Rank-Dependent Choice Equilibrium: A Non-Parametric Generalization of QRE

Jacob K. Goeree, Charles A. Holt, Philippos Louis,
Thomas R. Palfrey, and Brian Rogers[1]

April 19, 2018

Abstract

Quantal Response Equilibrium (QRE) builds the possibility of errors into an equilibrium analysis of games. One objection to QRE is that specific functional forms must be chosen to derive equilibrium predictions. As these can be chosen from an infinitely dimensional set, another concern is whether QRE is falsifiable. Finally, QRE can typically only be solved numerically. We address these concerns through the lens of a novel set-valued solution concept, Rank-Dependent Choice Equilibrium (RDCE), which imposes a simple ordinal monotonicity condition: equilibrium choice probabilities are ranked the same as their associated expected payoffs. We first discuss important differences between RDCE and QRE and then show that RDCE envelopes all QRE models. Finally, we show that RDCE (and, hence, QRE) is falsifiable since the measure of the RDCE set, relative to the set of all mixed-strategy profiles, converges to zero at factorial speed in the number of available actions.

Keywords: *rank-dependent choice, RDCE, QRE, falsifiability, empirical implications*

[1]Goeree: AGORA Center for Market Design, UNSW. Holt: Department of Economics, University of Virginia. Louis: Department of Economics, University of Cyprus. Palfrey: Division of the Humanities and Social Sciences, California Institute of Technology. Rogers: Department of Economics, Washington University St. Louis. The idea of Rank-Dependent Choice Equilibrium was introduced by Goeree in 2004 and was further explored by Goeree, Palfrey, and Rogers in 2005 and more recently by Goeree and Louis in 2017. These efforts have (finally) come to fruition in the form of this Handbook chapter and an independent paper by Goeree and Louis (2018). We gratefully acknowledge financial support from the Australian Research Council (DP150104491) and the National Science Foundation (SES-1426560).

1. Introduction

The central solution concept of game theory, Nash equilibrium, is based on two premises. First, players make *optimal* choices given their beliefs about others' behavior, and, second, these beliefs are *correct* in a probabilistic sense. Together these assumptions imply mutually consistent behavior in the form of an equilibrium distribution of actions that is stable, i.e. no player has an incentive to deviate from the planned action distribution.

While observed behavior conforms to Nash play in some settings, a large body of work in experimental game theory has documented systematic departures from Nash equilibrium. In some cases, the Nash equilibrium works well while in others it fails miserably. "Ten Treasures and Ten Contradictions of Game Theory" (Goeree and Holt, 2001), for instance, exhibits ten games in which Nash equilibrium predicts well but, under a simple change in parameters that preserves the game structure and Nash prediction, behavior shifts dramatically away from Nash. These anomalies have stirred interest in generalizations of the Nash equilibrium that relax the assumptions on which it relies.

One approach, pioneered by McKelvey and Palfrey (1995, 1998), is *Quantal Response Equilibrium (QRE)*. The methodological innovation of QRE is to incorporate the possibility of errors into an equilibrium framework. Players may make errors, i.e. choose actions that do not maximize expected payoffs, and are aware that other players may do so as well. This creates a feedback loop because errors of one player affect the expected payoffs of other players, which , in turn, affect their choice probabilities. Unlike prior approaches, QRE assumes errors are payoff sensitive, i.e. larger errors are committed less frequently. However, QRE retains the Nash premise that beliefs are correct on average. That is, a player's QRE choice probabilities are derived relative to the *true* expected payoffs of that player's possible actions, taking into account the distribution of errors across the other players.[1]

QRE has proven to be a remarkable tool for organizing data across a wide variety of games in experiments. For the "Treasures" studied by Goeree and Holt (2001), for instance, QRE predictions resemble those of the Nash equilibrium, while for the "Contradictions" they differ sharply. Importantly, QRE predictions match observed data from *both* the "Treasures" and the "Contradictions." Beyond fitting data, QRE has important implications for doing

[1] Goeree and Holt (2004) propose a model of "noisy introspection" that allows for surprises. The intuition is that players think through several steps of iterative reasoning but that each step becomes increasingly hard and, hence, noisier. See also Goeree, Louis, and Zhang (2017) for applications of the noisy introspection model to data from variants of the "11-20" game.

theory. In extensive-form games, for example, there is no play "off the equilibrium path" in QRE as all strategy profiles have a strictly positive chance of being used. QRE therefore avoids the subtle-yet-complex refinements necessary to address the probability-zero (out-of-equilibrium) events that standard theory has to deal with. Two decades of work on QRE, summarized in the recent book *Quantal Response Equilibrium: A Stochastic Theory of Games* by Goeree, Holt, and Palfrey (2016), has enhanced the relevance of game theory for the social sciences. Experimenters, econometricians, and game theorists now routinely use QRE for their data analyses and theoretical modeling.

Despite this success (or, perhaps, because of it), there has been criticism of QRE as well. One objection that has been raised by theorists is that QRE, unlike Nash equilibrium, is a parametric theory. For instance, in the *structural approach* originally proposed by McKelvey and Palfrey (1995), an error or noise term is added to expected payoffs, and each player chooses the action with highest perturbed payoff. In this formulation, QRE corresponds to the Bayesian equilibrium of the disturbed game where the error terms, typically assumed to be i.i.d., are private information (types). In the *regular approach* to QRE proposed by Goeree, Holt, and Palfrey (2005), best response functions are more directly replaced by smoothed-out "better response" functions, embodying the soft optimization of players' choices. With either the structural or regular approach, specific functional forms must be chosen (e.g. the error distributions) to derive equilibrium predictions. The model's predictions therefore depend on the ad hoc, if not arbitrary, choice of functional forms.

A related objection, raised by empiricists, is that the freedom to choose arbitrary functions might render QRE not falsifiable. For the structural approach, for instance, Haile, Hortacsu, and Kosenok (2008) show that without restrictions on the error structure, one can exactly match *any* profile of observed choices including those that violate payoff monotonicity.[2,3] The latter is not possible with regular QRE, which imposes monotonicity on the quantal response functions. However, also in this approach, the question remains what outcomes are possible when varying over all possible quantal response functions.

A final objection is that the QRE fixed-point equations can typically be solved only nume-

[2]This observation also applies to the non-game-theoretic additive random utility models frequently used in econometrics.

[3]Haile, Hortacsu, and Kosenok (2008) rely on error structures that violate interchangeability to construct non-monotonic choice probabilities. Since applications of structural QRE have typically assumed i.i.d. errors (which is stronger assumption than interchangeability), this part of their falsifiability critique has no apparent practical relevance for the application of QRE.

rically. For instance, the commonly used Logit and Probit models result in non-polynomial ("transcendental") fixed-point equations that cannot be solved analytically. The lack of closed-form solutions may hamper adoption of QRE in applied work, e.g. Industrial Organization.

We address these objections through the lens of a new equilibrium concept, *Rank-Dependent Choice Equilibrium* (RDCE), see also Goeree and Louis (2018). Rank-dependent choice rests on a simple ordinal monotonicity condition: if one action has a higher expected payoff than another, it is chosen with higher probability but there is no need to specify exactly by how much. All that is required is that choice probabilities are monotone in expected payoffs. The RDCE set then consists of those strategy profiles that are mutually consistent, i.e. those strategy profiles so that for all players the ranking of choice probabilities matches the ranking of expected payoffs calculated from other players' strategies.

