligned} P(x | \mathcal{I}) &= \int dq q(x)P(q | \mathcal{I}) \\ &= \prod_i \int dq_i q_i(x_i)P(q_i | \mathcal{I}) \end{aligned} \quad (56)$$

So our estimate of the joint distribution over moves, $P(x | \mathcal{I})$, is a product distribution. This contrasts with the coupled players scenario (see Sec. 4.3).

However just like in the coupled players scenario, in general the distribution $P(x | \mathcal{I})$ need not be a Nash equilibrium, even if the players are all fully rational. This can be the case even if the players all agree exactly on the counter-factual game, if that counter-factual game is one in which everyone is perfectly rational, and if that game has a unique Nash equilibrium q , one that all our (real game) players predict. In particular, this can be true in a game with a single Nash

equilibrium where it is *not* the case that at the Nash equilibrium each player’s mixed strategy is uniform over its support.³³

To see this, first recall that each player i considers the counterfactual game to predict what the other players are doing in the actual game. Doing this gives i a set of moves that she could make, all of which are best-response. Now by symmetry, *our* estimate of i ’s distribution, $P(x_i \mid \mathcal{I})$, is uniform over those counterfactual game best-response moves, and zero elsewhere. Intuitively, this reflects the fact that we have no reason to presume that i would prefer one over the other of the moves all of which she has concluded would maximize her expected utility.³⁴

This uniformity will hold for all players. Therefore the estimate we make of the distribution over joint moves, $P(x \mid \mathcal{I})$, is uniform over its support. This means it is *not* the Nash equilibrium of the game (under which some players have mixed strategies that are non-uniform over their support). Recall that for quadratic loss, our Bayes-optimal estimate of the joint mixed strategy q is $P(x \mid \mathcal{I})$. So our result is equivalent to saying that the Bayes-optimal q is not a Nash equilibrium.

Example: Consider a 2-player game where each player has move space $\{A, B\}$. Have the row player’s utility function be $u_R(A, A) = -1, u_R(A, B) = u_R(B, A) = 0, u_R(B, B) = -2$, and the column player’s utility function be $u_C(A, A) = 1, u_C(A, B) = u_C(B, A) = 0, u_C(B, B) = 2$. Then the only Nash equilibrium is the mixed strategy $q_R(A) = q_C(A) = 2/3$.

The row player would consider a counterfactual version of this game. Doing so, she would conclude (since she presumes all the players in the counterfactual game are perfectly rational) that in the counterfactual game only one q has non-zero probability, namely the mixed strategy $q_R(A) = q_C(A) = 2/3$. She would presume that the other player (the column player) in the real game uses the associated mixed strategy of that counterfactual Nash equilibrium. Given this presumption of hers for the actual q_C she faces, she would conclude that both of her possible moves (A and B) maximize her expected utility, since that is the case in the counterfactual game. However she has no reason to prefer one move over the other. So *we* would estimate that her distribution is uniform.

Doing this for both players, we would conclude that the distribution over joint moves is uniform over all possible x . Such a distribution is not a Nash equilibrium.

This result does not mean that we claim that the Nash equilibrium q is impossible. We assign non-zero $P(q \mid \mathcal{I})$ to that Nash equilibrium q in general.

³³Some have worried that this type of scenario calls into question the validity of the Nash equilibrium concept. The issue is why a player i should play a particular non-uniform mixed strategy over its best response pure strategies, when the only “advantage” of that mixed strategy is that it happens to make the mixed strategies of other players be best-response. See for example [62, 63, 64].

³⁴It is interesting to consider this result in light of experimental and theoretical work concerning risk-dominant Nash equilibria.

It is just that our estimate of the distribution over moves will not be that Nash equilibrium.

This contrasts with the coupled players scenario. In that scenario, if you are explicitly provided \mathcal{S} saying that all players are perfectly rational, then it is precisely that \mathcal{S} that tells you that player i must play the Nash equilibrium non-uniform distribution. If you are *not* provided that explicit prior information — as you are not in the independent players scenario — then in fact you should not assume that there is perfect rationality.

5.3 The MAP q for independent players

Since for independent players the posterior is a product distribution, the MAP q is also. So with some abuse of notation, we can write

$$\begin{aligned}
\text{MAP}(q) &\triangleq \operatorname{argmax}_q P(q \mid \mathcal{S}) \\
&= \operatorname{argmax}_q \prod_i e^{\alpha S(q_i)} \delta(q_i \cdot U_c^i - K(U_c^i, \beta_i)) \\
&= \prod_i \operatorname{argmax}_{q_i} P(q_i \mid \mathcal{S}) \\
&= \prod_i \text{MAP}(q_i)
\end{aligned} \tag{57}$$

where the index variable $x = (x_1, x_2, \dots)$ is implicit, as is the conditioning on the independent-players \mathcal{S} . For notational simplicity define

$$\text{MAP}(q_i) \triangleq \tilde{q}_i \tag{58}$$

for each i , So we can rewrite Eq. 57 as $\tilde{q} = \prod_i \tilde{q}_i$.

In the usual way, by maximizing entropy subject to the associated equality constraint, each \tilde{q}_i can be written as $e^{\beta_i U_c^i(x_i)}$ up to an overall proportionality constant. Recall that in writing \tilde{q}_i this way that β_i is the Lagrange parameter enforcing our constraint that $U_c^i \cdot \tilde{q}_i = K_i$, i.e., enforcing our restriction that the \tilde{q}_i be “well-consistent” with U_c^i . Writing it out,

$$\begin{aligned}
U_c^i \cdot \tilde{q}_i &= K(U_c^i, \beta_i) \\
&= \frac{\int dx_i e^{\beta_i U_c^i(x_i)} U_c^i(x_i)}{\int dx_i e^{\beta_i U_c^i(x_i)}}.
\end{aligned} \tag{59}$$

Given this form for each \tilde{q}_i , we can write the value of \tilde{q} for some arbitrary x as

$$\tilde{q}(x) \propto \prod_i e^{\beta_i U_c^i(x_i)}. \tag{60}$$

where the proportionality constant is independent of x . Plugging in, this becomes

$$\tilde{q}(x) \propto \prod_i \exp[\beta_i' \int dx_{-i} dq u^i(x_i, x_{-i}) q_{-i}(x_{-i}) \prod_{j \neq i} e^{\alpha^j S(q_j)} \delta(q_j \cdot U_{q_{-j}}^j - \epsilon_j^i(U_{q_{-j}}^j))], \tag{61}$$

where for simplicity we have absorbed all proportionality constants into β_i , writing that new value of β_i as β'_i .

As an example, say that our game has a single Nash equilibrium over pure strategies, x^* . Let the b_j^i (implicit in the ϵ_j^i) all go to infinity, keeping the a^i all finite, in such a way that the posterior distribution over q for the counterfactual coupled game approaches a single q which is a delta function about x^* . So $U_c^i(\cdot)$ approaches $u^i(\cdot, x_{-i}^*)$. Then \tilde{q} approaches a product of (independent) Boltzmann distributions:

$$\tilde{q}(x) \propto \prod_i e^{\beta_i u^i(x_i, x_{-i}^*)}. \quad (62)$$

This is a product of mixed strategies, each of the form of the Boltzmann distribution. As such it is similar to the QRE. Unlike the QRE though, there is no coupling between the different mixed strategies comprising \tilde{q} . This reflects the fact that, by hypothesis, the players are independent of each other in how they form their mixed strategies, as well as in the subsequent moves they make. Whenever there is such independence — which is the case in much of conventional noncooperative game theory, implicitly or otherwise — the QRE is not an appropriate choice for what kind of product of Boltzmann distributions to use to capture bounded rationality.

