Bounded Rationality and Correlated Equilibria

Fabrizio Germano
Peio Zuazo-Garin

February 2015

Barcelona GSE Working Paper Series

Working Paper nº 812

Bounded Rationality and Correlated Equilibria*

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

February 13, 2015

Abstract

We study an interactive framework that explicitly allows for non-rational behavior. We do not place any restrictions on how players can deviate from rational behavior. Instead we assume that there exists a lower bound $p \in [0, 1]$ such that all players play and are believed to play rationally with a probability p or more. 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 $p \in [0, 1]$. 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 compute a measure $\overline{p}$ that bounds from below the probability with which any given player is choosing 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 for further discussion. Germano acknowledges financial support from the Spanish Ministry of Science and Technology (Grants SEJ2007-64340 and ECO2011-28965), 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 also thank the Center for the Study of Rationality at the Hebrew University of Jerusalem for their 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`.

[‡]*BRiDGE*, Rovira i Virgili University, Department of Economics, 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. Instead we assume that, for any given $p \in [0, 1]$, a fraction at least p of any given player's chosen actions is consistent with rationality understood as payoff-maximizing behavior, while we put no constraint on what the remaining actions (occurring with frequency at most $1 - p$) may be. Among other things, we study the robustness of behavior with respect to p 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, 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 4, 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 is consistent with common knowledge of

[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 non-optimization-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.

payoff maximization. The second example, represented in Figure 2, is a game taken from Goeree and Holt (2001) due to David Kreps.

$G_{Kreps} \equiv$

	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 ($\overline{p} \approx 0.7$)

Here subjects playing the column player typically play strategy (NN) that is not in the support of the unique profile consistent with common knowledge of rationality, which is also the (unique) pure Nash equilibrium of the game ((T, L)), and which is played with a frequency of 0.178. Nonetheless, according to our measure, at least 70% of each player's (row player's and column player's) actions is consistent with common knowledge of payoff maximization.

In order to develop a theory that allows for boundedly rational behavior in normal form games, 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 at the interim level 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 *joint p-belief in 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. Characterizing the p-rational outcomes is tantamount to characterizing behavior that is robust to making mistakes, arbitrary mistakes, that for any given player can occur with frequency 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 an explicit 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 ($X \subset A$) of the overall action space, and where actions from this subset are believed to be played with certain minimum probability $p \in [0, 1]$. Our main characterization theorem uses the (X, p)-correlated equilibria to relate the p-rational outcomes to correlated equilibria and can thus be seen as giving a generalization of the main result of Aumann (1987). The p-rational outcomes are described by linear inequalities consisting of incentive and p-belief constraints that, for any p, always contain the correlated equilibria and such that, when $p = 1$, they coincide with the set of correlated equilibria, and when $p = 0$ they make up the whole space of distributions over action space A, $\Delta(A)$.

[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 non-rational.

In Theorem 2, we show that, besides being 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 strategy profiles involving strategies that do not survive the iterated elimination of strictly dominated strategies have probability at most $1-p$ under the common prior. These results confirm in a strong sense the robustness of the correlated equilibrium concept to bounded rationality. Theorem 4 shows that our main characterization result extends directly to the case of games of incomplete information using the notion Bayes correlated equilibrium of Bergemann and Morris (2013). 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 underlying bounded rationality of 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 obtain a unique number $\overline{p} \in [0,1]$ that in some sense quantifies proximity to common knowledge of rationality in a normal form strategic interaction. Indeed, 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 derive a unique number $\overline{p} \in [0,1]$, that gives the largest p such that each player plays actions that are consistent with payoff maximization and in fact 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. In those interactions, where the common prior assumption does not hold, then players may be acting rational, 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 easily 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. The value $\overline{p}$ can be useful as a measure of minimum amount of rationality or consistency of behavior in experimental data. We discuss this in more detail Section 4.

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

[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).

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 (2013) 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 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 defines the main epistemic concepts, including the notion of p-rational outcome. Section 3 contains all the main results as well as some simple examples. Section 4 shows how the theory implies a natural measure to quantify the degree of "rational" behavior in strategic interactions. Section 5 contains some extensions, including to games of incomplete information, and Section 6 provides some concluding remarks. All the proofs are in the Appendix.

2 p-Rationality in Games

In this section we define several basic concepts used throughout the paper, and we also introduce the notion of p-rational outcome, which is at the center of the paper. Properties and extensions of the p-rational outcomes are studied in subsequent sections.

A strategic-form *game* with complete information is defined as a tuple $G = \langle I, (A_i)_{i\in I}, (u_i)_{i\in I}\rangle$, 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*. We say that distribution $\pi \in \Delta(A)$ is a *correlated equilibrium* of G (Aumann, 1987) if for any player i and any action $a_i \in A_i$ the following incentive constraints are satisfied: $\sum_{a_{-i}\in A_{-i}} \pi(a_{-i}; a_i)[u_i(a_{-i}; a_i) - u_i(a_{-i}; a_i')] \geq 0$ for any action a_i'. We denote the set of correlated equilibria of G by $CE(G)$.

2.1 Knowledge, Rational Belief Systems and Correlated Equilibria

Following the epistemic framework by Aumann (1987), knowledge and beliefs are exogenously modeled by a *belief system*, which consists of a list $B = \langle \Omega, (\Pi_i)_{i \in I}, (\alpha_i)_{i \in I}, \mu \rangle$, where (i) Ω is a finite set of *states of the world*, and for any player i we have, (ii) Π_i, a *knowledge partition* of Ω where for any state ω we denote by $\Pi_i(\omega)$ the element of the partition containing ω, (iii) a *strategy map* $\alpha_i : \Omega \to A_i$ measurable w.r.t. Π_i, and (iv) a full support *common prior* $\mu \in \Delta(\Omega)$.[6] Note that belief system B induces a distribution $\pi_B \in \Delta(A)$ given by $\pi_B(a) = \mu(\alpha = a)$ for any action profile a.

An *event* is a collection of states $E \subseteq \Omega$. For any player i and any state ω, player i's differential information is then represented by $\Pi_i(\omega)$, and we say that player i *knows* event $E \subseteq \Omega$ 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 $E \subseteq \Omega$ 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 $CK(C)$. Additionally, the common prior and the knowledge partitions induce *interim beliefs*, conditional on differential information, that is, for any state ω,

$$\mu(E \,|\, \Pi_i(\omega)) = \frac{\mu(E \cap \Pi_i(\omega))}{\mu(\Pi_i(\omega))}, \text{ for any } E \subseteq \Omega.$$

For any state ω, we say that *player i is rational at ω* if her choice is optimal given her interim beliefs, *i.e.*, if

$$\alpha_i(\omega) \in \underset{a_i \in A_i}{\operatorname{argmax}} \, \mathbb{E}_B\left[u_i(\alpha_{-i}; a_i) \,|\, \Pi_i(\omega)\right],$$

where

$$\mathbb{E}_B\left[u_i(\alpha_{-i}; a_i) \,|\, \Pi_i(\omega)\right] = \sum_{a_{-i} \in A_{-i}} \mu(\alpha_{-i} = a_{-i} \,|\, \Pi_i(\omega))\, u_i(a_{-i}; a_i).$$

The set of states in which player i is rational is denoted by R_i and as usual, $R_{-i} = \bigcap_{j \neq i} R_j$ and $R = \bigcap_{i \in I} R_i$, respectively, denote the event that every player but i is rational, and the event that every player is rational. Note that in particular, a player knows at every state whether she is rational or not.

