Quantal Response Equilibrium with Symmetry: Representation and Applications*

Evan Friedman† and Felix Mauersberger‡

First version: 1st March 2022
Current version: 17th November 2022

Abstract

We study an axiomatic variant of quantal response equilibrium (QRE) for normal form games that augments the regularity axioms (Goeree et al., 2005) with various forms of "symmetry" across players and actions. The model refines regular QRE, generalizes logit QRE, and is tractable in many applications. The main result is a representation theorem that characterizes the model's set-valued predictions by taking unions and intersections of simple sets. We completely characterize the predictions for (almost) all 2×2 games, a corollary of which is to show, in coordination games, which Nash equilibrium is selected by the principal branch of the logit correspondence. As applications, we consider three classic games: public goods provision with heterogenous costs of participation, jury voting with unanimity, and the infinitely repeated prisoner's dilemma. For each, we characterize all equilibria within a particular large class. An analysis of existing experiments shows the model's potential for organizing experimental data.

Keywords: quantal response; symmetry; regular QRE; representation theorem

JEL Classification: C44, C72, C92

*We thank Alessandra Casella, Jacob Goeree, Philippos Louis, Ted Turocy, and seminar participants at Essex, Indiana, the briq institute, UCL, PSE, Bar-Ilan, JiLAEE-ESA 2021, RES 2022, BRIC 2022, KERI 2022, D-TEA 2022, EC 2022, SAET 2022, the CREST/LESSAC Experimental conference 2022, the Lancaster Game Theory Conference, and the 25+ Years of QRE conference for helpful comments.

†Email: evan.friedman@psemail.eu

‡Email: felix.mauersberger@utdallas.edu

The central concept of game theory is Nash equilibrium (NE). Despite its appeal and influence, however, a large body of literature, primarily based on controlled lab experiments, has documented systematic deviations from NE predictions (Camerer [2003]). This has inspired many alternative concepts to help explain the observed data.

One such concept, Quantal Response Equilibrium (QRE) (McKelvey and Palfrey [1995] and Chen et al. [1997]), generalizes NE by allowing players to make probabilistic mistakes in best responding to others' behavior, but otherwise maintains fixed-point consistency. QRE has had considerable success in organizing experimental datasets (Goeree et al. [2016]) and is often the benchmark against which other concepts are compared.

QRE is also an important theoretical benchmark as it deviates from NE in a minimal way by injecting noise, a realistic feature of human decision-making. The only modeling consideration is how to model this noise: one must select the admissible family of noise structures. The literature has proposed a number of such families, ranging from the very precise to the very flexible. At one extreme, noise is governed by a specific parametric family, whereas on the other, there are so many degrees of freedom that the model is difficult to reject (see, e.g., Haile et al. [2008]).

In this paper, we study a QRE model that is somewhere "in between". It allows for any noise structure that satisfies a number of common axioms as well as additional axioms imposing "symmetry" across players and actions. The symmetry axioms are natural; they ought to hold whenever players are drawn from the same (possibly heterogeneous) population[1] and action labels do not convey any information conditional on payoffs. The axioms are also satisfied by the common logit QRE and many other models,[2] accounting for the large majority of applications.

Our main result is a representation theorem which characterizes the entire set of these *Symmetric QRE* in terms of sets that are themselves easy to characterize. After reviewing other approaches below, we argue that the model is especially well-suited for applications as it makes fairly precise predictions and, due to our theorem, is tractable. The model also implies novel bounds on the common parametric models nested within

[1]In Section 1.4, we show that our notion of *player symmetry* is satisfied if each player is the representative agent for a population of individuals with heterogeneous noise-levels, provided each player is drawn from the same population. This would hold, for example, under the standard experimental paradigm of randomizing subjects into player roles. Our result is related to that of Golman [2011], which derives conditions under which heterogeneous quantal responders admit a representative agent.

[2]The axioms are satisfied by any *structural* QRE model (see, for example, Goeree et al. [2005]) with errors that are i.i.d. across players and actions. This family includes logit and has been of significant theoretical interest (see, for example, Haile et al. [2008]).

it, which has both theoretical and computational implications.

The most commonly used form of QRE is *logit QRE*, or *LQRE*. In an LQRE, players make mistakes according to the logit function where a free parameter, usually called λ, governs the degree of noisy behavior. In common practice, λ is estimated in-sample to best fit the data, with the resulting prediction compared to that of other models.

LQRE is an extremely useful tool with many virtues. It is broadly applicable and makes precise predictions: for any given λ, there are a finite number of LQRE, and theory suggests a natural way of making a unique selection among them.[3] Still, LQRE has drawbacks. First, the logit functional form is somewhat arbitrary.[4] Second, as solutions to systems of non-polynomial equations, LQRE can typically only be found via numerical methods. For this reason, LQRE is not particularly amenable to theoretical analysis, with theory often focusing on the limiting case as $\lambda \to \infty$. Solving for LQRE numerically can also be computationally demanding.

In part to address these concerns about LQRE, Goeree et al. [2005] introduce *regular QRE*, or *RQRE*, in which logit response is generalized to allow for any "quantal response function" that satisfies a number of axioms. The important axioms are *monotonicity*, which requires that actions with higher payoffs are played more often, and *responsiveness*, which requires that an an all-else-equal increase in the payoff to some action means it is played even more often. Goeree et al. [2005] show that these axioms impose testable restrictions on the data, without relying on any specific functional form.[5]

An important insight, due to Goeree and Louis [2021], is that, whereas it is difficult to solve for the QRE associated with any given quantal response function, it is often easy to solve for the *set of* QRE that are associated with *some* quantal response function satisfying the axioms. Goeree and Louis [2021] introduce the set-valued *M equilibrium*,[6] a theory that nests a number of existing concepts. In particular, they show that the union of M-equilibrium *choice sets* coincides with the set of RQRE. They further show that this set is *semi-algebraic*, i.e. characterized by a finite number of polynomial (in)equalities.

[3]For almost every game, there is a "principal branch" of the LQRE correspondence that gives a unique LQRE for each value of λ. See McKelvey and Palfrey [1995] for details.

[4]Stahl [1990] and Weibull et al. [2007] derive logit response as the solution to a control costs problem with entropic costs, but this choice of cost function is itself somewhat arbitrary.

[5]Hence, RQRE is not subject to the critique of Haile et al. [2008], which shows that *structural* QRE is non-falsifiable if the errors are not restricted to be i.i.d. across players' actions.

[6]Unlike QRE, M equilibrium is explicitly set-valued. It is defined by a set of mixed action profiles–an M-equilibrium choice set–and a set of supporting beliefs–an M-equilibrium belief set–such that (1) mixed actions in the choice set exhibit rank-ordered "better response" to any beliefs in the belief set and (2) beliefs in the belief set "induce the same ordering of expected payoffs" as any mixed actions in the choice set.

Hence, when viewed as a set-valued concept, RQRE can be quite tractable. What it gains in tractability, however, it loses in precision. RQRE makes fairly weak predictions in many games, and so is rarely used in applications.

Since LQRE is intractable and RQRE can be overly flexible, we study *Symmetric QRE*, or *SQRE*, which is intermediate between the two. The model builds upon RQRE by augmenting regularity with additional axioms that embed "symmetry" into the quantal response function. In an SQRE, players effectively use the same quantal response function (*player symmetry*) that depends only on payoff differences between actions (*translation invariance*) and does not favor any action a priori (*label independence*).

We find that SQRE resolves the tradeoff well, balancing the tractability of RQRE with the precision of LQRE. For several classes of games, we completely characterize the set of SQRE, and so the model is tractable when viewed as a set-valued concept. SQRE also makes more precise predictions than RQRE, with symmetry reducing the measure of predictions in our examples by more than half. Since SQRE imposes additional restrictions on the data relative to RQRE, we find new bounds on the models nested within it. Hence, we develop tools for analyzing LQRE and other workhorse models.

We first focus on binary-action games (with any number of players), for which we derive a representation theorem. To this end, we define a player's "stakes" to be the absolute expected utility difference between her actions, and we define the "extremeness" of a player's behavior as the maximum probability with which she takes one of her actions. The theorem states that the set of SQRE equals *the intersection of* (1) the set of RQRE and (2) the set of mixed action profiles that result in the same rankings of players by stakes and extremeness. Hence, to characterize the set of SQRE, one need only characterize the set of RQRE as well as a newly defined set of outcomes, and then take their intersection. Since the concept of stakes does not require that one keep track of which action yields the highest payoff and the concept of extremeness does not depend on payoffs at all, the set of SQRE is often easy to characterize. Our results also imply that, like M equilibrium, the set of SQRE is semi-algebraic and therefore computable by a finite algorithm.

To illustrate the theorem, we completely characterize the sets of SQRE for (almost) all 2×2 games.[7] This includes generalized versions of matching pennies, coordination games, prisoner's dilemma, and games that are dominance solvable in two steps.

[7] For simplicity, we exclude the measure zero set of games with actions that are either (1) weakly–but not strongly–dominant or (2) "safe" in the sense of yielding the same payoff independent of the opponent's behavior.

Whereas each game is defined by eight payoff parameters, the set of SQRE depends only on three transformed parameters. For each such game, the set of SQRE is a union of polytopes in the unit-square. The set of LQRE must be contained within these polytopes, and we show that, wherever two polytopes intersect at a single point, such points are LQRE. This implies that we are able to find certain special LQRE explicitly as solutions to systems of linear equations.

As a bi-product of our analysis of 2×2 games, we find that SQRE can be used for equilibrium selection in coordination games.[8] Under certain conditions on the game's payoffs parameters, we show that the centroid–whereby both players' behavior is uniformly random–is path-connected via the set of SQRE to exactly one NE. Since the set of LQRE is contained within the set of SQRE, this NE is necessarily that which is selected by LQRE in the manner of McKelvey and Palfrey [1995]. Hence, we are able to show which NE is selected, and the result is robust to noise processes beyond logit.[9]

The use of boundedly rational concepts in general, and QRE in particular, is often hindered by a lack of tractability. To illustrate the potential of SQRE in applied theory, we consider three classic games; and for each, we use our representation theorem to characterize all SQRE within a particular class that generalizes a class of NE that has been previously studied.

First, we consider two N-player games from the political economy literature: a public goods game with heterogeneous costs of participation (Diekmann [1986]) and a jury voting game with unanimity (Feddersen and Pesendorfer [1998]).[10] For these games, we characterize the set of semi-symmetric SQRE whereby players of a given type have the same behavior. As a third application, we study the infinitely repeated prisoner's dilemma for which we characterize all Markov-perfect SQRE for a particular specification of states. Hence, SQRE is tractable enough to be applied to games with many players, incomplete information, or dynamic strategies.

In all applications, we find that SQRE leads to a significant refinement of RQRE and imposes significant bounds on LQRE. We also find SQRE compelling in that symmetry rules out many implausible equilibria. For example, while there are RQRE of the par-

[8] A 2×2 coordination game is any that has two pure-strategy NE and one mixed-strategy NE.

[9] The result we obtain is essentially the same as that of Zhang and Hofbauer [2016], though our method is different, following directly from the SQRE characterization. Zhang and Hofbauer [2016] builds on earlier work by Turocy [2005] based on "homotopy methods."

[10] Participation game experiments include Franzen [1995] and Goeree et al. [2015]; jury voting experiments include Guarnaschelli et al. [2000]. These studies analyze the data using LQRE based on numerical approximation.

ticipation game in which players with higher costs contribute *more* often, there are no such SQRE.[11] Hence, implicit in our results is a novel critique of more flexible concepts.

To extend the representation theorem to arbitrary normal form games, we consider a strengthening of the SQRE axioms and give necessary conditions based on the cyclic monotonicity inequalities of Melo et al. [2018]. In the case where each player has two actions, the necessary conditions are also sufficient and the result coincides with the binary-action representation theorem. Through examples, we find that symmetry continues to impose significant restrictions in games with larger action spaces.

We show that the same tools developed in this paper can be used to characterize *asymmetric QRE* (AQRE) whereby some players are *noisier than* others (in the sense of having less extreme behavior despite facing higher stakes). This is the relevant model for situations in which there is reason to expect a particular ranking of players by noise, perhaps because some players are known to have more experience than others. The joint SQRE-AQRE framework can also be used to organize data by non-parametrically identifying which, if any, players appear to be noisier in a given dataset.

As an empirical application, we consider two studies of experimental matching pennies games (Selten and Chmura [2008] and McKelvey et al. [2000]). Overall, we find that SQRE performs nearly as well as RQRE, despite making much more precise predictions. There are, however, some deviations that the SQRE-AQRE framework allows us to interpret as asymmetries in noise across players. By relating these asymmetries to features of the games, we document an interesting regularity. Several explanations are possible, but the one we favor points to a role for *complexity* considerations typically absent from QRE models. Hence, the SQRE-AQRE framework provides a lens through which to interpret experimental data, and may help to suggest directions for alternative theories.

We consider various extensions. First, we give a representation theorem for *sets of games*,[12] which gives rise to comparative static predictions across games and stronger restrictions on the data than the product of individual-game restrictions. Second, we give a representation theorem for an SQRE variant based on *scale invariance* as opposed to *translation invariance*. This leads to a different refinement of RQRE and nests models such as the generalized Luce QRE (Luce [1959]), which is based on multiplicative errors. We derive similar lessons for the Luce model as we do for LQRE and provide intuition

[11] We also find that symmetry rules out RQRE of the infinitely repeated prisoner's dilemma in which players cooperate more often following previous defection.

[12] Melo et al. [2018] studies structural QRE in sets of games using results from convex analysis.

for why Luce QRE and LQRE give rise to different predictions.

The rest of the paper is organized as follows. Section 1 reviews the forms of QRE we consider and introduces the SQRE axioms, Section 2 gives the representation theorem for binary-action games, Section 3 characterizes the SQRE of 2×2 games, Section 4 explores applications, Section 5 shows how to characterize AQRE, Section 6 extends results to games with larger action spaces, Section 7 tests SQRE using data from existing experiments, and Section 8 discusses extensions and broader implications of our work.

1 Quantal response equilibrium

We introduce the QRE concepts we consider, discuss their containment relationship, define a general notion of QRE-based NE selection, and demonstrate how symmetry imposes restrictions.

1.1 General QRE for normal form games

QRE is defined for any finite normal form game $\Gamma = \{\mathcal{N}, A, u\}$ where $\mathcal{N} = \{1, ..., N\}$ is the set of players, $A = A_1 \times ... \times A_N$ is the action space, and $u : A \to \mathbb{R}^N$ is the payoff function. Player i's set of pure actions, A_i, has $J(i)$ elements and is enumerated as $\{a_{i1}, ..., a_{iJ(i)}\}$. Let Δ_i be the set of probability measures on A_i and let $\Delta = \Delta_1 \times ... \times \Delta_N$ be the set of independent probability measures on A with $\sigma = (\sigma_1, ..., \sigma_N)$ an arbitrary element. For simplicity, let $\sigma_{ij} = \sigma_i(a_{ij})$. Using expected utility, extend payoffs to the probability domain via $u_i(\sigma) = \sum_{a \in A} \sigma(a) u_i(a)$.

In addition to these standard objects, we follow the QRE literature in defining $\bar{u}_{ij}(\sigma_{-i}) = u_i(a_{ij}, \sigma_{-i})$ and $\bar{u}_i(\sigma_{-i}) = (\bar{u}_{i1}(\sigma_{-i}), ..., \bar{u}_{iJ(i)}(\sigma_{-i})) \in \mathbb{R}^{J(i)}$ for every $\sigma_{-i} \in \times_{j \neq i} \Delta_j$. These are the expected payoffs to each action given opponents' behavior.

Given an arbitrary vector of expected payoffs $\bar{u}_i = (\bar{u}_{i1}, ..., \bar{u}_{iJ(i)}) \in \mathbb{R}^{J(i)}$ (i.e. $\bar{u}_i = \bar{u}_i(\sigma_{-i})$ for some σ_{-i}), player i makes probabilistic mistakes in best responding according to a *quantal response function* $Q_i : \mathbb{R}^{J(i)} \to \Delta$ where $Q_{ij}(\bar{u}_i)$ is the probability with which player i takes action j. The quantal response function $Q = (Q_1, ..., Q_N)$ is the primitive of QRE and so must be specified exogenously. For a given Q, a QRE is defined as any fixed point of the composite function $Q \circ \bar{u}$, whereby players' behavior is consistent with their quantal response functions.

Definition 1. Fix $\{\Gamma, Q\}$. A QRE is any $\sigma \in \Delta$ such that $Q \circ \bar{u}(\sigma) = \sigma$ (for all $i \in 1, ..., N$ and all $j \in 1, ..., J(i)$, $\sigma_{ij} = Q_{ij}(\bar{u}_i(\sigma_{-i}))$).

The above model is completely vacuous: only by imposing restrictions on Q does the model place any testable restrictions on the data. Next, we discuss the three families of QRE models we consider, each of which implies (different) testable restrictions.

1.2 LQRE

The form of QRE that is by far most commonly applied to experimental data is logit QRE or *LQRE*. Here, the quantal response function takes the familiar logit form:

$$Q_{ij}(\bar{u}_i; \lambda) = \frac{e^{\lambda \bar{u}_{ij}}}{\sum_{k=1}^{J(i)} e^{\lambda \bar{u}_{ik}}}, \tag{1}$$

where the parameter $\lambda \in [0, \infty)$ indexes how much noise is embedded in the quantal response function.[13]

1.3 RQRE

Generalizing LQRE, Goeree et al. [2005] introduce regular QRE or *RQRE*. This "reduced-form" model is defined by any quantal response function satisfying the so-called regularity axioms:

(R1) **Interiority**: $Q_{ij}(\bar{u}_i) \in (0, 1)$ for all $j \in 1, ..., J(i)$ and for all $\bar{u}_i \in \mathbb{R}^{J(i)}$.