Now say the counterfactual game has two pure strategy Nash equilibria, $x^*(1)$ and $x^*(2)$, and that in evaluating the counterfactual game agent i gives them probabilities c_i and $1 - c_i$, respectively. Then rather than Eq. 62, we get

$$\tilde{q}(x) \propto \left[\prod_i e^{\beta_i c_i u^i(x_i, x_{-i}^*(1))} \right] \times \left[\prod_i e^{\beta_i (1-c_i) u^i(x_i, x_{-i}^*(2))} \right], \quad (63)$$

i.e., a product of the kind of equilibria arising for the two Nash equilibria taken separately. If $\beta_i \rightarrow \infty$, then agent i chooses the best response to either $x_{-i}^*(1)$ or $x_{-i}^*(2)$, depending on which gives i higher expected payoff (where the expectation is evaluated according to the distribution $(c_i, 1 - c_i)$).

6 Miscellaneous topics

This section presents some illustrative extensions of the basic PGT framework presented above.

6.1 Cost of computation

For a large range of games, the independent players scenario results in a tradeoff between how smart a player is and the cost of the computation they must engage in to determine their behavior. This relation between the cost of computation and bounded rationality emerges from the mathematics; it is not some *ad hoc* hypothesis we make to explain the observed (bounded rational) behavior of real human beings. In addition, using this mathematics, we can quantify the tradeoff and when it occurs, and more generally determine what characteristics of the

game are most intimately related to the tradeoff. (All of that analysis is the subject of future work.)

Say β_i increases while all other parameters are fixed, so U_c^i doesn't change. Then the set of q_i satisfying our invariant \mathcal{S} shifts (cf. Sec. 3.2). Typically such shifts in that set arising from increases in β_i also shrink that set (i.e., its measure decreases). Intuitively, the smarter player i is (for the counterfactual game), the more assured it is in assessment of the counterfactual game, and therefore the more assured it is in making its move. As an example, say that a^i and the values $\{b_j^i\}$ restrict $P(q | \mathcal{S}_c)$ to one q that is a Nash equilibrium of the game, an equilibrium which is a joint pure strategy of the players. So $P(x_{-i} | \mathcal{S}_c)$ is a delta function about the moves of the players other than i at that Nash equilibrium. Then for $\beta_i \rightarrow \infty$, q_i also becomes restricted to that equilibrium, i.e., the support of the likelihood, $P(\mathcal{S} | q_i)$ gets restricted to a single q_i (one that is a delta function about that Nash equilibrium's x_i). Accordingly the measure of q_i allowed by the likelihood goes to 0 as β_i approaches infinity.

When the set of q_i allowed by the likelihood shrinks this way, the set of q_i allowed by the associated posterior, $P(q_i | \mathcal{S})$ (i.e., the set of q_i in the support of that posterior) must also shrink. Typically this mean that the entropy of that posterior shrinks. Usually this in turn means that the integral of that posterior, $P(x_i | \mathcal{S}) = \int dq_i q_i(x_i)P(q_i | \mathcal{S})$, also get a smaller entropy as β_i increases. We can illustrate this by returning to our single pure strategy Nash equilibrium example. In that example, for $\beta_i \rightarrow \infty$, the support of $P(x_i | \mathcal{S})$ gets restricted to the Nash equilibrium x_i , and therefore its entropy goes to zero, the smallest possible value. As another example, recall from Sec. 3.2 that since we have fixed U_c^i , the entropy of the MAP q_i cannot increase as β_i increases.

In such situations, all these distributions with decreasing entropy have more and more information as β_i increases (recall that the amount of information in a distribution is the negative of its entropy). Now model agent i 's computational process (in deciding how to move) as starting with the assumption that $a^i, \{b_j^i\}$ accurately describes the other agents, so that the associated counterfactual game results in an accurate approximation of U^i . Under this model, we can interpret the amount of information in $P(x_i | \mathcal{S})$ as the amount of "computational effort" i expends to try to approximate $P(x_{-i} | \mathcal{S}_c)$ accurately and guess accordingly.

As just argued, typically that amount of information in $P(x_i | \mathcal{S}_c)$ — the negative of its entropy — increases as β_i does. So under this model, the larger β_i is, the more computational effort i expends. On the other hand, assume that the $a^i, \{b_j^i\}$ going into i 's counterfactual game calculation give an accurate approximation to the actual U^i . In this case, the expected payoff to i rises as β_i does. So when the $a^i, \{b_j^i\}$ give an accurate approximation to U^i (i.e., i 's modeling is accurate), rising β_i both means more expected payoff to i and more computational effort by i . Evidently β_i controls a tradeoff between how smart i is and how much computational effort it expends.³⁵

³⁵The analogous argument for the coupled players scenario is more problematic. This is because as i changes her distribution, for example by increasing her (coupled players value) b_i , the distribution of the other players must also change, due to the coupling between players. This means that the effect on the entropy of i 's distribution and to her expected payoff can

6.2 Rationality functions

In many situations it would be useful to have a way of quantifying the rationality of a player i , based purely on its behavior, without any model of its decision-making process (even as ill-specified a model as saying that the player “evaluates a counterfactual game to some given degree of accuracy”). We would like to be able to do this for any mixed strategy q_i and for any environment U^i (whether that mixed strategy is the choice of player i , as in type II games, or instead governs how i makes choice, as in type I games). We would like similar generality for judging potential moves x_i .

In particular, we do not want to *require* that the mixed strategy of real-world players has some *a priori*-specified parameterized form, e.g., a Boltzmann distribution over its environment. We do not want to assume that our data is a (perhaps noise-corrupted) stochastic realization of such a mixed strategy, and accordingly solve for the best-fit values of the associated parameters to some experimental data (as is done in much of the experimental work involving the QRE, e.g., [26]). After all, any requirement that the mixed strategy of a real-world player is *exactly* given by such a parametric function will almost always be in error, at least to a degree. This section presents such a broader quantification of rationality.

Consider the situation where players i has mixed strategy q_i and her environment is some fixed U^i . It is reasonable to say that two choices of q^i are equally rational if they have the same dot product with U^i . However we will often want to do more than simply say whether two q_i are equally rational for some particular U^i ; we will often want to say whether a q_i operating in environment U^i is more or less rational than a q'_i operating in environment $(U')^i$. To do this we need a scalar-valued function $R(V, p)$ that measures how rational an arbitrary distribution $p(y)$ is for an arbitrary utility function $V(y)$, i.e., that measures how peaked $p(y)$ is about the maximizers of $V(y)$, $\operatorname{argmax}_y V(y)$, and about the other y that have large $V(y)$ values.