Following Aumann and Dreze (2008), we say that a belief system B is *rational* if players are rational at every state, *i.e.*, if $\Omega = R$. Note that this last assumption is just a technical abbreviation for having common knowledge of rationality. Making use of all these elements, the main theorem

[6]Formally, a belief system B as defined here is analogous to a *correlating device* as defined by Dekel and Siniscalchi (2015), who interpret B as a model of a signal structure. Since we are interested in epistemic aspects, at this respect this paper follows Aumann's (1987) interpretation of B.

in Aumann (1987) shows that the distribution on A induced by a rational belief system B is a correlated equilibrium of G.

2.2 p-Belief, p-Rational Belief Systems and p-Rational Outcomes

The notion of p-belief, introduced by Monderer and Samet (1989) is a standard tool for relaxing knowledge. Formally, for probability p, player i's *p-belief operator* is defined as,[7]

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

For probability p, event E and state ω we say that *player i p-believes E at ω*, if $\omega \in B_i^p(E)$, and that E is *p-evident* if $E \subseteq \bigcap_{i \in I} B_i^p(E)$. An event C is *common p-belief at ω* if there exists some p-evident event E such that $\omega \in E \subseteq \bigcap_{i \in I} B_i^p(C)$. We denote the event that C *is common p-belief* by $CB^p(C)$.

We now introduce our p-belief counterpart of rational belief systems mentioned above, *p-rational belief systems*, whose central concept is *joint p-belief in rationality.* For reasons explained below in Lemma 1, joint p-belief in rationality serves as the notion of approximate knowledge of rationality that replaces in the present paper the notion of common knowledge of rationality assumed in Aumann (1987) and elsewhere.

Definition 1 (Joint p-belief in rationality, p-rational belief system, p-rational outcome) *Let G be a game, p a probability, B a belief system, and $\pi \in \Delta(A)$ a distribution. Then:*

(*i*) Joint p-belief in rationality *is the event that there is mutual p-belief in opponents' rationality, i.e. $J_pBR = \bigcap_{i \in I} B_i^p(R_{-i})$.*

(*ii*) *We say that B is a* p-rational *belief system if there is joint p-belief in rationality at every state, i.e. if $J_pBR = \Omega$.*

(*iii*) *We say that π is a* p-rational outcome *of G if there exists some p-rational belief system B that induces π, i.e. such that $\pi = \pi_B$. We denote the set of p-rational outcomes of G by p-$RO(G)$.*

The following lemma establishes some elementary relations between notions regarding joint p-belief in rationality, common knowledge of rationality, and common p-belief in rationality, and hence, the relation between rational belief systems as the ones studied by Aumann (1987) and the p-rational belief systems introduced above.

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

(*i*) *If $p > 0$, there is rationality at every state if and only if there is common p-belief in rationality at every state.*

[7] Note that in this context of finiteness and full support priors, 1-belief equals knowledge.

(*ii*) *If there is a common prior and there is joint* p*-belief in rationality at every state, then,* $\mu(CB^p(R)) \geq \mu(R) \geq p^2$.

(*iii*) *If* $p = 1$*, then there is common* p*-belief in rationality at every state if and only if there is joint* p*-belief in rationality at every state.*

If our aim is to replace the assumption that there is common knowledge of rationality at every state by a less restrictive one involving some p-belief in rationality, the lemma above suggests that joint p-belief in rationality at every state could be a sensible choice, since (*i*) common p-belief in rationality at every state would lead to a model analogous to the one already considered before (or without) introducing p-beliefs,[8] (*ii*) although joint p-belief in rationality at every state is defined as a *joint* and not necessarily *common* belief, it nonetheless implies common p-belief in rationality with probability p^2, and, (*iii*) as p converges to 1, joint p-belief in rationality at every state converges to *common* 1-belief in rationality at every state (or *common certainty* of rationality at every state) and therefore, to common knowledge of rationality. Note that in particular, in a p-rational belief system, there is common knowledge of joint p-belief in rationality at every state.

The interest of p-rational outcomes lies in the fact that they capture in a precise sense behavior that is robust to bounded rationality. Our aim is to provide a strategic characterization of p-rational outcomes, that is, of the distributions over strategy profiles that arise when there is joint p-belief in rationality at every state, as is done in Aumann (1987) for the case in which rationality holds at every state.[9]

3 The Set of p-Rational Outcomes

In this section, we characterize the p-rational outcomes introduced in the previous section. To do this, we first introduce the notion of (X, p)-correlated equilibrium that generalizes the notion of correlated equilibrium, and that may be of interest in its own right. After providing our main characterization result, we prove additional properties of the set of p-rational outcomes, and also present some simple examples.

3.1 Strategic Characterization of p-Rational Outcomes

The main objective in this paper is to characterize the set $p\text{-}RO(G)$. Before doing so we need to first introduce the following notion of (X, p)-correlated equilibrium that generalizes the notion of correlated equilibrium and that plays a key role in all our characterizations.

[8] The first result of Lemma 1 implies the equivalence between rationality and common p-belief in rationality and may seem striking at first, however, it should be noted that it follows from the fact that certain assumptions are applied to every state. Since the set of all possible states is structurally commonly known, so is any event assumed to hold in every state. In particular, if rationality holds at every state, rationality is commonly known, and indeed commonly p-believed.

[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 our notion of J_pBR 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, let $X = \prod_{i \in I} X_i \subseteq A$, and let p be a probability. We say that distribution $\pi \in \Delta(A)$ is a (X, p)-*correlated equilibrium *of G if the following conditions are satisfied:*

(i) $\sum_{a_{-i} \in A_{-i}} \pi(a_{-i}; a_i)[u_i(a_{-i}; a_i) - u_i(a_{-i}; a_i')] \geq 0$ *for any player i, any $a_i \in X_i$ and any action $a_i' \in A_i$* **(incentive constraints)**.

(ii) $\pi(X_{-i} \times \{a_i\}) \geq p\pi(A_{-i} \times \{a_i\})$ *for any player i and any $a_i \in A_i$* **(p-belief constraints)**.

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 for each player i a subset $X_i \subseteq A_i$ and a probability $p \in [0, 1]$ such that the distribution on the overall set of action profiles A satisfies (i) standard *incentive constraints* for all actions in X_i, and, (ii) p-*belief constraints*, meaning each player assigns probability at least p to the *other* players all choosing action profiles from $X_{-i} = \prod_{j \neq i} X_i$. In words, the (X, p)-correlated equilibria allow to *relax* incentive constraints on actions *not* in the X_i's, while restricting the probability with which this occurs to $1 - p$.[10] Note that $CE(G) \subseteq (X, p)$-$CE(G)$ for any $X \subseteq A$ and $p \in [0, 1]$, and that, if $X = A$ and $p = 1$, or just $X = A$, then $CE(G) = (X, p)$-$CE(G)$.