RDCE differs from QRE in a number of important ways. First, RDCE is a *non-parametric* theory. It does not rely in any way on a particular choice of functional form and can therefore not be criticized as generating predictions that rely on a modeler's choice of response functions or error distributions, as those concepts have no meaning in RDCE.

Second, RDCE is a set-valued equilibrium concept: since RDCE only imposes conditions on the *ranks* of choice probabilities, there is typically a continuous range of choice probabilities consistent with those ranks. Consider, for example, a symmetric dominant-strategy solvable game (e.g. Prisoner's Dilemma) with two possible actions, a_1 and a_2, and a_1 has the higher expected payoff. Then the RDCE set contains of all strategy profiles where both players choose a_1 more likely, i.e. $\{p_1(a_1), p_2(a_1)\} \in [\frac{1}{2}, 1]^2$. While this set may be relatively large for "small games" we show its size falls quickly in the number of available actions. Moreover, even for two-by-two games the RDCE set can be arbitrary small.

Third, computation of the RDCE *set* does not involve any fixed-*point* equations. Rather it follows from computing the choice probabilities for which the ranking of expected payoffs changes, i.e. for which there are payoff indifferences. We illustrate the construction of the RDCE set for a symmetric two-player game with three options in the next section. In this case, the RDCE set is easy to compute and can be illustrated graphically as a subset of the two-dimensional simplex. The idea is that expected payoffs are linear in (other players') choice probabilities so points of indifference, i.e. where the expected payoffs of two actions are the same, satisfy a linear equation. Either there are no solutions to the indifference equations in the simplex (when one option dominates another) or the solutions divide the probability simplex into two parts: a part where one expected payoff ranks higher and a complementary part where it ranks lower. Making all pairwise comparisons divides the probability simplex into at most 6 regions. Intersecting these regions with those that follow from ranking the probabilities themselves then yields the RDCE set, fully characterized by a collection of linear inequality conditions. This construction stands in contrast to QRE and Nash equilibrium, both of which require the solution of complex fixed-point equations. In fact, except in very special cases, no analytical solution exists for QRE and, instead, one typically resorts to numerical solutions of the model.

Despite these differences, there is an important connection between set-valued RDCE and various parametric theories, including QRE, that are based on fixed-points. Specifically, for the case of QRE, Theorem 1 establishes that for all finite games the RDCE set is essentially equivalent to the collection of fixed-points from all *different* QRE models. In other words, suppose one picks a set of quantal responses and solves the fixed-point equations, then picks a new set of quantal responses and solves the fixed-point equations again, and repeats this process *ad infinitum* until all possible quantal response functions have been sampled, then the collection of fixed-points so found fills up the RDCE set. Importantly, the same result holds for other (parametric) models, including those that do not fit the QRE framework (Goeree and Louis, 2018): RDCE is precisely the solution concept that nests *all* equilibrium theories that respect monotonicity of choice probabilities in expected payoffs.

Because RDCE generalizes all parametric theories that respect payoff monotonicity it can be used to address their falsifiability: what outcomes are possible if one does not commit to a particular set of parameters (e.g. quantal responses) upfront but instead tailors these to the data? Theorem 2 establishes an upper bound on the size of the RDCE set, which turns out to be quite strong. If K is the largest number of actions available to any player in the game, then the measure of the RDCE set relative to the measure of the full set of mixed strategy profiles is bounded above at $1/K!$. This shows tat RDCE is falsifiable. A fortiori, various parametric theories, including QRE, are falsifiable even when their parameters are estimated to best fit the data.

Section 4 illustrates the RDCE set for several commonly-studied two-by-two games and shows that the measure of the RDCE set is typically smaller than the $1/K!$ upper bound. This section also discusses comparative-statics properties of RDCE. Section 5 concludes.

2. An Example: the Lieberman Game

When confronted with data from experimental games, the Nash equilibrium suffers from two obvious problems – systematic contradictions to predicted behavior and overly restrictive point predictions. To illustrate, consider the symmetric two-player zero-sum game in Table 1, which is dominance solvable. Each player has three possible actions. The entries in the payoff matrix in Table 1 are Row's payoffs (and the negative of Column's payoffs). Action a_2 is strictly dominated by a_3 (for each player). Once these dominated actions are removed, a_1 is strictly dominated by a_3. Thus the unique rationalizable action, and so also the prediction from any standard equilibrium concept, is for both players to choose a_3. Yet, in the laboratory, one observes all three actions being chosen by both Row and Column players (Lieberman, 1960). Indeed, both Row and Column often choose a_1.

	a_1	a_2	a_3
a_1	0	15	−2
a_2	−15	0	−1
a_3	2	1	0

Table 1: A dominance solvable game (Lieberman, 1960).

In order to model behavior in games like this in a rigorous way, one needs a model that admits the statistical possibility for any strategy to be used. Indeed, without such a statistical model, for any game with a unique equilibrium in pure strategies, a single observation of a player using a different strategy (as is invariably the case in the laboratory) leads to immediate rejection of Nash equilibrium at any level of significance. This is sometimes referred to as the zero-likelihood problem, since the theoretical model assigns zero-likelihood to some data sets.[4] Call a model that assigns positive probability to all data sets *a statistical model.* Without a statistical model, standard maximum likelihood methods are virtually useless for analyzing most data sets generated by laboratory experiments, due to unspecified variation outside the model.

One commonly-used statistical model is the logit-QRE, which stipulates that choice probabilities are positively, but not perfectly, related to expected payoffs via a logit choice

[4]For the Lieberman game in Table 1, the Nash equilibrium assigns zero likelihood to nearly all data sets.

function. The probability of choosing a_k for $k = 1, \ldots, 3$ is given by:

$$p_k = \frac{\exp(\lambda \pi_k)}{\sum_{l=1}^{3} \exp(\lambda \pi_l)} \tag{1}$$

where the non-negative "precision" parameter λ determines the sensitivity of choices with respect to expected payoffs. When $\lambda = 0$, choice probabilities are uniform irrespective of expected payoffs, while only the action with the highest expected payoff is chosen as $\lambda \to \infty$. Importantly, the expected payoffs on the right side of (1)

$$\begin{aligned} \pi_1 &= 2p_1 + 17p_2 - 2 \\ \pi_2 &= -14p_1 + p_2 - 1 \\ \pi_3 &= 2p_1 + p_2 \end{aligned} \tag{2}$$

depend on the choice probabilities that appear on the left side of (1). In other words, (1) defines a set of fixed-point equations. These equations cannot be solved analytically but, for a given value of λ, they can be solved numerically. The red curve in the bottom triangle of Figure 1 shows the set of logit QRE: starting in the center when $\lambda = 0$ and ending at the unique Nash equilibrium choice a_3 when $\lambda = \infty$.

This example highlights three criticisms of QRE: (i) it relies on the numerical solution of a fixed-point equation, (ii) the equilibrium probabilities in (1) rely on the choice of a specific functional form, i.e. a choice rule different from logit, e.g., probit, would yield a different curve in the bottom triangle of Figure 1, which (iii) raises the falsifiability question of whether any point of the probability simplex can possibly be reached by using a suitably chosen choice rule.