Say that p is a Boltzmann distribution over $V(y)$, $p(y) \propto e^{\beta V(y)}$. Then we can use information theory in general, and effective invariants and the functions ϵ_i discussed above in particular, to motivate quantifying the rationality of p for V as the value β . The larger β is, the more peaked p is about the better mixed strategies, and therefore the more “rational” p is.

In addition, so long as p'' and p' are Boltzmann distributions for V'' and V' respectively, this measure of the associated β value can be used to compare the rationality of p'' for V'' with the rationality of p' for V' . We can do this even if the range of the function V' differs from that of V'' . This attribute of our measures differs from other naive choices for measuring rationality. In particular, it differs from the choice of measuring rationality as $p \cdot V$, which not only reflects how peaked p is about y that give large $V(y)$, but also reflects the range of values of $V(\cdot)$. (Indeed, simply translating the values of $V(\cdot)$ by a constant will modify the value of this alternative choice of rationality function.)

be more complicated.

In general though p will not be a Boltzmann distribution. So we need to extend our reasoning, to define an R that we can reasonably view as a quantifier of rationality for any p . Formally, we make two requirements of R :

1. If $p(y) \propto e^{\beta V(y)}$, for non-negative β , then the peakedness of the distribution — the value of $R(V, p)$ — is β .
2. Out of all p satisfying $R(V, p) = \beta$, the one that has maximal entropy is proportional to $e^{-\beta V(y)}$. In other words, we require that the Boltzmann distribution maximizes entropy subject to a provided value of the rationality/temperature.

We call any such R a **rationality function**.

Note that a rationality function can be applied to physical systems, where $V(y)$ is interpreted as the Hamiltonian over microstates y . Such a function is defined even for systems that are not at physical equilibrium (and therefore aren't described by Boltzmann distributions). In this, rationality functions are an extension of the conventional definition of temperature in statistical physics.

As an illustration, a natural choice is to define $R(V, p)$ to be the β of the Boltzmann distribution that “best fits” p . To formalize this we must quantify how well any given Boltzmann distribution “fits” any given p . Information theory provides many measures for how well a distribution p_1 is fit by a distribution p_2 . On such measure is the **Kullback-Leibler distance** [8, 65, 51]:

$$KL(p_1 \parallel p_2) \triangleq S(p_1 \parallel p_2) - S(p_1) \tag{64}$$

where $S(p_1 \parallel p_2) \triangleq - \int dy p_1(y) \ln[\frac{p_2(y)}{\mu(y)}]$ is known as the **cross entropy** from p_1 to p_2 (and as usual we implicitly choose uniform μ).

The KL distance is always non-negative, and equals zero iff its two arguments are identical. In addition, $KL(\alpha p^1 + (1 - \alpha)p^2 \parallel p^2)$ is an increasing function of $\alpha \in [0.0, 1.0]$, i.e., as one moves along the line from p^1 to p^2 , the KL distance from p^1 to p^2 shrinks.³⁶ The same is true for $KL(p^2 \parallel \alpha p^1 + (1 - \alpha)p^2)$. In addition, those two KL distances are identical to 2nd order about $\alpha = 0$. However they differ as one moves away from $\alpha = 0$ in general; KL distance is not a symmetric function of its arguments. In addition, it does not obey the triangle inequality, although it obeys a variant [8]. Despite these shortcomings, it is by far the most common way to measure the distance between two distributions.

Recall the definition of the partition function, $Z(V) \triangleq \int dy e^{V(y)}$ (the normalization constant for the distribution proportional to $e^{V(y)}$). Using the KL

³⁶This follows from the fact that the second derivative with respect to α is non-negative for all α , combined with the fact that KL distance is never negative and equals 0 when $\alpha = 0$.

distance and this definition, we arrive at the rationality function

$$\begin{aligned}
R_{KL}(V, p) &\triangleq \operatorname{argmin}_{\beta} KL(p \parallel \frac{e^{\beta V}}{Z(\beta V)}) \\
&= \operatorname{argmin}_{\beta} [-\beta \int dy p(y) V(y) + \ln(Z(\beta V)) - S(p)] \\
&= \operatorname{argmax}_{\beta} [\beta \int dy p(y) V(y) - \ln(Z(\beta V))]. \tag{65}
\end{aligned}$$

In [29] it is proven that R_{KL} respects the two requirements of rationality functions. Note that the argument of the argmin is globally convex (as a function of the minimizing variable β). In addition its second derivative is given by the variance (over y) of the Boltzmann distribution $e^{\beta V(y)}/Z(\beta V)$. This typically makes numerical evaluation of R_{KL} quite fast.

Comparing the definition of R_{KL} to Eq. 20, we see that the KL rationality of a distribution p is just the value of β for which p has minimal free utility gap. When p is a Boltzmann distribution over the states of a statistical physics systems, this β is (the reciprocal of) what is called temperature in the in App. 9. Systems described by such distributions are at physical equilibrium. In other words, the physical temperature of a physical system at physical equilibrium is (the reciprocal of) its KL rationality. KL rationality is also defined for off-equilibrium systems however, unlike physical temperature.

To help understand the intuitive meaning of the KL rationality function, consider fixing its value for agent i to some value ρ_i . Say q_{-i} is also fixed (and therefore so is player i 's environment, $U_{q_{-i}}^i$). Then there is a value a_i such that the set of all q_i having rationality value ρ_i is identical to the set of all q_i for which $E_{q_i}(U_{q_{-i}}^i) = a_i$. In fact, a_i is the expected value of U^i that would arise if $q_i(x_i)$ were a Boltzmann distribution (over $U_{q_{-i}}^i(x_i)$ values) with Boltzmann exponent $\beta_i = \rho_i$.³⁷

So knowing that player i has KL rationality ρ_i is equivalent to knowing that the actual expected value of U^i under q_i equals the ‘‘ideal expected value’’, in which q_i is replaced by the Boltzmann distribution over $U_{q_{-i}}^i(x_i)$ values with exponent $\beta_i = \rho_i$. (However note that such a constraint on the value of ρ_i does not actually specify q_{-i} , so it does not specify that ideal expected value of U^i .) The (loose) physical analog of this result is that all distributions over states of

³⁷To see all this, note that by definition of KL rationality function,

$$-\frac{\partial \ln(Z(\beta U_{q_{-i}}^i))}{\partial \beta} \Big|_{\beta=R_{KL}(U_{q_{-i}}^i, q_i)} = \int dx_i q_i(x_i) U_{q_{-i}}^i(x_i).$$

However by the discussion in Sec. 3.2, we know that the quantity on the left-hand side is just the Boltzmann utility evaluated at the specified value of β , $K_i(\beta)$. So $R_{KL}(U_{q_{-i}}^i, q_i) = R_{KL}(U_{q_{-i}}^i, q'_i) \Rightarrow K(R_{KL}(U_{q_{-i}}^i, q_i)) = K(R_{KL}(U_{q_{-i}}^i, q'_i)) \Leftrightarrow E_{q_i}(U_{q_{-i}}^i) = E_{q'_i}(U_{q_{-i}}^i)$. So any two q_i 's with the same rationality must have the same expected $U_{q_{-i}}^i$. To prove the other direction, recall that for fixed $U_{q_{-i}}^i$, the Boltzmann utility is a bijection from values of β into \mathbb{R} . **QED.**

a physical system having the same (potentially non-equilibrium) temperature also have the same expected value of the Hamiltonian.