(R2) **Continuity**: $Q_{ij}(\bar{u}_i)$ is a continuous and differentiable function for all $\bar{u}_i \in \mathbb{R}^{J(i)}$.

(R3) **Responsiveness**: $\frac{\partial Q_{ij}(\bar{u}_i)}{\partial \bar{u}_{ij}} > 0$ for all $j \in 1, ..., J(i)$ and for all $\bar{u}_i \in \mathbb{R}^{J(i)}$.

(R4) **Monotonicity**: $\bar{u}_{ij} > \bar{u}_{ik} \implies Q_{ij}(\bar{u}_i) > Q_{ik}(\bar{u}_i)$ for all $j, k \in 1, ..., J(i)$.

(R1) and (R2) are non-falsifiable technical axioms. Taken together, (R3) and (R4) are a stochastic generalization of "best response", requiring than an all-else-equal increase in the payoff to an action increases the probability it is played and that, given any belief, higher payoff actions are taken more often in a rank-order sense.

1.4 SQRE

In this paper, we introduce a model called symmetric QRE or *SQRE*. This model builds on RQRE by augmenting (R1)-(R4) with three additional axioms that capture various notions of symmetry:

[13] The case that $\lambda = 0$ involves uniformly mixing over all actions independently of payoffs, and, as $\lambda \to \infty$, the probability of taking (one of) the highest payoff action(s) approaches one.

(S1) **Player symmetry**: If $J(i) = J(j) = J$, there exists a bijection $\iota : \{1, ..., J\} \to \{1, ..., J\}$ such that if $\bar{u}_{ik} = \bar{u}_{j\iota(k)}$ for all k, then $Q_{ik}(\bar{u}_i) = Q_{j\iota(k)}(\bar{u}_j)$ for all k.

(S2) **Translation invariance**: $Q_i(\bar{u}_i + \gamma e_{J(i)}) = Q_i(\bar{u}_i)$ for any $\bar{u}_i \in \mathbb{R}^{J(i)}$ and $\gamma \in \mathbb{R}$, where $e_{J(i)} = (1, ..., 1)$.

(S3) **Label independence**: $Q_{ij}(\bar{u}_i) = Q_{ik}(\bar{u}_i')$ if $\bar{u}_{ij} = \bar{u}_{ik}'$ and $\bar{u}_{il} = \bar{u}_{i\iota(l)}'$ for all l, where $\iota : \{1, ..., J(i)\} \to \{1, ..., J(i)\}$ is a bijection.

Player symmetry (S1) ensures that if players have the same number of available actions, they have the same quantal response function relative to some ordering of their respective actions. *Translation invariance* (S2) implies that quantal response only depends on differences in payoffs between actions. *Label independence* (S3) implies that quantal response treats all actions symmetrically in the sense that only payoffs matter–not the actions' labels. We make a few remarks.

Remark 1. All *structural QRE* models (see, for example, Goeree et al. [2005]) with i.i.d. errors satisfy (R1)-(R4) and (S1)-(S3); and so SQRE generalizes this important class.

Remark 2. *Player symmetry* (S1) does not preclude *heterogeneity*; it holds whenever players are drawn from the same (possibly heterogeneous) population. Consider the partition $\mathcal{P}$ of players $\{1, ..., N\}$ such that within each element $P \in \mathcal{P}$, players have the same number of actions ($J(i) = J(j)$ for all $i, j \in P$). We suppose that each player $i \in \{1, ..., N\}$ is drawn from the same population of individuals. Individual s has quantal response function Q_i^s when in the role of player i, which satisfies (R1)-(R4), (S2), and (S3). Letting F^s denote the distribution of individuals, player i admits a representative agent with quantal response function $Q_i = \int Q_i^s dF^s$. If $Q_i^s = Q_j^s$ whenever $i, j \in P$ for some P, then we have that $Q = (Q_i)_{i \in \{1,...,N\}}$ satisfies (R1)-(R4) and (S1)-(S3).[14]

1.5 Nested sets of QRE

In this paper, we view each model as giving set-valued predictions–the set of QRE that can be supported for some quantal response function within the restricted class, be it symmetric, regular, or logit.

[14]For example, the heterogenous LQRE model of McKelvey et al. [2000] in which each player is drawn from a population of logit responders with heterogeneous noise levels, satisfies the SQRE axioms.

Definition 2. The sets of SQRE, RQRE, and LQRE, respectively, are defined as

$$S := \{\sigma \in \Delta : Q \circ \bar{u}(\sigma) = \sigma \text{ for some } Q \text{ satisfying (R1)-(R4) and (S1)-(S3)}\},$$
$$R := \{\sigma \in \Delta : Q \circ \bar{u}(\sigma) = \sigma \text{ for some } Q \text{ satisfying (R1)-(R4)}\}, \text{ and}$$
$$L := \{\sigma \in \Delta : Q(\bar{u}(\sigma); \lambda) = \sigma \text{ for some } \lambda \in [0, \infty) \text{ where } Q \text{ is as in (1)}\}.$$

We also say that, if $\sigma \in S$, then σ is an SQRE, and similarly for RQRE and LQRE.

It is easy to show that LQRE satisfies all of the previously mentioned axioms; and, by definition, SQRE satisfies the RQRE axioms (and more). This immediately implies the containment relationship $L \subset S \subset R$. The first containment is almost always strict. As we will show through examples, the second containment can also be strict, with the set of SQRE having a much smaller measure than the set of RQRE.

1.6 Nash equilibrium selection

In generic games, LQRE provides a theory of NE selection (McKelvey and Palfrey [1995]). For sufficiently low λ, there is a unique LQRE that is close to the centroid $\sigma^{centroid} = \times_i(\frac{1}{J(i)}, ..., \frac{1}{J(i)}) \in \Delta$ whereby all players are uniformly mixing. This unique LQRE is associated with the "principal branch" of the LQRE correspondence. By increasing λ, the principal branch can be "traced", limiting to a unique NE, the so-called *logit selection*. More generally, we say that a given QRE model selects a given NE if there is a path of QRE within the relevant family that connects it to the centroid.

Definition 3. A QRE model (SQRE, RQRE, or LQRE) *selects* NE $\sigma^* \in \Delta$ if σ^* is path-connected to the centroid $\sigma^{centroid} \in \Delta$ via its equilibrium set (S, R, or L, respectively). Moreover, if σ^* is the only such NE, then the QRE model *uniquely selects* σ^*.

Since $L \subset S$, the following result is immediate.

Lemma 1. *If SQRE uniquely selects NE σ^*, then LQRE uniquely selects σ^* (if it makes any selection at all).*[15]

The value of this result is that, as we show through examples, it is sometimes easy to find the NE uniquely selected by SQRE. In such cases, the selection is robust to many underlying noise processes, including logit whenever it is defined.

[15]Since LQRE does make a unique selection in generic games, we conjecture the stronger claim that drops the part in parenthesis.

1.7 How symmetry imposes restrictions: an example

A simple example shows the role of each axiom and gives much of the intuition behind our general results. Abstracting from the underlying game, consider the following dataset from a 2×2 game. Player 1 takes her first action with probability $\hat{p} \in (\frac{1}{2}, 1)$ and player 2 takes her first action with probability $\hat{q} \in (\frac{1}{2}, \hat{p})$. Hence, both players take their first action more often than not, and player 1 takes her first action more often than does player 2. These probabilities pin down both players' expected payoff vectors, which come out to $\bar{u}_1 = (3, 1)$ and $\bar{u}_2 = (5, 2)$. The question then is whether $(\hat{q}, \hat{p})$ is consistent with RQRE and SQRE.

The data *is* consistent with RQRE: $(\hat{q}, \hat{p}) \in R$. This is immediate after observing that only *monotonicity* (R3) imposes substantive restrictions on R.[16] This requires that each player take her action with the higher expected payoff more often than not, which holds in the example. We note that *responsiveness* (R4), which requires that an all-else-equal increase in the payoff to some action increases the probability that that action is played, imposes no restriction on R. This is because the axiom concerns *changes* in payoffs, and we only observe a single point on each player's quantal response function.[17]

More interestingly, the data is *not* consistent with SQRE: $(\hat{q}, \hat{p}) \notin S$. Whereas *responsiveness* does not impose any restriction under RQRE, the SQRE axioms (S1)-(S3) and *responsiveness* jointly imply a restriction *across players* which is not satisfied in this example. Since both players use the same quantal response function by *player symmetry* (S1), observing both players' data gives two points on their common quantal response function; and so *responsiveness* may imply a restriction. However, unless the two players' payoff vectors are ordered by an all-else-equal increase in the payoff to some action, *responsiveness* imposes no restriction: the "responsiveness order" is incomplete. *Responsiveness* on its own has no bite in the example since $\bar{u}_1$ features higher expected payoffs than $\bar{u}_2$ in *both* components. Luckily, *translation invariance* (S2) and *label independence* (S3) complete the order. Jointly, (S1)-(S3) and *responsiveness* imply that if a player has a higher expected payoff difference between her two actions than does her opponent, she should take her action that yields the higher payoff more often than does her opponent. In the example, the expected utility difference for player 1 is $3 - 1 = 2$, and the expected

[16] This point is made by Goeree and Louis [2021] who show that R equals the union of M-equilibrium choice sets. Goeree et al. [2018] make a similar point in the context of their rank-dependent choice equilibrium.

[17] *Responsiveness* does impose restrictions on the shape of quantal response functions, but since R is defined by taking the union over quantal response functions within the relevant class, the axiom has no bearing on R.

utility difference for player 2 is $5 - 2 = 3$. Hence, the axioms imply that $\hat{q} > \hat{p}$, which is violated in the data.

2 Binary-action games: a representation theorem

We begin by giving a representation theorem for N-player binary-action games, those in which each player has exactly two actions, i.e. $J(i) = 2$ for all i. This is an important class as it contains all 2×2 games as well as many games of theoretical interest (e.g. voting games). We refer to an arbitrary such game by $\Gamma^{N\times 2}$.

Relative to RQRE, the additional restrictions of SQRE depend on being able to make comparisons across players, as suggested by the example of Section 1.7. Hence, we introduce two player orders.

First, we say that a player *faces higher stakes than* another if the *absolute* expected utility difference between her actions is greater. Formally, defining player i's *stakes* in taking an action as $\delta(\bar{u}_i) \equiv |\bar{u}_{i1} - \bar{u}_{i2}|$, we have the following order:

Definition 4. Suppose players i and j have two actions each. Player i *faces higher stakes than* player j if $\delta(\bar{u}_i) \geq \delta(\bar{u}_j)$ (and *faces strictly higher stakes* if $\delta(\bar{u}_i) > \delta(\bar{u}_j)$).

Second, we say that a player is *more extreme than* another if she takes *some* action with a higher probability than the maximum probability with which the other player takes an action. Unlike the stakes order, this order does not depend at all on payoffs.

Definition 5. Suppose players i and j have two actions each. Player i is *more extreme than* player j if $max(\sigma_{i1}, 1 - \sigma_{i1}) \geq max(\sigma_{j1}, 1 - \sigma_{j1})$ (and *strictly more extreme* if $max(\sigma_{i1}, 1 - \sigma_{i1}) > max(\sigma_{j1}, 1 - \sigma_{j1})$).

Intuitively, an SQRE is an RQRE with additional restrictions across players. Specifically, each player has some action which, for her, yields a higher payoff; and an SQRE is an RQRE in which players who *face higher stakes* are more likely to take their higher-payoff actions. In any RQRE, since each player takes her higher-payoff action more often than not, a player is more likely than another to take her higher-payoff action if and only if she is *more extreme*. Combining these observations gives our representation theorem: "an action profile is an SQRE if and only if it is an RQRE in which players are ranked in the same way by stakes and extremeness". Formally, we define by X the set of action profiles such that the two rankings coincide:

$$X = \{\sigma \in \Delta : \delta(\bar{u}_i) > \ \delta(\bar{u}_j) \iff max(\sigma_{i1}, 1 - \sigma_{i1}) > \ max(\sigma_{j1}, 1 - \sigma_{j1})\}.$$

The set of SQRE is precisely this set *intersected with* the set of RQRE.

Theorem 1. *Fix $\Gamma^{N\times 2}$. $S = X \cap R$.*

Proof. Suppose $\sigma \in S$, i.e. that σ is a fixed point of $Q \circ \bar{u}$ where Q satisfies (R1)-(R4) and (S1)-(S3). By definition of R, $\sigma \in R$, so we need only check that $\sigma \in X$. Suppose not. Then there exists two players i and j such that i faces higher stakes than, but is less extreme than, j: $\delta(\bar{u}_i) \geq \delta(\bar{u}_j)$ and $max\{\sigma_{i1}, 1-\sigma_{i1}\} \leq max\{\sigma_{j1}, 1-\sigma_{j1}\}$ with at least one of these strict. But this cannot be because of the symmetry axioms (S1)-(S3) and *responsiveness* (R3) and *monotonicity* (R4).

Conversely, suppose $\sigma \in X \cap R$. We must construct quantal response function Q satisfying (R1)-(R4) and (S1)-(S3) such that σ is a fixed point of $Q \circ \bar{u}$. Since $\sigma \in X$, we have that $\delta(\bar{u}_{i_1}) \geq \delta(\bar{u}_{i_2}) \geq ... \geq \delta(\bar{u}_{i_N})$ and $max(\sigma_{i_1 1}, 1-\sigma_{i_1 1}) \geq max(\sigma_{i_2 1}, 1-\sigma_{i_2 1}) \geq ... \geq max(\sigma_{i_N 1}, 1-\sigma_{i_N 1})$ for some ordering of players $i_1, ..., i_n$ where $\delta(\bar{u}_{i_k}) > (=)\ \delta(\bar{u}_{i_{k+1}})$ if and only if $max(\sigma_{i_k 1}, 1-\sigma_{i_k 1}) > (=)\ max(\sigma_{i_{k+1} 1}, 1-\sigma_{i_{k+1} 1})$. Further, since $\sigma \in R$, we have that $max(\sigma_{i_N 1}, 1-\sigma_{i_N 1}) \geq \frac{1}{2}$. Set "pre-quantal response" function $\tilde{Q} : [0,\infty) \to [\frac{1}{2}, 1)$ to be any that is strictly increasing, differentiable, and satisfies $\tilde{Q}(0) = \frac{1}{2}$ and $\tilde{Q}(\delta(\bar{u}_{i_k})) = max(\sigma_{i_k 1}, 1-\sigma_{i_k 1})$ for all $k = 1, ..., N$. Because $\delta(\bar{u}_{i_1}) \geq \delta(\bar{u}_{i_2}) \geq ... \geq \delta(\bar{u}_{i_N})$ and $max(\sigma_{i_1 1}, 1-\sigma_{i_1 1}) \geq max(\sigma_{i_2 1}, 1-\sigma_{i_2 1}) \geq ... \geq max(\sigma_{i_N 1}, 1-\sigma_{i_N 1}) \geq \frac{1}{2}$, such a function exists. Define $Q_{i1} : \mathbb{R} \to (0,1)$ from $\tilde{Q}$ as $Q_{i1}(\bar{u}_{i1}, \bar{u}_{i2}) = \tilde{Q}(\bar{u}_{i1} - \bar{u}_{i2})\mathbf{1}_{\{\bar{u}_{i1} \geq \bar{u}_{i2}\}} + \Big(1 - \tilde{Q}(\bar{u}_{i2} - \bar{u}_{i1})\Big)\mathbf{1}_{\{\bar{u}_{i1} < \bar{u}_{i2}\}}$. Set player i's quantal response function as $Q_i = (Q_{i1}, 1-Q_{i1}) : \mathbb{R}^2 \to \Delta$. By construction, if the quantal response function is given by $Q = (Q_1, ..., Q_N)$, then σ is a fixed point of $Q \circ \bar{u}$. Finally, we must check that Q satisfies (R1)-(R4) and (S1)-(S3), but this is immediate from the construction. □

Hence, to characterize S, we need only characterize X and R independently, and then take their intersection. To characterize R, we rely on the results of Goeree and Louis [2021]. They show for arbitrary games that the union of M-equilibrium choice sets, which coincides with R, is *semi-algebraic.* In other words, R is characterized by finite polynomial (in)equalities and therefore computable by a finite algorithm:

Theorem (Goeree and Louis 2021)**.** [18]

$$R = \{\sigma \in \Delta : \bar{u}_{ij}(\sigma_{-i}) < \bar{u}_{ik}(\sigma_{-i}) \iff \sigma_{ij} < \sigma_{ik}\}.$$

[18] This is based on Proposition 3 of Goeree and Louis [2021], which establishes the relationship of M equilibrium to other concepts.

The new set X is also defined by (new) polynomial (in)equalities, and so S is itself semi-algebraic.[19] Hence, just like M equilibrium, S can be characterized without resorting to fixed-point calculations.

In addition to pointing out the semi-algebraic nature of S, we study the geometry of X. In each of our examples, we characterize it explicitly as the union and intersection of simple sets.

3 2×2 Games

We apply Theorem 1 to characterize the SQRE of (almost) all 2×2 games. Such games are widely used in experiments as they are the simplest that can capture a multitude of strategic features.

A general 2×2 game is given in Table 1. Parameters a_L, a_R, b_U, and b_D are the *base payoffs*, and parameters c_L, c_R, d_U, and d_D are the *payoff differences*. We specialize notation by using p for the probability player 1 chooses U, q for the probability player 2 chooses L, and $Q = (Q_U, Q_L)$ for quantal response.

