Bounded Rationality and Correlated Equilibria*

Fabrizio Germano[†] and Peio Zuazo-Garin[‡]

November 2, 2015

Abstract

We study an interactive framework that explicitly allows for nonrational behavior. We do not place any restrictions on how players' behavior deviates from rationality. Instead we assume that there exists a probability p such that all players play rationally with at least probability p, and all players believe, with at least probability p, that their opponents play rationally. This, together with the assumption of a common prior, leads to what we call the set of p-rational outcomes, which we define and characterize for arbitrary probability p. We then show that this set varies continuously in p and converges to the set of correlated equilibria as p approaches 1, thus establishing robustness of the correlated equilibrium concept to relaxing rationality and common knowledge of rationality. The p-rational outcomes are easy to compute, also for games of incomplete information, and they can be applied to observed frequencies of play to derive a measure $\overline{p}$ that bounds from below the probability with which any given player chooses actions consistent with payoff maximization and common knowledge of payoff maximization.

Keywords: strategic interaction, correlated equilibrium, robustness to bounded rationality, approximate knowledge, incomplete information, measure of rationality, experiments. *JEL Classification*: C72, D82, D83.

*Earlier versions of this paper circulated under the title "Approximate Knowledge of Rationality and Correlated Equilibria." We are indebted to Bob Aumann, Eddie Dekel, Sergiu Hart, and Dov Samet, for valuable comments and conversations. We also thank audiences in Barcelona, Bilbao, Chicago, Istanbul, Jerusalem, Paris, and York, as well as an anonymous referee for insightful comments. Germano acknowledges financial support from the Spanish Ministry of Economy and Competitiveness (Grant ECO2014-59225-P), as well as from the Barcelona GSE Research Network and the Generalitat de Catalunya. Zuazo-Garin acknowledges financial support from the Spanish Ministry of Science and Technology (Grant ECO2009-11213), and inestimable help from professors Elena Iñarra and Annick Laruelle. Both authors thank the Center for the Study of Rationality at the Hebrew University of Jerusalem for their generous hospitality; Germano also thanks AMSE-IMéRA at Université Aix-Marseille for generous hospitality. All errors are our own.

†Universitat Pompeu Fabra, Departament d'Economia i Empresa, and Barcelona GSE, Ramon Trias Fargas 25-27, E-08005 Barcelona, Spain; `fabrizio.germano@upf.edu`.

‡Universitat Rovira i Virgili, Department d'Economia, CREIP and *BRiDGE*, Avinguda de la Universitat 1, E-43204, Reus, Spain; `peio.zuazo@urv.cat`.

1 Introduction

Rationality, understood as consistency of behavior with stated objectives, information, and strategies available, naturally lies at the heart of game theory. Still today, most of game theory and its applications takes the rationality of the agents, knowledge and higher knowledge thereof as given and, in an important way, relies on these to make (ideally robust) predictions about behavior.[1] However, it is clear that, in practice, departures from full consistency or rationality not only occur and occur often, but they also occur in innumerable ways.[2] To address this, we develop a theory that relaxes the assumption of rationality and higher order knowledge of rationality. We assume that for any given probability $p \in [0, 1]$, players choose rationally with probability at least p and with respect to beliefs that assign probability at least p to their opponents' choosing rationally as well; we put no constraint on what the remaining actions or action profiles (occurring with frequency at most $1-p$) may be. Among other things, we study the robustness of behavior with respect to p, that is, its sensitivity to the introduction of nonrationality, and derive a measure of "rationality" for observed frequencies of play.

As a brief illustration, consider the following two examples. The first is a stylized penalty kick game, taken from Palacios-Huerta (2003), based on actual penalty kicks shot by professional soccer players in European leagues between 1995 and 2000.[3] Players' strategies are reduced to kick left or right (KL, KR) for the kicker (row player) and to jump left or right (JL, JR) for the goalkeeper (column player). The payoffs are described in Figure 1 in the left-hand matrix (and correspond to probability of scoring for the kicker and probability of saving for the goalkeeper); the right hand matrix describes the empirical frequencies with which the different strategies were played in the field.

$G_{PK} \approx$

	JL	JR
KL	58,42	95, 5
KR	93, 7	70,30

, $\pi^{emp}_{PK} \approx$

	JL	JR
KL	0.168	0.232
KR	0.252	0.348

,

Figure 1: Penalty kicks in professional European soccer leagues ($\overline{p} \approx 0.96$)

Although close, it can be checked that, if one assumes the existence of a common prior, strictly speaking, frequencies of play are inconsistent with common knowledge of rationality, as players are not playing best-responses to one another. As we will discuss in Section 5, where we introduce an empirical measure of rationality derived from our theory that allows to quantify discrepancies from equilibrium play, we compute that at least 96% of each player's (row player's and column player's) actions are consistent with payoff maximization and common knowledge of payoff maximization. The second example, represented in Figure 2, is a game taken from Goeree and Holt (2001) due to David Kreps. Here subjects playing the column player typically play strategy (NN) that is the only strategy that is not in the support of any of the Nash equilibria of the game. Nonetheless, according to our measure, at least 70% of each player's (row player's and column player's) actions are consistent with payoff-maximization and common knowledge of payoff maximization, whether or not the common prior assumption holds.

[1]This is true of almost any modern textbook in game theory, or most of the excellent chapters of the three volumes of the *Handbook of Game Theory with Economic Applications*, edited by Robert J. Aumann and S. Hart, and of the recent fourth volume edited by P. Young and S. Zamir, which meanwhile also has chapters specifically devoted to behavioral game theory and evolutionary game theory.

[2]See the exhaustive list of examples in the survey of Conlisk (1996); see also Rubinstein (1998) and Mallard (2011) on models of bounded rationality; Crawford (2013) and Harstad and Selten (2013) distinguish optimization-based from nonoptimization-based models of bounded rationality; Camerer *et al.* (2011) and Camerer and Ho (2015) contain surveys of behavioral game theory and economics. The experimental literature has played an important role in advancing the research on bounded rationality in game theory.

[3]We refer to that paper for discussions on the meaning of the strategies, payoffs, and the overall setup; see also Palacios-Huerta and Volij (2008), Chiappori *et al.* (2002) for further related results.

$G_{Kreps} \equiv$

	L	M	NN	R
T	200, 50	0, 45	10,30	20,-250
B	0,-250	10,-100	30,30	50, 40

, $\pi^{emp}_{Kreps} \approx$

	L	M	NN	R
T	0.178	0.054	0.462	0
B	0.082	0.026	0.218	0

.

Figure 2: Kreps game ($p \approx 0.7$)

In this paper, we are interested in a theory of strategic interaction that incorporates the following two aspects of bounded rationality: (i) some (possibly small) amount of nonrational behavior, and (ii) the capacity of players to expect, and optimally react, to nonrational behavior by their opponents. In order to develop such a theory, we relax the assumptions that agents are rational at all times and that there is common knowledge of rationality. We replace these with a substantially weaker assumption, namely, that there exists a lower bound on the probability that players assign to their opponents being rational, that is, to choosing actions that are payoff-maximizing given their own information. More specifically, we study the behavior that arises, if, in every state of the world, each player believes that the other players are rational with a probability p or more. This is what we call *common knowledge of mutual p-belief in opponents' rationality*. Together with the existence of a common prior it defines the notion of *p-rational outcome*, which is at the center of our paper. Thus characterizing the set of p-rational outcomes is tantamount to characterizing behavior that allows for mistakes, arbitrary mistakes, and optimization *w.r.t.* the possibility of mistakes, as long as the expected frequency of mistakes by others is bounded above by $1 - p$. Importantly, we put no restriction on what it means to be *non*-rational, except that the rules of the game implicitly require agents to select some action from the action space.[4]

After defining our central notion of p-rational outcomes, we give a strategic characterization in Theorem 1 in terms of what we call *(X, p)-correlated equilibria*. These are correlated equilibria, where incentive constraints hold on a subset ($X \subseteq A$) of the overall action space, and where actions from this subset are believed to be played with certain minimum probability p. The theorem provides a nonepistemic characterization that uses the (X, p)-correlated equilibria to relate the p-rational outcomes to correlated equilibria and can be seen as giving a generalization that extends the main result by Aumann (1987) to contexts of bounded rationality. The set of p-rational outcomes is described by linear inequalities consisting of incentive and p-belief constraints, and for any p; it always contains the set of correlated equilibria, with which it coincides when $p = 1$. When $p = 0$, the set of p-rational outcomes makes up the whole space of distributions over action space A, $\Delta(A)$.

In Theorem 2, we show that, besides being nonempty, convex and compact, the set of p-rational outcomes varies continuously in the underlying parameter p. We further show that when p is sufficiently close to 1, then rationalizable strategy profiles are played with probability at least p. These results confirm in some sense the robustness of the correlated equilibrium concept to bounded rationality. Proposition 2 further characterizes the p-rational equilibria as Bayesian Nash equilibria of incomplete information games, namely certain "canonical elaborations" as defined in Kajii and Morris (1997*a*,*b*), that are perturbations of the underlying (complete information) game G with some restrictions on the frequencies of "standard" and "committed" types. The proposition can be seen as also providing an epistemic foundation of Bayesian equilibria of such perturbed games with "standard" and "committed" types.

Theorem 4 then shows that our main characterization result extends directly to the case of games of

[4]As we discuss below, this is what separates this paper, from many other papers in the literature, that have looked at specific ways in which agents can deviate from "full rationality" such as the models of k-level reasoning, cognitive hierarchy or λ-quantal response, or theories of ϵ-equilibria and so on; see, e.g., Camerer and Ho (2015) for a discussion of some of these theories. By contrast, we remain agnostic about how players behave when they are nonrational.

incomplete information using the notion of Bayes correlated equilibrium of Bergemann and Morris (2015). To the extent that their results show robustness of the Bayes correlated equilibria to underlying private information structures, we can view our results as showing that the p-rational Bayes outcomes we characterize (which coincide with the Bayes correlated equilibria when $p = 1$) are robust to nonrational behavior by the players, provided it occurs with probability no more than $1 - p$.

As a further application of our theory, we use the p-rational outcomes to derive a unique number $\overline{p} \in [0, 1]$ that quantifies proximity to common knowledge of rationality in a normal-form strategic interaction. In interactions where the common prior assumption can be expected to hold, for any given distribution of play, say $\pi \in \Delta(A)$, we can define a unique number $\overline{p} \in [0, 1]$, that gives the largest p such that each player plays actions that are consistent with common knowledge of payoff maximization given π with frequency at least p. This gives a direct measure of the maximum possible amount of actions consistent with payoff maximization reflected in the distribution π. For interactions, where the common prior assumption does not hold, then players may be acting rationally, but their rationality is underestimated by $\overline{p}$ as it does not take into account possible inconsistencies in beliefs. Allowing for this, one can show that $\overline{p}$ is a lower bound for the maximum possible amount of actions consistent with payoff maximization and common knowledge of payoff maximization, reflected in the distribution π. Therefore, the value $\overline{p}$ can be useful as a measure of minimum amount of rationality in experimental data, whether or not there is a common prior. We discuss this in more detail Section 5.

At a theoretical level, our analysis builds on the epistemic literature, centered around the concepts of rationalizability, Bernheim (1984) and Pearce (1984), and correlated equilibrium, Aumann (1974, 1987), that characterizes behavior under varying assumptions on players' rationality and their reciprocal beliefs in each others' rationality. Tan and Werlang (1988) show that independent rationalizability characterizes rationality and common certainty of rationality; and Brandenburger and Dekel (1987) connect it to subjective correlated equilibria and correlated rationalizability. Using the notion of common p-belief (of Monderer and Samet (1989), who introduce the concept to study robustness of equilibria to incomplete information regarding payoffs, and thus, do not account for deviations from rationality), Hu (2007) introduces the notion of (correlated) p-rationalizability, and shows that it characterizes rationality and common p-belief in rationality, for general $p \leq 1$.[5] He also shows that as p converges to 1, the set of p-rationalizable actions approaches the set of rationalizable actions.

For incomplete information games and within an ex ante context, Forges (1993, 2006) introduces several notions of correlated equilibrium. Lehrer *et al.* (2010, 2013) clarify epistemically the role of different assumptions and information structures and study their effect on equilibrium behavior. Bergemann and Morris (2015) introduce a further broader notion of correlated equilibrium, which they call Bayes correlated equilibrium, and which they show characterizes behavior robust to varying information structures. Bayes correlated equilibrium is the equilibrium notion we use for our incomplete information analysis. At the interim stage, starting from hierarchies of beliefs, Dekel *et al.* (2007) introduce the notion interim correlated rationalizability and show that it characterizes common certainty of rationality. Germano and Zuazo-Garin (2015), introduce the notion of interim correlated p-rationalizability, and show that it characterizes common p-belief of rationality, for general $p \leq 1$. We view the ex ante and the interim approaches as providing complementary results; the ex ante approach used in this paper making epistemically speaking more restrictive assumptions (assuming besides a common prior also common knowledge of the model and the epistemic assumptions it entails). To the extent that the additional assumptions are satisfied, the ex ante approach

[5] To highlight an important difference with our approach, notice that, within our finite, ex ante context, assuming common p-belief of rationality and a common prior for any $p > 0$, amounts to the same as assuming common knowledge of rationality (see Lemma 1).

provides an effective tool for characterizing resulting behavior and, in our case, also yields sharper bounds and predictions as compared to the notions of p-rationalizability or interim correlated p-rationalizability.

The paper is structured as follows. Section 2 sets up both the game-theoretic and epistemic framework and recalls and discusses the main results by Aumann (1987). Section 3 is the main section, which introduces the notion of p-rational outcome and contains strategic and topological characterizations, as well as some simple examples. Section 4 contains some extensions, including to games of incomplete information. Section 5 shows how the theory implies a natural measure to quantify the degree of "rational" behavior in strategic interactions, and Section 6 provides some concluding remarks. All the proofs are in the Appendix.

2 Preliminaries

In this section, we recall some well-known concepts in game theory and epistemic game theory needed for the analysis of Section 3. In Section 2.1, we present the game-theoretic framework and the standard solution concepts later generalized: correlated and subjective correlated equilibrium (Aumann, 1974). Then, in Section 2.2, we introduce the epistemic framework, in which interactive knowledge and beliefs are formalized; this uses standard constructions by Aumann (1987), Brandenburger and Dekel (1987) or Monderer and Samet (1989), among others. Finally, in Section 2.3, we discuss the main result by Aumann (1987), which relates common knowledge assumptions regarding rationality with correlated and subjective correlated equilibria.