Comparing with the discussion in Sec. 4, we see that specifying the KL rationalities of all the players is exactly the same as specifying that they all obey the coupled players invariant, with the parameters of the functions ϵ_i given by those specified rationality values. An \mathcal{S} specifying the one scenario is identical to an \mathcal{S} specifying the other one. Accordingly, all the discussion in Sec.'s 4.1, 4.8 holds for making predictions based on specified rationalities of the players. In particular, as discussed in Sec. 4.1, the rationalities of the players in a game reflects the structure of that game, as much as it reflects the intrinsic characteristics of the players.

All of the foregoing was for quantifying the rationality of a particular q_i . However we can view the rationality of a particular x_i as a special case, where the “mixed strategy” q_i is a delta function about one of its moves. (*behaviorally*, it makes no difference if that x_i is a sample of some preceding q_i that i chose, or instead is i 's choice directly.) Plugging that in to the KL rationality function, we get the following definition of the rationality of a move x_i :

Say that a player i makes move x_i when there is an environment U^i . Then the KL rationality of that move is the β such that if i had instead chosen a Boltzmann mixed strategy with exponent β , the resultant expected value of u^i would have been the same as i 's actual expected utility. Formally, the KL rationality function is the mapping from (x_i, U^i) to the β such that

$$U^i(x_i) = \frac{\int dx'_i U^i(x'_i) e^{\beta U^i(x'_i)}}{\int dx'_i e^{\beta U^i(x'_i)}}. \quad (66)$$

6.3 Variable numbers of players

There are many statistical ensembles considered in statistical physics in addition to the CE. In particular, in the Grand Canonical Ensemble (GCE), the numbers of the particles of various types in the system is itself a stochastic quantity, in addition to the states of those particles. This is how one analyzes the statistics of physical systems involving chemical and/or particle physics interactions that change the particles of the system.

Recall that the CE can most cleanly be derived as an MAP distribution with an entropic prior and an appropriate expectation value constraint (App. 9). The GCE can be derived the same way. Whereas with the CE the expectation value constraint only concerns the expected energy, in the GCE it also concerns the expected numbers of particles of the various possible types [19].

As pointed out in [11, 29, 25], the same same approach used in the GCE can also be applied in a game theory context. In such a context, rather than “particles of various types”, one has “players of various types”. Broadly speaking, after this substitution, the ensuing analysis for the game theory context proceeds analogously to that of the statistical physics context.

To illustrate this we present a game theory scenario that roughly parallels the

GCE.³⁸ We postulate some pre-fixed set of **player types**. All players of a given type have the same move space and the same payoff function. At the beginning of each instance of our scenario, a set of players is randomly chosen, and each is assigned a rationality value randomly. Those players are then coupled as discussed above in Sec. 4, e.g., via a sequence of noncooperative games, and the instance ends with all of the players making a move.

We know that the expected number of players of any one of the player types is the same from one instance to the next, although we do not necessarily know that expectation value. We similarly assume the expected rationality for each player type (i.e., the expect value of b_i , in the terminology of Sec. 4.2) is the same from one instance to the next, without necessarily knowing those rationality values. These rationality values are statistically independent from each other.

We formalize this with an encoding of our variables into x modeled on the scheme used to derive the GCE [19]. For all player types i , x_i^N indicates the number of players of that type. For all integers $j > 0$, and all player types i , $x_{i,j}^M$ indicates the move of the j 'th player of type i , assuming there is such a player (i.e., assuming that $j \leq x_i^N$). The meaning of $x_{i,j}^M$ for larger j is undefined/irrelevant. Similarly $x_{i,j}^R$ indicates the rationality of the j 'th player of type i , assuming there is such a player, and is undefined otherwise.

We write x^N , x^M , and x^R , respectively to indicate the vector of all player-type cardinalities, the (countably infinite dimensional) vector of the moves by all possible players (including those that do not actually exist), and the (countably infinite dimensional) vector of the rationalities of all possible players (including those that do not actually exist). We also write the utility function of the type i players as $g_i(x^M, x^N)$, where $\forall i$, $g_i(x^M, x^N)$ is independent of $x_{k,j}^M \forall j > x_k^N$. Finally, we write \bar{N}_i and \bar{R}_i to indicate the (fixed but potentially unknown) expected number of players of type i and expected rationality of those players, respectively.

As in Sec. 4.2, the moves of our players are independent once the characteristics of the game are fixed (i.e., we are dealing with a conventional noncooperative game in which the moves are given by sampling an associated joint mixed strategy). However here the moves can be statistically dependent on those characteristics. For example, if the rationality $x_{i,j}^R = 0$ for some $j < x_i^N$, then we know that $q_{i,j}^M$ must be uniform, independent of the mixed strategies of the other players.

Reflecting this, we write

$$q(x) \triangleq \prod_{i,j} [q_i^N(x_i^N) q_{i,j}^R(x_{i,j}^R) \prod_{i',j'} q_{i',j'}^M(x_{i',j'}^M | x^N, x^R)] \quad (67)$$

³⁸One difference is that the GCE allows arbitrary statistical coupling between all variables. In contrast, here we impose numerous statistical independences among the variables, e.g., statistical independence between the moves of the players. Another difference is that there are multiple utility functions in games, whereas there is only analogous quantity (the Hamiltonian) in physical systems. This makes the formulas here more complicated than those in the GCE.

where the products over j and j' both run from 1 to ∞ . When the argument makes clear what the superscript $\{M, N, R\}$ should be, we will sometimes leave that superscript implicit. Note that in reflection of the statistical coupling of the components of x , q is *not* a product distribution. So in particular the entropy of q is not a sum of the entropies of its marginalizations, as it was above.

Writing it out,

$$q(x_{i,j}^M | x^N, x^R) = \frac{\int dx' q(x') \delta(x'_{i,j}^M - x_{i,j}^M) \delta(x'^R - x^R) \delta(x'^N - x^N)}{\int dx' q(x') \delta(x'^R - x^R) \delta(x'^N - x^N)}. \quad (68)$$

With some abuse of notation, we will write “ $q_{i,j}^M(\cdot | x^N, x^R)$ ” to mean the (infinite-dimensional) vector with component $x_{i,j}^M$ given by $q(x_{i,j}^M | x^N, x^R)$.