	L [q]	**R**
U [p]	$a_L + c_L$, b_U	a_R, $b_U + d_U$
D	a_L, $b_D + d_D$	$a_R + c_R$, b_D

Table 1: *General* 2×2 *game.* Parameters a_L, a_R, b_U, and b_D are the base payoffs, and c_L, c_R, d_U and d_D are the payoff differences.

Having two players with two actions each allows us to rewrite the set X in a more convenient form, which we use to give a specialized version of Theorem 1. We first define the sets of action profiles such that player 1 faces strictly higher stakes, player 2 faces strictly higher stakes, or both players face the same stakes:

$$\begin{aligned} U_1 &= \{(q, p) \in \square : \delta(\bar{u}_1(q)) > \delta(\bar{u}_2(p))\}, \\ U_2 &= \{(q, p) \in \square : \delta(\bar{u}_1(q)) < \delta(\bar{u}_2(p))\}, \text{ and} \\ U_0 &= \{(q, p) \in \square : \delta(\bar{u}_1(q)) = \delta(\bar{u}_2(p))\}, \end{aligned} \tag{2}$$

[19] We thank Jacob Goeree and Philippos Louis for this insight.

where for convenience we have defined $\square = [0,1]^2$ to be the unit square.

Next, we define the sets of action profiles such that player 1 is strictly more extreme, player 2 is strictly more extreme, or both players are equally extreme:

$$\begin{aligned} E_1 &= \{(q,p) \in \square : max(p, 1-p) > max(q, 1-q)\}, \\ E_2 &= \{(q,p) \in \square : max(p, 1-p) < max(q, 1-q)\}, \text{ and} \\ E_0 &= \{(q,p) \in \square : max(p, 1-p) = max(q, 1-q)\}. \end{aligned} \tag{3}$$

Unlike U_1, U_2, and U_0, which depend on the game's payoff parameters, E_1, E_2, and E_0 are completely invariant to payoffs and therefore the same for all games. It follows that, for 2×2 games, $X = (U_1 \cap E_1) \sqcup (U_2 \cap E_2) \sqcup (U_0 \cap E_0)$, where $A \sqcup B$ means the union of A and B and that A and B are disjoint.

For completeness, we restate the definition for the set of RQRE in 2×2 games: $R = \{(q,p) : Q_U(q) = p, Q_L(p) = q$ for some Q satisfying (R1)-(R4)$\}$. The SQRE representation for 2×2 games follows from Theorem 1 after substitution.

Corollary 1. *Fix* $\Gamma^{2\times 2}$. $S = (U_1 \cap E_1) \sqcup (U_2 \cap E_2) \sqcup (U_0 \cap E_0) \cap R.$

3.1 Analysis of four classes of games

We completely characterize SQRE for almost all 2×2 games. To this end, we consider the four classes of games in Table 2, which are defined by the signs of the payoff differences c_L, c_R, d_U, and d_D. The only games not represented are those non-generic games with either weakly (but not strongly) dominant actions or duplicate actions, precisely the games for which one or more payoff differences equal zero.[20]

These four classes of games differ in terms of their number and type of NE, as shown in the table. Referring to fully mixed NE as "MSNE" and pure strategy NE as "PSNE", the game can either have 1 MSNE (matching pennies); 1 MSNE and 2 PSNE (coordination); or 1 PSNE (prisoner's dilemma or two-step dominance solvable). Despite these differences, the SQRE of all such games are found in a similar fashion. Recalling that $S = X \cap R$, we first give a procedure for finding X that works in any game class.

[20] In order for some games to be classified as in Table 2, one must relabel actions and/or players. For example, a game in which $c_L < 0$, $c_R < 0$ and $d_U > 0$, $d_D > 0$ is a coordination game up to labelling.

[21] If ($q^* < 0$ and $p^* > 1$) or ($q^* > 1$ and $p^* < 0$), then $r > 0$. Otherwise, $r < 0$.

[22] If $q^* < 0$, then $r > 0$. Otherwise, $r < 0$.

Game	c_L	c_R	d_U	d_D	NE	q^*	p^*	r
Matching pennies	$+$	$+$	$+$	$+$	1 MSNE	$\in (0,1)$	$\in (0,1)$	$\in (0,\infty)$
Coordination	$+$	$+$	$-$	$-$	1 MSNE, 2 PSNE	$\in (0,1)$	$\in (0,1)$	$\in (-\infty,0)$
Prisoner's dilemma	$+$	$-$	$-$	$+$	1 PSNE	$\notin [0,1]$	$\notin [0,1]$	$\in (-\infty,\infty)$, $sign(r)$ depends on (q^*,p^*)[21]
Two-step dom. solvable	$+$	$-$	$+$	$+$	1 PSNE	$\notin [0,1]$	$\in (0,1)$	$\in (-\infty,\infty)$, $sign(r)$ depends on q^*[22]

Table 2: *Four classes of* 2×2 *games.* This table gives the four classes of 2×2 games that almost all such games can be classified into. They are defined by signs of the payoff difference parameters c_L, c_R, d_U, and d_D. We use "+" and "–" to indicate strictly positive and strictly negative values, respectively. The table also gives the number and type of NE and the values of the transformed parameters q^*, p^*, and r that define each game class.

3.1.1 Finding X

Our first step is to reparameterize the game by specifying four transformed parameters:

$$q^* = \frac{c_R}{c_L + c_R}, \; p^* = \frac{d_D}{d_U + d_D}, \; r = \frac{c_L + c_R}{d_U + d_D}, \text{ and } s = c_L + c_R.$$

If the game has an MSNE, as in the first two classes of games, q^* and p^* coincide with the NE mixing probabilities for players 2 and 1, respectively. If the game does not have an MSNE, as in the second two classes of games, q^* and p^* are still the relevant parameters for our analysis even though they do not necessarily belong to the $(0,1)$ interval. Parameter r is the *ratio* of the two players' sums of payoff differences, and s is a *scale* factor given here as the sum of player 1's payoff differences. Because SQRE is invariant to translation, only the payoff differences are relevant for S. From q^*, p^*, and r, all the payoff differences can be recovered up to a positive scaling, which is pinned down by s. In this paper, no theoretical results depend on the scale of the game, so we ignore s entirely. The last three columns of Table 2 give the ranges of transformed parameter values that uniquely identify each game class.

For simplicity, we assume that $c_L \neq -c_R$ and $d_U \neq -d_D$, which means that no player has a perfectly safe action yielding the same payoff independent of the opponent's behavior. This assumption is precisely that which makes the transformed parameters well-defined. It holds automatically for games within the first two classes and for almost all games within the second two classes.

We first characterize the sets U_1, U_2, and U_0 (2) whereby one player or the other faces higher stakes. The region whereby players face the same stakes is given by $U_0 = U_0^+ \cup U_0^- \cap \square$ where

$$U_0^+ = \{(q,p) \in \mathbb{R}^2 : p = rq + (p^* - rq^*)\} \text{ and}$$
$$U_0^- = \{(q,p) \in \mathbb{R}^2 : p = -rq + (p^* + rq^*)\}$$

are two lines in $\mathbb{R}^2$.[23] The lines U_0^+ and U_0^- intersect at (q^*, p^*) and have slopes r and $-r$, respectively. Hence, these lines divide the $\mathbb{R}^2$-plane into four regions, as shown in the top-left panel of Figure 1 for an example in which (q^*, p^*) is in the interior of the unit square. Importantly, however, depending on the game, (q^*, p^*) may fall outside of the unit square (see Table 2), and hence the unit square may intersect anywhere from one to four of these regions. Since $(q^*, p^*) = U_0^+ \cap U_0^-$, it is necessarily the case that $\delta(\bar{u}_1(q^*)) = \delta(\bar{u}_2(p^*))$ (though in some cases this requires we extend $\bar{u}_i$ to be defined over all of $\mathbb{R}$). In fact, it must be that $\delta(\bar{u}_1(q^*)) = \delta(\bar{u}_2(p^*)) = 0$. This is well-known for the first two game classes as (q^*, p^*) gives the MSNE in which each player mixes so as to keep her opponent indifferent. This is also true for the second two game classes in which (q^*, p^*) falls outside of the unit square and does not represent an MSNE.

The lines U_0^+ and U_0^- serve as boundaries of U_1 and U_2 in which one player faces strictly higher stakes than the other. Of the four regions defined by U_0^+ and U_0^-, it is easy to show that the regions west and east of (q^*, p^*) involve player 1 facing higher stakes, and hence their union (intersected with $\square$) is U_1.[24] Similarly, the regions north and south of (q^*, p^*) involve player 2 facing higher stakes, and hence their union (intersected

[23] To derive U_0^+ and U_0^-, write out each players' stakes: $\delta_1(q) = |c_L q - c_R(1-q)|$ and $\delta_2(p) = |d_U p - d_D(1-p)|$. It is easy to show that $\delta_1(q) = \delta_2(p)$ if and only if $c_L q - c_R(1-q) = d_U p - d_D(1-p)$ or $c_L q - c_R(1-q) = d_D(1-p) - d_U p$. Solving each equation for p as a function of q and then substituting p^*, q^*, and r gives U_0^+ and U_0^-.

[24] Since $U_0^+ \cup U_0^-$ is the set of points whereby players face the same stakes, and stakes are continuous in (q,p), there must be exactly one player who faces strictly higher stakes within each of the four regions. Therefore, to determine which player faces the higher stakes within a given region, it suffices to check a single point. Since players both face zero stakes at (q^*, p^*), player 1 necessarily faces strictly higher stakes at $(q^* + \epsilon, p^*)$ for all $\epsilon \neq 0$, and thus player 1 faces higher stakes in the west and east regions. Similarly, it can be shown that player 2 faces higher stakes in the north and south regions. We note that, depending on the game, r may be positive or negative, but U_1 and U_2 are still found in the same way: regardless of which of U_0^+ and U_0^- slopes upward, the west-east region gives U_1 and the north-south region gives U_2.

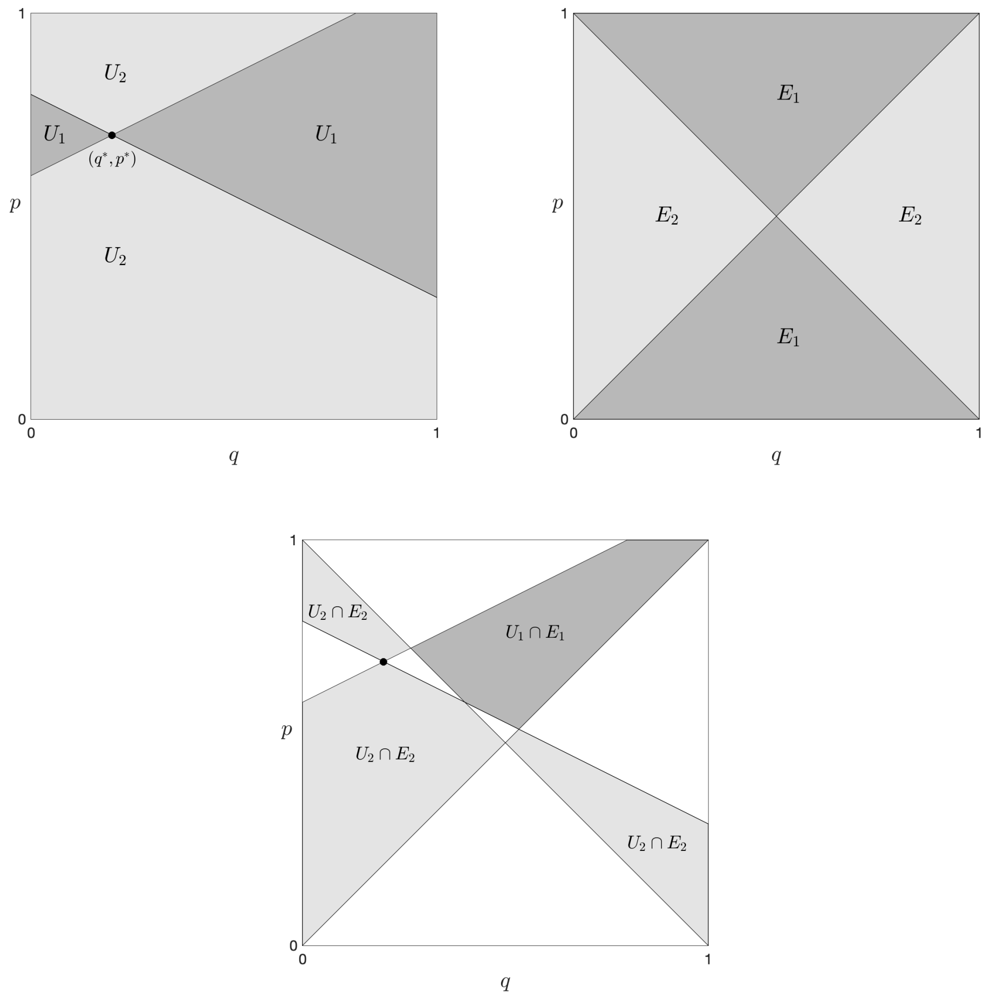

Figure 1: *Stakes- and extremeness-rankings in* 2×2 *games.* For a 2×2 game with $(q^*, p^*) \in (0,1)^2$ (i.e. either matching pennies or coordination), this figure plots the components of X. The top-left panel gives U_1, U_2, and U_0 whereby one player or the other faces higher stakes. The top-right panel gives E_1, E_2, and E_0 whereby one player or the other is more extreme. The bottom panel gives the set X whereby players have the same ranking by stakes and extremeness.

with $\square$) is U_2. Formally, recasting U_0^+ and U_0^- as functions of q,[25] we have

$$U_1 = \{(q,p) \in \mathbb{R}^2 | p < max\{U_0^-(q), U_0^+(q)\} \text{ and } p > min\{U_0^-(q), U_0^+(q)\}\} \cap \square \text{ and}$$
$$U_2 = \{(q,p) \in \mathbb{R}^2 | p < min\{U_0^-(q), U_0^+(q)\} \ \text{ or } \ p > max\{U_0^-(q), U_0^+(q)\}\} \cap \square.$$

These regions are plotted in the top-left panel of Figure 1 for the same example in which (q^*, p^*) is in the interior of the unit square.

The top-right panel of Figure 1 gives the sets E_1, E_2, and E_0 (3) whereby one player or the other has more extreme behavior. These sets do not depend on the game's payoff parameters, meaning they are the same for all games. They also have a simple geometrical description. E_0 is given by the diagonals of the unit square that pass through the centroid $(\frac{1}{2}, \frac{1}{2})$ and have slopes 1 and -1. These intersect, resulting in four regions. E_1 corresponds to the union of the north and south regions, and E_2 corresponds to the union of the west and east regions.

In the bottom panel of Figure 1, we plot the set X for the example game. The set, which looks like an "X", is constructed by intersecting U_1 with E_1 (dark gray regions) and U_2 with E_2 (light gray regions), and then taking the union. Since $S \subset X$, the size of S is bounded by the size of X. In this example, X has a measure of less than $\frac{1}{2}$. More generally, we find a uniform bound of $\frac{5}{8}$.

Proposition 1. *For any $\Gamma^{2\times2}$, the measure of X is no more than $\frac{5}{8}$.*

Proof. See Appendix 9.1. $\square$

Since $S = X \cap R$, additional restrictions on S come from regularity, and what remains in the characterization of S is to characterize R. Luckily, the qualitative shape of R depends only on which of the four classes the game belongs to. In the next sections, we give R and explore features of S for each game class.

3.1.2 Matching pennies

Matching pennies is any 2×2 game whose only NE is the MSNE $(q^*, p^*) \in (0,1)^2$. Such games have been particularly prominent in experimental economics due to their simplicity. They are also the prototypical game for illustrating how QRE deviates from NE. In the case that $q^* = p^* = \frac{1}{2}$, the unique RQRE (and thus SQRE) coincides with the NE. Otherwise, we impose $q^* \leq \frac{1}{2}$ and $p^* > \frac{1}{2}$, which is without loss.

[25] That is, define $U_0^+ : \mathbb{R} \to \mathbb{R}$ by $U_0^+(q) = \{p | (q,p) \in U_0^+\} = rq + (p^* - rq^*)$ and $U_0^- : \mathbb{R} \to \mathbb{R}$ by $U_0^-(q) = \{p | (q,p) \in U_0^-\} = -rq + (p^* + rq^*)$.