Rank-dependent choice, in contrast, is a non-parametric model that stipulates that actions with higher expected payoffs are more likely to be chosen without explicitly modeling by how much. It relies only on the ordinal content of expected utility calculations. For instance, from the expected payoffs in (2) it is clear that $\pi_3 > \pi_2$, which means that a_3 is always more likely to be chosen than a_2. The ranking of a_1 depends on the choice probabilities. To find the regions of the probability simplex where a_1 ranks highest, middle, or lowest we can simply solve for points of indifference, i.e where $\pi_1 = \pi_2$ or $\pi_1 = \pi_3$, which yields $p_1 + p_2 = \frac{1}{16}$ and $p_2 = \frac{1}{8}$ respectively. These solutions are shown by the dashed lines

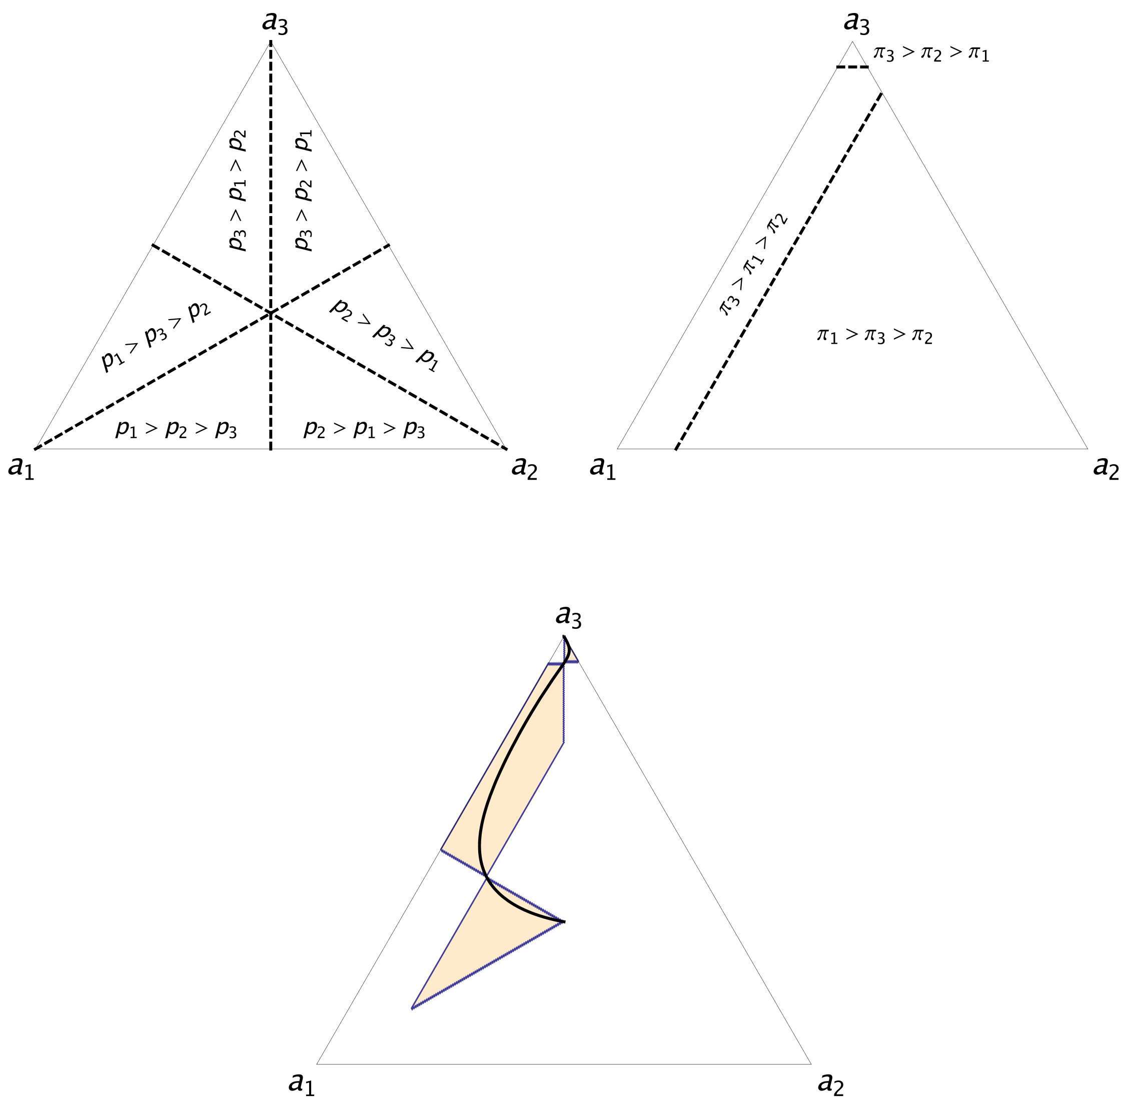

Figure 1: Construction of the set of RDCE for the Lieberman game. The top-left triangle shows the ranking of choice probabilities over the simplex and the top-right triangle shows the ranking of expected payoffs. Their "intersection," i.e. the points of the simplex where the two rankings match, define the set of RDCE and are indicated by the shaded areas in the bottom triangle. The curve that runs from the centroid to the Nash equilibrium shows the symmetric logit equilibria.

in the top-right triangle of Figure 1, which also indicates the ranking of expected payoffs in each of the three areas. A rank-dependent choice equilibrium (RDCE) now follows from the consistency condition, embedded in the equilibrium definition, that the ordering of payoffs matches the ordering of underlying choice probabilities as shown in the top-left triangle of Figure 1. The points of the simplex for which this matching holds constitute the RDCE set and are indicated by the shaded areas of the bottom triangle.

3. Rank-Dependent Choice Equilibrium

Consider a finite normal-form game $G = (N, \{A_i\}_{i=1}^n, \{\Pi_i\}_{i=1}^n)$, where $N = \{1, \ldots, n\}$ is the set of players, $A_i = \{a_{i1}, \ldots, a_{iK_i}\}$ is player i's set of K_i actions (or pure strategies), with $A = A_1 \times \cdots \times A_n$ the set of action profiles, and $\Pi_i : A \to \mathbb{R}$ is player i's payoff function. Let Σ_i be the set of probability distributions over A_i. An element $\sigma_i \in \Sigma_i$ is a mixed strategy for player i, which is a mapping from A_i to Σ_i, where $\sigma_i(a_{ik})$ is the probability that player i chooses pure strategy a_{ik}. Denote by $\Sigma = \prod_{i \in N} \Sigma_i$ the set of mixed strategy profiles. For each $k = 1, \ldots, K_i$ and any $\sigma \in \Sigma$, player i's expected payoff of choosing action a_{ik} is $\pi_{ik}(\sigma) = \sum_{a_{-i} \in A_{-i}} p(a_{-i}) \Pi_i(a_{ik}, a_{-i})$ where $p(a_{-i}) = \prod_{j \neq i} \sigma_j(a_j)$. Since player i has K_i possible actions, the vector of expected payoffs for these actions, denoted by $\pi_i = (\pi_{i1}, \ldots, \pi_{iK_i})$, is an element of $\mathbb{R}^{K_i}$.