Our invariant says that each q_i^N must result in an average of x_i^N that equals \bar{N}_i , and similarly for each q_i^R and \bar{R}_i . It also says that once x^N and x^R are fixed, q^M must be the joint mixed strategy appropriate for an associated coupled players type II game. To write out this latter condition, first define “ $-(i, j)$ ” to mean all players other than (i, j) (including players of type i other than the j 'th one of that type). Next as shorthand we will often take the distribution over all agents other than (i, j) implicit and write

$$\begin{aligned} U^{i,j}(x_{i,j}^M, x^R, x^N) &\triangleq U_{q_{-(i,j)}^M(\cdot | x^R, x^N)}^{i,j}(x_{i,j}^M, x^N) \\ &\triangleq \int dx_{-(i,j)} q(x_{-(i,j)}^M | x^R, x^N) g^i(x_{i,j}^M, x_{-(i,j)}^M, x^N) \end{aligned} \quad (69)$$

where we will write $U^{i,j}(\cdot, x^R, x^N)$ to mean the (infinite-dimensional) vector with component $x_{i,j}^M$ given by $U^{i,j}(x_{i,j}^M, x^R, x^N)$. So the coupled players portion of our invariant says that³⁹

$$q_{i,j}^M(\cdot | x^R, x^N) \cdot U^{i,j}(\cdot, x^N, x^R) = K(U^{i,j}(\cdot, x^N, x^R), x_{i,j}^R) \quad \forall i, j \quad (70)$$

Combining these three separate aspects of the invariant and explicitly expanding in full each instance that a component of q occurs, we get

$$\begin{aligned} P(\mathcal{S} | q) &\triangleq \prod_i \left[\delta(\bar{N}_i - \int dx_i^N q^N(x_i^N) x_i^N) \delta(\bar{R}_i - \int dx_i^R q^R(x_i^R) x_i^R) \times \right. \\ &\quad \prod_j \int dx^N dx^R q^N(x^N) q^R(x^R) \times \\ &\quad \left. \delta(q_{i,j}^M(\cdot | x^R, x^N) \cdot U_{q_{-(i,j)}^M(\cdot | x^R, x^N)}^{i,j}(\cdot, x^N) - \right. \\ &\quad \left. K(U_{q_{-(i,j)}^M(\cdot | x^R, x^N)}^{i,j}(\cdot, x^N), x_{i,j}^R)) \right]. \end{aligned} \quad (71)$$

³⁹Unfortunately, even with this abusive notation, book-keeping in the equations can get messy.

We then combine this likelihood with an entropic prior over q . This gives us the posterior $P(q | \mathcal{S})$. As usual, if we wish to we can consider the MAP q according to this posterior, various Bayes-optimal q 's according to this posterior, etc., thereby getting a single distribution over x 's.

Again just like in the usual analysis, as an alternative to these distributions, over x 's we can simply write $P(x | \mathcal{S})$ directly, getting the same answer as the Bayes-optimal q under quadratic loss:

$$P(x | \mathcal{S}) = \int dq P(q | \mathcal{S})q(x). \quad (72)$$

To evaluate this integral we must use Eq. 67 to plug in for $q(x)$, Eq. 71 for the likelihood, and then use the usual entropic prior. Also as usual we must be careful to calculate the normalization constant for the posterior $P(q | \mathcal{S})$ and divide that into the product of the likelihood and prior.

However arrived at, once we get a distribution over x , we can then marginalize over various components of x to get distributions over the associated quantities of interest. For example, we can do this to determine the typical move of a player of a particular type, the typical number of players of some type conditional on a particular move made by the first player of that type, etc..

7 Discussion and Future Work

It is worth comparing PGT to approaches based on models of actual humans beings, like those using models of agent learning [66] or models incorporating the mathematical structure of statistical physics [67, 68]. Broadly speaking, PGT's motivation is more like that of conventional game theory than that of model-based approaches. Like conventional game theory, PGT investigates what can be gleaned by careful consideration of the abstract problem of interacting goal-directed agents, before the introduction of experiment-based insight concerning the behavior of those agents.

An even closer analogy to PGT's motivation than that provided by conventional game theory is Bayesian statistics, and especially Bayesian statistics using invariance-based arguments to set the prior [1]. Like such Bayesian statistics, PGT is a first-principles-driven derivation of a framework for analyzing systems, a framework into which one can "slot in" any kind of experimental data as it becomes available."

While the extraordinary success of statistical physics has been used to choose the entropic prior for this paper, it is important to emphasize that many other priors can also be motivated using first-principles arguments, many of them also based on information-theoretic arguments. Similarly, many other choices of likelihood (the invariant) can be motivated (as discussed above). PGT is *not* restricted to the prior and likelihood considered in this paper, any more than conventional game theory is restricted to some particular refinement of the Nash equilibrium concept. The defining characteristics of PGT is the application of such priors and likelihoods to game outcomes rather than (or in addition to)

within games. The prior and likelihoods considered here are simply the examples worked out in this initial paper.

Obviously, if you happen to know what algorithm the players are using, then that should be reflected in the likelihood. PGT for various simple choices of such algorithms/likelihoods is the subject of future work. More generally, humans have lots of cognitive quirks presumably arising due to evolution. Accordingly the precise priors and likelihood investigated here may work best for computational agents involved in a game with no foreknowledge of the game. Important future work involves analysis with other priors and likelihoods incorporating behavioral economics results, prospect theory, etc.. These alternatives can be used for the external scientist’s assessment of the individual players and/or (in the independent players scenario) for the “models” the players have of each other.

Indeed, PGT can be seamlessly extended to encompass other kinds of \mathcal{I} , even kinds that do not involve utility functions. In particular, one or more observed samples of a mixed strategy q_i can naturally be incorporated into the likelihood term, $P(q_i | \mathcal{I})$. As another example, we can remove from \mathcal{I} the stipulation that our players’ choices of pure strategy are independent of one another, i.e., the stipulation that we use a product distribution. Doing so naturally results in correlated moves among the players, without any need for carefully designed ansatz’s like those behind correlated equilibria [20].

Similarly, there is a good deal of empirical evidence that human players do not prefer to maximize expected utility functions $\int dx_i q_i(x_i) U^i(x_i)$. Rather a long line of experiments starting with Allais’ paradox [69] indicate that what is invariant in the decision-making of a human i is some non-linear functional of its mixed strategy q_i . As more gets understood about such psychological phenomena [70] it should be straightforward to incorporate that understanding into (Bayesian) PGT. One simply changes what is considered invariant from one instance of the inference problem to the next, from being a linear functional of q_i to being some other type of functional.

Related future work will integrate behavior modeling (“user modeling”, belief nets, etc.) with PGT, to get an empirical science of human interactions. Such behavior modeling can run the gamut from knowledge concerning humans in general (e.g., behavioral economics) to knowledge concerning certain particular humans (psychological profiling, and in particular “games against nature”, i.e., the decision-making belief net of a particular human, in a non-game theory context [71]).

In addition to the foregoing, there is a huge amount of future work in PGT that carries over from conventional game theory. At the risk of being glib, almost every aspect of conventional game theory can be re-analyzed using PGT. This includes in particular cooperative game theory, in which context PGT should cut the Gordian knot of what equilibrium concept to adopt. Other broad topics that should be investigated using PGT — and therefore bounded rationality — are mechanism design, folk theorems, and signaling theory. It may also prove profitable to have such investigations be extended to allow varying number of players. Similarly, “bounded rational” evolutionary game theory, in particular for finite numbers of agents, can be investigated using the “GCE” (variable

number of players) variant of PGT illustrated above. All of this is in addition to more circumscribed game theory issues, like different types of noncooperative games (Bayesian games, correlated equilibrium games, differential games, etc.).

Other future work involves completing the analysis of the relationship between QRE and the coupled players MAP (and Bayes-optimal) q 's. This can also be extended to the independent players. Similarly, coverage issues like those presented in Prop.'s 1 and 2 for the coupled players scenario bears investigating for the independent players scenario.