The following definition generalizes the idea of *doubled game* in Aumann and Dreze (2008). For any $n \in \mathbb{N}$ we denote $N = \{0, 1 \ldots, n-1\}$; as we will see, it is the cases $n = 2$ and $n = 3$ that play a role in our epistemic characterization results.

Definition 3 (n-Game) *Let $G = \langle I, (A_i)_{i \in I}, (u_i)_{i \in I} \rangle$ be a game and $n \in \mathbb{N}$. The n-*game *is the tuple $nG = \langle I, (nA_i)_{i \in I}, (u_{n,i})_{i \in I} \rangle$, where for each player i,*

(i) $nA_i = \prod_{k \in N} A_{i,(k)}$, *where for each $k \in N$, $A_{i,(k)} = (A_i \times \{k\})$, is player i's set of pure actions. We denote a generic element of $nA = \prod_{i \in I} nA_i$ by (a, ν), where $\nu \in N^I$ specifies which copy of A_i in nA_i each player i's pure action belongs to.*

(ii) $u_{n,i}$ *is player i's payoff function, where for each $(a, \nu) \in nA$, $u_{n,i}(a, \nu) = u_i(a)$.*

In this context when writing the action spaces of the game nG as $nA_i = A_i \times N$ we mean that for each player there are n *copies* of the original action space A_i. Note that any distribution on the action profiles of nG, $\hat{\pi} \in \Delta(nA)$, induces a distribution on the action profiles of G in a natural way by taking the marginal on A, that is, $\pi = \text{marg}_A \hat{\pi}$. For any subset $Y \subseteq \Delta(nA)$ we denote by $\mathbf{marg}_A[Y] = \{\text{marg}_A \hat{\pi} \text{ where } \hat{\pi} \in Y\}$. These elements provide all the tools required to go on with the characterization of the set of p-rational outcomes of the game G. The following theorem shows that these can be expressed in terms of computationally simple (X, p)-correlated equilibria of the doubled game $2G$.

[10] 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). The belief constraints can reflect bounds or statistical regularities with which deviations from "rationality" are roughly known to occur.

Theorem 1 *Let G be a game and p a probability. Then, $\pi \in \Delta(A)$ is a p-rational outcome of G if and only if it is the distribution in $\Delta(A)$ induced by some $(A_{(0)}, p)$-correlated equilibrium of $2G$. Formally,*

$$p\text{-}RO\,(G) = \mathbf{marg}_A \left[\left(A_{(0)}, p \right) \text{-}CE\,(2G) \right],$$

This characterizes behavior in a game G under joint p-belief of rationality and a common prior, or, in other words, all behavior in G that is robust to bounded rationality as captured by the notion of p-rationality. 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 $A_{i,(0)}$) and by the irrational type (in $A_{i,(1)}$). The $(A_{(0)}, p)$-correlated equilibria with are distributions on $\Delta(2A)$ that by construction satisfy the incentive constraints just for the rational types, and where the p-belief constraints ensure that all players believe the others play rationally with probabilities p or more. Taking marginals, finally, ensures that the distributions are on $\Delta(A)$.[11]

3.2 Further Properties of p-Rational Outcomes

The next two results further characterize the structure and nature of the set of p-rational outcomes. The first result in this paragraph strengthens this by showing 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, 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 *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.*[12] *Moreover, for $p = 0$, we have $0\text{-}RO(G) = \Delta(A)$, for $p = 1$, we have $1\text{-}RO(G) = CE(G)$, 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 $\boldsymbol{p}$-rational outcomes whenever $p < 1$. The next result qualifies this by showing that if p is close enough 1, then strategy profiles involving strategies that do not survive the iterated elimination of strictly dominated strategies get a total weight of at most $1-p$. 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 and denote its complement in A by $(A^\infty)^c = A \setminus A^\infty$.

Proposition 1 *Let G be a game. Then there exists $\bar{p} \in (0, 1)$ such that, for any $p \in [\bar{p}, 1]$, we have that, if $\pi \in p\text{-}RO(G)$, then $\pi((A^\infty)^c) \leq 1 - p$.*

[11] Clearly, due to the symmetric role of the different copies of the action spaces of $2G$, the theorem would also hold for $X = A_{(1)}$, whereby only one of the two copies of players' actions satisfies the incentive constraints.

[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 above results confirm in a precise sense the robustness of the correlated equilibrium benchmark when weakening the underlying assumption of common knowledge of rationality to joint p-belief of rationality.

3.3 Examples

The following examples illustrate the p-rational outcomes for two simple 2×2 games.

3.3.1 Dominance Solvable Game

Consider 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

, $2G_D \equiv$

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

.

To compute the p-$RO(G_D)$ we compute the $(a_{(0)}, p)$-$CE(2G_D)$. For this notice that the strategies $(B,0)$ and $(T,1)$ of the row player and $(R,0)$ and $(L,1)$ of the column player are strictly dominated, so that the remaining constraints that need to be satisfied are the p-belief constraints, and one obtains,

$$p\text{-}RO(G_D) = \left\{ \pi \in \Delta(A) \;\middle|\; \begin{array}{l} \pi_{TL} \geq p_1(\pi_{TL} + \pi_{TR}), \pi_{BL} \geq p_1(\pi_{BL} + \pi_{BR}) \\ \pi_{TL} \geq p_2(\pi_{TL} + \pi_{BL}), \pi_{TR} \geq p_2(\pi_{TR} + \pi_{BR}) \end{array} \right\}.$$

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, $N_\epsilon(CE(G_D))$, and the set of ϵ-correlated equilibria, ϵ-$CE(G_D)$,[13] both with $\epsilon = 0.20$. Clearly the three sets are all distinct.

3.3.2 Matching Pennies Game

Consider the following version 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

, $2G_{MP} \equiv$

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

.

[13] 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, analagous to Radner's ϵ-Nash equilibria, formally, π is an *ϵ-correlated equilibrium* (ϵ-CE) if for any $i \in I$,

$$\sum_{a_i \in A_i} \max_{a'_i \in A_i} \sum_{a_{-i} \in A_{-i}} \pi(a_i, a_{-i}) \left(u_i(a'_i, a_{-i}) - u_i(a_i, a_{-i}) \right) \leq \epsilon.$$

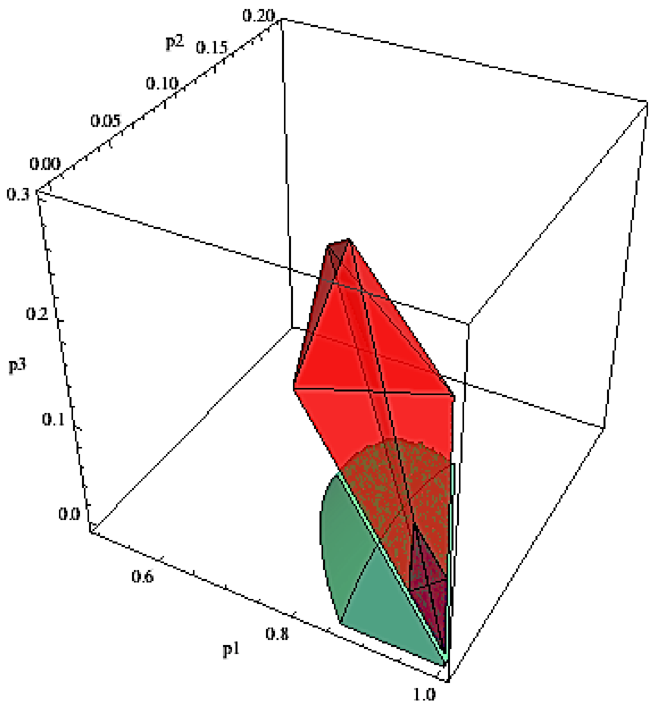

Figure 3: 0.95-$RO(\Gamma_D)$ (blue), 0.80-$RO(\Gamma_D)$ (red), $N_{0.10}CE(\Gamma_D)$ (green)