3.1. Regular QRE

The main idea behind regular QRE is to replace best responses with "better" responses, i.e. instead of players simply choosing the strategy with the highest expected payoff, they instead choose higher-payoff strategies with greater probability. The sensitivity of choices with respect to expected payoffs is determined by the better-response function, or quantal response function, which can be thought of as a "smooth" version of the best-response function that underlies Nash equilibrium. In particular, a quantal response function for player i, denoted R_i, maps the vector of expected payoffs for i's alternative actions into a probability distribution of choices over these strategies. To ensure some degree of (stochastic) rationality we require the quantal response functions satisfy the intuitive notions of responsiveness and monotonicity, as well as the technical conditions of interiority and continuity. Such functions are called *regular*.

Definition 1 *$R_i : \mathbb{R}^{K_i} \to \Sigma_i$ is a* ***regular quantal response function*** *if it satisfies the following four axioms.*

(R1) ***Interiority****: $R_{ik}(\pi_i) > 0$ for all $k = 1, \dots, K_i$ and for all $\pi_i \in \mathbb{R}^{K_i}$.*

(R2) ***Continuity****: $R_{ik}(\pi_i)$ is a continuous function for all $k = 1, \dots, K_i$, $\pi_i \in \mathbb{R}^{K_i}$.*

(R3) ***Responsiveness****: $R_{ik}(\pi_i)$ is strictly increasing in π_{ik} for all $k = 1, \dots, K_i$, $\pi_i \in \mathbb{R}^{K_i}$.*

(R4) ***Monotonicity****: $\pi_{ik} > \pi_{il} \Rightarrow R_{ik}(\pi_i) > R_{il}(\pi_i)$ for all $k, l = 1, \dots, K_i$.*

Interiority ensures a statistical model, i.e. one that is logically consistent with all possible data sets. As mentioned, avoiding the zero-likelihood problem is important for empirical applications. *Continuity* ensures that R_i is single-valued, and also seems natural from a behavioral standpoint, since arbitrarily small changes in expected payoffs should not lead to jumps in choice probabilities. *Responsiveness* requires that if the expected payoff of an action increases, ceteris paribus, the choice probability must also increase. *Monotonicity* is a stochastic form of rationality that involves binary comparisons of actions: an action with higher expected payoff is chosen more frequently than one with a lower expected payoff.

Recall that an element $\sigma \in \Sigma$ specifies a mixed strategy profile, i.e., one mixed strategy for each player. Since players' beliefs can also be represented by probability distributions over other players' decisions, one can think of σ in two ways: as representing a common set of beliefs regarding play and as stochastic responses to the expected payoffs determined by those beliefs. Loosely speaking, a QRE is a set of belief distributions for which the stochastic responses are consistent with the beliefs, i.e. it is a set of mixed strategies that gets mapped into itself via the response functions. More formally, define $R(\pi) = (R_1(\pi_1), \dots, R_n(\pi_n))$ to be regular if each R_i satisfies axioms (R1)–(R4). Since $R(\pi) \in \Sigma$ and $\pi = \pi(\sigma)$ is defined for any $\sigma \in \Sigma$, $R \circ \pi$ defines a mapping from Σ into itself.

Definition 2 *Let R be regular. A* ***regular Quantal Response Equilibrium*** *of the finite normal-form game G is a mixed-strategy profile σ^* such that $\sigma^* = R(\pi(\sigma^*))$.*

Since regularity of R includes continuity, $R \circ \pi$ is a continuous mapping. Existence of a regular QRE therefore follows directly from Brouwer's fixed-point theorem. Like the Nash equilibrium, this fixed-point is not necessarily unique. Let $QRE_R(G)$ denote the set of all fixed-points for the finite game, G, given the quantal responses $R = (R_1, \ldots, R_n)$. For generic games, $QRE_R(G)$ is a finite set.

Definition 2 makes explicit that regular QRE depends on the choice of R, the vector of quantal response functions, and it is in this sense that QRE can be called a parametric theory. A common choice, and the one used in Section 2, is to set the elements of R equal to

$$R_{ik} = \frac{\exp(\lambda \pi_{ik})}{\sum_{l=1}^{K_i} \exp(\lambda \pi_{il})}, \qquad i = 1, \ldots, n, \quad k = 1, \ldots, K_i, \tag{3}$$

which is known as the logit choice rule and the corresponding equilibrium is the logit QRE. Obviously, many other choice rules can be used, e.g. the Probit choice rule, see Goeree, Holt, and Palfrey (2016) for an overview. Since regular QRE is a parametric theory, it is not clear *a priori* that its predictions are falsifiable non-parametrically given that the parameter space is infinite dimensional, i.e. the set of all regular quantal response *functions*.

We want to understand the nature of the empirical restrictions of regular QRE, i.e., what is the full set of strategy profiles that can be generated from all *different* regular QRE. This appears to be a difficult question to address, since any given QRE is difficult to compute and the space of all possible quantal response functions is infinitely dimensional. It turns out, however, that understanding the "envelope theory" of all different QRE models is much simpler than computing any particular QRE.

In order to do so, we apply the Rank-Dependent Choice Equilibrium (RDCE) concept. RDCE is a non-parametric theory, which relies only on the ordinal ranking of expected payoffs and mutual consistency. We show an equivalence between the RDCE set and the envelope of all different QRE models. We further show that the RDCE set is easy to compute via payoff indifference conditions, and quickly becomes small as the number of actions in the game grows.

3.2. RDCE

Let $\text{rank}(v_k)$ denote the set-valued rank of the k-th element of a vector $v \in \mathbb{R}^K$. Then, for example, if action a_{ik} yields higher expected payoffs than all other actions $\text{rank}(\pi_{ik}) = \{1\}$. On the other hand, if action a_{ik} and a_{il} (with $l \neq k$) yield the same expected payoff, which is higher than that of all other strategies, then $\text{rank}(\pi_{ik}) = \text{rank}(\pi_{il}) = \{1, 2\}$.

Player i's best response set, $BR_i \subset \Sigma$, can be defined as follows:

$$BR_i(\Sigma) \;=\; \{\sigma \in \Sigma \,|\, \sigma_{ik} = 0 \text{ if } 1 \notin \text{rank}(\pi_{ik}) \text{ for } k = 1, \ldots, K_i\}\,.$$

In other words, a best response assigns zero probability to any action that yields less than the maximum expected payoff. Instead, consider a relaxation of the traditional best response based on the idea that players choose strategies with higher expected payoffs more frequently without specifying by how much.

Definition 3 *Player i's* ***rank-dependent response set*** *is given by*

$$RDR_i(\Sigma) \;=\; \{\sigma \in \Sigma \,|\, \text{rank}(\sigma_{ik}) \subset \text{rank}(\pi_{ik}) \text{ for } k = 1, \ldots, K_i\}\,.$$

Note that the rank-dependent response set is typically not closed. In particular, it does not include limit points where choice probabilities are equal but expected payoffs are strictly ranked.[5]

Just like the set of Nash equilibria for the game G can be defined as the intersection of players' best responses, the RDCE set is given by the intersection of players' rank-dependent response sets.

Definition 4 *For any normal-form game G, the Rank-Dependent Choice Equilibrium set is given by*

$$RDCE(G) \;=\; \text{cl}(\cap_i RDR_i).$$

The closure is necessary as the rank-dependent response set does not contain all its limit points.