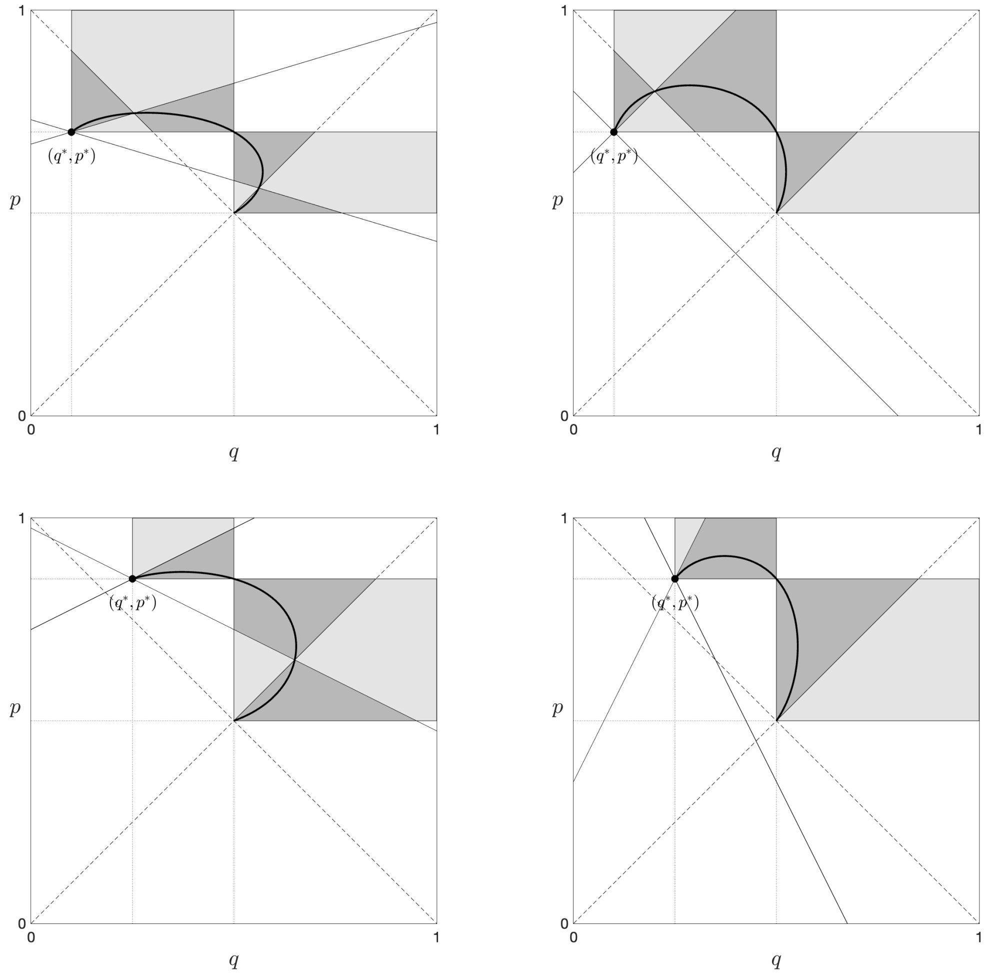

Figure 2: *Matching pennies.*

The set of RQRE is not closed,[26] and so we give the *closure* of R, which is easier to express. In an abuse of notation, we refer to this set also as R. In matching pennies, it is well-known that this set is given by $R = R^1 \cup R^2$ where $R^1 = [q^*, \frac{1}{2}] \times [p^*, 1]$ and $R^2 = [\frac{1}{2}, 1] \times [\frac{1}{2}, p^*]$ are two rectangular components.[27]

Figure 2 plots R, S, and L for four examples of matching pennies games. R is based on the above expression, S is based on Corollary 1, and L is based on numerical approximation. In the figure, U_0 is given by solid lines, E_0 is given by dashed lines, S is given by the dark gray region, $R \backslash S$ is given by the light gray region, and L is given

[26] The set of RQRE is not closed because the composite function $Q \circ \bar{u}$ of which RQRE is a fixed point must be interior and cannot have any jumps by *interiority* (R1) and *continuity* (R2).

[27] If $p^* = \frac{1}{2}$, the second component of R is degenerate: $R^2 = \emptyset$.

by the black curve. The four examples in the figure correspond to the four qualitatively distinct shapes of S. Here, S is the union of up to four polytopes, and each case is defined by how many polytopes (one or two) are contained in each component of R. In all cases, S has much smaller measure than R; and since L is contained in S, this implies significantly tighter bounds on L than implied by regularity alone.

While all LQRE are solutions to nonlinear systems that can only be solved numerically, we are able to find some special LQRE in closed form. For any matching pennies game, one polytope contains the centroid $(\frac{1}{2}, \frac{1}{2})$, one polytope contains the MSNE (q^*, p^*), and all polytopes are connected, with some polytopes intersecting at single points. Since L must be contained within these polytopes and connect the centroid to the MSNE, this implies that L must pass through all points where two polytopes intersect. Since the boundaries of these polytopes are linear, these special points are LQRE that can be solved for explicitly–along with the associated λ.[28] More generally, we are able to find similar LQRE explicitly in all of the binary-action games we consider.

3.1.3 Coordination

Coordination is any 2×2 game with three NE: two PSNE representing successful coordination, $(q, p) = (1, 1)$ and $(q, p) = (0, 0)$, and the MSNE $(q, p) = (q^*, p^*) \in (0, 1)^2$. Such games are the simplest with multiple NE and so have been used in experiments testing theories of equilibrium selection. These games include those that have been referred to as "battle of the sexes", "hawk-dove", "chicken", and "stag hunt", though the distinction between these particular variants is unimportant for our purposes. As with matching pennies, we restrict attention to $q^* \leq \frac{1}{2}$, which is without loss, but now we must also allow for any $p^* \in (0, 1)$ to cover all coordination games.

In coordination games, the (closure of the) set of RQRE is given by $R = R^1 \cup R^2 \cup R^3$ where $R^1 = [0, q^*] \times [0, min\{p^*, \frac{1}{2}\}]$, $R^2 = [q^*, \frac{1}{2}] \times [\frac{1}{2}, p^*]$, and $R^3 = [\frac{1}{2}, 1] \times [max\{p^*, \frac{1}{2}\}, 1]$ are three rectangular components.[29]

Figure 3 plots R, S, and L for six examples of coordination. These six examples correspond to qualitatively distinct shapes of S, which again appears as the union of some number of polytopes. Unlike for matching pennies, these polytopes are not necessarily connected, an observation that is useful for equilibrium selection.

[28]Take any (non-limiting) LQRE in which player i has two actions and is not uniformly mixing, i.e. $\sigma_{i1} \neq \frac{1}{2}$. Letting $d_i = \bar{u}_{i1} - \bar{u}_{i2}$ denote the (signed) difference in player i's equilibrium expected payoffs, we have that $\sigma_{i1} = \frac{e^{\lambda d_i}}{1+e^{\lambda d_i}} \iff \lambda = \frac{1}{d_i} ln(\frac{\sigma_{i1}}{1-\sigma_{i1}})$.

[29]If $p^* \leq \frac{1}{2}$ or $q^* = \frac{1}{2}$, the second component of R is degenerate: $R^2 = \emptyset$.

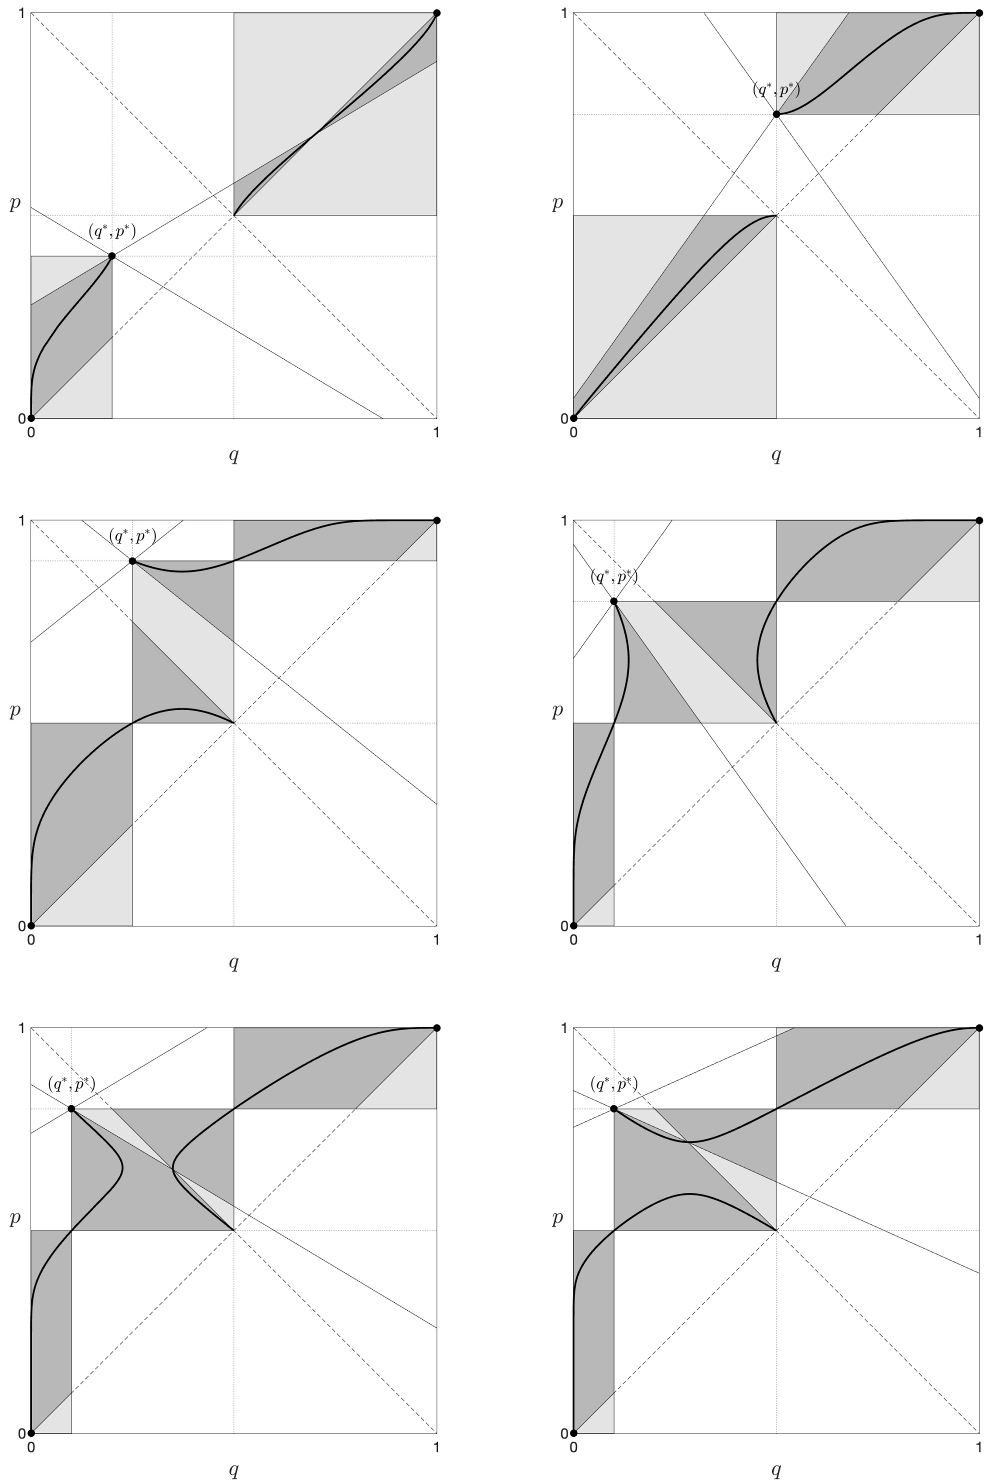

Figure 3: *Coordination.*

We find that, under certain conditions on parameters, there is only one NE that is path-connected via S to the centroid: SQRE uniquely selects this NE in the sense of Definition 3, which implies that LQRE makes the same selection (Lemma 1). The theorem below summarizes these cases. The proof follows directly from the geometry of S and so is omitted.

Theorem 2. *For coordination games, Table 3 gives conditions on p^*, q^*, and r for which SQRE uniquely selects the given NE ("Selected NE").*

Case	p^*	$q^* \leq \frac{1}{2}$	r	Selected NE
i	$p^* < \frac{1}{2}$	–	–	$(1, 1)$
ii	$p^* > \frac{1}{2}$	$q^* = \frac{1}{2}$	–	$(0, 0)$
iii	$\frac{1}{2} > p^* \geq 1 - q^*$	$q^* < \frac{1}{2}$	$r \geq \frac{p^* - \frac{1}{2}}{\frac{1}{2} - q^*}$	$(1, 1)$
iv	$p^* \geq 1 - q^*$	$q^* < \frac{1}{2}$	$r \leq \frac{p^* - \frac{1}{2}}{\frac{1}{2} - q^*}$	$(0, 0)$

Table 3: *Conditions for SQRE equilibrium selection in coordination games.*

Figure 2 illustrates the theorem. The top two panels correspond to cases *(i)* and *(ii)*. Here, the result does not hinge on symmetry in the sense that RQRE already makes a unique selection. The middle two panels give the more interesting cases, *(iii)* and *(iv)*. Here, while all three NE are selected by RQRE, SQRE makes a unique selection. The bottom two panels show cases where even SQRE does not make a unique selection, though symmetry still imposes significant bounds on L.

Relating our result to existing literature, we find that Theorem 2 gives the same conditions for NE selection as does Theorem 1 of Zhang and Hofbauer [2016] for a model similar to SQRE, though the results are proved in different ways. Our result follows from an exact SQRE characterization, which makes it possible to directly "observe" the selected equilibrium. Zhang and Hofbauer [2016], building upon earlier work by Turocy [2005], uses "homotopy methods". Such methods are very powerful, but do not give much insight into the overall shape of the equilibrium set. Finally, we note that all cases *(i)*-*(iv)* of Theorem 2 involve selecting the risk-dominant equilibrium (Harsanyi and Selten [1988]).[30]

[30] It is easy to show that $(q, p) = (1, 1)$ ($(q, p) = (0, 0)$) is risk-dominant if and only if $p^* + q^* < 1$ ($p^* + q^* > 1$).

3.1.4 Prisoner's dilemma and two-step dominance solvable games

The supplemental material (Section 10.1) gives the two other cases of 2×2 games: prisoner's dilemma and two-step dominance solvable games. An interesting feature of these games is that (q^*, p^*) falls outside of the unit square, which has implications for the shape of the set of SQRE.

4 Applications

To illustrate the potential of SQRE in applications, we study three classic games of both theoretical and experimental interest. We first consider two N-players games from political economy, a participation game with heterogenous costs of entry and a jury voting game with unanimity. We then consider the infinitely repeated prisoner's dilemma. For each game, we characterize all SQRE within a particular class–one that generalizes both a class of NE that has been analyzed theoretically and a class of LQRE that has been studied numerically. To explore the breadth of potential applications, we selected these games so as to involve many players, incomplete information, and dynamic strategies.

4.1 Participation game with heterogeneous costs

Following Diekmann [1986], there are N players, each of whom decides whether to enter or not. If at least 1 player enters, all receive $B > 0$.[31] As in Diekmann [1986], entering is costly, but to make the problem more interesting and general, we assume costs are heterogeneous. With probability c, the cost is low, $c_1 \in (0, B)$, and with probability $1 - c$, the cost is high, $c_2 \in (c_1, B)$.

We search for semi-symmetric SQRE in which all players of the same type have the same behavior, and so we let σ_t denote the probability with which a player of type t enters. Theorem 1 extends in the natural way to this case, and so we find the semi-symmetric part of S explicitly as the union and intersection of sets.

By entering, one ensures that the contribution threshold is reached, so a player of type t receives $B - c_t > 0$. By not entering, one avoids the cost and receives B if at least one of the $N - 1$ other players enters. The probability that at least one other player enters equals 1 minus the probability no other player enters, so the expected payoff from not entering is $B(1 - (c(1 - \sigma_1) + (1 - c)(1 - \sigma_2))^{N-1})$.

[31] When the threshold is exactly 1, as we assume here, the game is sometimes referred to as the "Volunteer's dilemma."

The difference in expected payoffs from entering versus not entering for a player of type t is $B(c(1-\sigma_1)+(1-c)(1-\sigma_2))^{N-1}-c_t$. Hence, entering yields a higher payoff if and only if

$$\sigma_1 \leq \underbrace{1-\frac{1}{c}\left(\frac{c_t}{B}\right)^{\frac{1}{N-1}}+\left(\frac{1-c}{c}\right)-\left(\frac{1-c}{c}\right)\sigma_2}_{\equiv f_t(\sigma_2;c_t)}. \tag{4}$$

In other words, it is optimal for type t to enter if σ_1 is less than a downward-sloping, linear function of σ_2, which for convenience we have defined as $f_t(\sigma_2;c_t)$ (suppressing most parameters in the notation). The slope of f_t depends only on c, but the intercept depends on c, B, N and c_t. The intercept is decreasing in c_t, which means that whenever entering is optimal for high cost-types, it is also optimal for low cost-types.

Using the inequality (4), R is found by collecting regions of the unit square that are consistent with best responding more often than not: $R=\{(\sigma_2,\sigma_1)|\sigma_t>\frac{1}{2} \iff \sigma_1 < f_t(\sigma_2;c_t) \text{ for } t=1,2\}$. The (closure of) set R can also be written explicitly.[32]

To complete the characterization of S, we need to find expressions for U_0, U_1, and U_2 whereby one type of player or the other faces higher stakes. U_0, where both types face the same stakes, is defined by (σ_2,σ_1) that satisfy $|B(c(1-\sigma_1)+(1-c)(1-\sigma_2))^{N-1}-c_1| = |B(c(1-\sigma_1)+(1-c)(1-\sigma_2))^{N-1}-c_2|$. Noting that $c_1<c_2$, this can only hold when

[32] $R=R^1\cup R^2$ where

$$\begin{aligned} R^1=\{(\sigma_2,\sigma_1):\ &\sigma_2\in[med\{0,\frac{1}{2},\{\sigma_2'|f_2(\sigma_2';c_2)=1\},\frac{1}{2}],\\ &\sigma_1\in[med\{f_1(\sigma_2;c_1),f_2(\sigma_2;c_2),\frac{1}{2}\},min\{max\{f_1(\sigma_2;c_1),\frac{1}{2}\},1\}]\} \text{ and}\\ R^2=\{(\sigma_2,\sigma_1):\ &\sigma_2\in[\frac{1}{2},med\{\frac{1}{2},1,\{\sigma_2'|f_2(\sigma_2';c_2)=\frac{1}{2}\}\}],\\ &\sigma_1\in[\frac{1}{2},min\{f_2(\sigma_2;c_2),1\}]\}\end{aligned}$$

are two *non*-rectangular components. If the high-cost type's cost-benefit ratio is sufficiently low $\left(\frac{c_2}{B}\leq\left(\frac{1-c}{2}\right)^{N-1} \iff \{\sigma_2'|f_2(\sigma_2';c_2)=1\}\geq\frac{1}{2}\right)$, then the first component is degenerate: $R^1=\emptyset$. If the high-cost type's cost-benefit ratio is sufficiently high $\left(\frac{c_2}{B}\geq\left(\frac{1}{2}\right)^{N-1} \iff \{\sigma_2'|f_2(\sigma_2';c_2)=\frac{1}{2}\}\leq\frac{1}{2}\right)$, then the second component is degenerate: $R^2=\emptyset$.