2.1 Correlated equilibria

A (finite, normal-form) *game* is defined as a tuple $G = \left\langle I, (A_i)_{i\in I}, (u_i)_{i\in I}\right\rangle$, where I is a finite set of *players*, and for any player i we have: a finite set of *actions* A_i and a *payoff function* $u_i : A \to \mathbb{R}$, where $A = \prod_{i\in I} A_i$ denotes the set of action *profiles*. Given distribution $\pi \in \Delta(A)$, for any player i, we say that action a_i is *optimal* w.r.t. π if,

$$\underset{a_i'\in A_i}{\operatorname{argmax}} \sum_{a_{-i}\in A_{-i}} \pi[(a_{-i}; a_i)] \cdot u_i\left(a_{-i}; a_i'\right),$$

where $A_{-i} = \prod_{j\neq i} A_j$. Then, following Aumann (1974): (i) a distribution $\pi \in \Delta(A)$ is a *correlated equilibrium* if for any player i every action a_i is optimal w.r.t. π, and (ii) a family of distributions $(\pi_i)_{i\in I} \subseteq \Delta(A)$ is a *subjective correlated equilibrium* if for any player i every action a_i is optimal w.r.t. π_i. We denote the sets of correlated equilibria and subjective correlated equilibria of G by $CE(G)$ and $SCE(G)$, respectively. It follows by definition that $(CE(G))^I \subseteq SCE(G)$. It is easy to see that every Nash equilibrium induces a correlated equilibrium; thus, $CE(G)$ and $SCE(G)$ are both nonempty.

2.2 Epistemic framework

Interactive knowledge and beliefs are exogenously modeled by a *belief system*, which consists of a list $B = \left\langle \Omega, (\Pi_i)_{i\in I}, (\alpha_i)_{i\in I}, (\mu_i)_{i\in I}\right\rangle$, where Ω is a finite set of *states* (of the world), and for each player i we have: (i) Π_i, a *knowledge partition* of Ω, where for any state ω we denote the cell containing ω by $\Pi_i(\omega)$, (ii) a *strategy map* $\alpha_i : \Omega \to A_i$ measurable w.r.t. Π_i, and (iii) a (subjective) *prior belief* $\mu_i \in \Delta(\Omega)$ with full support. An *event* is a collection of states $E \subseteq \Omega$. As mentioned above, belief systems formalize the following two epistemic aspects:

1. **Interim knowledge.** For any player i and any state ω, player i's knowledge at ω is represented by $\Pi_i(\omega)$: we say player i *knows* event E at state ω if $\Pi_i(\omega) \subseteq E$. Player i's *knowledge operator* is thus

defined as follows,

$$E \mapsto K_i(E) = \{\omega \in \Omega \,|\Pi_i(\omega) \subseteq E\}, \text{ for any } E \subseteq \Omega.$$

Note that the measurability of the strategy maps implies that each player knows at every state what action she chooses. We say that an event E is *evident knowledge* if $E \subseteq \bigcap_{i\in I} K_i(E)$, and for state ω, event C is *commonly known* at ω if there exists some evident knowledge E such that $\omega \in E \subseteq \bigcap_{i\in I} K_i(E)$. We denote the event that C *is commonly known* by $CK(C)$.

2. **Interim beliefs**. For each player i knowledge partition Π_i and prior belief μ_i induce *interim beliefs* at each state ω,

$$\mu_i(\omega)[E] = \frac{\mu_i[E \cap \Pi_i(\omega)]}{\mu_i[\Pi_i(\omega)]}, \text{ for any } E \subseteq \Omega.$$

Then, following Monderer and Samet (1989), for any probability $p \in [0,1]$ we say that player i p-*believes* event E at state ω if $\mu_i(\omega)[E] \geq p$. Player i's p-*belief operator* is thus defined as follows,

$$E \mapsto B_i^p(E) = \{\omega \in \Omega \,|\mu_i(\omega)[E] \geq p\}, \text{ for any } E \subseteq \Omega.$$

We say that an event E is p-*evident belief* if $E \subseteq \bigcap_{i\in I} B_i^p(E)$, and for state ω, event C is *common* p-*belief* at ω if there exists some p-evident belief E such that $\omega \in E \subseteq \bigcap_{i\in I} B_i^p(E)$. We denote the event that C *is common* p-*belief* by $CB^p(C)$. It is easy to see that in this framework, knowledge and 1-belief coincide, due to the fact that prior beliefs have full support.

2.3 'Correlated equilibria as an expression of Bayesian rationality'

Each belief system B induces, for each player i, the following interim beliefs on her opponents' behavior at each state ω: $\mu_i(\omega)[a_{-i}] = \mu_i(\omega)[\bigcap_{j\neq i} \alpha_j^{-1}(a_j)]$ for any $a_{-i} \in A_{-i}$, and thus, the following interim expected payoff,

$$\mathbb{E}_B(\omega)[u_i(\alpha_{-i}; a_i)] = \sum_{a_{-i}\in A_{-i}} \mu_i(\omega)[a_{-i}] \cdot u_i(a_{-i}; a_i),$$

for each $a_i \in A_i$. Then, we say that player i is (Bayesian) *rational* at state ω if her choice is optimal w.r.t. her interim beliefs, i.e., if $\alpha_i(\omega) \in \text{argmax}_{a_i\in A_i} \mathbb{E}_B(\omega)[u_i(\alpha_{-i}; a_i)]$, and denote the event that player i is rational by R_i. We denote the event that every player is rational by R. Finally, note that belief system B also induces, for each player i, a *subjective outcome* (distribution) $\pi_i^B \in \Delta(A)$ given by $\pi_i^B(a) = \mu_i[\bigcap_{j\in I} \alpha_j^{-1}(a_j)]$ for any $a \in A$. Aumann (1987) studies the impact of the following two properties in belief systems:

Common knowledge of rationality. Belief system B is *rational* if players are rational at every state, i.e., if $\Omega = R$. Note that since Ω is evident knowledge, it follows that rationality is commonly known at every state. Aumann (1987) shows that if belief system B is rational, then the family of subjective outcomes it induces is a subjective correlated equilibrium: $(\pi_i^B)_{i\in I} \in SCE(G)$.

Common prior assumption. Belief system B satisfies the *common prior assumption* if all the players hold the same prior belief, i.e., if $\mu_i = \mu_j$ for any $i, j \in I$. In case the common prior assumption is satisfied, we drop subscripts and denote the common prior belief, to which we refer as the *common prior*, by simply μ. In this case, B induces an *objective outcome* (distribution) $\pi^B \in \Delta(A)$ given by $\pi^B(a) = \mu[\bigcap_{i\in I} \alpha_i^{-1}(a_i)]$ for any $a \in A$.[6] The main result in Aumann (1987) shows that if B

[6] Or, alternatively, every player i's subjective outcome π_i^B happens to be identical.

is rational and satisfies the common prior assumption, then the objective outcome it induces is a correlated equilibrium: $\pi^B \in CE(G)$.

Two aspects of the formalization of belief systems are crucial for obtaining equilibrium behavior: (i) strategy maps $(\alpha_i)_{i\in I}$ are *structurally* commonly known,[7] and (ii) the common prior assumption implies that players hold correct beliefs about how information is distributed among their opponents.[8] The combination of this features yields that players hold, indeed, correct beliefs about how their opponents play, since: they hold correct beliefs about how information is distributed and they hold correct beliefs about how, specifically, choice is made contingent on information. For examples of epistemic frameworks that offer a more transparent distinction between equilibrium assumptions (as the two mentioned) and common knowledge of rationality see Tan and Werlang (1988), Dekel *et al.* (2007) or Battigalli *et al.* (2011).

3 Bounded rationality and correlated equilibria

In this section we characterize an extension of correlated equilibrium that incorporates the following two aspects of bounded rationality: (i) some (possibly small) amount of nonrational behavior, and (ii) the capacity of players to expect, and optimally react, to nonrational behavior by their opponents. Obviously, such behavior is at odds with common knowledge of rationality: a weaker epistemic assumption is required. In Section 3.1, we present our central epistemic assumption, discuss why it stands as a compelling relaxation, and define the probabilistic behavior it induces as p-rational outcomes (Definition 1). It remains unclear whether the p-rational outcomes can be characterized without having to go through a tedious process of epistemic modeling. This problem is solved in Section 3.2, where we present an easily computable generalization of correlated equilibria, (X, p)-correlated equilibria, and show in Theorem 1 that the set of p-rational outcomes can be characterized in terms of (X, p)-correlated equilibria of a related game, Aumann and Dreze's (2008) doubled game. In Section 3.3, we present some examples illustrating the p-rational outcomes, and in Section 3.4, we discuss geometric properties of the set of p-rational outcomes, and show that this set is robust to small variations in the amount of nonrational behavior.

3.1 p-Rational outcomes: Epistemic motivation

In order to capture the two features of bounded rationality mentioned above, we need to depart from common knowledge of rationality. A natural and standard way to relax knowledge assumptions is to resort to Monderer and Samet's (1989) notion of p-belief, as recalled in Section 2.2. However, it is not straightforward how p-beliefs should be applied in order to accommodate the kind of nonrational behavior we are looking for. Thus we propose to substitute common knowledge of *rationality* with common knowledge of *mutual p-belief in opponents' rationality*. That is, we assume that it is commonly known that there exists some probability p for which every player believes, with probability at least p, that the rest of players are rational. This idea is formalized by belief systems that satisfy:

Common knowledge of mutual p-belief in opponents' rationality (MB^{p}R). For fixed probability p, mutual p-belief in opponents' rationality holds at every state, i.e., $\Omega = \bigcap_{i\in I} B_i^p(R_{-i})$.

[7]That is, player i knows that, if player j gets the information corresponding to some state ω, then j plays action $\alpha_j(\omega)$; formally this is represented by the fact that $K_i(\neg\Pi_j(\omega) \cup [\alpha_j = \alpha_j(\omega)]) = \Omega$ for any $\omega \in \Omega$, and any $i, j \in I$.

[8]That is, the probability of player i receiving the information corresponding to state ω, $\mu[\Pi_i(\omega)]$ is exactly the probability that, *ex ante* each player j assigns to player i obtaining the information corresponding to state ω.

Note then that, under common knowledge of mutual p-belief in rationality (MB^{p}R), we have: (i) some (possibly small) amount of nonrational behavior, and (ii) capacity of players' to expect, and optimally react, to nonrational behavior by their opponents. Thus, MB^{p}R captures the two aspects of bounded rationality we are looking to characterize.[9] Lemma 1 below studies how common knowledge of rationality, common p-belief in rationality, and MB^{p}R relate to each other, and illustrates some properties of the latter, which motivate its suitability as a relaxation of common knowledge of rationality.

Lemma 1 *Let G be a game, p, a positive probability, and B, a belief system. Then:*

(i) $\Omega = CB^p(R)$ *if and only if* $\Omega = R$.

(ii) *If* $p = 1$, *then* $\Omega = R$ *if and only if* $\Omega = \bigcap_{i \in I} B_i^p(R_{-i})$.

(iii) *If B satisfies MB^{p}R and the common prior assumption, then* $\mu[R] \geq p^2$.

We provide some interpretations. The first result states that assuming common p-belief in rationality at every state is identical to assuming rationality at every state. Thus, for any $p > 0$, behavior induced by common p-belief in rationality corresponds exactly to correlated equilibria, and hence, we can conclude that common p-belief in rationality is not appropriate to capture the aspects of bounded rationality we are interested in. The second result shows that in the limit case of full rationality, when $p = 1$, MB^{p}R and common knowledge of rationality coincide. Thus, when nonrational behavior is excluded, MB^{p}R induces precisely, correlated equilibria. The third result studies the impact of the common prior assumption on MB^{p}R and shows that the fact that the common prior assumption entails correct belief implies that, for fixed p, MB^{p}R leads to rational behavior with at least probability p^2 so that nonrational behavior is bounded by probability $(1-p^2)$. Then, we define the outcome distributions induced by belief systems satisfying MB^{p}R and the common prior assumption as follows:

Definition 1 (p-Rational outcome) *Let G be a game, and p, a probability. Then, we say that distribution* $\pi \in \Delta(A)$ *is a* p-rational outcome *if it is induced by some belief system B that satisfies MB^{p}R and the common prior assumption. We denote the set of p-rational outcomes by p-RO(G).*

The interest of p-rational outcomes lies in the notion of bounded rationality that they capture via MB^{p}R. Yet, it seems problematic that, given a game G and a probability p, in order to characterize the set of p-rational outcomes, we need to consider all the possible belief systems satisfying MB^{p}R and the common prior assumption, and then, compute the outcomes they induce. We know that this can be circumvented for rational systems satisfying common knowledge of rationality: by Aumann's (1987) theorem, the behavior of these belief systems is captured by the set of correlated equilibria, which is easily computable through certain incentive constraints in the game G. In the next section, we explore whether it is possible to characterize the set of p-rational outcomes without the need of epistemic modeling.

3.2 p-Rational outcomes: Strategic characterization

We are interested in characterizing the set of p-rational outcomes in terms of the original game G, without evoking belief systems. Before doing so we need to first introduce the following notion of (X, p)-correlated equilibrium. This concept generalizes the notion of correlated equilibrium by explicitly allowing for some actions to be played nonoptimally, and plays a key role in our characterization result, Theorem 1 below.

[9] Another natural alternative, suggested to us by Dov Samet, would be to consider a common prior that satisfies $\mu[CB^p(R)] \geq 1 - \varepsilon$ for some $\varepsilon > 0$. This is a weakening of mutual p-belief in opponents' at every state in that it imposes less structure on players' beliefs, nonetheless it appears to be computationally less tractable; we return to this later; see Remark 1 in Section 6.

Definition 2 ((X, p)-Correlated equilibrium) *Let G be a game. Then, for any $X = \prod_{i\in I} X_i \subseteq A$ and any probability p we say that distribution $\pi \in \Delta(A)$ is a (X, p)-*correlated equilibrium *if for any player i the following hold:*

- (i) **Incentive constraints:** *For any action $a_i \in X_i$, $a_i \in \operatorname{argmax}_{a'_i\in A_i} \sum_{a_{-i}\in A_{-i}} \pi[(a_{-i}; a_i)] \cdot u_i(a_{-i}; a'_i)$.*
- (ii) **p-belief constraints:** *For any action $a_i \in A_i$, $\pi[X_{-i} \times \{a_i\}] \geq p \cdot \pi[A_{-i} \times \{a_i\}]$.*

We denote the set of (X, p)-correlated equilibria of G by (X, p)-$CE(G)$.