RDCE is thus an explicitly *set*-valued solution concept, unlike QRE and Nash, which correspond to specific strategy profiles that solve certain fixed-*point* equations. There is, however, a connection between the RDCE set and the collection of fixed-points from all *different* QRE models. Let $\mathcal{R}$ denote the (infinitely-dimensional) space of regular quantal

[5]For instance, given any σ, the centroid of Σ_i, where each action in A_i is played with equal probability $\frac{1}{K_i}$ ("purely random behavior"), can be approximated arbitrarily closely by elements of $RDR_i(\sigma)$, but for most games and most σ_{-i} the centroid of Σ_i is not an element of $RDR_i(\sigma)$, as expected payoffs are typically not all equal.

functions with typical element $R \in \mathcal{R}$. Recall that, $QRE_R(G)$ denotes the finite set of QRE fixed-points for the particular quantal response, R.

Theorem 1 *For any normal-form game G, the RDCE set is nonempty, closed, and is essentially the same as the union of fixed-points for all* **different** *QRE models, i.e.*

$$RDCE(G) \,=\, \mathrm{cl}(\bigcup_{R \in \mathcal{R}} QRE_R(G))$$

In particular, $RDCE(G)$ contains all Nash equilibria of G and perfectly random behavior.

Proof. Condition (R4) of Definition 1 implies that in any regular QRE, choice probabilities satisfy payoff monotonicity. Hence, $QRE_R(G)$ is an element of $RDCE(G)$ for any R. Since the RDCE set is closed by Definition 4, it contains all its limit points so we have

$$RDCE(G) \,\supset\, \mathrm{cl}(\bigcup_{R \in \mathcal{R}} QRE_R(G))$$

The right side is non-empty, hence, so is the RDCE set. Since both Nash equilibria and perfectly random behavior can be seen as limit points of QRE (e.g. the logit-QRE in (3) with $\lambda \to \infty$ and $\lambda \to 0$ respectively), they are contained in the RDCE set as well.

To establish the reverse, i.e.

$$RDCE(G) \,\subset\, \mathrm{cl}(\bigcup_{R \in \mathcal{R}} QRE_R(G))$$

we next show that for any strategy profile, σ, in the interior of the RDCE set one can construct a set of quantal response functions, $R = (R_1, \ldots, R_n)$, so that $\sigma \in QRE_R(G)$.

Denote the set of totally mixed strategies by $\Sigma^o = \{\sigma \in \Sigma \,|\, \sigma_{ik} \in (0,1) \text{ for all } i\,\{1,\ldots,n\}, k \in \{1,\ldots,K_i\}\}$ and denote the interior of the RDCE set by $RDCE^o(G) = RDCE(G) \cap \Sigma^o$. Let $\sigma \in RDCE^o(G)$ then the K_i components of $\pi_i(\sigma)$ are ordered the same as the K_i components of σ_i. Without loss of generality, relabel i's actions so that $j < k$ implies $\pi_{ij}(\sigma) \leq \pi_{ik}(\sigma)$ and $\sigma_{ij} \leq \sigma_{ik}$. Construct a strictly positive, continuous, strictly increasing function, $h_i : \mathbb{R} \to \mathbb{R}^+$ as follows. For each $j \in \{1,\ldots,K_i\}$ let $h_i(\pi_{ij}(\sigma)) = \sigma_{ij}$. This pins down the value of h_i at K_i ordered points[6] and these values are increasing. Next let

$$h_i(x) = \begin{cases} \frac{\sigma_{i1}}{1+\pi_{i1}(\sigma)-x} & \text{for} \quad x \leq \pi_{i1}(\sigma) \\ \frac{\sigma_{ij+1}(x-\pi_{ij}(\sigma))+\sigma_{ij}(\pi_{ij+1}(\sigma)-x)}{\pi_{ij+1}(\sigma)-\pi_{ij}(\sigma)} & \text{for} \quad \pi_{ij}(\sigma) \leq x \leq \pi_{ij+1}(\sigma), \quad j = 1,\ldots,K_i-1 \\ \sigma_{iK_i} + x - \pi_{iK_i}(\sigma) & \text{for} \quad x \geq \pi_{iK_i}(\sigma) \end{cases}$$

which extends h_i to the real line such that it is strictly increasing and strictly positive everywhere. Now, for each $i \in \{1,\ldots,n\}$, define i's quantal response function R_i as

$$R_{ij}(\pi_i) = \frac{h_i(\pi_{ij})}{\sum_{k=1}^{K_i} h_i(\pi_{ik})}$$

It is readily verified R_i satisfies conditions (R1)–(R4) of Definition 1 and, hence, is a regular quantal response function. Furthermore, by construction we have $\sigma_{ij} = R_{ij}(\pi_i(\sigma))$ for all $i \in \{1,\ldots,n\}$ and for all $j \in \{1,\ldots,K_i\}$. Hence, σ is a regular QRE for the regular quantal response functions $R = (R_1,\ldots,R_n)$. ■

Remark 1 While Theorem 1 focuses on QRE, it should be noted that a similar result holds for other (parametric) equilibrium concepts that respect payoff monotonicity, see Goeree and Louis (2018).

Remark 2 Usual applications of QRE, for reasons of parsimony, assume that each player has the same quantal response function and estimate a single parameter, such as λ in the case of logit-QRE. While Theorem 1 does not impose symmetry, the empirical restrictions on the data that are implied by requiring all players to have the same quantal response function would be more severe. It is an open question as to exactly how much stronger the empirical implications would be for single-parameter models.

Remark 3 The result that all QRE lie inside the RDCE set also applies to Heterogeneous QRE (HQRE), which is characterized and analyzed in Rogers et al. (2009). That is, suppose that each player behaves according to a quantal response function that is drawn from a distribution of possible quantal response functions. For example, each player has a logit response function with precision parameters that are independent draws from a log normal

[6] There may be fewer than K_i points if some of i's actions yield the same expected payoffs at σ.

distribution over λ. Then the prediction of HQRE behavior in the aggregate will be restricted to the RDCE set, for any distribution of quantal response functions from which player types are drawn.

One might worry that RDCE lacks predictive power as it is far less restrictive than the Nash equilibrium or logit-QRE, both of which restrict choice probabilities to a finite set of (mixed) strategy profiles for almost all games. However, we next show that as the size of action sets increase, the RDCE set converges quickly to a sharp prediction in almost all games. To show this result, we constrain attention to the set of games $\mathcal{G}$ with the property that for each player i, the set of strategy profiles σ_{-i} that make i indifferent between two actions has measure zero in Σ_{-i}. This is a generic class of games, i.e. $\mathcal{G}$ is open and dense in the space of all finite normal games.

Theorem 2 *For every game $G \in \mathcal{G}$, $\mu(RDCE(G)) \leq 1/K!$ where $K = \max_i K_i$.*

Proof. Let i^* be such that $K_{i^*} = K$. By the definition of $\mathcal{G}$, at almost every σ_{-i^*}, i^* has a strict ranking of expected payoffs associated with the K possible actions. So the measure of player i^*'s rank-dependent response set is $1/K!$ for almost every σ_{-i^*}. In other words, even without imposing an equilibrium consistency condition (i.e. intersecting the rank-dependent response sets across players), the measure of the RDCE set is already limited to $1/K!$. ■

The proof of Theorem 2 makes clear that the theoretical upper bound is attained only when players' rank-dependent response sets are identical (and, hence, also equal to their intersection). As we show in the next section, this is a rather special case. Typically, the measure of the RDCE set is smaller than the upper bound in Theorem 2.