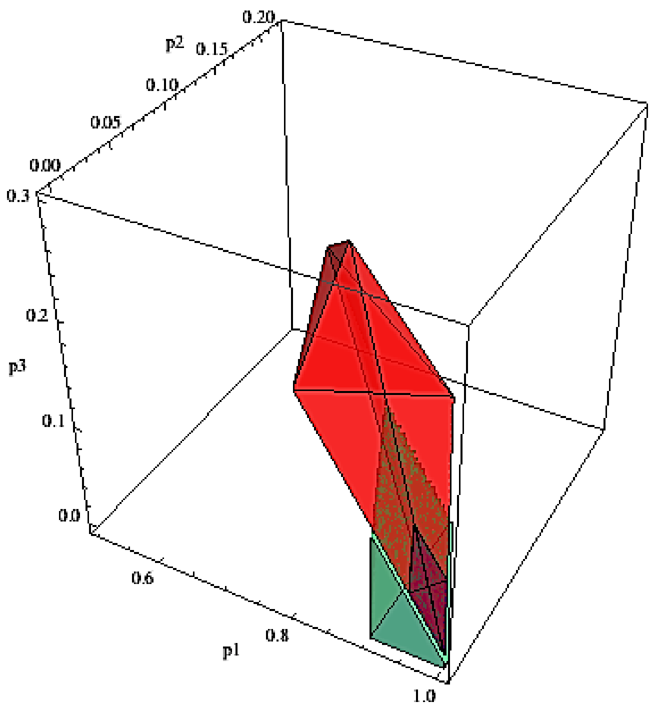

Figure 4: 0.95-$RO(G_D)$ (blue), 0.80-$RO(G_D)$ (red), 0.10-$CE(G_D)$ (green)

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-neighbourhood of the set of correlated equilibria, $N_\epsilon(CE(G_{MP}))$, and the set of ϵ-correlated equilibria, ϵ-$CE(G_{MP})$, both with $\epsilon = 0.10$, respectively. Again, the sets p-$RO(G_{MP})$ and $N_\epsilon(CE(G_{MP}))$ and ϵ-$CE(\Gamma_{MP})$ are visibly distinct.

4 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 ourcome. 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:

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

By definition of the p-rational outcomes, $\overline{p}$ is also the largest value of p consistent with common knowledge of joint p-belief of rationality for the distribution π. In particular, this means that at the distribution π, every player chooses actions that are consistent with 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 2 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.

Threrefore the unique value $\overline{p} \in [0, 1]$ defined above can be interpreted loosely, and in a sense to be further qualified below, as the largest level of rationality in the sense of behavior compatible with payoff-maximization given the distribution of play π. This can be used to derive a unique and easily comparable number for any observed finite strategic interaction. In particular, it can be applied to various games played in experimental settings, also incomplete information games.

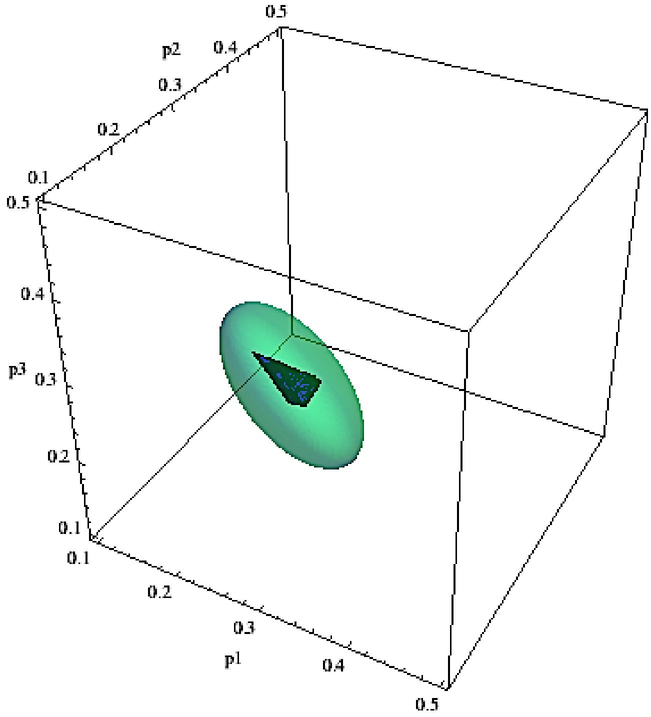

Figure 5: 0.95-$RO(G_{MP})$ (blue), $N_{0.10}CE(G_{MP})$ (green)

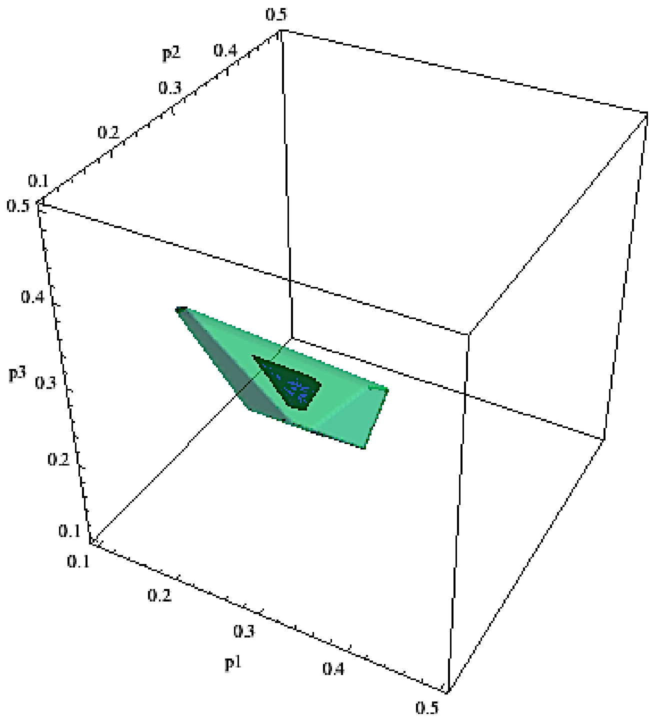

Figure 6: 0.95-$RO(G_{MP})$ (blue), 0.10-$CE(G_{MP})$ (green)

This differs from other notions in the literature such as the level λ of fitted λ-quantal response models, which is *not* comparable generally speaking across games, without normalizing the payoffs; similarly, the mode level k of a fitted k-level model also does not always give a clear-cut value, as it generally depends on the specified level 0; the approach also typically consists of a distribution of levels k within the population.[14] We discuss some experimental data and make an important caveat to our measure $\overline{p}$ that stems from the assumption of a common prior that is implicit in our theory.