This notion weakens the usual incentive constraints in the following sense. Fix a probability $p \in [0, 1]$ and, for each player i, a subset $X_i \subseteq A_i$ such that the distribution on the overall set of action profiles $\pi \in \Delta(A)$ satisfies two kind of constraints:

- (i) Standard *incentive constraints* for not all actions, but rather, only for those in X_i. Thus, strictly smaller sets $X_i \subsetneq A_i$ can reflect agents that follow a mediator's advice without questioning the "rationality" of doing so (*e.g.*, in the sense of a social norm), or also agents that simply act irrationally in the sense of making mistakes (for whatever reason and in whichever way). It is easy to see that, when $X = A$, the incentive constraints imply that (A, p)-$CE(G) = CE(G)$ regardless of the value of p.
- (ii) *p-belief constraints*, meaning that at interim level, each player assigns probability at least p to the *rest* of players all choosing action profiles from $X_{-i} = \prod_{j\neq i} X_i$. The belief constraints can reflect bounds or statistical regularities with which deviations from "rationality" are roughly known to occur, by restricting the probability of this occurrence to $1 - p$. Note that in case the possibility of irrational behavior is excluded ($p = 1$), then $\bigcup_{X\in C_A} (X, 1)$-$CE(G) = CE(G)$, where $C_A = \prod_{i\in I} 2^{A_i}$.

Finally, note that the computation of (X, p)-correlated equilibria is similar to that of correlated equilibria; it only involves $\sum_{i\in I} (|X_i| (|A_i| - 1) + |A_i|) \leq \sum_{i\in I} |A_i| (|A_i| + 1)$ linear incentive and p-belief constraints. Now, in order to complete our characterization result in Theorem 1 we need to recall the notion of *doubled game* due to Aumann and Dreze (2008):

Definition 3 (Doubled game, cf. Aumann and Dreze (2008)) *Let G be a game. Then, the* doubled game *of G is defined as the tuple $2G = \langle I, (A'_i)_{i\in I}, (u'_i)_{i\in I}\rangle$, where for each player i:*

- (i) *$A'_i = A_i \times \{1, 2\}$ is player i's set of pure actions. With some abuse of notation, we denote a generic element of $A' = \prod_{i\in I} A'_i$ by (a, ν), where, for $\nu \in \{1, 2\}^I$, ν_i specifies which copy of A_i in A'_i player i's pure action belongs to.*
- (ii) *$u'_i : A' \to \mathbb{R}$, given by $(a, \nu) \mapsto u_i(a)$ is player i's payoff function,*

Thus, in this context, when writing the action spaces of the game $2G$ as $A'_i = A_i \times \{1, 2\}$ we mean that for each player there are two *copies* of the original action space A_i, with the same payoffs as in G. Note that any distribution on the action profiles of $2G$, $\hat{\pi} \in \Delta(A')$, induces a distribution on the action profiles of G in a natural way by taking the marginal on A, that is, $\pi = \operatorname{marg}_A \hat{\pi}$. For any subset $Y \subseteq \Delta(A')$ we denote $\mathbf{marg}_A(Y) = \bigcup_{\hat{\pi}\in Y} \{\operatorname{marg}_A \hat{\pi}\}$. These elements provide all the tools required to go on with the characterization of the set of p-rational outcomes of the game G.

Our next theorem shows that these can be expressed in terms of computationally simple (X, p)-correlated equilibria of the doubled game $2G$. The intuition for the proof is as follows. A doubled game can be seen as splitting players' actions into ones chosen by the rational type (in $A_i \times \{1\}$) and by the irrational type (in

$A_i \times \{2\})$. Then, for each $X = \prod_{i \in I}(A_i \times \{1\})$, the (X, p)-correlated equilibria are distributions on $\Delta(A')$ that by construction satisfy the incentive constraints just for the rational types, and where the p-belief constraints ensure that all players believe at interim level that others play rationally with probability p or more. Finally, taking marginals ensures that the distributions are on $\Delta(A)$.[10]

Theorem 1 (Strategic characterization of p-rational outcomes) *Let G be a game and p, a probability. Then, distribution $\pi \in \Delta(A)$ is a p-rational outcome of G if and only if it is the distribution in $\Delta(A)$ induced by some $(A_{(1)}, p)$-correlated equilibrium of $2G$, where $A_{(1)} = \prod_{i \in I}(A_i \times \{1\})$. Formally,*

$$p\text{-}RO\,(G) = \mathbf{marg}_A \left(\left(A_{(1)}, p\right) \text{-}CE\,(2G) \right),$$

This characterizes behavior in a game G under MB^{p}R and the common prior assumption, or, in other words, all behavior in G representing the following two aspects of bounded rationality: (i) some (possibly small) amount of nonrational behavior, and (ii) the capacity of players to expect, and optimally react, to nonrational behavior by their opponents. In particular, Theorem 1 implies that, given the structure of the doubled game and of its $(A_{(1)}, p)-$correlated equilibria, the set of p-rational outcomes of G as subset of $\Delta(A)$ is characterized by $\sum_{i \in I} |A_i| \left(|A_i| + 1\right)$ linear inequalities, (of which $\sum_{i \in I} |A_i| \left(|A_i| - 1\right)$ are incentive constraints and $\sum_{i \in I} 2|A_i|$ are belief constraints), which in turn are all linear functions of the payoffs of the original game G and the probability p. Thus, the characterization of the set of p-rational outcomes is similar, in computational terms, to that of the set of correlated equilibria, which is defined by $\sum_{i \in I}(|A_i|(|A_i)| - 1)$ linear inequalities.

3.3 Examples

The following examples illustrate the p-rational outcomes for two simple 2×2 games. Consider first the following game G_D, solvable by strict dominance with corresponding augmented game $2G_D$,

$G_D \equiv$

	L	R
T	2,2	1,1
B	1,1	0,0

, $2G_D \equiv$

	$(L,1)$	$(R,1)$	$(L,2)$	$(R,2)$
$(T,1)$	2,2	1,1	2,2	1,1
$(B,1)$	1,1	0,0	1,1	0,0
$(T,2)$	2,2	1,1	2,2	1,1
$(B,2)$	1,1	0,0	1,1	0,0

.

To compute p-$RO(G_D)$ we compute first $(A_{(1)}, p)$-$CE(2G_D)$ and apply Theorem 1. Notice that the strategies $(B, 1)$ and $(T, 2)$ of the row player and $(R, 1)$ and $(L, 2)$ of the column player are strictly dominated, so that the remaining constraints that need to be satisfied are the p-belief constraints. This gives:

$$p\text{-}RO(G_D) = \left\{ \pi \in \Delta(A) \,\middle|\, \begin{array}{l} \pi_{TL} \geq p \cdot (\pi_{TL} + \pi_{TR}), \pi_{BL} \geq p \cdot (\pi_{BL} + \pi_{BR}) \\ \pi_{TL} \geq p \cdot (\pi_{TL} + \pi_{BL}), \pi_{TR} \geq p \cdot (\pi_{TR} + \pi_{BR}) \end{array} \right\}.$$

Figures 3 and 4 show the set p-$RO(G_D)$ for $p = 0.95$ and $p = 0.80$ together with respectively the ε-neighborhood of the set of correlated equilibria of G_D, $N_\varepsilon(CE(G_D))$, and the set of ε-correlated equilibria, ε-$CE(G_D)$,[11] both with $\varepsilon = 0.20$. Clearly the three sets are all distinct. Consider now the following version

[10] Clearly, due to the symmetric role of the different copies of the action spaces of $2G$, the theorem would also hold for any $X = A \times \{\nu\}$, whereby only one of the two copies of players' actions satisfies the incentive constraints.

[11] In general, this is the set of probability distributions $\pi \in \Delta(A)$ that satisfy the incentive constraints for correlated equilibria with a slack of ε, analogous to Radner's ε-Nash equilibria, formally, π is an *ε-correlated equilibrium* (ε-CE) if for any $i \in I$,

$$\sum_{a_i \in A_i} \max_{a'_i \in A_i} \sum_{a_{-i} \in A_{-i}} \pi[(a_{-i}; a_i)] \cdot \left(u_i(a_{-i}; a'_i) - u_i(a_{-i}; a_i)\right) \leq \varepsilon.$$

G_{MP} of matching pennies, with corresponding doubled game $2G_{MP}$,

$G_{MP} \equiv$

	L	R
T	1,0	0,1
B	0,1	1,0

, $2G_{MP} \equiv$

	$(L,1)$	$(R,1)$	$(L,2)$	$(R,2)$
$(T,1)$	1,0	0,1	1,0	0,1
$(B,1)$	0,1	1,0	0,1	1,0
$(T,2)$	1,0	0,1	1,0	0,1
$(B,2)$	0,1	1,0	0,1	1,0

.

The set p-$RO(G_{MP})$ is now somewhat more tedious to characterize, nonetheless we know it is a compact, convex polyhedron around $\bar{\pi} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, which converges to $\bar{\pi}$ as p converges to 1. In particular it contains profiles that do not yield the agents their value of the game, but rather something in a neighborhood thereof. Figures 5 and 6 show the set p-$RO(G_{MP})$ for $p = 0.95$ together with the ε-neighborhood of the set of correlated equilibria, $N_\varepsilon(CE(G_{MP}))$, and the set of ε-correlated equilibria, ε-$CE(G_{MP})$, both with $\varepsilon = 0.10$, respectively. Again, the sets p-$RO(G_{MP})$ and $N_\varepsilon(CE(G_{MP}))$ and ε-$CE(G_{MP})$ are visibly distinct.

3.4 Further properties of p-rational outcomes

The next results further characterize the structure and nature of the set of p-rational outcomes. The first shows that as p converges to 1 the p-rational outcomes converge to the set of correlated equilibria. But more generally it also shows that the p-rational outcomes *always* vary continuously in p,[12] at any $p \in [0,1]$; and go from being the entire set $\Delta(A)$ when $p = 0$ to the set of correlated equilibria when $p = 1$.

Theorem 2 (Topological properties of the set of p-rational outcomes) *Let G be a game and p a probability. Then the set of p-rational outcomes of the game G is a nonempty, convex, compact set that varies continuously in p. Moreover, for $p = 0$, we have 0-$RO(G) = \Delta(A)$, for $p = 1$, we have 1-$RO(G) = CE(G)$, and for any $p \in [0,1)$, we have $dim[p\text{-}RO(G)] = dim[\Delta(A)]$.*

The very last statement further shows that *all* strategies can be in the support of p-rational outcomes whenever $p < 1$. The next result qualifies this by showing that if p is close enough 1, then rationalizable strategy profiles or ones that survive the iterated elimination of strictly dominated strategies get a total weight of at least p. This can be interpreted as the p-rationality counterpart of the fact that strategies that do not survive the iterated elimination of strictly dominated strategies are not in the support of correlated equilibria. In what follows we denote by A^∞ the set of all strategy profiles that survive the iterated elimination of strictly dominated strategies.

Proposition 1 (Rationalizability and p-rational outcomes) *Let G be a game. Then there exists $\bar{p} < 1$ such that $\pi[A^\infty] \geq p$ for any $\pi \in p$-$RO(G)$ and any $p \geq \bar{p}$.*

This shows that if the probability p with which the opponents' rationality is believed is sufficiently high, then the probability with which players play rationalizable strategies is also high. In other words, besides being close to the correlated equilibria in a topological sense, the p-rational outcomes, for p close to one, will be close also in terms of their support.

The next result relates the p-rational outcomes to Bayesian Nash equilibria of incomplete information games, where there is uncertainty not about the rationality of the players but about their payoffs. To make the connection precise, we recall some definitions from the work of Kajii and Morris (1997*a*,*b*). They

[12] A correspondence is *continuous* if it is both upper- and lower hemicontinuous; see, e.g., Ch. 17 in Aliprantis and Border (2006) for further details and related definitions. The topology is the standard one, inherited from Euclidean spaces.

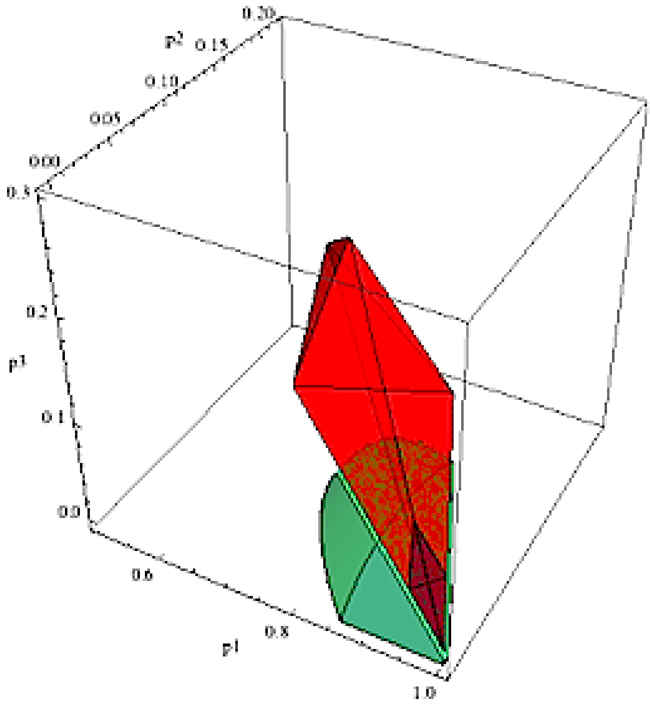

Figure 3: 0.95-$RO(G_D)$ (blue), 0.80-$RO(G_D)$ (red), $N_{0.10}CE(G_D)$ (green)

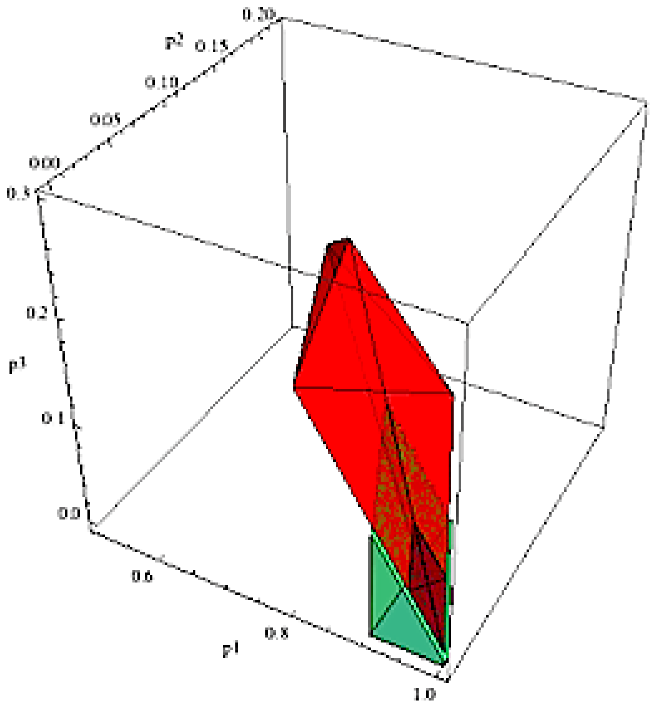

Figure 4: 0.95-$RO(G_D)$ (blue), 0.80-$RO(G_D)$ (red), 0.10-$CE(G_D)$ (green)