Other future work will investigate what happens in the variable number of players scenario if the random variable of the number of player of type i is not independent of the random variable of the total utility accrued by all players of that type. One aspect of such an investigation would see what happens if that random variable is statistically coupled to \bar{U}^i/x_i^N , the average, of players of type i , of the expected utility of those players. In particular, it is interesting to see what happens if that variable is coupled to $\frac{\bar{U}^i}{x_i^N \sum_j \bar{U}^j}$, the ratio of total expected utility that is earned by players of type i , divided out among those players.

All of this is in addition to the future work mentioned in the preceding sections.

8 Appendix 1 — Historical context of PGT

Despite its widespread and profound usefulness in other fields, attempts to use Shannon entropy in game theory, psychology, and economics has proven controversial (see for example [72] and references therein). By and large though those attempts have considered Shannon entropy as a physical quantity occurring *within* the system under study, and then tried to relate that physical quantity to other aspects of the system. In contrast, where Shannon entropy has proven so successful in statistical physics, statistics, signal processing, etc., is in guiding the external scientist in his *inference about* the system under study. It is in this latter sense that Shannon entropy is used in PGT.

The results in [11, 29, 25] can be viewed as the first derivation of bounded rational equilibria using full probabilistic reasoning. (The arguments in [27] concerned equilibrium concepts rather than distributions over the space of all possible mixed strategies.) It should be noted though that the maxent Lagrangian has a history far predating both the work in [11, 29, 25] and that in [27]. As the free energy of the CE it has been explored in statistical physics for well over a century. Indeed, the QRE is essentially identical to the “mean field approximation” of statistical physics. (See also [67].)

In the context of game theory, the maxent Lagrangian was given an *ad hoc* justification and investigated in [31, 32, 73] and related work. The first attempt to derive it in that context using first principles reasoning occurred in [30].

The use of the Boltzmann distribution mixed strategies also has a long history in the Reinforcement Learning (RL) community, i.e., for the design of computer algorithms for a player involved in an iterated game with Nature [74, 75].

Related work has considered multiple computational players [76, 77]. In particular, some of that work has been done in the context of “mechanism design” of many computational players, i.e., in the context of designing the utility functions of the players to induce them to maximize social welfare [78, 54, 58, 57]. In all of this RL work the Boltzmann distribution is usually motivated either as an *a priori* reasonable way to trade off exploration and exploitation, as part of Markov Chain Monte Carlo procedure, or by its asymptotic convergence properties [79].

The work in [80, 81, 82] in particular, and econophysics in general, also concerns the relation between statistical physics and the social sciences. In particular, much of that work considers the relation between equilibrium distributions of statistical physics and notions of equilibrium in social science settings. None of it concerns game theory though. To relate that domain to statistical physics one must drill deeper into statistical physics, into its information-theoretic foundations as elaborated by Jaynes. The first relatively simple-minded work relating information theory, statistical physics, and bounded rational game theory this way was [51].

9 Appendix 2 - Using the Entropic Prior to Derive the CE

This appendix elaborates — in a very detailed manner — the application of the entropic prior to statistical physics that results in the **canonical ensemble** (CE). The level of detail presented borders on overkill. However it turns out not to be trivial how to set the analogous details arising in the application of the entropic prior to game theory. In addition, the subtleties of how to use the entropic prior to derive the CE are invariably slighted in the literature. Hence first working through the well-understood statistical physics case can help hone intuition.

On the other hand though, it turns out that in the CE, the temperature T — our prior knowledge concerning the system — equals the Lagrange parameter of a constrained optimization problem, rather than the value of the constraint associated with that Lagrange parameter. This is not the case with PGT, and introduces some subtleties that are mostly absent from PGT. In addition, the PGT analogue of what in the CE is the system’s Hamiltonian function are the players’ utility functions. While there is a single Hamiltonian in the CE, there are multiple utility functions in noncooperative games. Associated with this, in PGT there are multiple analogues of what in the CE is (global) temperature. All of this introduces complications into PGT that are absent from the CE. Due to all this, readers already comfortable with the entropic prior and how to apply it in the CE may want to skip this appendix.

Write the precise microscopic state of a physical system under consideration as y . So for example, in classical (non-quantum) statistical physics, y is the set of positions and momenta of all the particles in the system. Arguments from

physics are typically invoked to justify a claim that the temperature T of the system “determines the expected energy of the system” for the (known) energy function of the system, $H(y)$.⁴⁰

Note that for this conclusion of those arguments to be a falsifiable statement, expectation values (“expected energy”) must be meaningful. So T must be associated with a (falsifiable) physical distribution over multiple y ’s, $q(y)$, rather than with a single y . Typically this distribution arises by not fully specifying the starting y and by allowing unknown stochastic external influences to perturb the system between when we acquire the value T and any subsequent observation of a property of the system. (It is implicitly required that those external influences do not change the value that a repeated temperature measurement would give.) All that is fixed (in addition to T) are some high-level aspects of how the system is set up, and of how it is opened to the external world. (For example, it may be that the identity of the person performing the experiment, how they physically hold the experimental instruments, etc., fixes those physical details.) Accordingly, the state y at that subsequent observation can vary.

So physically, to falsify a prediction of what q is associated with a particular T , we can imagine repeatedly setting up our system in the way specified and measuring the temperature, then opening it to (unknown) external influences in the way specified, and after that recording the resultant state y ; the associated distribution across y ’s is the (falsifiable) q . What we are interested in is the relation between that q and the measured T .

The aforementioned “arguments from physics” tell us that for any specification of how the system is set up and then opened, there is the same single-valued function of the measured T to the expected energy of the system under q . So via that single-valued function, a particular value of T picks out a unique set of q ’s (namely those with the associated expected energy).

Formally then, what we know is that the value of the temperature T uniquely fixes the expected value under q of a measurement of the system energy $H(y)$, independent of the details of how the system is set up and then opened to external influences, i.e., T fixes $E_q(H) \triangleq \int dy H(y)q(y)$. In general, for the same T , different choices of how the system is set up and then opened to external influences will result in a different one of the possible q consistent with the T -specified value of $E_q(H)$. However we don’t know *a priori* how the specification of the system’s setup and opening chooses among the set of all q that are all consistent with a particular value of $E_q(H)$. So even if the precise value of $E_q(H)$ were given to us, and even if how the system is set up and then opened were also specified, *for us*, it is as though nothing is specified concerning which of the q consistent with $E_q(H)$ has been picked out by what we know. (It is in encapsulating this ignorance of the distribution across q ’s that the entropic prior will arise.)

Moreover, while for a particular choice of how the system is set up and then opened we can ascertain the expected energy of the system by repeated experiments, often we cannot do this directly from a *single* one of those experi-

⁴⁰Properly speaking, H is the system’s “Hamiltonian”.

ments. This means in particular that while typically we can measure T in such a single experiment, often we do not know how that value T fixes the expected energy under q .⁴¹ So observing T does not always allow *us* to write down the expected energy, only to know that it has been fixed. In such instances, having observed T , we do not know what set of q 's that value of T has picked out, only that there is some such set.