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 Introcuction. 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.[15] 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.[16] 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 ($\overline{p} \approx 0.8$)

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

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

[15] 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 $\overline{p} \approx 0.96$:

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

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

.

Essentially, since the entry at $((KL,0),(JR,1))$ is positive, this indicates that the goalkeepers jumped right a little more than optimal, and since the entry $((KL,1),(JR,0))$ is (barely) positive, this indicates that the kickers kicked left very slightly more than optimal. The remaining entries with $(KL,0)$, $(KR,0)$, $(JL,0)$, and $(JR,0)$ are all consistent with rationality.

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

(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 ($\overline{p} \approx 0.25$).[17]

As a third example, consider the games depicted in Figure 8.[18] 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.

Finally, consider again the Kreps game G_{Kreps} represented in Figure 2 in the Introduction. Here players typically play a strategy (NN) that is not in the support of the (unique) Nash equilibrium of the game. Although the strategy is not part of any Nash equilibrium, it nonetheless can be "rationalized" to some extent, such that by our measure, although the Nash equilibrium profile is played with a frequency of 0.178, the interaction has a level of "rationality" of $\overline{p} \approx 0.7$.

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 the maximum level p such that every player plays actions consistent

[17] 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 equilibrim 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.

[18] 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.

with payoff maximization with probability at least p, at the empirical distribution of play π^{emp}.[19]

On the other hand, in games with multiple equilibria, such as the coordination game 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:

$$G_{BS} \equiv \begin{array}{c|c|c|} & L & R \\ \hline T & 2,1 & 0,0 \\ \hline B & 0,0 & 1,2 \\ \hline \end{array}\ , \quad \pi_{BS} = \begin{array}{c|c|c|} & L & R \\ \hline T & 0 & 0 \\ \hline B & 1 & 0 \\ \hline \end{array}\ .$$

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 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 "non-rationality," namely, non-payoff-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 probability of payoff-maximizing behavior 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.

5 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 relate them to the notion of rational expectations of Aumann and Dreze (2008) and compute what payoffs to expect at p-rational outcomes. Then, 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 (2013).

[19]The degree of approximation improves the closer the empirical distribution is to the unique equilibrium and the higher the value $\overline{p}$ is. The Kreps game is a case, where the unique equilibrium is played with low probability and some players may anticipate opponents playing the NN action. Here the prior may be further from the one implied by common belief in rationality.

[20]Kneeland (2013) 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.

5.1 p-Rational Expectations

Following Aumann and Dreze (2008), we can analyze expected payoffs or *expectations in a game* from the point of view of a fixed player who already gained differential information and has therefore, *interim* beliefs.

Definition 4 (p-Rational Expectation) *Let G be a game, let p be a probability, and let B be a p-rational belief system for G. Then a* p-rational expectation *in G is an interim expected payoff of some player. We denote the set of all such p-rational expectations of G by p-RE (G).*

The following result provides the joint p-belief of rationality counterpart of Aumann and Dreze's characterization:

Theorem 3 *Let G be a game and p a probability. Then the p-rational expectations in G are the expected payoffs of the $\left(A_{(0)}, p\right)$-correlated equilibria of the tripled game $3G$, conditional on playing an action in $3A_i$. Moreover, the p-rational expectations of players acting rationally at the interim stage are the expected payoffs of the $\left(A_{(0)}, p\right)$-correlated equilibria of the tripled game $3G$, conditional on playing an action in $A_{i,(0)}$.*

5.2 Games with Incomplete Information

We now extend the characterization of p-rational behavior to games of incomplete information.

5.2.1 Preliminaries

We follow the formalization in Lehrer *et al.* (2010, 2013) and Bergemann and Morris (2013) that splits the game with incomplete information in two components so that strategic and informational aspects can be studied separately. The first one is a *basic game* $G = \left\langle I, \Theta, \psi, (A_i)_{i\in I}, (u_i)_{i\in I}\right\rangle$, 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*. The game proceeds then as follows:

(i) A states of nature θ is randomly drawn with probability $\psi(\theta)$.

(ii) A profile of types t is randomly drawn with conditional probability $\sigma(t|\theta) = \sigma(\theta)(t)$.

(iii) Each player i, who privately receives signal t_i, chooses an action a_i and gets payoff $u_i((a_{-i}; a_i), \theta)$.

In this context, the concept of *belief system* is extended in order to be able to include payoff-uncertainty and signal structures, so that we have a list $B = \left\langle \Omega, (\Pi_i)_{i\in I}, \mu, \kappa, (\alpha_i)_{i\in I}, (\tau_i)_{i\in I}\right\rangle$, 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. Following Bergemann and Morris (2013), we assume that a belief model always satisfies the following *consistency* condition:

$$\mu(\tau = t, \kappa = \theta) = \psi(\theta)\sigma(t|\theta) \text{ for any } (t, \theta) \in T \times \Theta.$$

Rationality is defined as the event that players choose optimally given their private information. Thus,

$$R_i = \{\omega \in \Omega \,|\vee_{a_i \in A_i} \mathbb{E}\left[u_i\left((\alpha_{-i}, a_i), \kappa\right)|\Pi_i(\omega)\right] \leq \mathbb{E}\left[u_i\left((\alpha_{-i}, \alpha_i(\omega)), \kappa\right)|\Pi_i(\omega)\right]\}$$

represents the event that player i is rational. Again, we denote $R_{-i} = \bigcap_{j \neq i} R_j$ and $R = \bigcap_{i \in I} R_i$. For each player i, both p-belief operators and joint p-belief in rationality are defined exactly as done in Subsection 2.2, and again, for any $p \in [0, 1]$ we say that a belief system is p-rational, if joint p-belief in rationality holds at every state.

Then, for game with incomplete information (G, S) and probability p, distribution $\pi \in \Delta(T \times A \times \Theta)$ is a *p-rational Bayes outcome* of G if there is some belief system B satisfying consistency and joint p-belief in rationality such that $\pi(t, a, \theta) = \mu(\tau = t, \alpha = a, \kappa = \theta)$ for any $(t, a, \theta) \in T \times A \times \Theta$. We denote the set of p-rational Bayes outcomes of (G, S) by $\pi \in p\text{-}RBO(G, S)$.

5.2.2 Strategic Characterization of p-Rational Bayes Outcomes

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: we first define a notion analogous to that of a correlated equilibrium for games with incomplete information, and second, we define the n-basic game.

Definition 5 ((X, p)-Bayes Correlated Equilibrium) *Let (G, S) be a game with incomplete information, let $X = \prod_{i \in I} X_i \subseteq A$, and let p be a probability. We say that distribution $\pi \in \Delta(T \times A \times \Theta)$ is a* (X, p)-Bayes correlated equilibrium *of (G, S) if the following conditions are satisfied:*

(i) $\pi(A \times \{(t, \theta)\}) = \psi(\theta)\sigma(t|\theta)$ *for any* $(t, \theta) \in T \times \Theta$ **(consistency constraints)**.