$0 < c_1 < B(c(1-\sigma_1)+(1-c)(1-\sigma_2))^{N-1} < c_2$, in which case we have

$$B(c(1-\sigma_1)+(1-c)(1-\sigma_2))^{N-1} - c_1 = c_2 - B(c(1-\sigma_1)+(1-c)(1-\sigma_2))^{N-1} \iff$$

$$\sigma_1 = \underbrace{1 - \frac{1}{c}\left(\frac{c_1+c_2}{2B}\right)^{\frac{1}{N-1}} + \left(\frac{1-c}{c}\right) - \left(\frac{1-c}{c}\right)\sigma_2}_{\equiv g(\sigma_2)},$$

where for convenience we have defined $g(\sigma_2)$. We note that $g(\sigma_2)$ is a downward sloping line that is parallel to $f_1(\sigma_2; c_1)$ and $f_2(\sigma_2; c_2)$ that falls between them: $f_2(\sigma_2; c_2) < g(\sigma_2) < f_1(\sigma_2; c_1)$ for all σ_2. It is now easy to show that

$$\begin{aligned} U_0 &= \{(\sigma_2, \sigma_1) : \sigma_1 = g(\sigma_2)\}, \\ U_1 &= \{(\sigma_2, \sigma_1) : \sigma_1 < g(\sigma_2)\}, \text{ and} \\ U_2 &= \{(\sigma_2, \sigma_1) : \sigma_1 > g(\sigma_2)\}. \end{aligned}$$

Hence, U_0 is the line defined by $g(\sigma_2)$, U_1 is the half-space below it, and U_2 is the half-space above it.

With E_0, E_1, and E_2 defined as in (3), but grouping players of the same type, we characterize S as in Corollary 1. Figure 4 plots R, S, and L for two parameterizations of the participation game (parameters given in the figure). We see that S is given as the union of polytopes, that it significantly refines R, and that it imposes significant bounds on L–including giving some LQRE exactly as the solution to linear systems.

In the left panel of Figure 4, parameters are such that, while most SQRE and RQRE involve only the low-cost player tending to enter, there do exist QRE in which both types enter more often than not. In the right panel, parameters are such that, in all QRE, low-cost players tend to enter whereas high-cost players tend to stay out.

In the next result, we argue that RQRE gives implausible predictions. If the high-cost player's cost-benefit ratio is sufficiently low, as in the left panel of Figure 4, there are RQRE in which the high-cost player actually participates *more* often than the low-cost player. Interestingly, such implausible equilibria are ruled out by symmetry as no SQRE involves this feature.

Proposition 2. *In the participation game: (i) If the high-cost type's cost-benefit ratio is sufficiently low ($\frac{c_2}{B} < \left(\frac{1}{2}\right)^{N-1}$), there are RQRE in which the high-cost type participates more often than the low-cost type ($\sigma_2 > \sigma_1$). (ii) For all parameters, all SQRE are such that the low-cost type participates more often than the high-cost type ($\sigma_1 > \sigma_2$).*

Proof. See Appendix 9.1. □

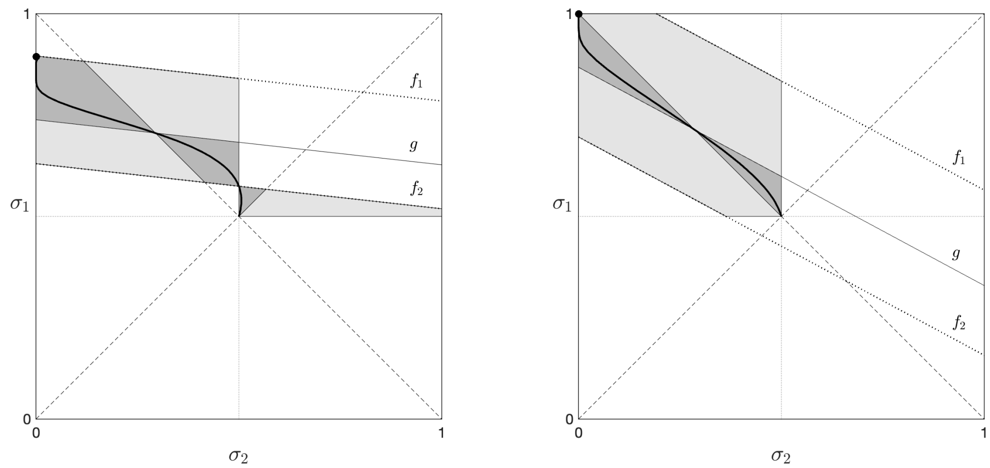

Figure 4: *Participation with heterogeneous costs.* Left panel: $c = 0.9$, $B = 8$, $N = 3$, $c_1 = 0.3$, and $c_2 = 1.5$. Right panel: $c = 0.65$, $B = 5$, $N = 3$, $c_1 = 0.4$, and $c_2 = 1.5$.

4.1.1 Jury voting and the infinitely repeated prisoner's dilemma

In the supplemental material (Section 10.2), we explore two additional applications: a jury voting game with unanimity and the infinitely repeated prisoner's dilemma. These show how SQRE can be applied to games with incomplete information and dynamic strategies.

5 Asymmetric QRE

The same tools developed in this paper can be used to characterize *asymmetric QRE* (AQRE) whereby some players are noisier than others. We view the value of AQRE as two-fold. First, it is the relevant model whenever there is some reason to expect a particular ranking of players by noise, perhaps because some players are known to have more experience than others. Second, as we illustrate in Section 7, AQRE can be used to organize data by non-parametrically identifying which players appear to be noisier. To the extent there is a relationship between features of games and the types of asymmetries observed, this may point to alternative models that are more consistent with the data.

To characterize AQRE, first consider any $\sigma \in R$ in a binary-action game. We say that player i is *strictly less noisy than* player j if player i *faces smaller stakes than* and is *more extreme than* player j (with at least one of these strict). This order is incomplete: for any given $\sigma \in R$, it may be that two players are unranked by noise. In fact, the set of SQRE equals the subset of R for which the order is maximally incomplete. More generally, the order implies a partition of R whereby all action profiles within a partition element yield the same ranking and incompleteness relation between players.

In 2×2 games, for example, either player 1 is strictly less noisy than player 2, player 2 is strictly less noisy than player 1, or neither player is strictly noisier than the other. Hence, we partition the set of RQRE into three parts: $R = S \cup S^1 \cup S^2$ where $S = (U_1 \cap E_1) \sqcup (U_2 \cap E_2) \sqcup (U_0 \cap E_0) \cap R$ is the set of SQRE and $S^i = (U_j \cap E_i) \sqcup (U_0 \cap E_i) \sqcup (U_j \cap E_0) \cap R$ is the component of the set of AQRE whereby player i is strictly less noisy than player j. In Figure 5, we plot this decomposition for a matching pennies game, with each component of the set of AQRE in a different color. More generally, these components can be solved for explicitly in all 2×2 games and all applications considered in this paper.

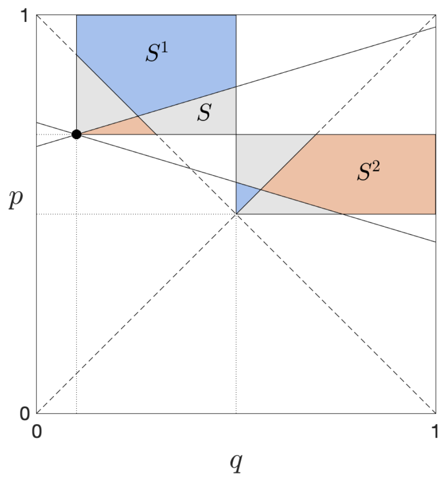

Figure 5: *Asymmetric QRE.* For a matching pennies game, we plot the set of SQRE S (gray) and the two components of the set of AQRE: S^1 for which player 1 is less noisy (blue) and S^2 for which player 2 is less noisy (red).

6 Larger action spaces

As suggested by the example of Section 1.7, the restrictions of SQRE depend on being able to make comparisons across players so that *responsiveness* implies restrictions. In binary-action games, the symmetry axioms complete the "responsiveness order". For games with larger action spaces, however, the symmetry axioms still leave the order incomplete. Hence, we consider a model with additional structure. This model satisfies the symmetry axioms and is equivalent to SQRE in binary-action games, so the results in this section extend the previous approach.

We consider the framework of *structural QRE* (see, for example, Goeree et al. [2005]). We follow the literature in defining this to be any QRE model in which player i's quantal response is induced by additive errors $\varepsilon_i = (\varepsilon_{i1}..., \varepsilon_{iJ(i)}) \in \mathbb{R}^{J(i)}$ such that

$$Q_{ij}(\bar{u}_i) = \int_{\{\varepsilon_i \in \mathbb{R}^i | \bar{u}_{ij} + \varepsilon_{ij} \geq \bar{u}_{ik} + \varepsilon_{ik} \ \forall k \in \{1,...,J(i)\}\}} dF_i(\varepsilon_i),$$

where distribution F_i is full-support, absolutely continuous, and mean-zero. In addition, we make the standard assumptions that the error distributions are *invariant to payoffs* ($F_i(\varepsilon_i|\bar{u}_i) = F_i(\varepsilon_i)$) and independent across players ($F_i \perp F_j$) but not necessarily actions.

We define *symmetric structural QRE* (SSQRE) to be any structural QRE model with additional symmetry conditions across players and actions that mirror the symmetry axioms. In particular, we require that (1) if $J(i) = J(j) = J$, then $F_i = F_j$ for some ordering of players' actions, i.e. there exists a bijection $\iota : \{1, ..., J\} \to \{1, ..., J\}$ such that, for every realization ε_i, $F_i(\varepsilon_{i1}, ..., \varepsilon_{iJ}) = F_j(\varepsilon_{i\iota(1)}, ..., \varepsilon_{i\iota(J)})$; and (2) F_i is exchangeable. Hence, if players have the same number of actions, then they have the same error distribution, which does not favor some actions a priori. For an arbitrary game Γ, we refer to the *set of* SSQRE as S^* and make two remarks:

Remark 3. Whereas structural QRE can violate *monotonicity* (R4) (as well as (S1) and (S3)), it is immediate that SSQRE satisfies (R1)-(R4) and (S1)-(S3); and thus $S^* \subset S$.

Remark 4. SSQRE is more general than structural QRE with i.i.d. errors since error distributions can differ across players with different numbers of actions, and, for each player, errors can be correlated across actions.

Using results from convex analysis, Melo et al. [2018] characterize when the data from *sets of games* are consistent with structural QRE using the so called "cyclic monotonicity" (CM) inequalities. The key assumption is that player i's error distribution is the same across games, so the data from M games gives M observations on player i's quantal response function; and the CM inequalities give the necessary and sufficient restrictions for each player i. In SSQRE, by imposing symmetry across players, if N' players have the same number of actions, SSQRE allows us to treat the N' players' data as coming from the same quantal response function. And since SSQRE imposes exchangeability, we can actually treat every permutation of a player's data as an independent observation.

With this perspective, our result is a straightforward adaptation of the CM inequalities.

For each player i, we define her set of *action orderings* by $\mathcal{A}_i = \{(a_{i\iota(1)}, a_{i\iota(2)}, ..., a_{i\iota(J(i))})|\iota : \{1, ..., J(i)\} \to \{1, ..., J(i)\}$ is a bijection$\}$, which simply gives every possible permutation of her actions.

For a given number of actions J, we consider J*-cycles* of action orderings for players with exactly J actions. Such a cycle is given by $C^J = \{A^0, A^1, ..., A^{L-1}, A^L\}$ where $A^L = A^0$ and $L \geq 2$ is the length of the cycle. Each term A^m in the cycle corresponds to an action ordering for some player with J actions: $A^m \in \mathcal{A}_i$ for some i with $J(i) = J$. It can be that $A^m, A^n \in \mathcal{A}_i$, meaning that multiple orderings are for the same player. However, with the exception of the endpoints, any two orderings are distinct: $A^m \neq A^n$ unless $m = 0$ and $n = L$.

For a given action ordering A^m, we define u_j^m and σ_j^m to be the equilibrium expected payoff to the player's jth action in the ordering and the probability this action is taken, respectively. For a given J-cycle C^J, we define $CM^J(C^J)$ as the set of action profiles (for *all* players) such that the behavior of the players with exactly J actions satisfies a version of the cyclic monotonicity inequalities:

$$CM^J(C^J) = \left\{ \sigma \in \Delta | \sum_{m=0}^{L-1} \sum_{j=1}^{J} (u_j^{m+1}(\sigma) - u_j^m(\sigma))\sigma_j^m \leq 0 \right\}. \tag{5}$$

We define $CM^J = \bigcap_{C^J} CM^J(C^J)$ as the intersection over all J-cycles, and we define CM^* as the intersection of CM^J over all J:

$$CM^* = \bigcap_J CM^J.$$

Admittedly, there are many cycles to consider. However, the inequalities in (5) are invariant to the choice of starting index, e.g. the inequalities emerging from (A^m,A^n,A^l,A^m) and (A^n,A^l,A^m,A^n) are the same. Melo et al. [2018] make the same point in the context of their problem. More specific to our problem, we make the following remarks:

Remark 5. If some player i satisfies *monotonicity* in the sense that $\bar{u}_{ij} > \bar{u}_{ik} \iff \sigma_{ij} > \sigma_{ik}$ for all j, k, then the inequality in (5) is automatically satisfied for all cycles involving only that one player. Hence, to determine if $\sigma \in CM^*$, one need not consider any player-specific cycles for players who satisfy *monotonicity*. Conversely, if any player violates *monotonicity* in the sense that $\bar{u}_{ij} > \bar{u}_{ik}$ and $\sigma_{ij} < \sigma_{ik}$ for some j, k, then $\sigma \notin CM^*$.

Remark 6. Any two cycles $C_1^J = \{A^0, A^1, ..., A^{L-1}, A^L\}$ and $C_2^J = \{B^0, B^1, ..., B^{L-1}, B^L\}$

lead to the same inequalities if (1) A^m and B^m correspond to the same player ($A^m, B^m \in \mathcal{A}_i$ for some i) for all m and (2) there exists a bijection $\gamma : \{1, ..., J\} \to \{1, ..., J\}$ such that $A^m = (a_1^m, ..., a_J^m)$ and $B^m = (a_{\gamma(1)}^m, ..., a_{\gamma(J)}^m)$ for all m.

Our first result is that the set of SSQRE S^* is contained in the intersection of CM^* and R. This follows from adapting the results of Melo et al. [2018] and the easy-to-show fact that SSQRE is regular.

Theorem 3. *Fix* Γ. $S^* \subset CM^* \cap R$.

Proof. See Appendix 9.1. □

Hence, to find necessary conditions on S^*, one need only characterize CM^* and R, and then take their intersection. In light of Remark 5, "most" of CM^* is contained within R: only when $\sigma \in CM^*$ is such $\bar{u}_{ij} = \bar{u}_{ik}$ and $\sigma_{ij} \neq \sigma_{ik}$ for some player i and actions j, k is it the case that $\sigma \notin R$. It is still useful to express the result as in Theorem 3 because of the comparative simplicity of R. For numerical analyses, this form of the result suggests to first express R and then search within R for all $\sigma \in CM^*$.

We conjecture the stronger claim that $S^* = CM^* \cap R$, but we are unable to show it for general games. In binary-action games, however, we find both that $S^* = CM^* \cap R$ and that $S^* = S = X \cap R$. Hence, Theorem 1 is necessary and sufficient for representing the set of SSQRE.

Proposition 3. *Fix* $\Gamma^{N \times 2}$. $S^* = CM^* \cap R = S = X \cap R$.

Proof. See Appendix 9.1. □

Player 1:

	R	**B**	**Y**
R	9	6	4
B	8	8	2
Y	4	4	5

Player 2 ($\times\beta$):

	R	**B**	**Y**
R	9	6	4
B	8	8	2
Y	4	4	5

Table 4: *The Mondrian Game.*

To illustrate the restrictions of SSQRE for a game in which each player has more than two actions, we consider the 3×3 "Mondrian" game whose payoffs are given in Table 4.[33] We consider a game that, at its base, is symmetric, but we then scale player 2's payoffs by constant $\beta > 0$. This form of asymmetry is special in that it still allows us

[33] This particular variant was considered by Goeree and Louis [2021], and Figure 4 is adapted from a figure in their paper.

to easily draw equilibria in a single simplex. Figure 4 shows the construction of the set of RQRE. The top left-panel gives ordinal rankings of actions by the probability each is played and the top-right panel gives rankings of actions by expected payoffs. By taking intersections, the bottom panel gives the set of RQRE, revealing that it is the union of three "colorable" sets: red, blue, and yellow. Hence, $\sigma \in R$ if and only if σ_1 and σ_2 are the same color.