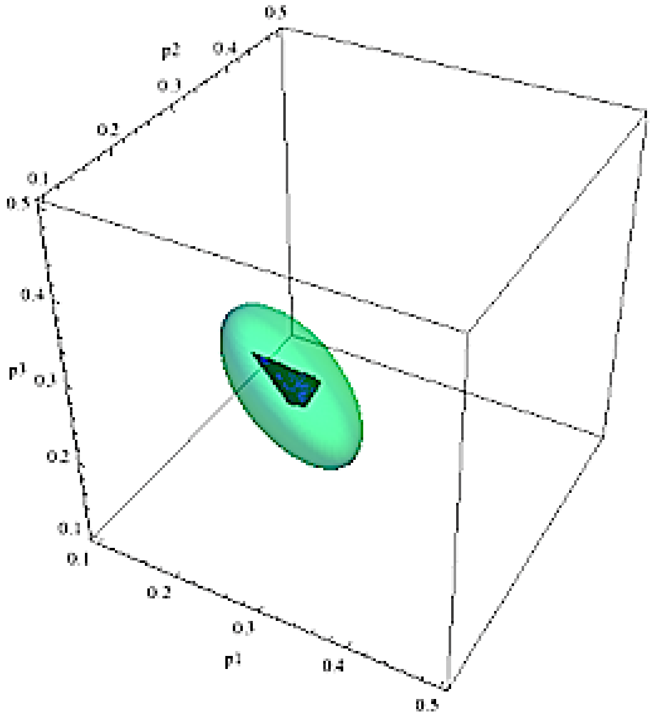

Figure 5: 0.95-$RO(G_{MP})$ (blue), $N_{0.10}CE(G_{MP})$ (green)

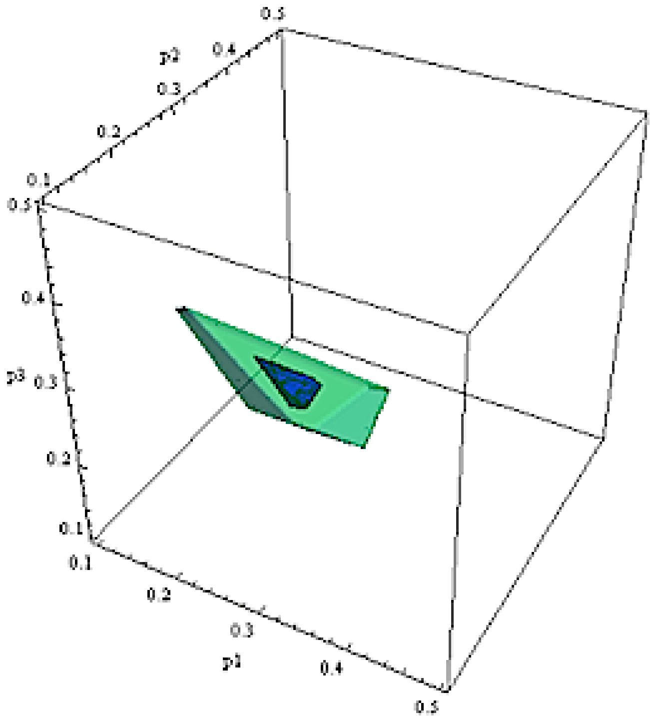

Figure 6: 0.95-$RO(G_{MP})$ (blue), 0.10-$CE(G_{MP})$ (green)

define incomplete information games they call elaborations, to be viewed as perturbations of the underlying complete information game G, in order to study notions of robust equilibrium. We follow their approach, especially Kajii and Morris (1997*a*), in defining the notion of (canonical) elaboration. A player's set of possible types is written as the union $T_i = T_i^s \cup T_i^c$, where T_i^s is a countable set of *standard* types whose payoffs coincide with the ones of the game G, while $T_i^c \equiv A_i$ is a set of *committed* types who have a strictly dominant action to play the strategy corresponding to their type. The set of all type profiles is then $T = \Pi_{i \in I} T_i$. We can then define an *elaboration* of the game G as an incomplete information game (G, P) with type space T, probability distribution $P \in \Delta(T)$, and payoff functions,

$$\widetilde{u}_i(a_i, a_{-i}; t) = \begin{cases} u_i(a_i, a_{-i}) & \text{if } t_i \in T_i^s \\ 1 & \text{if } a_i = t_i \in T_i^c \\ 0 & \text{if } a_i \neq t_i \text{ and } t_i \in T_i^c, \end{cases}$$

where u_i is the payoff function of the original game G. One can define Bayesian Nash equilibria as profiles of strategies $\alpha_i : T \to \Delta(A_i)$ as usual, and, using the above form for the payoff function $\widetilde{u}_i$, it is easy to see that a strategy profile α is an equilibrium if, for any player i and any type t_i with $P[t_i] > 0$, we have

(*i*) $\alpha_i(t_i)[a_i] = 1$ if $a_i = t_i \in T_i^c$, and

(*ii*) $a_i \in \operatorname{argmax}_{a_i' \in A_i} \sum_{t_{-i} \in T_{-i}} P[t_{-i} \,|\, t_i] \cdot \sum_{a_{-i} \in A_{-i}} \alpha_{-i}(t_{-i})[a_{-i}] \cdot u_i(a_i', a_{-i}) t$ if $t_i \in T_i^s$ and $\alpha_i(t_i)[a_i] > 0$, where $\alpha_{-i}(t_{-i})[a_{-i}] = \prod_{j \neq i} \alpha_j(t_j)[a_j]$.

We then say distribution $\mu \in \Delta(A)$ is an *equilibrium action distribution* (*EAD*) of (G, P) if there is an equilibrium α of (G, P) with $\mu[a] = \sum_{t \in T} P(t) \cdot \alpha(t)[a]$ for all $a \in A$, where $\alpha(t)[a] = \prod_{i \in I} \alpha_i(t_i)[a_i]$. We are then in a position to relate our p-rational outcomes of G to *EAD*'s of elaborations (G, P) putting sufficient probability on standard types.

Proposition 2 (Bayesian Nash equilibria and p-rational outcomes) *Let G be a game and p a probability. Then the distribution $\pi \in \Delta(A)$ is a p-rational outcome of G if and only if π is an equilibrium action distribution of an elaboration (G, P), where $P \in \Delta(T)$ is a probability measure on the set of types, that satisfies $P[T_{-i}^s \,|\, t_i] \geq p$, for any $t_i \in T_i$.*

Perhaps not surprisingly, the p-rational outcomes of a game G can be expressed as Bayesian Nash equilibria of an incomplete information game, where players believe their opponents are "standard" types (i.e., payoff-maximizing in the original game G) with probability at least p and can include "committed" types (i.e., committed to playing fixed, not necessarily payoff-maximizing strategies in G) with the remaining probability at most $1 - p$. Essentially, the "standard" types of the elaborations correspond to our rational types, whereas the "committed" types correspond to our nonrational types; the difference is that, in the elaborations, the "committed" types' strategies become "rational" given the "fictitious" payoffs assigned through the functions $\widetilde{u}_i$. Put differently, the proposition provides an epistemic foundation for Bayesian Nash equilibria of incomplete information games that are "elaborations" of the original game G with "committed" types. In particular, these equilibria assume a common prior and common knowledge of rationality.

4 Some Extensions

In this section, we extend our analysis of p-rational outcomes and relate them to further concepts studied in the literature. First, we present the bounded rationality counterpart of Aumann and Dreze's (2008) notion of rational expectations by providing, in Theorem 3, a strategic characterization of the interim expected payoffs of p-rational outcomes. Then, in Section 4.2, we extend our basic framework to games of incomplete information and relate the resulting p-rational Bayes outcomes to the notion of Bayes correlated equilibrium of Bergemann and Morris (2015).

4.1 p-Rational expectations

Aumann and Dreze (2008) define the *rational expectations* of game G as the set of interim expected payoffs of each player, given some belief system that satisfies common knowledge of rationality and the common prior assumption. Thus, according to these authors, rational expectations can be identified with the *value* of game G; that is, with the set of payoffs players can *reasonably* expect when taking part in G (being 'reasonably' understood, in their context, as consistent both with common knowledge of rationality and the common prior assumption). Following Aumann and Dreze's (2008) criteria, it is straightforward to extend their definition of rational expectation beyond common knowledge of rationality, so that it covers the aspects of bounded rationality we are interested in:

Definition 4 (p-Rational expectation) *Let G be a game, p a probability, and B a belief system satisfying MB^pR and the common prior assumption. Then, a p-*rational expectation *in G is the interim expected payoff of some player. We denote the set of p-rational expectations of G by $p\text{-}RE(G)$.*

We can interpret the set of p-rational expectations of game G as the set of payoffs players can *reasonably* expect when taking part in G, where now 'reasonably' is to be understood as consistent with MB^pR and the common prior assumption. The main result by Aumann and Dreze (2008) provides a characterization of the set of rational expectations of game G in terms of the correlated equilibria of the doubled game, $2G$. Specifically, they show that each player i's set of rational expectations is, exactly, the set of interim expected payoffs of some correlated equilibrium of $2G$. In order to provide a similar characterization for p-rational expectations, we need to invoke the notion of *tripled game*, which is analogous to that of doubled game in Section 3.2 but implies adding an additional payoff-irrelevant *copy* of each player's set of actions. Thus, the tripled game consists then on the list $3G = \langle I, (A'_i)_{i\in I}, (u'_i)_{i\in I}\rangle$, where for each player i we have set of actions $A'_i = A_i \times \{1, 2, 3\}$, and playoff map u'_i is again indifferent to which copy of each of their action players are choosing. Then, the set of p-rational expectations is characterized as follows:

Theorem 3 (Characterization of p-rational expectations) *Let G be a game, p, a probability, and let $A_{(1)} = \prod_{i\in I}(A_i \times \{1\})$. Then:*

- *(i) The p-rational expectations in G are the interim expected payoffs of the $(A_{(1)}, p)$-correlated equilibria of the tripled game $3G$ when playing an action in $A_i \times \{1, 2, 3\}$.*
- *(ii) The p-rational expectations of rational players are the interim expected payoffs of the $(A_{(1)}, p)$-correlated equilibria of the tripled game $3G$ when playing an action in $A_i \times \{1\}$.*

Theorem 3 shows that each player i's set of p-rational expectations is, exactly, the set of interim expected payoffs of some $(A_{(1)}, p)$-correlated equilibrium of $3G$. Furthermore, it also shows that, if we are only interested in the p-rational expectation of those 'types' of player i that act rationally, then this set is, exactly, the set of interim expected payoffs of some $(A_{(1)}, p)$-correlated equilibrium of $3G$, conditional on playing, not any arbitrary action, but rather, only those in $A_i \times \{1\}$.

4.2 Incomplete information

4.2.1 Preliminaries

Bayesian games. We follow the formalization in Lehrer *et al.* (2010, 2013) and Bergemann and Morris (2015) that splits the Bayesian game in two components, so that strategic and informational aspects can be studied separately. First, we have a *game with incomplete information* $G = \left\langle I, \Theta, \psi, (A_i)_{i\in I}, (u_i)_{i\in I}\right\rangle$, where: I is a finite set of players, Θ is a finite set of *states of nature*, $\psi \in \Delta(\Theta)$ is a *common prior* with full support, and, for any player i, we have a finite set of *actions* A_i, and a *payoff function* $u_i : A \times \Theta \to \mathbb{R}$, where $A = \prod_{i\in I} A_i$ is the set of *action profiles*. The second component is an *information structure* $S = \left\langle (T_i)_{i\in I}, \sigma\right\rangle$, where each T_i is a finite set of *signals* (or *types*) for player i, and we have *signal distribution* $\sigma : \Theta \to \Delta(T)$, where $T = \prod_{i\in I} T_i$ is the set of *signal profiles*. A *Bayesian game* consists then on a pair (G, S) in which the interaction proceeds as follows:

- A state of nature θ is randomly drawn with probability $\psi[\theta]$.
- A profile of types t is randomly drawn with conditional probability $\sigma[t|\theta] = \sigma(\theta)[t]$.
- Each player i, who privately receives signal t_i, chooses an action a_i and gets payoff $u_i((a_{-i}; a_i), \theta)$.

Belief systems (the Bayesian game case). The concept of *belief system* is extended in order to be able to include payoff-uncertainty and information structures. This way, in the present context a belief systems consists on a list $B = \left\langle \Omega, (\Pi_i)_{i\in I}, \mu_i, \kappa, (\alpha_i)_{i\in I}, (\tau_i)_{i\in I}\right\rangle$, where (i) Ω is a finite set of *states of the world*, (ii) each Π_i is a partition of Ω, (iii) $\kappa : \Omega \to \Theta$ is a random variable that assigns a state of nature to each state of the world, and (iv) for any player i we have random variables $\alpha_i : \Omega \to A_i$ and $\tau_i : \Omega \to T_i$, both measurable w.r.t. Π_i, that respectively determine the action and signal corresponding to player i at each state of the world. Finally, $\mu_i \in \Delta(\Omega)$ is a prior belief with full support. Again, we say that B satisfies the common prior assumption if all the players have the same prior belief. Following Bergemann and Morris (2015), we assume that a belief model always satisfies the following standard condition that excludes informational incompatibilities between the information structure and the belief system:

> **Consistency.** For any player i, $\mu_i[\tau = t, \kappa = \theta] = \psi(\theta) \cdot \sigma[t\,|\theta]$, for any type profile t and any state of nature θ.

Interim beliefs are defined exactly as in Section 2.2, and thus, at state ω, each player i's interim beliefs about opponents' behavior and the state of nature are given by $\mu_i(\omega)[(a_{-i}, \theta)] = \mu_i(\omega)[\kappa^{-1}(\theta) \cap \bigcap_{j\neq i} \alpha_j^{-1}(a_{-i})]$, for any $(a_{-i}, \theta) \in A_{-i} \times \Theta$. These beliefs induce an interim expected payoff, for each action a_i,

$$\mathbb{E}_B(\omega)\left[u_i\left((\alpha_{-i}; a_i); \kappa\right)\right] = \sum_{a_{-i}\in A_{-i}} \mu_i(\omega)[a_{-i}] \cdot u_i\left(a_{-i}; a_i\right),$$

Again, we say that player i is rational at state ω, if $\alpha_i(\omega) \in \operatorname{argmax}_{a_i\in A_i} \mathbb{E}(\omega)\left[u_i\left((\alpha_{-i}, a_i), \kappa\right)\right]$, and denote the event that player i is rational by R_i. The rest of epistemic notions defined in sections 2.2 and 3.1, in particular that of MB^{p}R are straightforwardly adapted to the incomplete information case studied here. Then, similarly as in Section 3.1, in the present context, the aspects of bounded rationality we are interested in are represented by the following notion:

Definition 5 (p-Rational Bayes outcome) *Let (G, S) be a Bayesian game, and p, a probability. Then, we say that distribution $\pi \in \Delta(T \times A \times \Theta)$ is a* p-rational Bayes outcome *if it is induced by some belief system*

B that satisfies consistency, MB^{p}R and the common prior assumption. We denote the set of p-rational Bayes outcomes by p-RBO(G,S).