4. Applications to 2×2 Games

Theorem 2 shows that the RDCE set shrinks quickly as the number of possible actions grows. RDCE seems least restrictive in small games, which is one reason why we discuss 2×2 games in some detail. Another reason is that it conveniently allows us to graph the RDCE set, even for asymmetric games, as part of the set of mixed-strategy profiles $\Sigma = [0, 1]^2$ (i.e. the unit square).

4.1. A Prisoner's Dilemma Game

In the prisoner's dilemma game, two players face the choice to stay silent (S) and do jail time or talk (T), i.e. provide evidence incriminating the other player, to get a reduced sentence. The game payoffs are shown in Table 2. Note that irrespective of the other's choice, a player is better off talking: $\pi(T) = \pi(S) + 1$. In other words, T is the dominant strategy.

	S	T
S	-1, -1	-3, 0
T	0, -3	-2, -2

Table 2: A prisoner's dilemma game.

Let $p_R(S)$ and $p_C(S)$ denote the probabilities with which S is chosen by Row and Column respectively. In terms of rank-dependent response sets the dominant-strategy property implies $p_R(S) < \frac{1}{2}$ for all $p_C(S)$ and, likewise, $p_C(S) < \frac{1}{2}$ for all $p_R(S)$. These sets are shown as the light-shaded areas in Figure 2. Their intersection, the dark-shaded area, corresponds to the RDCE set.

The filled square at $(\frac{1}{2}, \frac{1}{2})$ corresponds to perfectly random behavior while the filled circle at the origin indicates the unique Nash equilibrium. Furthermore, let $F(x) = \exp(x)/(1 + \exp(x))$ denote the logistic distribution function, then the logit QRE probabilities satisfy

$$\begin{aligned} p_R(S) &= F(\lambda_R(\pi_R(S) - \pi_R(T))) \\ p_C(S) &= F(\lambda_C(\pi_C(S) - \pi_C(T))) \end{aligned} \tag{4}$$

The solid red line connecting the center and origin shows the set of symmetric logit QRE, i.e. when $\lambda = \lambda_R = \lambda_C$ ranges from $\lambda = 0$ (random) to $\lambda = \infty$ (rational). It may seem intuitive that the predicted choice probabilities are symmetric as the game is symmetric. However, this result also relies on the assumption that both players use the same logit quantal response functions, e.g. with the same precision parameter. The low / middle / high dashed red lines in the bottom-left quadrant of Figure 2 show the set of asymmetric logit QRE when Column's precision parameter is half / thrice / ten times that of Row.

A priori there is no reason to assume that players' responsiveness to expected payoff

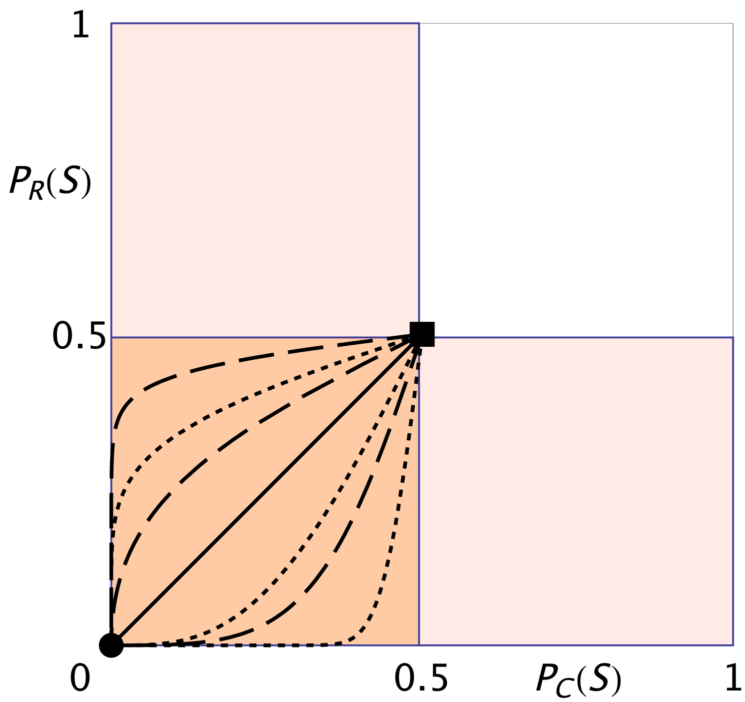

Figure 2: The rank-dependent response sets (light) and the rank-dependent choice equilibria (dark) for the Prisoner's Dilemma game in Table 2. The filled square at $(\frac{1}{2}, \frac{1}{2})$ corresponds to random behavior and the filled circle at $(0, 0)$ to the unique Nash equilibrium. The solid (long-dashed) curves show the set of symmetric (asymmetric) logit QRE while the short-dashed curves show the set of asymmetric probit QRE.

differences is the same – the RDCE set shows the possible outcomes for any combination of precision parameters and, in fact, for any set of quantal response functions. For example, probit QRE choice probabilities follow by replacing the logistic distribution function $F(\cdot)$ in (4) with the cumulative distribution for a standard normal variable, $\Phi(\cdot)$. The dotted blue lines in the bottom-left quadrant of Figure 2 show the set of asymmetric probit QRE when Row's precision parameter is half / thrice / ten times that of Column.

Note that the different dashed (dotted) lines suggest that any point in the RDCE set can be seen as a logit (probit) QRE for some asymmetric combination of precision parameters. For the prisoner's dilemma game this can easily be made more precise since $\pi(S) = \pi(T) - 1$. So (4) is equivalent to

$$\begin{aligned} \lambda_R &= F^{(-1)}(1 - p_R(S)) \\ \lambda_C &= F^{(-1)}(1 - p_C(S)) \end{aligned} \tag{5}$$

for any $p_R(S) \leq \frac{1}{2}$ and $p_C(S) \leq \frac{1}{2}$, where we used $F(-x) = 1 - F(x)$. In other words, any point in the RDCE set can be seen as a logit QRE with positive parameters (λ_R, λ_C) that satisfy (5). Equivalently, any point in the RDCE set can be seen as probit QRE with parameters $\tilde{\lambda}_R = \Phi^{(-1)}(F(\lambda_R))$ and $\tilde{\lambda}_C = \Phi^{(-1)}(F(\lambda_C))$. These findings suggest a conjecture that all regular QRE can be obtained via an asymmetric logit QRE, i.e., that the infinite dimensional space of all regular QRE reduces to the two-dimensional choice of logit precision parameters (λ_R, λ_C) in the case of two players. Finally, note that the measure of the RDCE set is only $\frac{1}{4}$, i.e. half the upper bound of Theorem 2.

4.2. A Minimum-Effort Coordination Game

Consider two students who work together on a project and who simultaneously have to decide whether to exert low effort, $L = 1$, or high effort, $H = 2$. The grade or payoff from the project to both is determined by the lowest of the efforts exerted: if both students work hard the payoff is 2 for each, but if one or both students slack off, the payoff is only 1. Moreover, every unit of effort costs $0 < c < 1$.

	L	H
L	$1-c,\ 1-c$	$1-c,\ 1-2c$
H	$1-2c,\ 1-c$	$2-2c,\ 2-2c$

Table 3: A 2×2 coordination game.