The invariant \mathcal{S} for this situation fixes T , and therefore specifies that the distribution $q(y)$ must lie on a hyperplane of the form $E_q(H) = h$. But it does not specify the value h . Nor does it specify anything concerning which q goes with any particular h , i.e., it tells us nothing concerning the distribution of q 's across that hyperplane. Our inference problem is to circumvent this handicap: \mathcal{S} is the value T , together with the knowledge that it fixes $E_q(H)$, and we must use this to say something concerning q , the quantity we wish to infer. Formally, we wish to evaluate $P(q | \mathcal{S})$ for this choice of \mathcal{S} . (Note that this is a distribution across distributions.)

Note that the distribution q concerns the physical world. So in particular, it is experimentally falsifiable. In contrast, a distribution $P(q | \mathcal{S})$ reflects *us* (the researcher), and our (in)ability to infer q from T and the specification of how the system is set up and then opened. Although such a perspective is not required, one can interpret $P(q | \mathcal{S})$ as a subjective “degree of belief” in the objective (i.e., falsifiable) distribution q . Alternatively, one can view \mathcal{S} as picking out a set of physical instances of our system that are consistent with \mathcal{S} , and then interpret $P(q | \mathcal{S})$ in terms of frequencies of those instances.

For the reasons elucidated above, it makes sense to use an entropic prior over q 's for this \mathcal{S} . With such a prior, the MAP q is the one that maximizes $S(q)$ subject to the constraint $E_q(H) = h$. We just happen not to know h .

This is a constrained optimization problem with unknown constraint value. The associated Lagrangian is

$$\mathcal{L}(\beta, q) \triangleq \beta[E_q(H) - h] - S(q). \quad (73)$$

Removing the additive constant βh and dividing by the constant β gives $E_q(H) - \frac{S(q)}{\beta}$. This is known in statistical physics as the **free energy** of the system. β is the Lagrange parameter of our constraint. To solve our constrained optimization problem, q and β are jointly set so that the partial derivatives of $\mathcal{L}(\beta, q)$ are all zero.⁴² The minimizer of the free energy — the MAP q — is given by the Boltzmann distribution,

$$q(y) \propto \exp(-\beta H(y)). \quad (74)$$

⁴¹Observationally, that expected energy is defined in terms of a set of multiple experiments in addition to the current one. Mathematically, even if we know the Hamiltonian function, often we cannot evaluate its expected value for a particular T .

⁴²Throughout this paper the terms in any Lagrangian that restrict distributions to the unit simplices are implicit. The other constraint needed for a Euclidean vector to be a valid probability distribution is that none of its components are negative. This will not need to be explicitly enforced in the Lagrangian here, since this constraint is always obeyed for the optimizing $\mathcal{L}(\beta, q)$.

For macroscopically large systems, the posterior over q is in essence a delta function about the MAP solution, so the Bayes-optimal solution for almost any loss function is given by Eq. [refeq:statphysex](#).

β turns out to be the (inverse of) the temperature of the physical system (measured in units where Boltzmann's constant equals 1). In other words, the invariant of our problem is the value of the Lagrange parameter, not of the associated constraint constant. (The precise relationship between β and h depends on the function H in general.)

This scenario and its solution q is exactly the CE discussed previously. It is the simplest of all scenarios considered in statistical physics (hence its name).

ACKNOWLEDGEMENTS: I would like to thank George Judge, Duncan Luce, Herb Gintis, Larry Karp, Dan Friedman, Mike George, Wendy Cho, Arnold Zellner, and Tom Vincent for helpful discussions.

References

- [1] E. T. Jaynes and G. Larry Bretthorst. *Probability Theory : The Logic of Science*. Cambridge University Press, 2003.
- [2] S. F. Gull. Bayesian inductive inference and maximum entropy. In *Maximum Entropy and Bayesian Methods*, pages 53–74. Kluwer Academic Publishers, 1988.
- [3] T.J. Loredo. From laplace to sn 1987a: Bayesian inference in astrophysics. In *Maximum Entropy and Bayesian Methods*, pages 81–142. Kluwer Academic Publishers, 1990.
- [4] A. Zellner. Some aspects of the history of bayesian information processing. *Journal of Econometrics*. to appear.
- [5] D.H. Wolpert. The bootstrap is inconsistent with probability theory. In *Maximum entropy and Bayesian Methods*. Kluwer publishers, 1995.
- [6] L. Breiman. Stacked regression. *Machine Learning*, 24, 1996.
- [7] D. H. Wolpert. Reconciling Bayesian and non-Bayesian analysis. In *Maximum Entropy and Bayesian Methods*, pages 79–86. Kluwer Academic Publishers, 1996.
- [8] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, 1991.
- [9] D. Mackay. *Information theory, inference, and learning algorithms*. Cambridge University Press, 2003.
- [10] C.E Strauss, D.H. D.H. Wolpert, and D.R. Wolf. Alpha, evidence, and the entropic prior. In A. Mohammed-Djafari, editor, *Maximum Entropy and Bayesian Methods 1992*. Kluwer, 1994.

- [11] D. H. Wolpert. Product distribution field theory. 2003. preprint cond-mat/0307630.
- [12] D. H. Wolpert. Generalizing mean field theory for distributed optimization and control. 2004. submitted.
- [13] W. Macready, S. Bieniawski, and D.H. Wolpert. Adaptive multi-agent systems for constrained optimization. Technical report IC-04-123, 2004.
- [14] C. Fan Lee and D. H. Wolpert. Product distribution theory for control of multi-agent systems. In *Proceedings of AAMAS 04*, 2004.
- [15] D. H. Wolpert and S. Bieniawski. Distributed Control by Lagrangian Steepest Descent. pages 1562–1567, 2004.
- [16] S. Bieniawski and D. H. Wolpert. Adaptive, distributed control of constrained multi-agent systems. In *Proc. of the Third Intl. Conf. on Autonomous Agents and Multiagent Systems*, 1230-1231.
- [17] S. Bieniawski, D. H. Wolpert, and I. Kroo. Discrete, continuous, and constrained optimization using collectives. In *Proceedings of 10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Albany, New York*, 2004.
- [18] N. Antoine, S. Bieniawski, I. Kroo, and D. H. Wolpert. Fleet assignment using Collective Intelligence. In *Proceedings of 42nd Aerospace Sciences Meeting*, 2004. AIAA-2004-0622.
- [19] E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106:620, 1957.
- [20] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, 1987.
- [21] K. E. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2003.
- [22] D. McFadden. Econometric models of probabilistic choice. In *Structural analysis of discrete data and econometric applications*. MIT Press, 1981.
- [23] C. A. Holt J. K. Goeree. A model of noisy introspection. *Games and Economic Behavior*, 46:365–382, 2004.
- [24] T. R. Palfrey J. K. Goeree, C. A. Holt. Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104:247–272, 2002.
- [25] D. H. Wolpert. What Information theory says about best response, binding contracts, and Collective Intelligence. In A. Namatame, editor, *Proceedings of WEHIA04*. Springer Verlag, 2004.