4.2.2 p-Rational Bayes outcomes: Strategic characterization

The characterization of the set of p-rational Bayes outcomes follows a similar pattern as the characterization of p-rational outcomes of a game with complete information: first, we need to generalize the notion of correlated equilibria (Bergemann and Morris's (2015) version, in this case of incomplete information) so that it accounts for nonoptimal behavior; second, we need to introduce the counterpart of doubled game corresponding to the game with incomplete information.

Definition 6 ((X,p)-Bayes correlated equilibrium) *Let (G,S) be a Bayesian game. Then, for any $X = \prod_{i\in I} X_i \subseteq A$ and any probability p we say that distribution $\pi \in \Delta(T\times A\times\Theta)$ is a (X,p)-*Bayes correlated equilibrium *if the following hold:*

- (i) **Consistency constraints:** *For any type profile t and any state of nature θ, $\pi[(t,\theta)] = \psi[\theta]\,\sigma[t|\theta]$.*
- (ii) **Incentive constraints:** *For any player i, any type $t_i\in T_i$ and any action $a_i\in X_i$,*

$$a_i \in \operatorname{argmax}_{a_i'\in A_i} \sum_{a_{-i}\in A_{-i}} \sum_{\theta\in\Theta} \pi[(t_i,(a_{-i};a_i),\theta)]\cdot u_i\left((a_{-i};a_i'),\theta\right).$$

- (iii) **p-Belief constraints:** *For any player i, any type $t_i\in T_i$ and any action $a_i\in A_i$,*

$$\pi\left[X_{-i}\times\{(t_i,a_i)\}\right] \geq p\cdot\pi\left[A_{-i}\times\{(t_i,a_i)\}\right].$$

We denote the set of (X,p)-Bayes correlated equilibria of (G,S) by (X,p)-BCE(G,S).

It is easy to see that, if the amount of bounded rationality vanishes, $p=1$, then every (X,p)-Bayes correlated equilibrium is also a Bayes correlated equilibrium as defined by Bergemann and Morris (2015).[13] The only remaining element in order to present the characterization result is then an appropriate version of the doubled game:

Definition 7 (Doubled game, the incomplete information case) *Let G be a game with incomplete information. Then, the* doubled game *of game G is defined as the tuple $2G = \left\langle I,\Theta,\psi,(A_i')_{i\in I},(u_i')_{i\in I}\right\rangle$, where for each player i:*

- (i) *$A_i' = A_i\times\{1,2\}$ is player i's set of pure actions. With some abuse of notation, we denote a generic element of $A' = \prod_{i\in I} A_i'$ by (a,ν), where, for $\nu\in\{1,2\}^I$, ν_i specifies which copy of A_i in A_i' player i's pure action belongs to.*
- (ii) *$u_i' : A'\Theta\to\mathbb{R}$, given by $((a,\nu),\theta)\mapsto u_i(a,\theta)$ is player i's payoff function,*

Finally, given a Bayesian game $(2G,S)$, we can project outcome distributions of the doubled game into outcome distributions of the original game. Let $\mathbf{marg}_{T\times A\times\Theta}(Y) = \left\{\text{marg}_{T\times A\times\Theta}\hat\pi \,|\, \hat\pi\in Y\right\}$ for any subset $Y\subseteq\Delta(T\times A'\times\Theta)$. Then, the characterization result in this case becomes:

[13] In such case $p=1$, the p-beliefs constraints only hold if $X = \text{supp}\,(\text{marg}_A\pi)$, and thus, the incentive constraints are satisfied if and only if π satisfies what Bergemann and Morris call *obedience*.

Theorem 4 (Strategic characterization of p-rational Bayes outcomes) *Let (G, S) be a Bayesian game and p, a probability. Then, distribution $\pi \in \Delta(T \times A \times \Theta)$ is a p-rational Bayes outcome of (G, S) if and only if it is the distribution in $\Delta(T \times A \times \Theta)$ induced by some $(A_{(1)}, p)$-Bayes correlated equilibrium of $(2G, S)$, where $A_{(1)} = \prod_{i \in I}(A_i \times \{1\})$. Formally,*

$$p\text{-}RBO\,(G, S) = \mathbf{marg}_{T \times A \times \Theta}\left(\left(A_{(1)}, p\right)\text{-}BCE\,(2G, S)\right),$$

This is parallel to our characterization result for the complete information case. Finally, consider a game with incomplete information (G, S) and belief system B. Besides consistency, MB^{p}R and the common prior assumption, and following Forges (1993, 2006), we can impose the following additional conditions on B:

(1) *Informational sufficiency of the joint type:* $\mu\left[\kappa = \theta \,|\vee_{i \in I} \Pi_i\right] = \mu\left[\kappa = \theta \,|\tau\right]$ for any state of nature θ.

(2) *Informational sufficiency of individual types:* $\mu\left[\tau_{-i} = t_{-i}, \kappa = \theta \,|\Pi_i\right] = \mu\left[\tau_{-i} = t_{-i}, \kappa = \theta \,|\tau_i\right]$ for any any player i, any partial profile of types t_{-i} and any state of nature θ.

Then it is easy to see that, if we impose only (1), or both (1) and (2), the respective distributions induced in $A \times \Theta$ are a bounded rationality counterpart of Forges' *Bayesian solution* and *belief invariant Bayesian solution.*

5 On p as an empirical measure of rationality

In the previous section we computed, for a given game G and for a given value $p \in [0, 1]$, the set of all distributions of play, $\pi \in \Delta(A)$, making up the p-rational outcomes. In this section, we go the other way around and compute for a game G and for a given distribution of play π, the unique largest value of p, say $\overline{p}$, that is compatible with π being a p-rational outcome. We then look again at games played in the field or in experimental settings and compute, for the observed distributions of play, the unique largest value $\overline{p}$ that is consistent with the empirical distribution of play π^{emp}. We argue that $\overline{p}$ can be interpreted as a lower bound measure for the degree of "rationality" understood as possible payoff-maximizing behavior that is compatible with the empirical frequency of play π^{emp}. We now make this more precise.

Recall that from Theorem 2 it follows that the set of p-rational outcomes is always compact and that it varies continuously in p. Moreover, since it goes from being the set of correlated equilibria (when $p = 1$) to being the entire set $\Delta(S)$ (when $p = 0$), it immediately follows that, for any given distribution of play $\pi \in \Delta(A)$, for any finite normal form game G, it is possible to compute a unique $\overline{p} \in [0, 1]$ such that:

$$\overline{p} = \max\left\{p \in [0, 1] \,|\, \pi \in p\text{-RO}(G)\right\}.$$

By definition of the p-rational outcomes, $\overline{p}$ is also the largest value of p consistent with common knowledge of mutual p-belief of opponents' rationality (MB^{p}R) for the distribution π. In particular, assuming the payoffs are the ones given in G, this means that at the distribution π, every player chooses actions that are consistent with common knowledge of payoff maximization with probability at least $\overline{p}$. (Notice that payoff-maximizing here is relative to some $\overline{p}$-rational belief system B deduced from π, see Section 3 for the definitions.) Moreover, given Theorem 1, the p-rational outcomes are defined by finitely many linear inequalities so that the value $\overline{p}$ is relatively easy to compute.

Therefore, the unique value $\overline{p} \in [0, 1]$ defined above can be interpreted as the largest level of rationality in a given observers data point π such that, for each player, a fraction $\overline{p}$ of his or her actions are consistent with common knowledge of payoff-maximization given the distribution of play π. This provides a unique value that can be computed for any observed finite strategic interaction or game played in an experimental setting, including incomplete information games. Moreover, the obtained measure $\overline{p}$ has the same interpretation and is thus comparable across games.

Especially the recent literature on behavioral game theory has provided many models of bounded rationality in games. Some of the most successful ones include the quantal response equilibria of McKelvey and Palfrey (1995), and the k-level reasoning models of Stahl and Wilson (1994, 1995), Costa-Gomes *et al.* (2001), and Camerer (2003).[14] These models indirectly provide measures of nonrational behavior that can be applied to experimental or field data.[15] Without questioning the models' success at explaining and predicting strategic behavior in different experimental settings, we believe the corresponding measures, as summary indicators for the level of rationality of a given interaction, are not as suitable as as our measure $\overline{p}$ for the following reasons. On one hand, the level λ of fitted λ-quantal response models is not comparable without renormalization of the payoffs of the game. On the other hand, the estimated levels k from k-level reasoning models need not give a unique or clear-cut value, as they typically consist of a distribution of levels k within the population, and, moreover, the estimated levels k generally depend on an assumed level 0. We now discuss some experimental data, including games, where the common prior assumption is unlikely to hold. In such cases, our measure $\overline{p}$ is a *lower* bound on the maximum frequency of actions consistent with payoff maximization and common knowledge of payoff maximization.

Consider again the penalty kick game (G_{PK}) based on penalty kicks shot by professional soccer players in European leagues, represented in Figure 1 from the Introduction. For the empirical frequencies provided, we compute a value of $\overline{p} \approx 0.96$ confirming its closeness to the unique equilibrium of the game.[16] As a second, closely related example, consider the following two matching pennies games with similar strategic characteristics as the penalty kicks game, and that were played in a lab.[17] The first is a standard (symmetric) matching pennies games (G_{MP}) and the second is an asymmetric version (G_{AMP}).

$G_{MP} \equiv$

	L	R
T	80,40	40,80
B	40,80	80,40

, $\pi^{emp}_{MP} \approx$

	L	R
T	0.230	0.250
B	0.250	0.270

,

$G_{AMP} \equiv$

	L	R
T	320,40	40,80
B	40,80	80,40

, $\pi^{emp}_{AMP} \approx$

	L	R
T	0.154	0.806
B	0.006	0.034

,

Figure 7: Matching pennies ($\overline{p} \approx 0.96$) and asymmetric matching pennies ($\overline{p} \approx 0.8$)

[14] Camerer and Ho (2015) contains a recent discussion of these different models; Wright and Leyton-Brown (2014) measure and compare quantitatively the predictive performance of various models (including k-level and λ-quantal response models) across a large sample of experiments.

[15] Aumann (1992) proposes a measure of "irrationality" using both probabilities and forgone payoffs; implicitly, our measure also takes into account forgone payoffs.

[16] To give a sense of what the number means in this case, we provide the underlying probability distribution $\overline{\pi}^{2G}_{PK} \in \Delta(2A)$ of the doubled game, that supports the value $\overline{p} \approx 0.96$:

$\overline{\pi}^{2G}_{PK} \approx$

	$(JL,1)$	$(JR,1)$	$(JL,2)$	$(JR,2)$
$(KL,1)$	0.166	0.216	0	0.016
$(KR,1)$	0.252	0.348	0	0
$(KL,2)$	0.002	0	0	0
$(KR,2)$	0	0	0	0

.

Essentially, since the entry at $((KL,1),(JR,2))$ is positive, this indicates that the goalkeepers jumped right a little more than optimal, and since the entry $((KL,2),(JR,1))$ is (barely) positive, this indicates that the kickers kicked left very slightly more than optimal. The remaining entries with $(KL,1)$, $(KR,1)$, $(JL,1)$, and $(JR,1)$ are all consistent with rationality.

[17] The games and frequencies of play are taken from Goeree and Holt (2001).

As Goeree and Holt (2001) explain, the games are chosen such that while the original game "conforms nicely to predictions of Nash equilibrium or relevant refinement," a change in the payoff structure produces a "large inconsistency between theoretical predictions and observed behavior." Therefore, while behavior is close to the predicted (unique) Nash equilibrium in the basic game (G_{MP}), it is less close in the asymmetric version (G_{AMP}). Again, our theory allows to quantify the level of "rationality" and obtains values of $\overline{p} \approx 0.96$ for the first interaction (G_{MP}) and a level of $\overline{p} \approx 0.80$ for the second one (G_{AMP}). Notice that while the asymmetric version (G_{AMP}) was "designed" to generate behavior visibly inconsistent with Nash behavior, the level of "rationality" we find ($\overline{p} \approx 0.80$) is significantly above what we would obtain if players had been choosing their strategies uniformly at random ($\overline{p} \approx 0.25$).[18]

As a third example, consider the games depicted in Figure 8.[19] The first one is solvable in two rounds of strict dominance, whereas the second one is solvable in three rounds.

$G_{DS_2} \equiv$

	L	R
T	75,51	42,27
B	48,80	89,68

, $\pi^{emp}_{DS_2} \approx$

	L	R
T	0.791	0.066
B	0.132	0.011

,

$G_{DS_3} \equiv$

	L	R
T	53,86	24,19
M	79,57	42,73
B	28,23	71,50

, $\pi^{emp}_{DS_3} \approx$

	L	R
T	0	0
M	0.181	0.604
B	0.050	0.165

.

Figure 8: Games solvable by two rounds ($\overline{p} \approx 0.86$) and three rounds ($\overline{p} \approx 0.79$) of strict dominance

In particular, both games have a unique outcome consistent with common knowledge of rationality, which are (T, L) for G_{SD_2}, played with frequency 0.79, and (B, R) for G_{SD_3}, played with frequency 0.165. Our computed level of "rationality" is $\overline{p} \approx 0.86$ for the first and $\overline{p} \approx 0.79$. The lower value of $\overline{p}$ in G_{SD_3} compared with that of G_{SD_2} is consistent with the intuition that coordination that requires higher levels of beliefs (in this case third order beliefs versus second order) is also more difficult to obtain.

In the above games, the assumption of rationality and higher order beliefs in rationality imply a unique outcome, so that the assumption of a common prior is implicit in predicting the equilibrium outcome. For such games, our measure $\overline{p}$ is indeed likely to approximately pick up the degree of "rationality" in the sense of a maximum level p such that every player plays actions consistent with payoff maximization with probability at least p, at the empirical distribution of play π^{emp}.

On the other hand, in games with multiple equilibria, such as the coordination and the Kreps games below, the assumption of a common prior becomes crucial in interpreting the value $\overline{p}$. Consider the following simple (battle of the sexes) coordination game:

$G_{BS} \equiv$

	L	R
T	2,1	0,0
B	0,0	1,2

, $\pi_{BS} =$

	L	R
T	0	0
B	1	0

.