(ii) $\sum_{(a_{-i}, \theta) \in A_{-i} \times \Theta} \pi(T_{-i} \times \{(t_i, a_{-i}, a_i, \theta)\})\left[u_i(a_{-i}, a_i, \theta) - u_i(a_{-i}, a_i', \theta)\right] \geq 0$ *for any player i, any $a_i \in X_i$, any $a_i' \in A_i$ and any $t_i \in T_i$* **(incentive constraints)**.

(iii) $\pi(T_{-i} \times X_{-i} \times \Theta \times \{(t_i, a_i)\}) \geq p\pi_i(T_{-i} \times A_{-i} \times \Theta \times \{(t_i, a_i)\})$ *for any player i, any $a_i \in A_i$ and any $t_i \in T_i$* **(p-belief constraints)**.

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 for any game with incomplete information (G,S), (X,p)-$BCE\,(G,S)$ where $p=1$ is precisely the set of Bayes correlated equilibria as defined by Bergemann and Morris (2013).[22] The only remaining element to proceed with the characterization result is a generalization of the n-game:

Definition 6 (n-Basic Game) *Let $G=\langle I,\Theta,\psi,(A_i)_{i\in I},(u_i)_{i\in I}\rangle$ be a basic game and let $n\in\mathbb{N}$, then the n-*basic game* is the tuple $nG=\langle I,\Theta,\psi,(nA_i)_{i\in I},(u_{n,i})_{i\in I}\rangle$, where for each player i,*

(i) *$nA_i=\prod_{k\in N}A_{i,(k)}$, where for each $k\in N$, $A_{i,(k)}=(A_i\times\{k\})$, is player i's set of pure actions. We denote a generic element of $nA=\prod_{i\in I}nA_i$ by (a,ν), where $\nu\in N^I$ specifies which copy of A_i in nA_i each player i's pure action belongs to.*

(i) *$nA_i=\prod_{k\in N}(A_i\times\{k\})$, is player i's set of pure actions. We denote a generic element of $nA=\prod_{i\in I}nA_i$ by (a,ν), where $\nu\in N^I$ specifies which copy of A_i in nA_i each player i's pure action belongs to.*

(ii) *$u_{n,i}$ is player i's payoff function, where for each $(a,\nu,\theta)\in nA\times\Theta$, $u_{n,i}(a,\nu,\theta)=u_i(a,\theta)$.*

Note that a distribution $\pi^*\in\Delta(T\times nA\times\Theta)$ naturally induces a distribution $\pi\in\Delta(T\times A\times\Theta)$ by setting $\pi(t,a,\theta)=\pi^*\left(\prod_{i\in I}(\{a_i\}\times N)\times\{(t,\theta)\}\right)$ for any $(t,a,\theta)\in T\times A\times\Theta$. For any subset $Y\subseteq\Delta(T\times nA\times\Theta)$, define $\mathbf{marg}_{T\times A\times\Theta}[Y]=\left\{\mathrm{marg}_{T\times A\times\Theta}\hat{\pi}\text{ where }\hat{\pi}\in Y\right\}$. The characterization result in this case becomes:

Theorem 4 *Let (G,S) be a game with incomplete information and let $p\in[0,1]^I$. Then, $\pi\in\Delta(T\times A\times\Theta)$ is a p-rational Bayesian outcome of (G,S) if and only if it is the distribution in $\Delta(T\times A\times\Theta)$ induced by some $(A_{(0)},p)$-Bayes correlated equilibrium of $(2G,S)$. Formally,*

$$p\text{-}RBO\,(G,S)=\mathbf{marg}_{T\times A\times\Theta}\left[\left(A_{(0)},p\right)\text{-}BCE\,(2G,S)\right],$$

This is parallel to our characterization result for the complete information case. Finally, consider a game with incomplete information (G,S) and epistemic model E. Besides consistency and joint p-incentive compatibility, and following Forges (1993, 2006), we can additionally impose the following conditions on E: (1) $\mu(\kappa=\theta\,|\vee_{i\in I}\Pi_i)=\mu(\kappa=\theta\,|\tau)$ for any $\theta\in\Theta$, and (2) $\mu(\tau_{-i}=t_{-i},\kappa=\theta\,|\Pi_i)=\mu(\tau_{-i}=t_{-i},\kappa=\theta\,|\tau_i)$ for any $t_{-i}\in T_{-i}$, and $\theta\in\Theta$ and any $i\in I$. It is then easy to see that if we impose only (1), or both (1) and (2), the respective distributions induced in $A\times\Theta$ are p-belief versions of Forges' *Bayesian solution* and *belief invariant Bayesian solution*.

6 Concluding Remarks

We conclude with a few remarks.

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

Remark 1 ($(\boldsymbol{p},\boldsymbol{q})$-Rational Outcomes) An important objective of the paper was to put as few restrictions on non-rational 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 J_pBR – that players always believe the other players are rational with probability p or more; thus we indirectly assumed that all players and all player types 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 non-rational 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 non-rational at a given state). In particular, we can assume that each player i believes the other players are rational with probability p_i or more when rational and believes others are rational with probability q_i or more when not rational; importantly, one can drop any restriction on the non-rational types and directly set $q_i = 0$ for all $i \in I$, which would allow to not impose *any* belief constraints on non-rational types.

This can be formalized by assuming a vector $(\boldsymbol{p},\boldsymbol{q}) \in [0,1]^{2I}$, where the components associated to states in which the agents are rational are represented by the vector of probabilities $\boldsymbol{p} \in [0,1]^I$, while the components associated to the non-rational states are represented by the vector of probabilities $\boldsymbol{q} \in [0,1]^I$. This leads to the more general $(\boldsymbol{p},\boldsymbol{q})$-*rational outcomes* of G ($(\boldsymbol{p},\boldsymbol{q})$-$RO(G)$). These are again marginals of $\left(A_{(0)},(\boldsymbol{p},\boldsymbol{q})\right)$-correlated equilibria of $2G$, in that they are distributions satisfying the same conditions as the $\left(A_{(0)},p\right)$-$CE(2G)$ except that the p-belief constraints now hold with probabilities p_i (instead of the fixed p) for all rational types of player i, and hold with probability q_i (instead of the same p) for all non-rational types of player i. That is we replace the original p-belief constraints ((ii)) with the more general $(\boldsymbol{p},\boldsymbol{q})$-beliefs constraints of the form,

(ii') $\pi\left(X_{-i}\times\{a_i\}\right) \geq p_i\pi\left(A_{-i}\times\{a_i\}\right)$ for any player i and any $a_i \in A_{i,(0)}$.

(ii'') $\pi\left(X_{-i}\times\{a_i\}\right) \geq q_i\pi\left(A_{-i}\times\{a_i\}\right)$ for any player i and any $a_i \in A_{i,(1)}$.