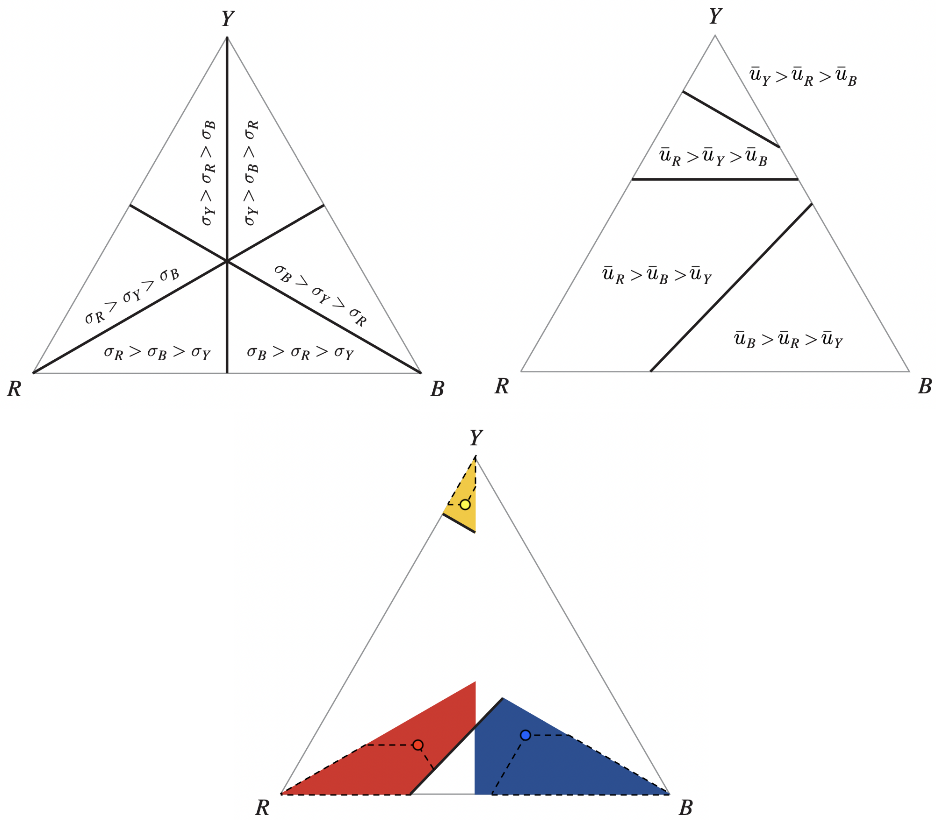

Figure 6: *QRE in the Mondrian Game.* The top-left panel gives regions of the simplex consistent with different rankings of actions by the probability each is played. The top-right panel gives regions of the simplex consistent with different rankings of actions by expected payoffs. By taking intersections, the bottom panel gives the sets of RQRE: $(\sigma_1, \sigma_2) \in R$ if and only if σ_1 and σ_2 are the same color. Fixing the behavior of player 1 at one of the colored circles, and choosing β to be sufficiently large, player 2's behavior must be in the more extreme dashed region of the same color to be part of an SSQRE.

Now we consider σ_1 to be the red circle in Figure 4. As long as σ_2 is in the red region, (σ_1, σ_2) is an RQRE, which is true for any $\beta > 0$. Now choose β sufficiently large so that, as long as σ_2 is in the red region, we have that $\bar{u}_{1R} > \bar{u}_{1B} > \bar{u}_{1Y}$, $\bar{u}_{2R} > \bar{u}_{2B} > \bar{u}_{2Y}$, $\bar{u}_{2R} - \bar{u}_{2B} > \bar{u}_{1R} - \bar{u}_{1B}$, and $\bar{u}_{2B} - \bar{u}_{2Y} > \bar{u}_{1B} - \bar{u}_{1Y}$. Hence, both players' actions are ordered by expected payoffs in the same way, but player 2 faces higher stakes in the sense that the expected payoff differences between any two of her actions is greater than it is for player 1. An easy-to-show implication of symmetry is that player 2's behavior must be more extreme in the sense that $\sigma_{2B} > \sigma_{1B}$ and $\sigma_{2Y} < \sigma_{1Y}$. Hence, σ_2 must lie in the dashed subset of the red region, consisting of all red points more extreme than σ_1. The figure also depicts, for arbitrary σ_1 in the yellow and blue regions the corresponding colored sets of σ_2 that are more extreme.

This example shows that symmetry imposes significant restrictions in games with larger action spaces. The implication of symmetry that we explore here, based on generalized stakes and extremeness, extends what we observed in the binary-action case. The inequalities implicit in Theorem 3 give additional restrictions.

7 Analysis of experimental data

We evaluate model performance using data from two studies on experimental matching pennies games, Selten and Chmura [2008] and McKelvey et al. [2000]. We focus on these studies because each played several games, subjects were randomly matched with opponents so that they could not use dynamic strategies, and there were sufficient rounds for some degree of learning to take place.

We first visualize the data using familiar graphs. Figure 7 corresponds to the 12 games of Selten and Chmura [2008], and Figure 8 corresponds to the 4 games of McKelvey et al. [2000]. Within each graph, as before, S is given by the dark gray region, $R\backslash S$ is given by the light gray region, and L is given by the black curve. In addition, we give the empirical frequency (black dot) and the prediction of each theory that minimizes the Euclidean distance to the data: the best-fit SQRE (white circle), best-fit RQRE (red diamond), and best-fit LQRE (blue square).

The three models we consider, LQRE, SQRE, and RQRE, are nested, but since the latter two are non-parametric "area theories", we cannot rely on standard measures of fit that penalize parameters. Instead, we statistically test (using an F-test, clustered by independent subject group) whether the best-fit prediction of each model generates the data in a given game. To the extent we reject the best-fit prediction of a theory, the data is not plausibly generated by any prediction of that theory. While this procedure does not "select a model", it does give some sense of what is the most precise theory that is consistent with the data.

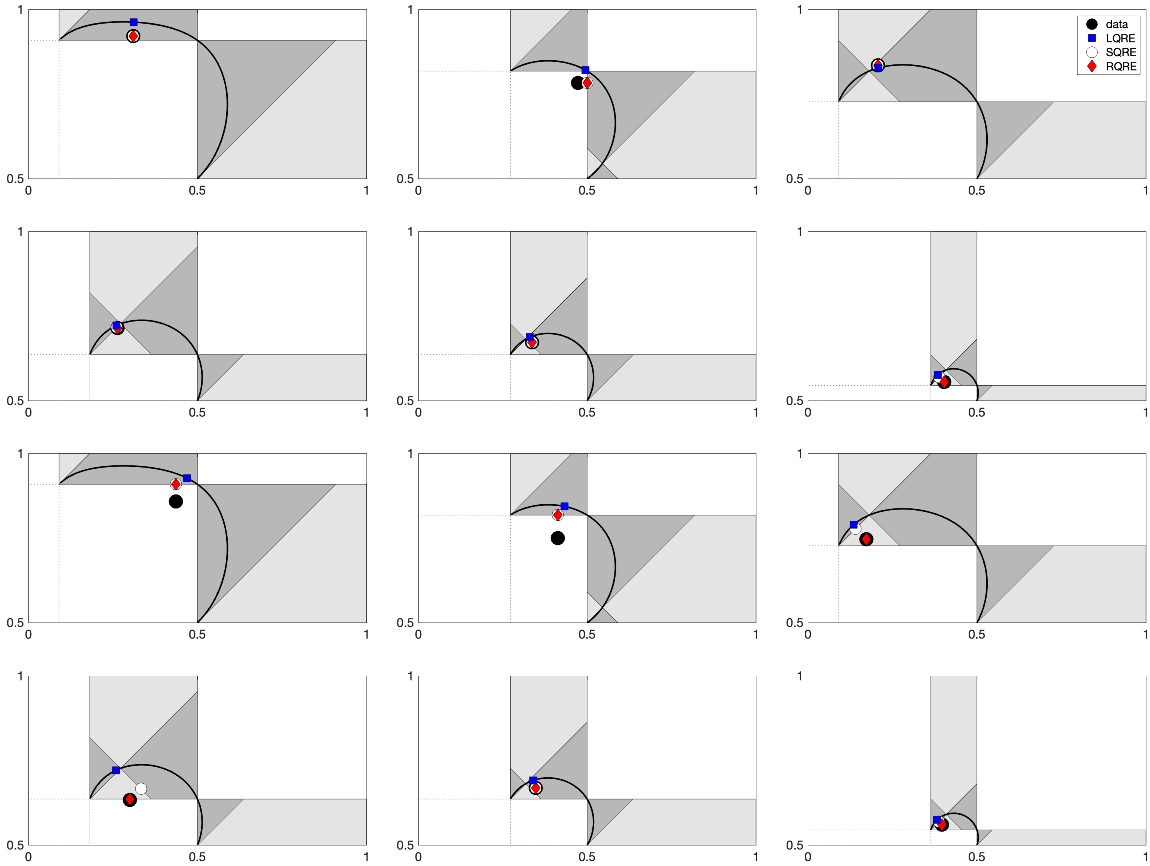

Figure 7: *Selten and Chmura (2008).*

Inspecting the figures, it appears visually that all three models perform rather well. This is confirmed in Table 5, which gives distances to the data and the results of the statistical tests. Ordering the models as LQRE, SQRE, and RQRE, the distances must be weakly decreasing. We find average distances of 0.064, 0.036, and 0.009, respectively, which are all quite low and much lower than the average distance for NE of 0.147. In terms of statistical significance at the 5%-level, we find that LQRE, SQRE, and RQRE are rejected in 11, 6, and 2 games, respectively, but again, the distances are typically small. Broadly speaking, the models are capable of capturing patterns in the data.

Focusing on the two area theories, SQRE and RQRE, Table 5 gives the area (i.e. Lebesgue measure) for each game. On average, the SQRE area of 0.076 is much smaller than the RQRE area of 0.183. Hence, while the best-fit RQRE comes closer to the data than the best-fit SQRE on average, SQRE makes much more precise predictions.

Looking at the figures more carefully, we see two interesting qualitative patterns. First, for all 16 games, the data never falls in the "extreme" components of $R\backslash S$, the

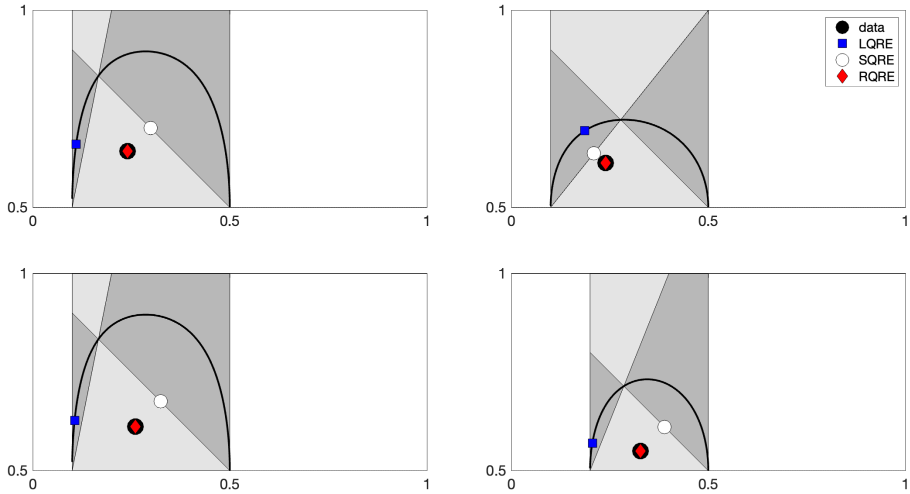

Figure 8: *McKelvey, Palfrey, and Weber (2000).*

light gray regions that are bordered by the boundary of the unit square. Hence, there are large areas of R that predictably do not improve fit. That this holds across all games suggests the value of SQRE as a refinement of RQRE. Second, there are interesting differences across the two studies, as we discuss below.

In Selten and Chmura [2008], the average distances of SQRE and RQRE are 0.023 and 0.012, respectively, and the average areas are 0.066 and 0.181, respectively. In this dataset, the gap in distance is relatively small and the gap in area is relatively big, meaning SQRE performs especially well. Furthermore, deviations from SQRE do not appear to be systematic.

In McKelvey et al. [2000], the data consistently falls in R/S, meaning it can be rationalized exactly by RQRE but not SQRE. SQRE also performs more poorly on average than it does in Selten and Chmura [2008]: the average distance is 0.074 (as opposed to 0 for RQRE) and the average area is 0.104 (as opposed to 0.188 for RQRE).

A natural question is if there is any consistent pattern in the way the data deviates from SQRE predictions in McKelvey et al. [2000]. Inspecting Figure 8, we see that the data for all 4 games falls in the region $U_1 \cap E_2 \cap R$, which is precisely the component of AQRE (see Section 5) in which player 2 is *less noisy than* player 1.

This is interesting because the most obvious difference between the Selten and Chmura [2008] and McKelvey et al. [2000] games is that the latter involve player 2 facing "symmetric" payoffs in the sense that she is made indifferent when player 1 uniformly mixes. Because of this, it is especially easy for player 2 to best respond: for any belief other than $p = 0.5$, one of her actions first-order stochastically dominates

Study	Game	LQRE		SQRE			RQRE		
		dist.	p	dist.	p	area	dist.	p	area
SC 2008	1	0.04	0.001***	0.00	1.000	0.12	0.00	1.000	0.24
	2	0.04	0.005***	0.03	0.462	0.08	0.03	0.462	0.20
	3	0.01	0.849	0.00	1.000	0.10	0.00	1.000	0.23
	4	0.01	0.848	0.00	0.978	0.06	0.00	1.000	0.18
	5	0.02	0.655	0.00	1.000	0.04	0.00	1.000	0.15
	6	0.03	0.046**	0.02	0.179	0.01	0.00	1.000	0.08
	7	0.08	0.000***	0.05	0.000***	0.12	0.05	0.000***	0.24
	8	0.10	0.001***	0.07	0.034**	0.08	0.07	0.034**	0.20
	9	0.06	0.000***	0.04	0.006***	0.10	0.00	1.000	0.23
	10	0.10	0.000***	0.05	0.099*	0.06	0.00	0.994	0.18
	11	0.02	0.417	0.00	1.000	0.04	0.00	1.000	0.15
	12	0.02	0.134	0.01	0.496	0.01	0.00	1.000	0.08
	avg.	0.043	–	0.023	–	0.066	0.012	–	0.181
MPW 2000	1	0.13	0.000***	0.08	0.000***	0.12	0.00	1.000	0.20
	2	0.10	0.000***	0.04	0.398	0.09	0.00	1.000	0.20
	3	0.16	0.001***	0.09	0.000***	0.12	0.00	1.000	0.20
	4	0.12	0.000***	0.09	0.000***	0.08	0.00	1.000	0.15
	avg.	0.127	–	0.074	–	0.104	0.000	–	0.188
Average	–	0.064	–	0.036	–	0.076	0.009	–	0.183

Table 5: *Model performance.*

the other. Hence, it would not be surprising if player 2-subjects are less noisy. This is both capable of rationalizing the observed data and consistent with previous studies that document rates of best response to stated beliefs.[34]

We interpret these results as suggesting a relationship between noise and *complexity* that is absent from traditional QRE models, which assume that stochastic choice depends only on expected payoffs. SQRE also does not embed such complexity considerations, but the SQRE-AQRE framework allowed us to identify the asymmetries that are their hallmark. More generally, by uncovering such asymmetries and their relationship to features of games, the framework may help to stimulate the development of new theories.

[34] Friedman and Ward [2022] show that subjects facing symmetric payoffs best respond at much higher rates, even conditional on subjective expected payoff differences between actions.

8 Discussion

8.1 Extensions

We consider two extensions in the supplemental material (Section 10.3). First, we derive a representation theorem analogous to Theorem 1 for *sets of games*, where the key assumption is that the quantal response function is held fixed across games. This gives rise to comparative static predictions across games and stronger restrictions on the data than the product of individual-game restrictions. Second, we give a similar representation theorem for an SQRE variant based on *scale invariance* as opposed to *translation invariance*. This leads to a different refinement of RQRE and nests models such as the generalized Luce QRE (Luce [1959]), which is based on multiplicative errors.

8.2 SQRE as a meta-theory

SQRE can be viewed as a meta-theory in that it nests a number of existing parametric models. Figure 9 plots, for a particular matching pennies game played by Selten and Chmura [2008], the set of SQRE along with the predictions from four parametric QRE models: LQRE, probit QRE, Heterogenous LQRE (McKelvey et al. [2000]), and Endogenous QRE (Friedman [2020] and McKelvey et al. [1997]).[35]

The four parametric models give rise to set-predictions that are curves connecting the centroid to the Nash equilibrium. Since all set-predictions are contained within the set of SQRE, the degree of freedom of being able to choose the point on each curve that best fits the data implies that the empirical performance of all models must be fairly similar. In particular, plotting the data for this example (black dot), we see that it happens to fall near the union of two of the polytopes that define the set of SQRE. Since this point is necessarily a prediction of all four models (based on a path-connectedness argument), the models perform just about equally well. This example suggests caution in introducing new models and designing experiments for distinguishing between them.

[35]Probit QRE is a structural QRE model with normally distributed errors whose parameter is the standard deviation. Heterogeneous LQRE is a two-parameter model in which each player represents a population of individuals with values of λ drawn i.i.d. from a log-normal distribution. We assume, somewhat arbitrarily, that the standard deviation of the log-normal distribution τ depends on the mean μ according to the increasing function $\tau = 3log(1+\mu)$; and we plot the set of predictions indexed by $\mu \in (0,\infty)$. Endogenous QRE is a model in which each player i chooses λ_i optimally subject to cost function $c(\lambda_i) = \theta\lambda_i^2$; and we plot the set of predictions indexed by $\theta \in (0,\infty)$.