For the extreme case where players play (L, B) with probability 1, this corresponds to a value $\overline{p} = 0$. At the same time, if we do not assume a common prior, the profile (L, B) is consistent with common knowledge of

[18] In a recent article, Martin *et al.* (2014) study chimpanzee behavior in matching pennies games and compare it with human behavior. They suggest that the chimpanzees' choices are closer to Nash equilibrium than humans', by calculating standard deviations of observed choices from the Nash prediction. This is a case where our measure provides a reasonable alternative for assessing and comparing levels of rationality. Notice that equal Euclidean distance from the equilibrium distribution within the strategy simplex need not at all imply equal $\overline{p}$ and vice versa.

[19] These are taken from Costa-Gomes *et al.* (2001). We are grateful to Miguel Costa-Gomes for kindly providing us the data for these experiments.

rationality. (Player 1 believes player 2 will play R, and player 2 believes player 1 plays T; it is a rationalizable profile). In this case, our measure, *confounds* the two possible sources of "nonrationality," namely, nonpayoff-maximizing behavior that is due to lack of rationality and higher order beliefs in rationality or behavior that is due to lack of a common prior. Without knowing whether or not the assumption of a common prior is met, we cannot separate the two, and so the resulting measure $\overline{p}$ cannot be interpreted as a measure of "rationality" in the sense of an approximate maximum probability of payoff-maximizing behavior at the distribution of play π^{emp}.[20] However, and this is important for many cases of empirical relevance, the value $\overline{p}$ can nonetheless be interpreted as a measure of "rationality" in the sense of a *lower bound* on the maximum frequency of behavior that is consistent with common knowledge of payoff maximization at the distribution of play π^{emp}. In other words, it remains true that a computed value $\overline{p}$ implies that, at π, every player chooses actions that are consistent with payoff maximization with probability at least $\overline{p}$, whether or not there is a common prior.[21] The only difference is that without a common prior this need no longer be the maximal such value. As the above example shows, the amount of payoff-maximizing behavior may be above $\overline{p}$ for all players; this cannot happen if there is a common prior.

Finally, consider again the Kreps game G_{Kreps} represented in Figure 2 in the Introduction. Here players typically play a strategy (NN) that is the only strategy that is not in the support of a Nash equilibrium of the game (these are (T, L), (B, R) and a mixed equilibrium $((\frac{30}{31}, \frac{1}{31}); (\frac{1}{21}, \frac{20}{21}, 0, 0)))$. Although the strategy is not part of any Nash equilibrium, it is both rationalizable and in the support of the set of correlated equilibria, and it is played with frequency 0.68. By our measure, the overall frequency of play has a level of "rationality" of at least $\overline{p} \approx 0.7$, whether or not there is a common prior.

6 Concluding remarks

We conclude with a few remarks.

Remark 1 ((p, q)-Rational outcomes) An important objective of the paper was to put as few restrictions on nonrational behavior as possible, so as to cover all sorts of departures from rationality. However, throughout the paper we implicitly assumed – as part of the notion of MB^pR – that players always believe the other players are rational with probability p or more; thus we indirectly assumed that all players have the same p whether or not they are rational at a given state. This is consistent with all players making mistakes with same lower bound probabilities and always being aware of others making mistakes with these lower bound probabilities. Strictly speaking though, it restricts behavior of rational and nonrational types.

A more general benchmark – also in line with our motivation – is to allow for different beliefs in rationality for different players and for different types (whether rational or nonrational at a given state). In particular, we can assume that each player i believes the other players are rational with probability p or more when rational and believes others are rational with probability q or more when not rational; importantly, one can drop any restriction on the nonrational types and directly set $q = 0$ for all $i \in I$, which would allow to not impose *any* belief constraints on nonrational types. This can be formalized assuming a pair of probabilities (p, q), where the components associated to states in which the agents are rational are represented by p, while the components associated to the nonrational states are represented by q This leads to the more general (p, q)-*rational outcomes* of G, (p, q)-$RO(G)$. These are again marginals of $(A_{(1)}, (p, q))$-correlated equilibria

[20] Kneeland (2015) estimates for an interesting class of "ring games" the degrees to which agents are rational, hold beliefs of opponents being rational, and consistency of beliefs, and deduces that deviations from "equilibrium behavior" are largely due to inconsistency of beliefs.

[21] To see this, notice that the computation of the largest $\overline{p}$ consistent with payoff maximization and no further constraint must yield a no smaller $\overline{p}$ than the same computation with the additional constraint of the common prior assumption.

of $2G$, in that they are distributions satisfying the same conditions as the $(A_{(1)}, p)$-$CE(2G)$ except that the p-belief constraints now hold with probabilities p for all rational types, and hold with probability q for all nonrational types. That is, we replace the original p-belief constraints ((ii)) with the more general (p, q)-beliefs constraints of the form,

(ii') For any player i and any $a_i \in X_i$, $\pi\left[X_{-i} \times \{a_i\}\right] \geq p \cdot \pi\left[A_{-i} \times \{a_i\}\right]$

(ii'') For any player i and any $a_i \in A_i \setminus X_i$, $\pi\left[X_{-i} \times \{a_i\}\right] \geq q \cdot \pi\left[A_{-i} \times \{a_i\}\right]$.

Thus, for any player i we have that $(A_{(1)}, p)$-$CE(2G) \subseteq (A_{(1)}, (p, q))$-$CE(2G) \subseteq (A_{(1)}, (p, 0))$-$CE(2G)$ and therefore, p-$RO(G) \subseteq (p, q)$-$RO(G) \subseteq (p, 0)$-$RO(G)$.

While the correspondence (p, q)-$RO(G)$ maintains the basic topological properties of the correspondence p-$RO(G)$, it need not converge to the set of correlated equilibria of G as $(p, q) \to (1, 0)$, i.e., if only the rational types believe opponents are rational with probability 1, but does so if one also requires $(p, q) \to (1, 1)$. This can be seen already in Example 1. A $(1, 0)$-rational belief system can be very far from a $(1, 1)$-rational belief system in that the former need not put any restriction on the total mass of states where all players are rational, $\mu[R]$.[22]

The alternative notion of approximate knowledge of rationality requiring $\mu[CB^p(R)] > 1 - \varepsilon$, for $\varepsilon > 0$, (instead of MB^{p}R), is more flexible with respect to the players' beliefs in that it only restricts the total mass of common p-belief and hence does not specify directly what interim beliefs individual players have. A characterization of p-rational outcomes with this definition is possible along the lines of our Theorem 1, but involves more complicated incentive and p-belief constraints that are imposed over all possible subsets and permutations of players. We leave such a characterization for future work.

Remark 2 (Noncommon priors) Throughout the paper we assumed the existence of common prior beliefs. This together with the notion of MB^{p}R allowed us to derive relatively stringent restrictions on behavior. It is natural to ask, what happens if the common prior assumption is relaxed. As it turns out, under *subjective* or *noncommon* prior beliefs, MB^{p}R puts *no* restrictions on possible behavior – even when $p = 1$.[23] This provides a stark contrast with the behavior under common knowledge of rationality and also common p-belief in rationality as in, respectively, Aumann (1974), Bernheim (1984), Brandenburger and Dekel (1987), Pearce (1984), Tan and Werlang (1988) and Börgers (1994), Hu (2007), Germano and Zuazo-Garin (2015), and in a sense further highlights the stringency of the common prior assumption.[24]

[22] To see that in a $(1, 0)$-rational belief system the total mass of states where agents are nonrational is unrestricted, take the game in Example 1 and consider the belief system B, where $\Omega = A$, $\alpha_i(a_{-i}; a_i) = a_i$, for any player i and any $(a_{-i}; a_i) \in A$, and where $\mu \in \Delta(A)$ is given by $\mu_{TL} = \mu_{TR} = \mu_{BL} = 0$ and $\mu_{BR} = 1$. It can be checked that it is $(1, 0)$-rational and clearly $\mu[R] = 0$. At the same time, in a (p, q)-rational belief system it is always the case that, for any player i and any state ω, $\mu(\omega)\left[R_{-i}\right] \geq q$, hence

$$\mu[R_{-i} \cap \Pi_i(\omega)] \geq q \cdot \mu[\Pi_i(\omega)] \Longrightarrow \sum_{\Pi_i(\omega) \in \Pi_i} \mu[R_{-i} \cap \Pi_i(\omega)] = q \cdot \sum_{\Pi_i(\omega) \in \Pi_i} \mu[\Pi_i(\omega)] \Longrightarrow \mu[R_{-i}] \geq q$$

which besides confirming the expected convergence to the correlated equilibria as $(p, q) \to (1, 1)$, also shows that positive q's do put restrictions on the total mass of states where agents are rational $\mu[R]$.

[23] To see the noncommon prior case, given game G and probability p we say that a family of distributions $(\pi_i)_{i \in I} \in (\Delta(A))^I$ is a *p-subjectively rational outcome* of G (p-$SRO(G)$) if there exists some belief system B that satisfies MB^{p}R, does not satisfy the common prior assumption and induces subjective outcomes $(\pi_i)_{i \in I}$. As shown in the Appendix, it is easy to see that, for any $p \in [0, 1]$, the whole space is obtained, namely:

$$p\text{-}SRO(G) = (\Delta(A))^I.$$

In particular, a player who is certain that all the other agents are rational may still select a nonrational action, and thus, *any* pure strategy profile in A is consistent with MB^{p}R, even when $p = 1$.

[24] Recall that the result of part (i) of Lemma 1 also holds with noncommon priors.

Remark 3 (Comparison with further solution concepts) Our sets of p-rational outcomes define sets of distributions of play that are broader than the correlated equilibria. As the examples show, they are distinct from ε-neighborhoods of the correlated equilibria, and put further structure on the deviations from the set $CE(G)$ that occur as p departs from 1. At the same time, they are distinct from the ε-correlated equilibria, reflecting the fact that they impose no constraints on the *type* of departure from rationality assumed – unlike the ε-optimizers of the ε-correlated equilibria. A similar remark applies to the quantal response equilibria of McKelvey and Palfrey (1995) or other models such as the level-k reasoning models (Stahl and Wilson (1994, 1995), Costa-Gomes *et al.* (2001), Camerer (2003)) that put specific restrictions on how players can deviate from rationality. More closely related are the rationalizable and the p-rationalizable strategy profiles (see respectively Bernheim (1984); Pearce (1984); Dekel *et al.* (2007) and Hu (2007); Germano and Zuazo-Garin (2015)), which are derived at the interim stage and without appealing to priors. Unlike the p-rational outcomes, whose set of distributions is fully supported on A, whenever $p < 1$, both the rationalizable and the p-rationalizable profiles may be strict subsets of A. It remains an empirical question to what extent the p-rational outcomes bound observed behavior in a robust and useful manner.

Remark 4 (Learning to play p-rational outcomes) Clearly, all learning dynamics that lead to correlated equilibria (see e.g., Hart (2005)) will also lead to p-rational outcomes, which includes dynamics that converge in polynomial time (see e.g., Hart and Mansour (2010)). The question arises as to what further dynamics (not necessarily converging to correlated equilibria) may converge to p-rational outcomes and whether they include interesting dynamics that for example allow for faster or more robust convergence.

References

Aliprantis, Charambolos D. and Kim C. Border (2006). *Infinite Dimensional Analysis: A Hitchhiker's Guide, Third Edition*. Springer, Berlin.

Aumann, Robert J. (1974). "Subjectivity and Correlation in Randomized Strategies". *Journal of Mathematical Economics* **1**, 67–96.

Aumann, Robert J. (1987). "Correlated Equilibria as an Expression of Bayesian Rationality". *Econometrica* **55**(1), 1–18.

Aumann, Robert J. (1992). "Irrationality in Game Theory". In: *Economic Analysis of Markets and Games: Essays in Honor of Frank Hahn* (P. Dasgupta, D. Gale, O. Hart and E. Maskin, Eds.). MIT Press, Cambridge.

Aumann, Robert J. and Jacques H. Dreze (2008). "Rational Expectations in Games". *American Economic Review* **98**, 72–86.

Battigalli, Pierpaolo, Alfredo Di Tillio, Eduardo Grillo and Antonio Penta (2011). "Interactive epistemology and solution concepts for games with asymmetric information". *The B.E. Journal of Theoretical Economics* **11**, Article 6.

Bergemann, Drik and Stephen Morris (2015). "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games". forthcoming *Theoretical Economics*.

Bernheim, B. Douglas (1984). "Rationalizable Strategic Behaviour". *Econometrica* **52**, 1007–1028.

Börgers, Tilman (1994). "Weak Dominance and Approximate Common Knowledge". *Jornal of Economic Theory* **64**, 265–276.

Brandenburger, Adam and Eddie Dekel (1987). "Rationalizability and Correlated Equilibria". *Econometrica* **55**, 1391–1402.

Camerer, Colin F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. New York, Princeton University Press.

Camerer, Colin F. and Teck-Hua Ho (2015). "Behavioral Game Theory, Experiments and Modeling". In: *Handbook of Game Theory, Volume 4* (H.P. Young and S. Zamir, Eds.). North-Holland, Elsevier, Amsterdam.

Camerer, Colin F., George Loewenstein and Matthew Rabin (2011). *Advances in Behavioral Economics*. Princeton University Press.

Chiappori, Pierre-Andre, Steven D. Levitt and Timothy Groseclose (2002). "Testing Mixed-Strategy Equilibria When Players are Heterogeneous: The Case of Penalty Kicks in Soccer". *American Economic Review* **92**, 1138–1151.

Conlisk, John (1996). "Why Bounded Rationality?". *Journal of Economic Literature* **34**, 669–700.

Costa-Gomes, Miguel, Vincent P. Crawford and Bruno Broseta (2001). "Cognition and Behavior in Normal-Form Games: An Experimental Study". *Econometrica* **69**, 1193–1235.

Crawford, Vincent P. (2013). "Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games". *Journal of Economic Literature* **51**, 512–527.

Dekel, Eddie, Drew Fudenberg and Stephen Morris (2007). "Interim Correlated Rationalizability". *Theoretical Economics* **2**, 15–40.

Forges, Françoise (1993). "Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information". *Theory and Decision* **35**, 277–310.

Forges, Françoise (2006). "Correlated Equilibrium in Games with Incomplete Information Revisited". *Theory and Decision* **61**, 329–344.

Germano, Fabrizio and Peio Zuazo-Garin (2015). "Uncertain Rationality and Robustness in Games with Incomplete Information". Mimeo.

Goeree, Jacob K. and Charles A. Holt (2001). "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions". *American Economic Review* **91**, 1402–1422.

Harstad, Ronald M. and Reinhard Selten (2013). "Bounded-Rationality Models: Tasks to Become Intellectually Competitive". *Journal of Economic Literature* **51**, 496–511.