- [26] T. R. Palfrey J. K. Goeree, C. A. Holt. Quantal response equilibrium and overbidding in private-value auctions. 1999.
- [27] R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38, 1994.
- [28] H. C. Chen and J. W. Friedman. *Games and Economic Behavior*, 18:32–54, 1997.
- [29] D. H. Wolpert. Information theory - the bridge connecting bounded rational game theory and statistical physics. In D. Braha and Y. Bar-Yam, editors, *Complex Engineering Systems*, 2004.
- [30] J. R. Meginniss. A new class of symmetric utility rules for gambles, subjective marginal probability functions, and a generalized Bayes' rule. *Proc. of the American Statistical Association, Business and Economics Statistics Section*, pages 471 – 476, 1976.
- [31] D. Fudenberg and D. Kreps. Learning mixed equilibria. *Game and Economic Behavior*, 5:320–367, 1993.
- [32] D. Fudenberg and D. K. Levine. Steady state learning and Nash equilibrium. *Econometrica*, 61(3):547–573, 1993.
- [33] D. Luce. *Individual Choice Behavior*. Wesley, 1959.
- [34] R. McAfee and J. McMillan. Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43:1–19, 1987.
- [35] R.J. Aumann and S. Hart. *Handbook of Game Theory with Economic Applications*. North-Holland Press, 1992.
- [36] D. H. Wolpert and S. Bieniawski. Adaptive distributed control: beyond single-instant categorical variables. In A. Skowron, editor, *Proceedings of MSRAS04*. Springer Verlag, 2004.
- [37] W. Macready and D. H. Wolpert. Distributed constrained optimization with semicoordinate transformations. submitted to *Journal of Operations Research*, 2005.
- [38] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.
- [39] T. Basar and G.J. Olsder. *Dynamic Noncooperative Game Theory*. Siam, Philadelphia, PA, 1999. Second Edition.
- [40] K. Binmore. *Fun and Games: A Text on Game Theory*. D. C. Heath and Company, Lexington, MA, 1992.
- [41] R. D. Luce and H. Raiffa. *Games and Decisions*. Dover Press, 1985.

- [42] A. Greif. Economic history and game theory: A survey. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 3. North Holland, Amsterdam, 1999.
- [43] J. Bernardo and A. Smith. *Bayesian Theory*. Wiley and Sons, 2000.
- [44] J. M. Berger. *Statistical Decision theory and Bayesian Analysis*. Springer-Verlag, 1985.
- [45] D. H. Wolpert. On bias plus variance. *Machine Learning*, 9:1211–1244, 1997.
- [46] J. B. Paris. *The Uncertain Reasoner’s Companion: A Mathematical Perspective*. Cambridge University Press, 1994.
- [47] K. S. Van Horn. Constructing a logic of plausible inference: a guide to cox’s theorem. *International Journal of Approximate Reasoning*, 34(1):3–24, 2003.
- [48] C.F. Camerer. *Behavioral Game theory: experiments in strategic interaction*. Princeton University Press, 2003.
- [49] D. Kahneman. A psychological perspective on economics. *American Economic Review (Proceedings)*, 93:2:162–168, 2003.
- [50] D. Kahneman. Maps of bounded rationality: Psychology of behavioral economics. *American Economic Review*, 93:5:1449–1475, 2003.
- [51] D. H. Wolpert. Bounded rational games, information theory, and statistical physics. In D. Braha and Y. Bar-Yam, editors, *Complex Engineering Systems*, 2004.
- [52] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2003.
- [53] D. Wolpert and K. Tumer. Beyond mechanism design. In H. Gao et al., editor, *International Congress of Mathematicians 2002 Proceedings*. Qingdao Publishing, 2002.
- [54] D. H. Wolpert and K. Tumer. Optimal payoff functions for members of collectives. *Advances in Complex Systems*, 4(2/3):265–279, 2001.
- [55] K. Tumer and D. H. Wolpert, editors. *Collectives and the Design of Complex Systems*. Springer Verlag, 2003.
- [56] A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5:297–323, 1992.
- [57] D. H. Wolpert. Theory of Collective Intelligence. In K. Tumer and D. H. Wolpert, editors, *Collectives and the Design of Complex Systems*, New York, 2003. Springer.

- [58] D. H. Wolpert and K. Tumer. Collective Intelligence, Data Routing and Braess' Paradox. *Journal of Artificial Intelligence Research*, 16:359–387, 2002.
- [59] R. J. Aumann and A. Brandenburger. Epistemic conditions for nash equilibrium. *Econometrica*, 63(5):1161–1180, 1995.
- [60] B. Kadane and P. Larkey. Subject probability and the theory of games. *Management Science*, 28(2):113–125, 1982. Same issue contains discussion of the article.
- [61] R. J. Aumann. Interactive epistemology ii: Probability. *Int. J. Game Theory*, 28:301–314, 1999.
- [62] J.C. Harsanyi. Games with randomly distributed payoffs: A new rationale for mixed-strategy equilibrium points. *Int. J. Game Theory*, 2:1–23, 1973.
- [63] J.C. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press, 1988.
- [64] R. J. Aumann. *Economic Decision Making: Games, Econometrics and Optimization*. Elsevier, 1990. The chapter “Nash Equilibria are not Self-Enforcing”.
- [65] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification (2nd ed.)*. Wiley and Sons, 2000.
- [66] D. Fudenberg and D. K. Levine. *The Theory of Learning in Games*. MIT Press, Cambridge, MA, 1998.
- [67] S. Durlauf. How can statistical mechanics contribute to social science? *Proc. Natl. Acad. Sci. USA*, 96:10582–10584, 1999.
- [68] W. A. Brock and S. N. Durlauf. Interaction-based models. In *Handbook of Econometrics: Volume 5*, pages 3297–3380. Elsevier, 2001. Chapter 54.
- [69] M. Allais. *Econometrica*, 21:503–546, 1953.
- [70] J. A. List and M. S. Haigh. A simple test of expected utility theory using professional traders. *Proceedings of the National Academy of Sciences*, 102:945–948, 2005.
- [71] R. Kurzban and D. Houser. Experiments investigating cooperative types in humans. *Proceedings of the National Academy of Sciences*, 102(5):1803–1807, 2005.
- [72] D. Luce. Whatever happened to information theory in psychology? *Review of General Psychology*, pages 183–188, 2003.
- [73] J.S. Shamma and G. Arslan. Dynamic fictitious play, dynamic gradient play, and distributed convergence to nash equilibria. *IEEE Trans. on Automatic Control*, 50(3):312–327, 2004.

- [74] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
- [75] L. P. Kaelbling, M. L. Littman, and A. W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- [76] R. H. Crites and A. G. Barto. Improving elevator performance using reinforcement learning. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems - 8*, pages 1017–1023. MIT Press, 1996.
- [77] J. Hu and M. P. Wellman. Multiagent reinforcement learning: Theoretical framework and an algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning*, pages 242–250, June 1998.
- [78] D. H. Wolpert, K. Tumer, and J. Frank. Using collective intelligence to route internet traffic. In *Advances in Neural Information Processing Systems - 11*, pages 952–958. MIT Press, 1999.
- [79] C. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3/4):279–292, 1992.
- [80] A. Dragulescu and V.M. Yakovenko. Statistical mechanics of money. *Eur. Phys. J. B*, 17:723–729, 2000.
- [81] M. Aoki. *Modeling Aggregate Behavior and Fluctuations in Economics : Stochastic Views of Interacting Agents*. Cambridge University Press, 2004.
- [82] J.D. Farmer, M. Shubik, and D. E. Smith. Economics: The next physical science? SFI working paper 05-06-027.