The normal-form representation of this game, see Table 3, represents a coordination game in the following sense. If Row chooses L, Column wants to choose L (and vice versa) and if Row chooses H, Column wants to choose H (and vice versa). These are the two pure-strategy Nash equilibria of the game. In addition, there exists a mixed-strategy Nash equilibrium of the game where both players choose L with probability $1 - c$ and H with probability c.

The theoretical construct most commonly used to "select" among the equilibria in 2×2 coordination games is Harsanyi and Selten's (1988) notion of risk dominance. Risk dominance can, in this case, be thought of as maximizing the loss incurred when deviating from equilibrium play and, in turn, this depends on the parameter c, as follows. When both players are choosing low efforts, the cost of a unilateral deviation to high effort is just the

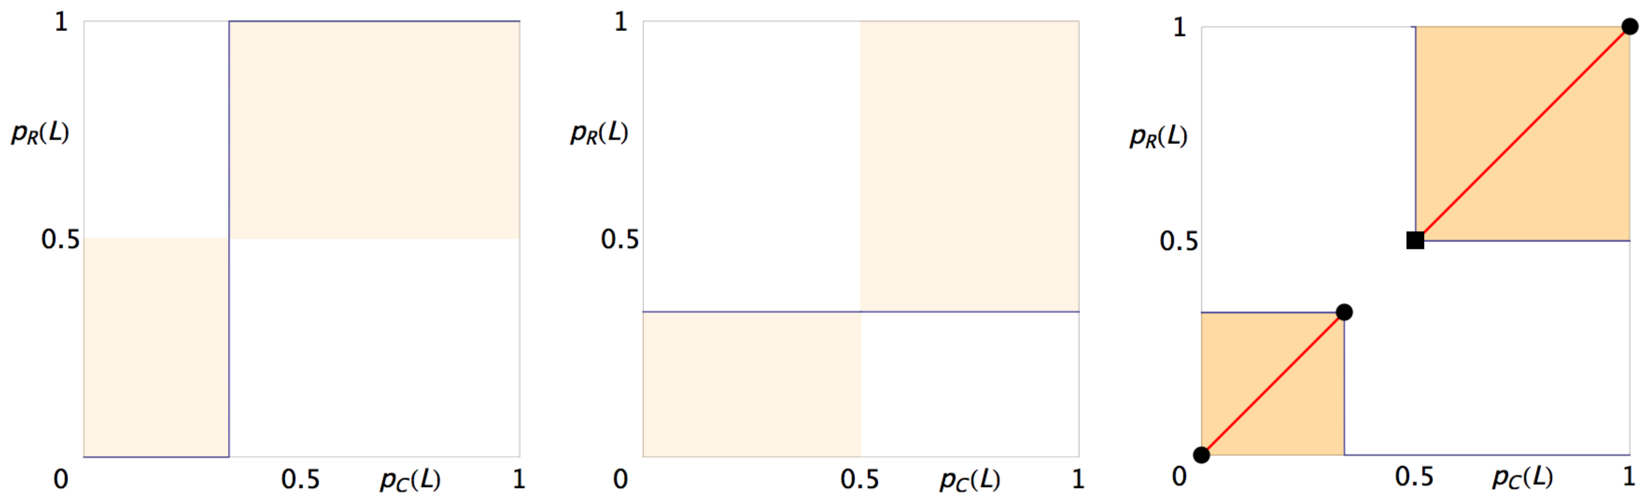

Figure 3: The rank-dependent response sets for Column (left panel) and Row (middle panel) and the RDCE set (right panel) for the Minimum-Effort Coordination game in Table 3. In the right panel, the filled square at $(\frac{1}{2}, \frac{1}{2})$ corresponds to random behavior and the filled circles at $(0, 0)$, $(1 - c, 1 - c)$, and $(1, 1)$ correspond to Nash equilibria. The lines show the set of symmetric logit QRE.

cost of the extra effort, c, which will be referred to as the "deviation loss." Similarly, the deviation loss at the (H, H) equilibrium is $1-c$, since a unilateral reduction in effort reduces the minimum by 1 but saves the marginal effort cost c. In this case, risk dominance selects the equilibrium that has the highest deviation loss. The deviation loss from the low-effort equilibrium is greater than that from the high-effort equilibrium if $c > 1-c$, or equivalently, if $c > \frac{1}{2}$, in which case the low-effort equilibrium is risk dominant. Risk dominance, therefore, has the desirable property that it selects the low-effort outcome if the cost of effort is sufficiently high.

The RDCE set is also sensitive to effort costs, but of course does not make point predictions. The left panel of Figure 3 shows Column's rank-dependent response set when $c > \frac{1}{2}$. When Column's probability of choosing L is less than $1-c$, Row is more likely to choose H, i.e. $p_R(L) < \frac{1}{2}$, and otherwise Row is more likely to choose L, i.e. $p_R(L) > \frac{1}{2}$. Similarly, the middle panel shows the rank-dependent response set for Row. The right panel shows their intersection, i.e. the RDCE set, as well as the Nash equilibria (three filled circles), random behavior (filled square), and the symmetric logit QRE (red line segments).

Figure 3 assumes $c > \frac{1}{2}$, in which case the measure of the RDCE set with L being the more likely choice, $\mu_L = \frac{1}{4}$, exceeds the measure of the RDCE set in which H is the more likely choice, $\mu_H = (1-c)^2$ and in this sense makes qualitatively similar predictions to risk dominance. The reverse is true when $c < \frac{1}{2}$, in which case these measures are $\mu_L = c^2$ and $\mu_H = \frac{1}{4}$ respectively. To summarize the total measure of the RDCE set is

$$\mu_L + \mu_H = \begin{cases} \frac{1}{4} + c^2 & \text{if } \; 0 < c \leq \frac{1}{2} \\ \frac{1}{4} + (1-c)^2 & \text{if } \; \frac{1}{2} \leq c < 1 \end{cases}$$

and the relative measure is given by

$$\frac{\mu_H}{\mu_L} = \begin{cases} (2c)^{-2} & \text{if } \; 0 < c \leq \frac{1}{2} \\ (2-2c)^2 & \text{if } \; \frac{1}{2} \leq c < 1 \end{cases}$$

which is strictly decreasing in the effort cost, c. In other words, RDCE makes the comparative statics prediction that an increase (decrease) in effort cost enlarges the set of outcomes in which L (H) is more frequently chosen.