Hart, Sergiu (2005). "Adaptive Heuristics". *Econometrica* **73**, 1401–1430.

Hart, Sergiu and Yishay Mansour (2010). "How long to equilibrium? The communication complexity of uncoupled equilibrium procedures". *Games and Economic Behavior* **69**, 107–126.

Hu, Tai-Wei (2007). "On p-rationalizability and approximate common certainty of rationality". *Journal of Economic Theory* **136**, 379–391.

Kajii, Atsushi and Stephen Morris (1997*a*). "Refinements and Higher Order Beliefs: A Unified Survey". Mimeo.

Kajii, Atsushi and Stephen Morris (1997*b*). "The Robustness of Equilibria to Incomplete Information Games". *Econometrica* **65**, 1283–1309.

Kneeland, Terri (2015). "Identifying Higher-Order Rationality". *Econometrica* **83**, 2065–2079.

Lehrer, Ehud, Dinah Rosenberg and Eran Shmaya (2010). "Signaling and Mediation in Games with Common Interest". *Games and Economic Behavior* **68**, 670–682.

Lehrer, Ehud, Dinah Rosenberg and Eran Shmaya (2013). "Garbling of Signals and Outcome Equivalence". *Games and Economic Behavior* **81**, 179–191.

Mallard, Graham (2011). "Modeling Cognitively Bounded Rationality: An Evaluative Taxonomy". *Journal of Economic Surveys* **26**, 674–704.

Martin, Christopher Flynn, Rahul Bhui, Peter Bossaerts, Tetsuro Matsuzawa and Colin Camerer (2014). "Chimpanzee Choice Rates in Competitive Games Match Equilibrium Game Theory Predictions". *Scientific Reports* **4**, 5182.1–6.

McKelvey, Richard D. and Thomas R. Palfrey (1995). "Quantal Response Equilibria for Normal Form Games". *Games and Economic Behavior* **10**, 6–38.

Monderer, Dov and Dov Samet (1989). "Approximating Common Knowledge with Common Beliefs". *Games and Economic Behavior* **1**, 179–190.

Palacios-Huerta, Ignacio (2003). "Professionals Play Minimax". *Review of Economic Studies* **70**, 395–415.

Palacios-Huerta, Ignacio and Oscar Volij (2008). "Experiencia Docet: Professionals Play Minimax in Laboratory Experiments". *Econometrica* **76**, 71–115.

Pearce, David G. (1984). "Rationalizable Strategic Behaviour and the Problem of Perfection". *Econometrica* **52**, 1029–1050.

Rubinstein, Ariel (1998). *Modelling Bounded Rationality*. Cambridge, MA, MIT Press.

Stahl, Dale O. II and Paul W. Wilson (1994). "Experimental evidence on players' models of other players". *Journal of Economic Behavior & Organization* **25**, 309–327.

Stahl, Dale O. II and Paul W. Wilson (1995). "On Players' Models of Other Players: Theory and Experimental Evidence". *Games and Economic Behavior* **10**, 218–254.

Tan, Tommy Chin-Chiu and Sérgio Ribeiro da Costa Werlang (1988). "The Bayesian Foundation of Solution Concepts of Games". *Journal of Economic Theory* **45**, 370–391.

Wright, James R. and Kevin Leyton-Brown (2014). "Level-0 Meta-Models for Predicting Human Behavior in Games". *Proceedings of the 15th Conference on Economics and Computation (EC'14)*.

APPENDIX

A Proof of Lemma 1

(*i*) The left implication is obvious, so let's focus on the right one. By definition, $CB^p(R) \subseteq \bigcap_{i\in I} B_i^p(R)$, and therefore, for any $i \in I$, $CB^p(R) \subseteq B_i^p(R_i)$. Then, since $p > 0$, and $\omega \in B_i^p(R_i)$, we have that $R_i \cap \Pi_i(\omega) \neq \emptyset$, and therefore, that $\Pi_i(\omega) \subseteq R_i$ and, in particular, $\omega \in R_i$. Thus, $B_i^p(R_i) \subseteq R_i$.

(*ii*) It is immediate that $CB^1(R) = \Omega$ if and only of $\bigcap_{i\in I} B_i^1(R) = \Omega$, so it suffices to prove that $\bigcap_{i\in I} B_i^1(R) = \Omega$ if and only if $\bigcap_{i\in I} B_i^1(R_{-i}) = \Omega$. The right implication is immediate. For the proof of the left one, from part (*i*) of the lemma, it is enough to check that if $\bigcap_{i\in I} B_i^1(R_{-i}) = \Omega$ then $R = \Omega$. But this is immediate: take $i, j \in I$, $i \neq j$, then $B_j^{p_j}(R_{-j}) \subseteq B_j^{p_j}(R_i)$, and therefore, $B_j^{p_j}(R_i) = \Omega$. Hence $\mu(R_i) = \sum_{\omega\in\Omega} \mu(R_i \cap \Pi_j(\omega)) = \sum_{\omega\in\Omega} \mu(\Pi_j(\omega)) = 1$. Since μ has full support on Ω, the latter implies that $R_i = \Omega$. As the proof applies for any $i \in I$, we obtain that $R = \Omega$.

(*iii*) If $\bigcap_{i\in I} B_i^p(R_{-i}) = \Omega$, then $R = \bigcap_{i\in I}(B_i^p(R_{-i}) \cap R_i) = \bigcap_{i\in I} B_i^p(R)$. Thus, R is p-evident belief and therefore, $R \subseteq CB^p(R)$. Now, since $\bigcap_{i\in I} B_i^p(R_{-i}) = \Omega$, we have both that $\mu(R_{-i}) \geq p$, and $\mu(R_{-i} \cap R_i) \geq p\mu(R_i)$. The fact that for any $j \neq i$, $\mu(R) = \mu(R_{-i}|R_i)\mu(R_i) \geq p\mu(R_i) \geq p\mu(R_{-j}) \geq p^2$ completes the proof.

B Proofs of the characterization results

In this section we first prove a technical lemma. Then we prove Theorem 4, and then Theorem 1 as a special case of Theorem 4. The technical lemma is the following:

Lemma 2 *Let (G, S) be a Bayesian game, p a probability, $n \in \{2, 3\}$, and $\pi^* \in (A_{(k)}, p)$-$BCE(nG, S)$ for some $k \leq n$. For game nG, let belief system $B = \langle \Omega, (\Pi_i)_{i\in I}, \mu, \kappa, (\alpha_i)_{i\in I}, (\tau_i)_{i\in I} \rangle$ be given by: (i) $\Omega = \{(t, a, \nu, \theta) \in T \times A' \times \Theta \,|\pi^*[(a, \nu, \theta)] > 0\}$, and for any $(t, a, \nu, \theta) \in \Omega$, (ii) $\mu[(t, a, \nu, \theta)] = \pi^*[(t, a, \nu, \theta)]$, (iii) $\kappa(t, a, \nu, \theta) = \theta$, and for any $i \in I$, (iv) we have cells $\Pi_i(t, a, \nu, \theta) = T_{-i} \times A'_{-i} \times \Theta \times \{(t_i, a_i, \nu_i)\}$, (v) $\alpha_i(t, a, \nu, \theta) = a_i$, and (vi) $\tau_i(t, a, \nu, \theta) = t_i$. Then, B is a belief system for (G, S) satisfying consistency, MB^pR and the common prior assumption, and induces π^* in $T \times A' \times \Theta$.*

Proof. It is immediate that α_i and τ_i are measurable w.r.t. Π_i for any $i \in I$. Take $(t, \theta) \in T \times \Theta$; then, we have $\mu[\tau = t, \kappa = \theta] = \pi^*[(t, \theta)] = \psi[\theta] \cdot \sigma[t|\theta]$, and therefore, B satisfies consistency. Now, note first the fact that for any $\omega = (t, a, \nu, \theta) \in \Omega$ and any $a'_i \in A_i$ it holds that,

$$\mathbb{E}(\omega)[u_i(\alpha_{-i}, a'_i, \kappa)] = \sum_{(t'_{-i},(a'_{-i},\nu'_{-i}),\theta')} \pi^*\left[(t'_{-i}, t_i, (a'_{-i}, \nu'_{-i}), (a_i, \nu_i), \theta')\right] \cdot u_i\left(((a'_{-i}, \nu'_{-i}); (a_i, \nu_i)), \theta'\right),$$

together with the incentive constraints, implies that for any $i \in I$ we have that $T \times A'_{-i} \times (A_i \times \{k\}) \times \Theta \subseteq R_i$, and therefore,[25] that $\mu\left[\prod_{j\neq i}(A_j \times \{k\}) \cap \Pi_i(\omega)\right] \leq \mu[R_{-i} \cap \Pi_i(\omega)]$ for any $i \in I$ and any $\omega \in \Omega$. Then, take $i \in I$ and $\omega = (t, a, \nu, \theta) \in \Omega$ and note that:

$$\mu[\prod_{j\neq i}(A_j \times \{k\}) \cap \Pi_i(\omega)] = \pi^*[\prod_{j\neq i}(A_j \times \{k\}) \times \{(t_i, a_i, \nu_i)\}],$$

$$\mu[\Pi_i(\omega)] = \pi^*[A'_{-i} \times \{(t_i, a_i, \nu_i)\}].$$

[25] We abbreviate, $\mu\left[\prod_{j\neq i}(A_j \times \{k\}) \cap \Pi_i(\omega)\right] = \mu\left[\left(T \times \prod_{j\neq i}(A_j \times \{k\}) \times A'_i \times \Theta\right) \cap \Pi_i(\omega)\right]$, with some abuse of notation.

In consequence, due to the p-belief constraints, we have $\mu\left[R_{-i}\cap\Pi_i(\omega)\right]\geq p\cdot\mu\left[\Pi_i(\omega)\right]$, and therefore, that B satisfies MB^{p}R. ■

B.1 Proof of Theorem 4

For the right inclusion, pick $\pi\in p\text{-}RBO\,(G,S)$ and B, a belief model that induces π. Take distribution $\pi^*\in\Delta\left(T\times A'\times\Theta\right)$ given by,

$$\pi^*\left[(t,a,\nu,\theta)\right]=\mu\left[\left[\tau=t,\alpha=a,\kappa=\theta\right]\cap\bigcap_{i:v_i=0}R_i\cap\bigcap_{i:v_i=1}\neg R_i\right]$$

for any $(t,a,\nu,\theta)\in T\times A'\times\Theta$. Then, $\pi^*\in\Delta\left(T\times A'\times\Theta\right)$. The consistency constraint is satisfied, since,

$$\pi^*\left[(t,\theta)\right]=\mu\left[\left[\tau=t,\kappa=\theta\right]\cap\bigcup_{\nu\in N^I}\left(\bigcap_{i:v_i=0}R_i\cap\bigcap_{i:v_i=1}\neg R_i\right)\right]=\mu\left[\tau=t,\kappa=\theta\right]=\psi\left[\theta\right]\cdot\sigma[t|\theta]$$

for any $(t,\theta)\in T\times\Theta$. Now, note that for any $i\in I$, any $a_i,a_i'\in A_i$, any $t_i\in T_i$ and any $\nu_i\in N$ we have that,

$$\begin{aligned}\sum_{(t_{-i},a_{-i},\nu_{-i},\theta)}&\pi^*[(t_{-i},t_i,(a_{-i},\nu_{-i}),(a_i,1),\theta)]\cdot u_i\left((a_{-i},\nu_{-i}),(a_i,1),\theta\right)\\ &-\sum_{(t_{-i},a_{-i},\nu_{-i},\theta)}\pi^*\left[(t_{-i},t_i,(a_{-i},\nu_{-i}),(a_i,1),\theta)\right]\cdot u_i\left((a_{-i},\nu_{-i}),(a_i',\nu_i),\theta\right)=\\ &=\sum_{\omega\in R_i\cap[\tau_i=t_i,\alpha_i=a_i]}\left(\mathbb{E}(\omega)\left[u_i\left(\left(\alpha_{-i},\alpha_i\left(\omega\right)\right),\kappa\right)\right]-\mathbb{E}(\omega)\left[u_i\left(\left(\alpha_{-i},a_i'\right),\kappa\right)\right]\right)\geq 0.\end{aligned}$$

In addition, note that for any $(t_i,a_i,\nu_i)\in T_i\times A_i'$,

$$\begin{aligned}\pi\left[X_{-i}\times\{(t_i,a_i,\nu_i)\}\right]&=\mu\left[R_{-i}\cap\left[\tau_i=t_i,\alpha_i=a_i\right]\cap\left[1_{R_i}=2-\nu_i\right]\right]\\ &=\sum_{\omega\in[\tau_i=t_i,\alpha_i=a_i]\cap\left[1_{R_i}=2-\nu_i\right]}\mu\left[R_{-i}\cap\Pi_i\left(\omega\right)\right]\\ &\geq\sum_{\omega\in[\tau_i=t_i,\alpha_i=a_i]\cap\left[1_{R_i}=2-\nu_i\right]}p\cdot\mu\left[\Pi_i\left(\omega\right)\right]=p\cdot\pi^*\left[A_{-i}'\times\{(t_i,a_i,\nu_i)\}\right].\end{aligned}$$

Thus, both the incentive constraints and the p-belief constraints are satisfied. For the left inclusion, just apply Lemma 2 to $A_i'=A_i\times\{1,2\}$ and $k=1$.

B.2 Proof of Theorem 1

This theorem can be seen as a corollary of Theorem 4. To see this, note that if G is a game with complete information, we can define Bayesian game (G',S), where we have (i) $G'=\left\langle I,\Theta,\psi,\left(A_i\right)_{i\in I},\left(u_i'\right)_{i\in I}\right\rangle$, with $\Theta=\{\theta\}$, $\psi=1_{\{\theta\}}$ and $u_i'(a,\theta)=u_i(a)$ for any $a\in A$ and $i\in I$, and (ii) $S=\left\langle\left(T_i\right)_{i\in I},\sigma\right\rangle$ with $T_i=\{t_i\}$ for any $i\in I$ and $\sigma(t|\theta)=1$. It is immediate that, for any $X=\prod_{i\in I}X_i\subseteq A$ and any $p\in[0,1]$, we have that $(X,p)\text{-}CE\,(G)=(X,p)\text{-}BCE\,(G',S)$. But note also that if we take some list $B'=\left\langle\Omega,\left(\Pi_i\right)_{i\in I},\left(\alpha_i\right)_{i\in I},\left(\tau_i\right)_{i\in I},\mu\right\rangle$, which is a candidate to be a belief system for G', for any $i\in I$, forcefully $\tau_i=1_{\{t_i\}}$; and therefore, B' is a belief system for G' if and only if $B=\left\langle\Omega,\left(\Pi_i\right)_{i\in I},\left(\alpha_i\right)_{i\in I},\mu\right\rangle$ is a belief system for G. Thus, it is immediate that, for any $p\in[0,1]$, we have $p\text{-}RO\,(G)=p\text{-}RBO\,(G',S)$.