Thus, whenever $p_i \leq p$ and $q_i \leq p$ for all i, we have that $\left(A_{(0)},p\right)$-$CE(2G) \subset \left(A_{(0)},(\boldsymbol{p},\boldsymbol{q})\right)$-$CE(2G) \subset \left(A_{(0)},(\boldsymbol{p},\boldsymbol{0})\right)$-$CE(2G)$ and therefore,

$$p\text{-}RO(G) \subset (\boldsymbol{p},\boldsymbol{q})\text{-}RO(G) \subset (\boldsymbol{p},\boldsymbol{0})\text{-}RO(G) \text{ for } \max\{p_i,q_i : i \in I\} \leq p.$$

While the correspondence $(\boldsymbol{p},\boldsymbol{q})$-$RO(G)$ maintains the basic topological properties of the correspondence p-$RO(G)$, it need not converge to the set of correlated equilibria of G as $(\boldsymbol{p},\boldsymbol{q}) \to (\mathbf{1},\mathbf{0})$, i.e., if only the rational types believe opponents are rational with probability 1, but does so if one also requires $(\boldsymbol{p},\boldsymbol{q}) \to (\mathbf{1},\mathbf{1})$. This can be seen already in Example 1. A $(\mathbf{1},\mathbf{0})$-rational belief system can be very far from a $(\mathbf{1},\mathbf{1})$-rational belief system in that the former need not put any restriction on the total mass of states where all players are rational $\mu(R)$.[23]

[23]To see that in a $(\mathbf{1},\mathbf{0})$-rational belief system the total mass of states where agents are non-rational is unrestricted,

The alternative notion of approximate knowledge of rationality requiring $\mu(CB^p(R)) > 1-\epsilon$, for $\epsilon > 0$, (instead of J_pBR), 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 beliefs individual players and types 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 (Non-Common Priors) Throughout the paper we assumed the existence of a common prior (CP). This together with the notion of joint p-belief of rationality 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 *non-common* priors, joint p-belief of rationality puts *no* restrictions on possible behavior – even when $p = 1$.[24] This provides a stark contrast with the behavior under common certainty of rationality and also common p-belief of 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.[25]

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 (e.g, Camerer (2003)) that put specific

take the game in Example 1 and consider the belief system $B = \langle \Omega, (\Pi_i)_{i\in I}, (\alpha_i)_{i\in I}, \mu \rangle$, where $\Omega = A$, $s_i(a_{-i}; a_i) = a_i$, for all $(a_{-i}; a_i) \in A_i$, $i \in I$, 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 $(\mathbf{1}, \mathbf{0})$-rational and clearly $\mu(R) = 0$. At the same time, in a $(\boldsymbol{p}, \boldsymbol{q})$-rational belief system it is always the case that, for any $i \in I, \omega \in \Omega$, $\mu(R_{-i} \mid \Pi_i(\omega)) \geq q_i$, hence

$$\mu(R_{-i} \cap \Pi_i(\omega)) \geq q_i \mu(\Pi_i(\omega)) \Longrightarrow \sum_{\Pi_i(\omega)\in\Pi_i} \mu(R_{-i} \cap \Pi_i(\omega)) = q_i \sum_{\Pi_i(\omega)\in\Pi_i} \mu(\Pi_i(\omega)) \Longrightarrow \mu(R_{-i}) \geq q_i,$$

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

[24] To see the non-common prior case, define belief system $B = \langle \Omega, (\Pi_i)_{i\in I}, (\alpha_i)_{i\in I}, (\mu_i)_{i\in I} \rangle$ to be *subjectively p-rational* if $\Omega = \bigcap_{i\in I} B_i^p(R_{-i})$. Given a finite game in strategic form G with set of players I and set of action profiles A, and given $p \in [0,1]$, 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 subjectively p-rational belief system $B = B = \langle \Omega, (\Pi_i)_{i\in I}, (\alpha_i)_{i\in I}, (\mu_i)_{i\in I} \rangle$ for G such that for any $i \in I$, we have $\pi_i = \mu_i \circ \alpha^{-1}$. 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, an agent that is certain that all the other agents are rational given his priors may still select a non-rational action given, and so *any* pure strategy profile in A is consistent with subjective p-rationality, even when $p = 1$.

[25] Recall that the result of part (a) of Lemma 1 also holds with non-common priors.

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.

Bergemann, Drik and Stephen Morris (2013). "The Comparison of Information Structures in Games: Bayes Correlated Equilibrium and Individual Sufficiency". Mimeo.

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 and Marciano Siniscalchi (2015). "Epistemic Game Theory". In: *Handbook of Game Theory, Volume 4* (H.P. Young and S. Zamir, Eds.). North-Holland, Elsevier, Amsterdam.

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 Decission* **35**, 277–310.

Forges, Françoise (2006). "Correlated Equilibrium in Games with Incomplete Information Revisited". CORE Discussion Paper 2006/41.

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". *Theory and Decission* **91**, 1402–1422.

Harstad, Ronald M. and Reinhard Selten (2013). "Boundedly Rational 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.

Kneeland, Terri (2013). "Testing Behavioral Game Theory: Higher-Order Rationality and Consistent Beliefs". Mimeo.

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 Behavior and the Problem of Perfection". *Econometrica* **52**, 1029–1050.

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

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*) 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*) 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.

(*iii*) 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 $R = \Omega$.

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 game with incomplete information and let $p \in [0, 1]$, $n \in \mathbb{N}$, and $\pi^* \in (A_{(k)}, p)$-BCE(nG, S), for some $k \in N$. Let belief system $B = \langle \Omega, (\Pi_i)_{i\in I}, \mu, \kappa, (\alpha_i)_{i\in I}, (\tau_i)_{i\in I} \rangle$ defined as follows: (i) $\Omega = \{(t, a, \nu, \theta) \in T \times nA \times \Theta \,|\pi^*[(a, \nu, \theta)) > 0\}$, 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 $\Pi_i(t, a, \nu, \theta) = T_{-i} \times nA_{-i} \times \Theta \times \{(t_i, a_i, \nu_i)\}$, (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 and joint p-belief in rationality that 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 $\mu(\tau = t, \kappa = \theta) = \pi^*(nA \times \{(t, \theta)\}) = \psi(\theta)\sigma(t|\theta)$, and therefore, B satisfies consistency. Now, note first the fact that for any $\omega = (t, a, \nu, \theta) \in \Omega$ and any $a_i' \in A_i$,

$$\mathbb{E}\left[u_i\left(\alpha_{-i}, a_i', \kappa\right)|\Pi_i(\omega)\right] = \sum_{(t_{-i}', (a_{-i}', \nu_{-i}'), \theta')} \pi^*\left(\left(t_{-i}', t_i, (a_{-i}', \nu_{-i}'), (a_i, \nu_i), \theta'\right)\right) u_i\left((a_{-i}', \nu_{-i}'); (a_i, \nu_i), \theta'\right),$$

together with the incentive constraints, implies that for any $i \in I$ we have that $T \times nA_{-i} \times A_{i,(k)} \times \Theta \subseteq R_i$, and therefore,[26] that $\mu\left(A_{-i,(k)} \cap \Pi_i(\omega)\right) \leq \mu\left(R_{-i} \cap \Pi_i(\omega)\right)$ for any $i \in I$ and any $\omega \in \Omega$. Then, take $i \in I$ and $\omega = (t, a, \nu, \theta) \in \Omega$; note that:

$$\mu\left(A_{-i,(k)} \cap \Pi_i(\omega)\right) = \sum_{t'_{-i} \in T_{-i}} \pi^*\left(A_{-i,(k)} \times \Theta \times \{(t_i, a_i, \nu_i)\}\right),$$
$$\mu\left(\Pi_i(\omega)\right) = \sum_{t'_{-i} \in T_{-i}} \pi^*\left(nA_{-i} \times \Theta \times \{(t_i, a_i, \nu_i)\}\right).$$