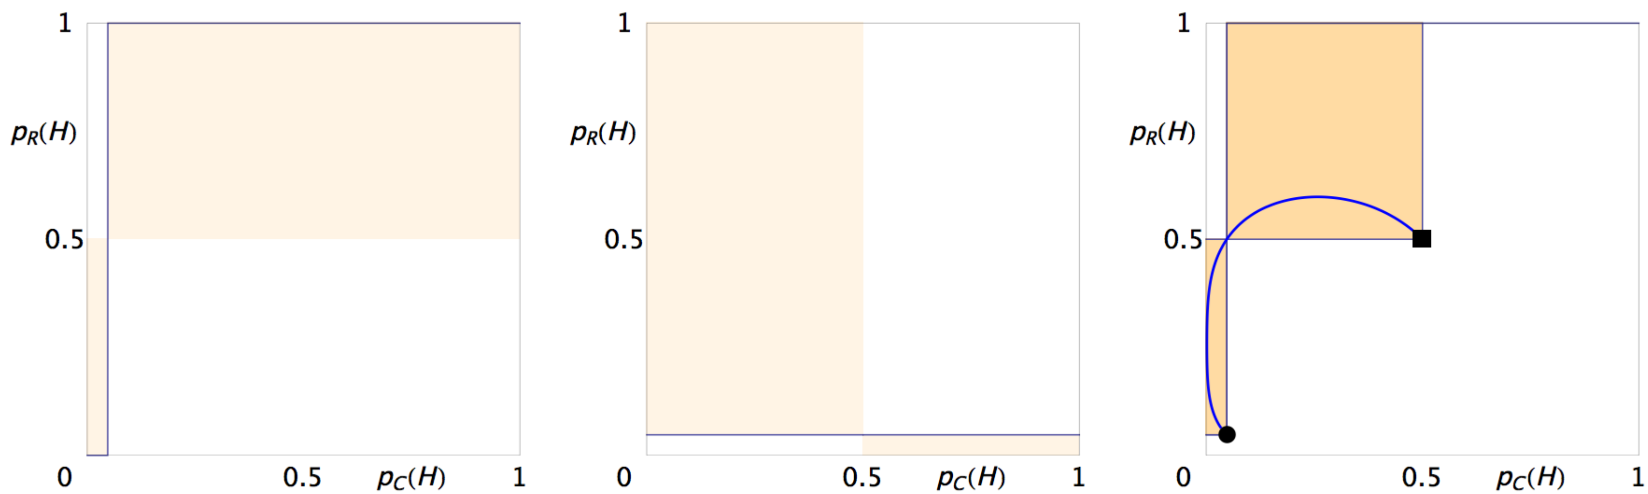

Figure 4: The RDCE set for the asymmetric matching-pennies game in Table 4 for $X = 20$. The filled square at $(\frac{1}{2}, \frac{1}{2})$ corresponds to random behavior and the filled circle at $(\frac{1}{X+1}, \frac{1}{X+1})$ corresponds to the unique mixed-strategy Nash equilibrium. The curve shows the set of probit QRE when $\lambda = \lambda_R = \lambda_C$ varies from $\lambda = 0$ (square) to $\lambda = \infty$ (circle).

4.3. An Asymmetric Matching-Pennies Game

A final example is the asymmetric matching-pennies game shown in Table 4 with $X > 1$. This game has a unique mixed-strategy Nash equilibrium in which both the Row and Column select Heads (H) with probability $1/(X+1)$ and Tails (T) with complementary probability $X/(X+1)$. For large X, the Nash equilibrium prediction puts almost all weight on the (T, T) outcome, which seems unintuitive from Row's point of view.

	H	T
H	X, 0	0, X
T	0, 1	1, 0

Table 4: An asymmetric matching-pennies game.

Figure 4 shows the RDCE set, the Nash equilibrium (filled circle) and random behavior (filled square) as well as the set of symmetric logit QRE for $X = 20$. The measure of the RDCE is readily computed as

$$\mu = \frac{1}{4} - \left(\frac{1}{X+1}\right)^2$$

and highlights that the theoretical upper bound of $\frac{1}{2}$ (see Theorem 2) is again loose: the measure of the RDCE set is less than $\frac{1}{4}$ for all $X > 1$ and tends to zero when X limits to 1.

Finally, note that for large X, RDCE makes the intuitive prediction that both players are more likely to choose the action that might yield X, i.e. $p_R(H) \geq \frac{1}{2}$ and $p_C(T) \geq \frac{1}{2}$.

5. Conclusion

Quantal Response Equilibrium (QRE) has proven an extremely useful tool for experimenters and theorists alike. It serves as a statistical model that can be applied to any data set from the laboratory or the field without running into zero-likelihood problems. It provides intuitive comparative statics predictions that are often absent in standard theoretical analyses. Yet, QRE has met some criticism as well. Its predictions rely on making specific parametric assumptions about the quantal response functions that map expected payoffs into choice probabilities. Moreover, the parameter space (i.e. the space of possible quantal response functions) is infinite dimensional, which raises the question of whether QRE is falsifiable. Finally, the fixed-point equations that define QRE can typically only be solved numerically.

We address these criticisms through the lens of a novel set-valued equilibrium concept called *Rank-Dependent Choice Equilibrium* (RDCE) (see also Goeree and Louis, 2018) . Rank-dependent choice relies on the idea that actions with higher expected payoffs are more likely to be chosen *without* explicitly specifying by how much. The RDCE consists of all strategy profiles that are mutually consistent, i.e. for each player, choice probabilities are ranked the same as their associated expected payoffs (given others' strategies). As we have shown, the RDCE set is often simple to compute, see e.g. the Lieberman example of Section 2. For multi-person games with large action sets the computation of the RDCE set is not as straightforward but its conceptual simplicity remains.

While the characterization of RDCE is achieved through completely different – and much simpler – means than QRE, it can be seen as an "envelope theory" for all QRE models. The RDCE set of a normal-form game G is essentially the same as the union of fixed-points from all *different* QRE models, see Theorem 1. Furthermore, for generic games, the size of the RDCE set decreases at factorial speed in the number of available actions, see Theorem 2. This demonstrates that RDCE (and, a fortiori, QRE) is a falsifiable theory.

Like the Nash equilibrium, RDCE is parameter free, but unlike the Nash equilibrium, it provides intuitive comparative statics. In coordination games, for instance, RDCE predicts that an increase in effort cost shrinks the set of predicted outcomes in which high efforts are more likely to be chosen. And in games with an asymmetric mixed-strategy Nash equilibrium, RDCE predicts that own payoff effects matter, as robustly observed. We leave it to future research whether RDCE can be as successfully applied in other contexts.

References

Goeree, J.K. and Holt, C.A. (2001) "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5), 1402–1422.

Goeree, J.K. and Holt, C.A. (2004) "A Model of Noisy Introspection," *Games and Economic Behavior*, 46(2), 365–382.

Goeree, J.K., Holt, C.A., and Palfrey, T.R. (2005) "Regular Quantal Response Equilibrium," *Experimental Economics*, 8, 347–367.

Goeree, J.K., Holt, C.A., and Palfrey, T.R. (2016) *Quantal Response Equilibrim: A Stochastic Theory of Games*, Princeton University Press.

Goeree, J.K., Louis, P., and Zhang, J. (2017) "Noisy Introspection in the 11-20 Game," *Economic Journal*, forthcoming.

Goeree, J.K. and Louis, P. (2018) "Rank-Dependent Choice Equilibrium," UNSW working paper.

Haile, P. A., Hortacsu, A., and Kosenok, G. (2008) "On the Empirical Content of Quantal Response Equilibrium," *American Economic Review*, **98**, 180-200.

Lieberman, B. (1960) "Human Behavior in a Strictly Determined 3×3 Matrix Game," *Behavioral Science*, **5**, 317-322.

McKelvey, R.D. and Palfrey, T.R. (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 10(1), 6–38.

McKelvey, R.D. and Palfrey, T.R. (1998) "Quantal Response Equilibria for Extensive Form Games," *Experimental Economics*, **1**, 9-41.

Rogers, B., Palfrey, T.R., and C. Camerer (2009) "Hetergeneous Quantal Response Equilibrium," *Journal of Economic Theory*, 144(4), 1440-67.