So, let G be a game, and $p \in [0,1]$. Then, we just checked above that both $p\text{-}RO\,(G) = p\text{-}RBO\,(G', S)$ and $\left(A_{(1)}, p\right)\text{-}CE\,(2G) = \left(A_{(1)}, p\right)\text{-}BCE\,(2G', S)$, hold, so from Theorem 4 we conclude that $p\text{-}RO\,(G) = \left(A_{(1)}, p\right)\text{-}CE\,(2G)$.

B.3 Proof of Theorem 3

We prove the first statement; the next one concerning the p-rational expectations of rational types then follows directly. We suppose we are taking some player i's expectation. For right inclusion, take p-rational belief system $B = \left\langle \Omega, (\Pi_i)_{i\in I}, (\alpha_i)_{i\in I}, \mu \right\rangle$ and $\omega \in \Omega$. We define:

$$\pi^*_{i,\omega}\left[(a,\nu)\right] = \mu\left[[\alpha = a] \cap W_i \cap \bigcap_{j\neq i, \nu_j = 1} R_j \cap \bigcap_{j\neq i, v_j \neq 1} \neg R_j\right]$$

for any $(a,\nu) \in A'$, where $W_i = R_i \setminus \Pi_i(\omega)$ if $\nu_i = 0$, $W_i = \neg\left(R_i \setminus \Pi_i(\omega)\right)$ if $\nu_i = 1$, and $W_i = \Pi_i(\omega)$ if $\nu_i = 2$. It is immediate that $\pi^*_{i,\omega} \in \Delta(A')$. By an argument similar to the one in the first part of the proof of Theorem 4, reduced to the degenerate case where $|\Theta| = 1$, we can conclude that $\pi_{i,\omega}$ is a $\left(A_{(1)}, p\right)$-correlated equilibrium of G; moreover, it is immediate that player i's expectation conditional on playing $(\alpha_i(\omega), 3)$ induced by $\pi_{i,\omega}$ is exactly $\mathbb{E}(\omega)\left[u_i\left(\alpha_{-i}, \alpha_i(\omega)\right)\right]$. For the left inclusion, take Lemma 2 for the case $A'_i = A_i \times \{1,2,3\}$, $k = 1$, and $|\Theta| = 1$.

C Proof of Theorem 2

Nonemptiness follows from the fact that correlated equilibria always exist for any finite game G and constitute p-rational outcomes for any $p \in [0,1]$. Given that the set of p-rational outcomes is a projection of the (X,p)-correlated equilibria of $2G$, with $X = A_{(k)}$ a copy of the action space of the original game G, the remaining properties follow once they have been shown for the (X,p)-correlated equilibria of $2G$. This is what we do next. For the given game G, define the (X,p)-correlated equilibrium correspondence, where $X = A_{(k)}$ with $k \in \{0,1\}$, is fixed:

$$\begin{array}{rccc} \rho: & [0,1] & \longrightarrow & \Delta(2A) \\ & p & \rightarrow & (X,p)\text{-}CE\,(2G)\ . \end{array}$$

Clearly ρ is convex- and compact-valued; it remains to be shown that it is also continuous. We do this by showing that it is upper- and lower-hemicontinuous (respectively, *uhc* and *lhc*) as a correspondence of p.

uhc

Since $2A$ is finite, $\Delta(2A)$ is compact, and hence upper-hemicontinuity is equivalent to showing that ρ has a closed graph. But this is immediate from inspection of the inequalities defining the sets $(X,p)\text{-}CE\,(2G)$. In particular, all the inequalities are all weak inequalities, linear in p. Moreover, the domain $[0,1]$ is compact.

lhc

Denote by $\Gamma_\rho \subset [0,1] \times \Delta(2A)$ the graph of the correspondence ρ. Fix $(p,\hat{\pi}) \in \Gamma_\rho$ and let $(p^n)_n \subset [0,1]$ be a sequence converging to p. We need to show that there exists a sequence $((\hat{\pi})^n)_n$ converging to $\hat{\pi}$ such that $(\hat{\pi})^n \in \rho(p^n)$ for sufficiently large n. Take the point $(p,\hat{\pi})$. Clearly this satisfies all inequalities defining $\rho(p)$, in particular also the p-rationality constraints. Consider the following sequence $(p^n,\hat{\pi})_n \subset [0,1] \times \Delta(2A)$. If for sufficiently large n the elements are contained in Γ_ρ we are done. So consider the case where they are not. Consider the family of projections $\Pi_\rho : [0,1] \times \Delta(2A) \longrightarrow [0,1] \times \Delta(2A)$ that map, for fixed $\bar{p} \in [0,1]$, any element $(\bar{p}, \bar{\pi}) \in [0,1] \times \Delta(2A)$ to the point in the set $\{\bar{p}\} \times \rho(\bar{p})$ that is closest to $(\bar{p}, \bar{\pi})$. Since the sets $\rho(\cdot)$ are always nonempty, convex, compact polyhedra, we have that $\Pi_\rho(p^n, \hat{\pi})$ is uniquely defined and moreover, $\Pi_\rho(p^n, \hat{\pi}) \in \Gamma_\rho$ for any point in the sequence $(p^n, \hat{\pi})_n$. It remains to be shown that the sequence $(\Pi_\rho(p^n, \hat{\pi}))_n$ converges to the point $(p, \hat{\pi})$. Apart from the p-belief constraints all other constraints defining $\rho(p)$ are independent of p. Hence, if $(p, \hat{\pi})$ satisfies those constraints, then so must any other point in the sequence $(p^n, \hat{\pi})_n$. Therefore the only constraints that can be violated by elements of the sequence $(p^n, \hat{\pi})_n$ are the p-belief constraints. Consequently, any point in the sequence $(\Pi_\rho(p^n, \hat{\pi}))_n$ lies on the boundary of the polyhedra defined by the p-belief constraints. As mentioned, these constraints are linear in p, and since they also define nonempty, convex, compact polyhedra, the sequence $(\Pi_\rho(p^n, \hat{\pi}))_n$ indeed converges to $(p, \hat{\pi})$. This shows the continuity of ρ and hence also of p-$RO(G)$ in p.

Finally, the claims that, for $p = 0$, we have 0-$RO(G) = \Delta(A)$, and for $p = 1$, we have 1-$RO(G) = CE(G)$, are immediate. To see that for any $p \in [0,1)$, we have $\dim[p\text{-}RO(G)] = \dim[\Delta(A)]$, notice that the (X,p)-correlated equilibria with $X = A^1$ and $p < 1$ entail distributions that put strictly positive weight on all strategies in A^2 as well as all convex combinations of such distributions. Projecting onto the original space $\Delta(A)$ implies distributions with strictly positive weights on all strategies in A as well as all possible convex combinations. This concludes the proof.

D Proof of Proposition 1

Fix G and let $A^n = \Pi_{i \in I} A_i^n$ denote the space of all pure strategy profiles that survive n rounds of iterated elimination of strictly dominated strategies in G, and similarly for the individual sets A_i^n. Let G^n denote the subgame of G with strategies restricted to A^n. Because G is finite, the limit sets A_i^∞, A^∞, and G^∞ are well defined (and are obtained after finitely many iterations). Also, for any subset $Y \subset A$, let $Y^c = A \setminus Y$ denote the complement of Y in A. For any given $p \in [0,1]$, take $p' \geq p$. We show that for p sufficiently close to 1, behavior is supported with high probability in A^∞. Specifically, we construct a $\bar{p} < 1$ such that for any $p \in [\bar{p}, 1]$, if $\pi \in p$-$RO(G)$, then $\pi[(A^\infty)^c] \leq 1 - p$. Consider the game $G^0 = G$ and pick some $p^1 < 1$. It immediately follows from p-rationality that for $p \in [p^1, 1]$, if $\pi \in p$-$RO(G)$, then we have $\pi\left[\left(A^1\right)^c\right] \leq 1 - p$. Suppose now that the above statement is true for $n-1$, namely there exists $p^{n-1} < 1$ such that for $p \in [p^{n-1}, 1]$, if $\pi \in p$-$RO(G)$, then we have $\pi\left[\left(A^{n-1}\right)^c\right] \leq 1 - p$. We show that the statement also holds for n. Fix game G^{n-1}. It follows from finiteness of G and continuity of the payoffs that there exists $p^n \in [p^{n-1}, 1)$ such a strategy in $A^{n-1} \setminus A^n$ that is strictly dominated in G^{n-1} (by some strategy in G^{n-1} and hence in G) is also strictly dominated in G (by the same strategy) given a p-rational belief system with $p \geq p^n$.[26] This implies that for any $p \in [p^n, 1]$ and any $\pi \in p$-$RO(G)$, we also have $\pi[(A^n)^c] \leq 1 - p$. Finiteness of the game implies that the process ends after finitely many steps implying that indeed there exists $p^\infty < 1$ such that for $p \in [p^\infty, 1]$ and any $\pi \in p$-$RO(G)$, we have $\pi[(A^\infty)^c] \leq 1 - p$. Taking $\bar{p} = p^\infty$ proves the claim.

[26]This follows from $p^n \geq p^{n-1}$, and because $\pi \in p$-$RO(G)$ with $p \geq p^{n-1}$ implies $\pi\left[\left(A^{n-1}\right)^c\right] \leq 1 - p$

E Proof of Proposition 2

To see the if part, take $\pi \in \Delta(A)$ a p-rational outcome of G. By Theorem 1, there exists an $(A_{(1)}, p)$-correlated equilibrium $\hat{\pi} \in \Delta(A')$ of the doubled game $2G$. Define $P \equiv \hat{\pi}$ and set

$$\alpha_i(t_i)[a_i] = \begin{cases} 1 & \text{if } t_i = a_i \\ 0 & \text{else ,} \end{cases}$$

for all $t_i \in T_i$. Then, by definition of the type spaces $T_i = T_i^s \cup T_i^c$, we have:

- $P[T_{-i}^s \,|\, t_i] \geq p$, for any $t_i \in T_i$;
- $\alpha_i(t_i)[a_i] = 1$ if $t_i \in T_i^c$ and $t_i = a_i$
- $a_i \in \text{argmax}_{a_i' \in A_i} \sum_{t_{-i} \in T_{-i}} P[t_{-i} \,|\, t_i] \cdot \sum_{a_{-i} \in A_{-i}} \alpha_{-i}(t_{-i})[a_{-i}] \cdot u_i(a_{-i}; a_i')$, where $\alpha_{-i}(t_{-i})[a_{-i}] = \prod_{j \neq i} \alpha_j(t_j)[a_j]$.

Let $\mu[a] = \sum_{t \in T} P(t) \cdot \alpha(t)[a]$, where $\alpha(t)[a] = \prod_{i \in I} \alpha_i(t_i)[a_i]$; then, again by construction, and using Theorem 1, we have $\mu = \pi$.

To see the only if part, take $\mu \in \Delta(A)$ an *EAD* of (G, P) where P satisfies $P[T_{-i}^s \,|\, t_i] \geq p$, for any $t_i \in T_i$, and for some equilibrium profile α satisfying conditions (i) and (ii). By definition, together with P, α can be seen as inducing a probability distribution on the outcomes of the doubled game $\hat{\mu} \in \Delta(2A)$, where types in T_i^s and corresponding strategies chosen are associated to $A_{(2)}$ and types in T_i^c are associated to strategies in $A_{(2)}$. As a consequence of (i) and (ii) and the definition of P, $\hat{\mu}$ will also satisfy the incentive constraints and the p-belief constraints for $2G$. By Theorem 1, computing the *EAD* μ of $\hat{\mu}$ gives us, a p-rational outcome of G.

F Proof of result in Remark 2

Adapting the original definition by Aumann (1974, 1987), for any $X = \prod_{i \in I} X_i \subseteq A$, we say that the family $(\pi_i)_{i \in I} \subseteq (\Delta(A))^I$ is a (X, p)-*subjective correlated equilibrium* of G, if, for any $i \in I$ the following are satisfied:

- **Incentive constraints.** For any $a_i \in X_i$, $a_i \in \text{argmax}_{a_i' \in A_i} \sum_{a_{-i} \in A_{-i}} \pi_i[(a_{-i}; a_i)] \cdot u_i(a_{-i}; a_i')$.
- ***p*-Belief constraints.** For any $a_i \in A_i$, $\pi_i[X_{-i} \times \{a_i\}] \geq p \cdot \pi_i[A_{-i} \times \{a_i\}]$.

We denote the set of (X, p)-subjective correlated equilibria of game G by (X, p)-$SCE(G)$. Given $n \in \{2, 3\}$ and nG, we have a map $\mathbf{marg}_{A^I} : (\Delta(A'))^I \to (\Delta(A))^I$, where for any $(\hat{\pi}_i)_{i \in I} \in (\Delta(A'))^I$, we have that $\mathbf{marg}_{A^I}\left((\hat{\pi}_i)_{i \in I}\right) = (\text{marg}_A(\hat{\pi}_i))_{i \in I}$. Then, the proof of the identity,

$$p\text{-}SRO(G) = \mathbf{marg}_{A^I}((X, p)\text{-}SCE(2G)),$$

where $X = \prod_{i \in I}(A_i \times \{1\})$, is the same as the one for Theorem 1 after slight modifications (just add sub-indices where needed). To see that the above marginals constitute the whole space, take $\left(a^i\right)_{i \in I} \subseteq A$, and for any $i \in I$, $\pi_i = 1_{\{a^i\}}$. Fix $k \in \{1, 2\}$, and define, for any $i \in I$, $\hat{\pi}_i = 1_{\{((a_j^i, k)_{j \neq i}; (a_i^i, 2-k))\}}$. It is immediate that $\mathbf{marg}_{A^I}\left((\hat{\pi}_i)_{i \in I}\right) = (\pi_i)_{i \in I}$. Now, take $i \in I$, then the incentive constraints are trivially satisfied, since $\hat{\pi}_i[A_i \times \{1\}] = 0$. Moreover, the p-belief constraint is also satisfied, because, regardless of i's action, the sums are on both sides 0 or 1. We conclude that, again for $X = \prod_{i \in I}(A_i \times \{1\})$, we have $\left(1_{\{a^i\}}\right)_{i \in I} \in \mathbf{marg}_{A^I}((X, p)\text{-}SCE(2G))$ for any $(a^i)_{i \in I} \subseteq A$; thus, by convexity, $\mathbf{marg}_{A^I}((X, p)\text{-}SCE(2G)) = ((\Delta(A))^I$.