In consequence, due to the p-belief constraints, we have $\mu\left(R_{-i} \cap \Pi_i(\omega)\right) \geq p\mu\left(\Pi_i(\omega)\right)$, and therefore, B satisfies joint p-incentive compatibility. ■

B.1 Proof of Theorem 4

For the right inclusion, take distribution $\pi \in p\text{-}RBO\,(G, S)$ and B, a belief model that induces π. Take distribution $\pi^* \in \Delta\left(T \times 2A \times \Theta\right)$ defined as follows:

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

for any $(t, a, \nu, \theta) \in T \times nA \times \Theta$. Then, $\pi^* \in \Delta\left(T \times nA \times \Theta\right)$. The consistency constraint is satisfied, since we have that,

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

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 $\nu_i \in N$,

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

[26] We abbreviate, $\mu\left(A_{-i,(k)} \cap \Pi_i(\omega)\right) = \mu\left(\left(T \times A_{-i,(k)} \times nA_i \times \Theta\right) \cap \Pi_i(\omega)\right)$, with some abuse of notation.

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

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

Thus, both the incentive constraints and the p-belief constraints are satisfied. For the left inclusion, just apply Lemma 2 to $n = 2$ and $k = 0$.

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 a game with incomplete information (G', S), where we have (i) $G' = \left\langle I, \Theta, \psi, (A_i)_{i \in I}, (u'_i)_{i \in I} \right\rangle$, 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) $S = \left\langle (T_i)_{i \in I}, \sigma \right\rangle$ 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 $B' = \left\langle \Omega, (\Pi_i)_{i \in I}, (\alpha_i)_{i \in I}, (\tau_i)_{i \in I}, \mu \right\rangle$, 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 $B = \left\langle \Omega, (\Pi_i)_{i \in I}, (\alpha_i)_{i \in I}, \mu \right\rangle$ 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 $(X, p)\text{-}CE(2G) = \left(A_{(0)}, p\right)\text{-}BCE(2G', S)$, hold, so from Theorem 4 we obtain $p\text{-}RO(G) = \left(A_{(0)}, p\right)\text{-}CE(2G)$.

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:

$$\pi^*_{i,\omega}(a, \nu) = \mu\left([\alpha = a] \cap W_i \cap \bigcap_{j \neq i, \nu_j = 0} R_j \cap \bigcap_{j \neq i, v_j \neq 0} \neg R_j\right)$$

for any $(a, \nu) \in 3A$, where $W_i = R_i \backslash \Pi_i(\omega)$ 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 $\pi^*_{i,\omega} \in \Delta(3A)$. 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 $\left(A_{(0)}, p\right)$-correlated equilibrium of G; moreover, it is immediate that player i's expectation conditional on playing $(\alpha_i(\omega), 2)$ induced by $\pi_{i,\omega}$ is exactly $\mathbb{E}_B\left[u_i(\alpha_{-i}, \alpha_i(\omega))\,|\Pi_i(\omega)\right]$. For the left inclusion, take Lemma 2 for the case of $n = 3$, $k = 0$, and $|\Theta| = 1$.

C Proof of Theorem 2

Non-emptiness 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)-$CE\,(2G)$. 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 non-empty, 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 non-empty, convex, compact polyhedra, the sequence $\left(\Pi_{\rho}\left(p^{n}, \hat{\pi}\right)\right)_{n}$ 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 $A_{i}^{\infty}, A^{\infty}$, 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 $p^{\prime} \geq p$. We show that for p sufficiently close to 1, behaviour 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 $\pi\left(\left(A^{\infty}\right)^{c}\right) \leq 1-p$. Consider the game $G^{0}=G$ and pick some $p^{1}<1$. It immediately follows from p-rationality that for $p \in\left[p^{1}, 1\right]$, 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 $p \in\left[p^{n-1}, 1\right]$, 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 $p^{n} \in\left[p^{n-1}, 1\right)$ 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 $p \geq p^{n}$.[27] This implies that for any $p \in\left[p^{n}, 1\right]$ and any $\pi \in p$-$RO(G)$, we also have $\pi\left(\left(A^{n}\right)^{c}\right) \leq 1-p$. Finiteness of the game implies that the process ends after finitely many steps implying that indeed there exists $p^{\infty}<1$ such that for $p \in\left[p^{\infty}, 1\right]$ and any $\pi \in p$-$RO(G)$, we have $\pi\left(\left(A^{\infty}\right)^{c}\right) \leq 1-p$. Taking $\bar{p}=p^{\infty}$ proves the claim.

E Proof of Result in Remark 2

Following Aumann (1974, 1987), for any $X=\prod_{i \in I} X_{i} \subseteq A$, we say that the family $\left(\pi_{i}\right)_{i \in I} \subseteq(\Delta(A))^{I}$ is a (X, p)-*subjective correlated equilibrium* of G, if, for any $i \in I$,

- For any $a_{i} \in X_{i}$, the following incentive constraints are satisfied,

$$\sum_{a_{-i} \in A_{-i}} \pi_{i}\left(a_{-i} ; a_{i}\right)\left[u_{i}\left(a_{-i} ; a_{i}\right)-u_{i}\left(a_{-i} ; a_{i}^{\prime}\right)\right] \geq 0,$$

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

for any $a_i' \in A_i$.

- For any $a_i \in A_i$ the following p-belief constraint is satisfied,

$$\sum_{a'_{-i} \in X_{-i}} \pi_i \left(a'_{-i}; a_i\right) \geq p \sum_{a_{-i} \in A_{-i}} \pi_i \left(a_{-i}; a_i\right).$$

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

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

where $X = \prod_{i \in I}(A_i \times \{0\})$, 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 $(a_i)^{i \in I} \subseteq A$, and for any $i \in I$, $\pi_i = 1_{\{a^i\}}$. Fix $k \in \{0, 1\}$, and define, for any $i \in I$, $\hat{\pi}_i = 1_{\left\{\left(\left(a^i_j, k\right) j \neq i; \left(a^i_i, 1-k\right)\right)\right\}}$. It is immediate that $\mathbf{marg}_A\left((\hat{\pi}_i)_{i \in I}\right) = (\pi_i)_{i \in I}$. Now, take $i \in I$, then the incentive constraints are trivially satisfied, since $\hat{\pi}_i\,(2A_{-i} \times (A_i \times \{0\})) = 0$. 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 $X = \prod_{i \in I}(A_i \times \{0\})$, we have $\left(1_{\{a^i\}}\right)_{i \in I} \in \mathbf{marg}_A\,((X, p)\text{-}SCE\,(2G))$ for any $\left(a^i\right)_{i \in I} \subseteq A$, so by convexity, $\mathbf{marg}_A\,((X, p)\text{-}SCE\,(2G)) = (\Delta\,(A))^I$.