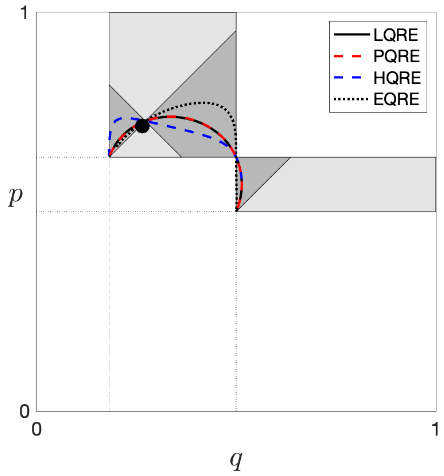

Figure 9: *SQRE as a meta-theory.*

8.3 SQRE for computation[36]

For numerically approximating LQRE, one can restrict attention to the set of SQRE to perform a more targeted and computationally efficient grid search. In some cases, one can do better. Consider the left panel of Figure 10, which depicts the sets of SQRE and LQRE for a matching pennies game. Following the discussion in Section 3.1.2 and Footnote 28, we solve explicitly for the values of λ associated with the LQRE at the intersections of the SQRE polytopes.[37] Since λ is monotonic across the arc of LQRE, each polytope is associated with a specific range of λ values. This allows one to refine the grid search further and gives information about the speed of convergence as $\lambda \to \infty$.

Information about convergence speed, intuitively, can also be useful for equilibrium selection. Consider the right panel of Figure 10, which depicts the sets of SQRE and LQRE for a coordination game. This is a case where SQRE does not make a unique selection. However, we can determine the logit selection by finding the LQRE associated with the intersections of the SQRE polytopes. In order for the principal branch to approach the NE $(1, 1)$, it must first pass through the point labelled λ^3 and then the point labelled λ^2. However, this cannot be since $\lambda^3 > \lambda^2$, and so the principal branch can only approach NE $(0, 0)$.

Finally, in order to use the powerful homotopy methods of Turocy [2005] to efficiently trace out a given branch of the LQRE correspondence, one must be able to find specific points on that branch. A general method for finding points on *non*-principal branches is

[36]We credit our conversations with Ted Turocy for inspiring this section.

[37]The game is scaled so that λ^1, the smallest non-zero λ that can be solved for, equals 1.

not known. However, we are able to find such points in coordination games (right panel of Figure 10) as well as in our application of the infinitely repeated prisoner's dilemma (supplemental material, Section 10.2.2), suggesting our methods may be more broadly applicable.

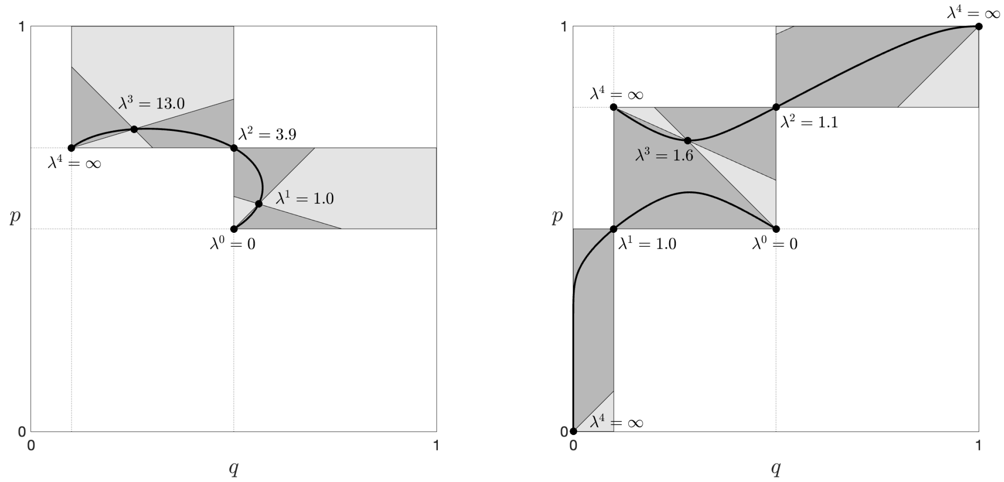

Figure 10: *Computational uses of SQRE.*

8.4 Symmetry in other models

The main insight of this paper is that symmetry implies restrictions in games without restricting individual behavior. Hence, one can sharpen predictions without having access to individual-level information, predictions that will hold in any population. This insight is very general, so it may be fruitful to explore the implications of symmetry in other models. This includes other QRE-type models, such as M equilibrium (Goeree and Louis [2021]), empirical equilibrium (Brown and Velez [2018]), and team equilibrium (Palfrey et al. [2021]); models of sampling (e.g. Rubinstein and Osborne [2003] and Goncalves [2020]) and random belief equilibrium (e.g. Friedman and Mezzetti [2005] and Friedman [2022]); and non-equilibrium models, such as level k and its successors (e.g. Camerer et al. [2004]).

8.5 Conclusion

We offer an axiomatic QRE model that imposes "symmetry" across players and actions, whose axioms are microfounded and satisfied in the large majority of applications. The model gives a significant refinement of regular QRE, with symmetry reducing the measure of predictions in our examples by more than half. By the same token, the theory implies much tighter bounds on logit QRE than previously known. Our main result, a representation theorem, makes the theory practical, showing how to characterize its set-predictions without having to solve for fixed points. Hence, *symmetric QRE*, or *SQRE*, balances the precision of logit QRE with the tractability of regular QRE.

In application to experimental data, we envision the SQRE framework being used alongside logit QRE, both to aid in its computation and as additional non-parametric benchmark. As illustrated in Section 7, the framework is also useful for organizing data and the identification of which players appear to be noisier, which may help to suggest directions for new theories of behavior in games.

As theoretical applications, we completely characterize SQRE's set-predictions for (almost) all 2×2 games. We also consider two classic N-player games from political economy as well as the infinitely repeated prisoner's dilemma. For each, we characterize all SQRE within a large class that generalizes a class of NE that has been previously studied. We find these solutions compelling as they exclude implausible equilibria, such as equilibria of a participation game in which higher-cost types participate more often. As a bi-product of our analysis, we also find that the SQRE framework is useful for NE selection. These results suggest that SQRE may be practical in theoretical applications in addition to being a tool for understanding experimental data.

References

Alexander Brown and Rodrigo Velez. Empirical equilibrium. *Working paper*, 2018. 8.4

Colin Camerer. Behavioral game theory: Experiments in strategic interaction. *Princeton University Press*, 2003. (document)

Colin Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 2004. 8.4

Hsiao-Chi Chen, James W. Friedman, and Jacques-Francois Thisse. Boundedly rational nash equilibrium: A probabilistic choice approach. *Games and Economic Behavior*, 1997. (document)

Andreas Diekmann. Volunteer's dilemma: A social trap without a dominant strategy and some empirical results. *Procedural Effects of Social Behavior: Essays in Honor of Anatol Rapoport*, 1986. (document), 4.1

Timothy Feddersen and Wolfgang Pesendorfer. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *American Political Science Review*, 1998. (document), 10.2.1

Axel Franzen. Group size and one shot collective action. *Rationality and Society*, 1995. 10

Evan Friedman. Endogenous quantal response equilibrium. *Games and Economic Behavior*, 2020. 8.2

Evan Friedman. Stochastic equilibria: Noise in actions or beliefs? *American Economic Journal: Microeconomics*, 2022. 8.4, 10.3.2

Evan Friedman and Jeremy Ward. Stochastic choice and noisy beliefs in games: an experiment. *Working Paper*, 2022. 34

James Friedman and Claudio Mezzetti. Random belief equilibrium in normal form games. *Games and Economic Behavior*, 2005. 8.4

Jacob Goeree and Philippos Louis. M equilibrium. *American Economic Review*, 2021. (document), 16, 2, 18, 33, 8.4

Jacob Goeree, Charles Holt, and Thomas Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 2003. 10.3.1

Jacob Goeree, Charles Holt, and Thomas Palfrey. Regular quantal response equilibrium. *Experimental Economics*, 2005. 2, (document), 1.3, 1, 6, 10.3.1, 45

Jacob Goeree, Charles Holt, and Angela Smith. An experimental examination of the volunteer's dilemma. *Games and Economic Behavior*, 2015. 10

Jacob Goeree, Charles Holt, and Thomas Palfrey. Quantal response equilibrium. *Princeton University Press*, 2016. (document), 10.2.2

Jacob Goeree, Charles Holt, Philippos Louis, Thomas Palfrey, and Brian Rogers. Rank-dependent choice equilibrium: A non-parametric generalization of qre. *Handbook of Research Methods and Applications in Experimental Economics*, 2018. 16

Russell Golman. Quantal response equilibria with heterogeneous agents. *Journal of Economic Theory*, 2011. 1

Duarte Goncalves. Sequential sampling and equilibrium. *Working paper*, 2020. 8.4

Serena Guarnaschelli, Richard McKelvey, and Thomas Palfrey. An experimental study of jury decision rules. *American Political Science Review*, 2000. 10

Philip Haile, Ali Hortacsu, and Grigory Kosenok. On the empirical content of quantal response equilibrium. *American Economic Review*, 2008. (document), 2, 5

John Harsanyi and Reinhard Selten. A general theory of equilibrium selection in games. *Massachusetts Institute of Technology*, 1988. 3.1.3

Robert Luce. Individual choice behavior. *Wiley*, 1959. (document), 8.1

Richard McKelvey and Thomas Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 1995. (document), 3, 1.6

Richard McKelvey, Thomas Palfrey, and Roberto Weber. Endogenous rationality equilibrium. *Working Paper*, 1997. 8.2

Richard McKelvey, Thomas Palfrey, and Roberto Weber. The effects of payoff magnitude and heterogeneity on behavior in 2x2 games with unique mixed strategy equilibria. *Journal of Economic Behavior and Organization*, 2000. (document), 14, 7, 7, 8.2, 44

Emerson Melo, Kiril Pogorelskiy, and Matthew Shum. Testing the quantal response hypothesis. *International Economic Review*, 2018. (document), 12, 6, 6, 6, 9.1

Rosemarie Nagel. Unraveling in guessing games: an experimental study. *American Economic Review*, 1995. 10.1.2

Jack Ochs. An experimental study of games with unique mixed strategy equilibria. *Games and Economic Behavior*, 1995. 10.3.1, 44

Thomas Palfrey, Jeongbin Kim, and Jeffrey Zeidel. Games played by teams of players. *Working paper*, 2021. 8.4

Ariel Rubinstein and Martin Osborne. Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 2003. 8.4

Reinhard Selten and Thorsten Chmura. Stationary concepts for experimental 2x2 games. *American Economic Review*, 2008. (document), 7, 7, 8.2

Dale Stahl. Entropy control costs and entropic equilibrium. *International Journal of Game Theory*, 1990. 4

Dale Stahl and Paul Wilson. On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 1995. 10.1.2

Theodore Turocy. A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 2005. 9, 3.1.3, 8.3

Jorgen Weibull, Lars-Goran Mattsson, and Mark Voorneveld. Better may be worse: some monotonicity results and paradoxes in discrete choice under uncertainty. *Theory and Decision*, 2007. 4

Boyu Zhang and Josef Hofbauer. Quantal response methods for equilibrium selection in 2×2 coordination games. *Games and Economic Behavior*, 2016. 9, 3.1.3

9 Appendix

9.1 Omitted proofs

Proof of Proposition 1. We define $X(q^*, p^*, r)$ to be the set X for any 2×2 game as a function of parameters (q^*, p^*, r), and we use $|\cdot|$ to denote the measure of a set. Our proof proceeds by taking arbitrary parameters (q_0^*, p_0^*, r_0) and constructing (q_1^*, p_1^*, r_1), (q_2^*, p_2^*, r_2), and (q_3^*, p_3^*, r_3) such that $|X(q_0^*, p_0^*, r_0)| \leq |X(q_1^*, p_1^*, r_1)| \leq |X(q_2^*, p_2^*, r_2)| \leq |X(q_3^*, p_3^*, r_3)|$ and $|X(q_3^*, p_3^*, r_3)|$ is maximal. Let (q_0^*, p_0^*, r_0) be an arbitrary set of parameters such that X looks qualitatively like the top-left panel of Figure 11: $q_0^* \in (-\infty, \frac{1}{2}]$, $p_0^* \in (\frac{1}{2}, \infty)$, and r_0 is such that the downward-sloping branch of $U_0^+ \cap U_0^-$ (whichever of U_0^+ or U_0^- is downward-sloping) crosses vertical segment $\{(q,p) | q = \frac{1}{2}, p \in (\frac{1}{2}, 1)\}$. This is without loss.[38] Find some (q_1^*, p_1^*, r_1) such that $r_1 = r_0$ and (q_1^*, p_1^*) is to the northwest of (q_0^*, p_0^*) on the original downward-sloping branch of $U_0^+ \cap U_0^-$ such that the upward sloping branch of $U_0^+ \cap U_0^-$ no longer intersects with $\square$. The resulting picture looks like

[38]Either (q_0^*, p_0^*, r_0) satisfies these conditions for some relabelling of players and actions or (q_0^*, p_0^*, r_0) resembles one of (q_1^*, p_1^*, r_1), (q_2^*, p_2^*, r_2), or (q_3^*, p_3^*, r_3) that we construct, in which case there are fewer steps required.

the top-right panel. It is clear from the geometry that $|X(q_1^*, p_1^*, r_1)| > |X(q_0^*, p_0^*, r_0)|$ as the area by which X increases for each small change is proportional to the total length of the blue segments minus that of the red segment in the top-left panel. Next, find some (q_2^*, p_2^*, r_2) such that the downward sloping branch of $U_0^+ \cap U_0^-$ rotates counter-clockwise about where it crosses vertical segment $\{(q,p)|q = \frac{1}{2}, p \in (\frac{1}{2}, 1)\}$ to the point where it is approximately horizontal (and the upward sloping branch of $U_0^+ \cap U_0^-$ continues not to intersect with $\square$). The resulting picture looks like the bottom-left panel. It is clear from the geometry that $|X(q_2^*, p_2^*, r_2)| > |X(q_1^*, p_1^*, r_1)|$ as the area by which X increases for each small change is proportional to the total length of the blue segments minus that of the red segments in the top-right panel. Since we started from arbitrary (q_0^*, p_0^*, r_0) that involves a downward-sloping branch of $U_0^+ \cap U_0^-$ and found an improvement associated with a downward-sloping branch of $U_0^+ \cap U_0^-$ that is approximately horizontal, this implies that it is optimal for the downward-sloping branch of $U_0^+ \cap U_0^-$ to be approximately horizontal. It remains to find the optimal level of this branch. It is easy to show that this is obtained when the level is exactly $\frac{3}{4}$. To see this, consider the picture in the bottom-left panel where the level is less than $\frac{3}{4}$. By increasing the level by a small amount, the increase in the area of X is proportional to the total length of the blue segments minus that of the red segment in the bottom-left panel; and only at the level of $\frac{3}{4}$ does the length of the red segment begin to exceed that of the blue segments. Calling the resulting parameters (q_3^*, p_3^*, r_3), we have shown that $|X(q_3^*, p_3^*, r_3)|$ is maximal. We plot $X(q_3^*, p_3^*, r_3)$ in the bottom-right panel, and it is easy to show that $|X(q_3^*, p_3^*, r_3)| = \frac{5}{8}$. $\square$

Proof of Proposition 2. (i): Referencing the left panel of Figure 4, if $\frac{c_2}{B} < \left(\frac{1}{2}\right)^{N-1}$, then $f_2(\sigma_2; c_2)$ crosses the upward-sloping diagonal at some $\sigma_2 = \sigma_2^{'} \in (\frac{1}{2}, 1)$ (or passes over the unit-square entirely). This means that R^2 is non-degenerate and involves some area below the diagonal for which $\sigma_2 > \frac{1}{2}$, which implies the existence of RQRE for which $\sigma_2 > \sigma_1$. *(ii)*: For any (σ_2, σ_1), the low-cost type has a larger net-benefit to entering. Hence, by symmetry, the low-cost type must enter more often than the high-cost type in any SQRE. $\square$

Proof of Theorem 3. It is straightforward to show that the quantal response function Q_i defined by $Q_{ij}(\bar{u}_i) = \int_{\{\varepsilon_i \in \mathbb{R}^i | \bar{u}_{ij} + \varepsilon_{ij} \geq \bar{u}_{ik} + \varepsilon_{ik} \ \forall k \in \{1,...,J(i)\}\}} dF_i(\varepsilon_i)$ satisfies (R1)-(R4) under the assumptions on the error distributions. Hence, it must be that $S^* \subset R$. To show that $S^* \subset CM^*$, we adapt the result of Melo et al. [2018]. They consider the same model without our two additional assumptions: (1) if $J(i) = J(j) = J$, then $F_i = F_j$ for some ordering of players' actions, i.e. there exists a bijection $\iota : \{1, ..., J\} \to \{1, ..., J\}$

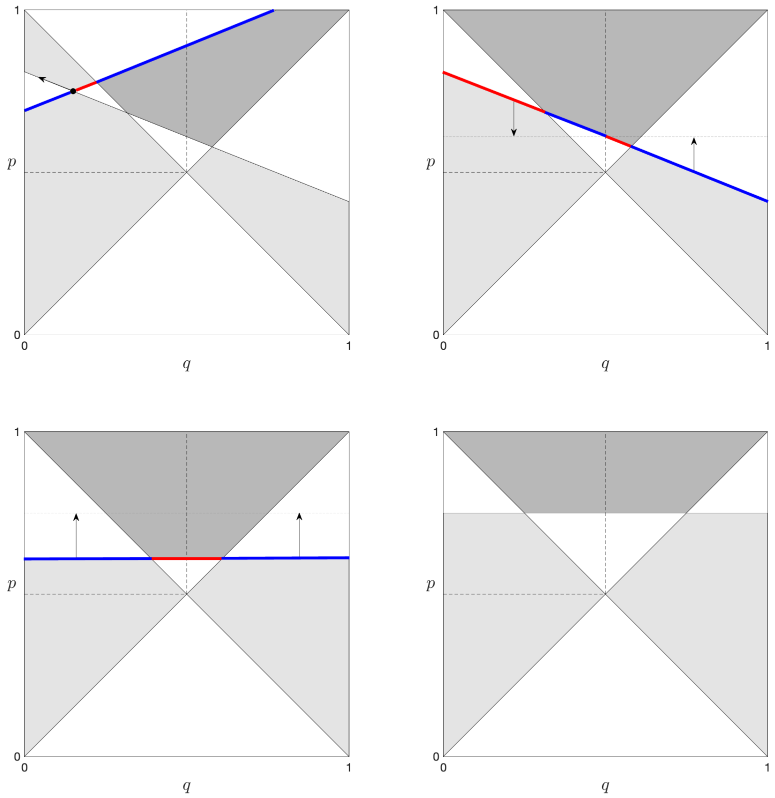

Figure 11: *Proof of Proposition 1.* This figure plots the objects used in the proof of Proposition 1: $X(q_0^*, p_0^*, r_0)$ (top-left panel), $X(q_1^*, p_1^*, r_1)$ (top-right panel), $X(q_2^*, p_2^*, r_2)$ (bottom-left panel), and $X(q_3^*, p_3^*, r_3)$ (bottom-right panel).

such that, for every realization ε_i, $F_i(\varepsilon_{i1}, ..., \varepsilon_{iJ}) = F_j(\varepsilon_{i\iota(1)}, ..., \varepsilon_{i\iota(J)})$; and (2) F_i is exchangeable. For a player with J actions, their result is as follows: it is necessary and sufficient that, for any cycle of *games* $\{G^0, G^1, ..., G^{L-1}, G^L\}$ where $G^L = G^0$ and $L \geq 2$, $\sum_{m=0}^{L-1} \sum_{j=1}^{J} (u_j^{m+1} - u_j^m)\sigma_j^m \leq 0$, where u_j^m and σ_j^m are the payoff to action j in game G^m and the corresponding probability this action is played, respectively. With the additional structure of (1), there exists a set of action-orderings, one ordering for each player, so that we may treat all players with J actions *within a game* as the same player. Thus, under the action-orderings that make players comparable, it is necessary that, for any cycle of *players* with J actions $\{P^0, P^1, ..., P^{L-1}, P^L\}$ where $P^L = P^0$ and $L \geq 2$, the above inequality holds where we reinterpret u_j^m as the payoff to action j of player m under the given ordering. Since we impose exchangeability (2), we must consider cycles with all possible action-orderings both within and across players with J actions. Doing this for all J and taking intersections gives rise to the inequalities in CM^*. $\square$

Proof of Proposition 3. We first show that $S^* = S$. It is obvious that SSQRE satisfies (R1)-(R4) and (S1)-(S3), and thus $S^* \subset S$. To show that $S \subset S^*$, suppose that $\sigma \in S$. We have (along the lines of the proof of Theorem 1) that $\delta(\bar{u}_{i_1}) \geq \delta(\bar{u}_{i_2}) \geq ... \geq \delta(\bar{u}_{i_N})$ and $max(\sigma_{i_1 1}, 1-\sigma_{i_1 1}) \geq max(\sigma_{i_2 1}, 1-\sigma_{i_2 1}) \geq ... \geq max(\sigma_{i_N 1}, 1-\sigma_{i_N 1})$ for some ordering of players $i_1, ..., i_N$ where $\delta(\bar{u}_{i_k}) > (=)\ \delta(\bar{u}_{i_{k+1}})$ if and only if $max(\sigma_{i_k 1}, 1-\sigma_{i_k 1}) > (=)\ max(\sigma_{i_{k+1} 1}, 1-\sigma_{i_{k+1} 1})$. Further, we have that $max(\sigma_{i_N 1}, 1-\sigma_{i_N 1}) \geq \frac{1}{2}$. We need to show that σ can be supported via a quantal response function based on structural errors with symmetry restrictions. To this end, if player i's quantal response function is derived from errors $\varepsilon_i = (\varepsilon_{i1}, \varepsilon_{i2})$, the probability she takes action 1 given payoff difference $\bar{u}_{i1} - \bar{u}_{i2}$ is given by $\tilde{F}(\bar{u}_{i1} - \bar{u}_{i2})$ where $\tilde{F} : (-\infty, \infty) \to (0, 1)$ is the CDF of difference $\varepsilon_{i2} - \varepsilon_{i1}$. Given that the distribution of ε_i must be full-support, absolutely continuous, mean zero, and exchangeable, the only restrictions on $\tilde{F} : (-\infty, \infty) \to (0, 1)$ are that: $\tilde{F}$ is continuous and strictly increasing, $\tilde{F}(0) = \frac{1}{2}$, and $\tilde{F}(x) = 1 - \tilde{F}(-x)$ for all x. Because $\delta(\bar{u}_{i_1}) \geq \delta(\bar{u}_{i_2}) \geq ... \geq \delta(\bar{u}_{i_N})$ and $max(\sigma_{i_1 1}, 1 - \sigma_{i_1 1}) \geq max(\sigma_{i_2 1}, 1 - \sigma_{i_2 1}) \geq ... \geq max(\sigma_{i_N 1}, 1 - \sigma_{i_N 1}) \geq \frac{1}{2}$, there exists such an $\tilde{F}$ that also satisfies $\tilde{F}(\delta(\bar{u}_{i_k})) = max(\sigma_{i_k 1}, 1 - \sigma_{i_k 1})$ for all $k = 1, ..., N$. Set $Q_{i1}(\bar{u}_{i1}, \bar{u}_{i2}) = \tilde{F}(\bar{u}_{i1} - \bar{u}_{i2})\mathbf{1}_{\{\bar{u}_{i1} \geq \bar{u}_{i2}\}} + \Big(1 - \tilde{F}(\bar{u}_{i2} - \bar{u}_{i1})\Big)\mathbf{1}_{\{\bar{u}_{i1} < \bar{u}_{i2}\}}$, and set player i's quantal response function as $Q_i = (Q_{i1}, 1 - Q_{i1}) : \mathbb{R}^2 \to \Delta$. By construction, if the quantal response function is given by $Q = (Q_1, ..., Q_N)$, then σ is a fixed point of $Q \circ \bar{u}$, and hence $S \subset S^*$.

Having established that $S^* = S$, we next show that $CM^* \cap R = X \cap R$, and thus $S^* = CM^* \cap R = S = X \cap R$ (as $S = X \cap R$ by Theorem 1). Following the argument in Footnote 11 of Melo et al. [2018], for $\Gamma^{N\times 2}$, we need only consider cycles of length 2 for determining if $\sigma \in CM^*$, and thus $CM^* = \{\sigma \in \Delta | u_1^m(\sigma) - u_2^m(\sigma) > (\geq) u_1^n(\sigma) - u_2^n(\sigma) \iff \sigma_1^m > (=)\sigma_1^n\ \forall \text{ action orderings } m, n\}$. Intersecting with R sharpens this to: $CM^* \cap R = \{\sigma \in \Delta | u_1^m(\sigma) - u_2^m(\sigma) > (=) u_1^n(\sigma) - u_2^n(\sigma) \iff \sigma_1^m > (=)\sigma_1^n\ \forall \text{ action orderings } m, n\}$. Notice that, when m and n correspond to different orderings of the same player, $u_1^m(\sigma) - u_2^m(\sigma) > (=) u_1^n(\sigma) - u_2^n(\sigma) \iff \sigma_1^m > (=)\sigma_1^n$ is equivalent to $\bar{u}_{i1}(\sigma) - \bar{u}_{i2}(\sigma) > (=) 0 \iff \sigma_{i1} > (=)\frac{1}{2}$. Given that this is true, when m and n correspond to different players, $u_1^m(\sigma) - u_2^m(\sigma) > (=) u_1^n(\sigma) - u_2^n(\sigma) \iff \sigma_1^m > (=)\sigma_1^n$ is equivalent to $|\bar{u}_{i1}(\sigma) - \bar{u}_{i2}(\sigma)| > (=)\ |\bar{u}_{j1}(\sigma) - \bar{u}_{j2}(\sigma)| \iff max(\sigma_{i1}, 1 - \sigma_{i1}) > (=)\ max(\sigma_{j1}, 1 - \sigma_{j1})$. Using notation $\delta(\bar{u}_i) = |\bar{u}_{i1}(\sigma) - \bar{u}_{i2}(\sigma)|$, we thus have that

$$\begin{aligned} CM^* \cap R = \{\sigma \in \Delta | &\bar{u}_{i1}(\sigma) - \bar{u}_{i2}(\sigma) > (=) 0 \iff \sigma_{i1} > (=)\frac{1}{2} \text{ and} \\ &\delta(\bar{u}_i) > (=)\ \delta(\bar{u}_j) \iff max(\sigma_{i1}, 1 - \sigma_{i1}) > (=)\ max(\sigma_{j1}, 1 - \sigma_{j1})\}, \end{aligned}$$

which equals $X \cap R$ by the fact that $\bar{u}_{i1}(\sigma) - \bar{u}_{i2}(\sigma) > (=) 0 \iff \sigma_{i1} > (=) \frac{1}{2}$ in any RQRE and the definition of X. $\square$

10 Supplemental Material

10.1 Additional 2×2 games

10.1.1 Prisoner's dilemma

Prisoner's dilemma is any 2×2 game in which both players have a strictly dominant action. The case in which it is efficient, by total payoffs, for each player to take her dominat*ed* action has been the focus of much theoretical and experimental work. Our analysis, however, does not depend on this feature. The only restriction on (q^*, p^*) is that $q^* \notin [0, 1]$ and $p^* \notin [0, 1]$, and so (q^*, p^*) must fall outside of the unit square.

The (closure of the) set of RQRE is given by $R = [\frac{1}{2}, 1] \times [\frac{1}{2}, 1]$, which is simply the north-east quarter of the unit square and corresponds to each player taking her dominant action more often than not.

Figure 12 plots R, S, and L for two examples of prisoner's dilemma. In the left panel, U_0 intersects E_0 within R, which implies that, for each player, there are SQRE in which she faces higher stakes and is the more extreme. In the right panel, R resides entirely within the south portion of U_2, meaning that it is always player 2 who faces higher stakes and is the more extreme.

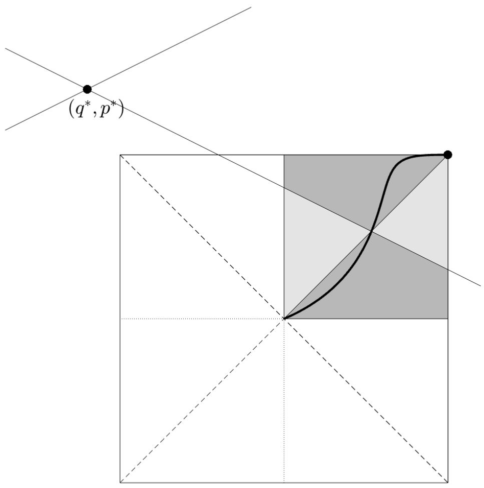

Figure 12: *Prisoner's dilemma.*

10.1.2 Two-step dominance solvable

A two-step dominance solvable 2×2 game is any in which one player has a strictly dominant action and the other player has undominated actions. Such games are interesting because for only one player does her optimal action depend on beliefs over the rationality of the opponent. Hence, such games have been used to test models such as level k (Nagel [1995] and Stahl and Wilson [1995]). We assume, without loss, that U is strictly dominant for player 1 and R is player 2's strict best response to U so that the unique NE is $(q,p) = (0,1)$. The only restriction on (q^*, p^*) is that $p^* \in (0,1)$ and q^* is bounded away from $(0,1)$. Hence, (q^*, p^*) must again fall outside of the unit square.

The (closure of the) set of RQRE is given by $R = R^1 \cup R^2$, where $R^1 = [0, \frac{1}{2}] \times [max\{p^*, \frac{1}{2}\}, 1]$ and $R^2 = [\frac{1}{2}, 1] \times [\frac{1}{2}, p^*]$ are two rectangular components.[39]

Figure 13 plots R, S, and L for two examples of dominance solvable games. As with matching pennies, we see that S can be the union of up to four polytopes. As with the prisoner's dilemma, depending on payoffs, it may be that each player has SQRE in which she faces higher stakes and is more extreme; or it may be that all SQRE involve the same player facing higher stakes and being more extreme (though we do not depict such a case here).

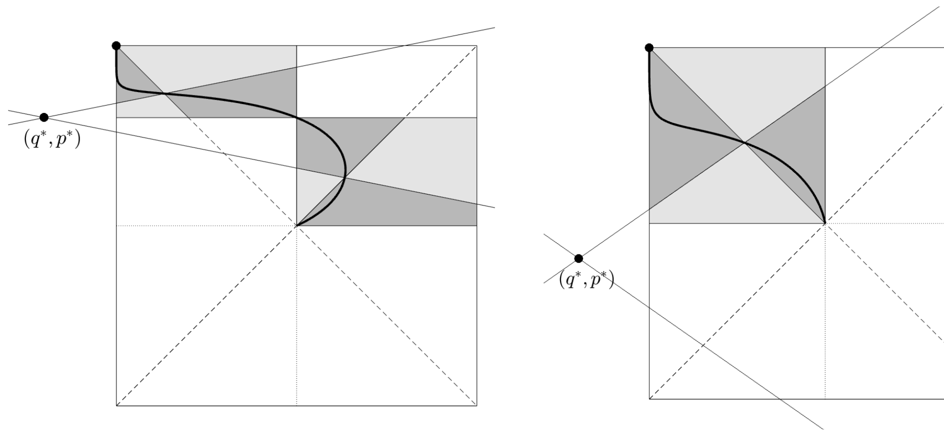

Figure 13: *Two-step dominance solvable.*

10.2 Additional applications

10.2.1 Jury voting with unanimity

Following Feddersen and Pesendorfer [1998], there are N jurors, who simultaneously vote whether to convict or acquit a defendant. If *all* jurors vote to convict, the defendant is convicted. Otherwise, if *any* juror votes to acquit, the defendant is acquitted. With

[39]If $p^* \leq \frac{1}{2}$, the second component of R is degenerate: $R^2 = \emptyset$.

equal probability, the defendant is guilty (G) or innocent (I). Jurors prefer to convict the defendant if guilty and acquit if innocent, with the following payoffs:

$$\begin{aligned} u(A,I) &= u(C,G) = 0 \\ u(C,I) &= -\mu \\ u(A,G) &= -(1-\mu), \end{aligned}$$

where A and C are the outcomes "acquit" and "convict", and $\mu \in (0,1)$ is parameter indexing how bad it is to convict the innocent relative to acquitting the guilty.

Prior to voting, each juror receives a signal that is drawn independently, conditional on the state. In state G, the signal is γ with probability $\pi > max\{\frac{1}{2},\mu\}$ and ι with probability $1-\pi$. In state I, the signal is ι with probability π and γ with probability $1-\pi$. Hence, γ signals "guilty", ι signals "innocent", and the signal structure is symmetric.

We search for semi-symmetric SQRE in which all players of the same type have the same behavior, and so we characterize the semi-symmetric part of S. To this end, we let σ_γ and σ_ι denote the probabilities with which γ- and ι-jurors, respectively, vote to convict.

The game is standard, so we give, without derivation, the expected payoff differences from voting to convict versus voting to acquit for each type. For γ-jurors, the expected payoff difference is $(1-\mu)\pi(\pi\sigma_\gamma+(1-\pi)\sigma_\iota)^{N-1}-\mu(1-\pi)((1-\pi)\sigma_\gamma+\pi\sigma_\iota)^{N-1}$ and, for ι-jurors, the expected payoff difference is $(1-\mu)(1-\pi)(\pi\sigma_\gamma+(1-\pi)\sigma_\iota)^{N-1}-\mu\pi((1-\pi)\sigma_\gamma+\pi\sigma_\iota)^{N-1}$. Hence, voting to convict yields a higher payoff for γ- and ι-jurors, respectively, if

$$\sigma_\gamma \geq \underbrace{\frac{(\pi(1-\mu))^{1/(N-1)}(1-\pi)-((1-\pi)\mu)^{1/(N-1)}\pi}{((1-\pi)\mu)^{1/(N-1)}(1-\pi)-(\pi(1-\pi))^{1/(N-1)}\pi}\cdot\sigma_\iota}_{\equiv f_\gamma(\sigma_\iota)}$$

and

$$\sigma_\gamma \geq \underbrace{\frac{((1-\pi)(1-\mu))^{1/(N-1)}(1-\pi)-(\pi\mu)^{1/(N-1)}\pi}{(\pi\mu)^{1/(N-1)}(1-\pi)-((1-\pi)(1-\mu))^{1/(N-1)}\pi}\cdot\sigma_\iota}_{\equiv f_\iota(\sigma_\iota)}.$$

In other words, it is optimal for γ-jurors to vote to convict if σ_γ is greater than an upward-sloping, linear function of σ_ι, which for convenience we have defined as $f_\gamma(\sigma_\iota)$. Similarly, it is optimal ι-jurors to vote to convict if σ_γ is greater than the upward-sloping, linear function $f_\iota(\sigma_\iota)$. We see immediately that both $f_\gamma(\sigma_\iota)$ and $f_\iota(\sigma_\iota)$ have a zero-intercept, meaning they are rays from the origin, and $f_\iota(\sigma_\iota) > f_\gamma(\sigma_\iota)$ for all $\sigma_\iota > 0$, meaning that whenever voting to convict is optimal for ι-jurors, it is also optimal for γ-jurors.

Using the above inequalities, R is found by collecting regions of the unit square that are consistent with best responding more often than not: $R = \{(\sigma_\iota, \sigma_\gamma) | \sigma_t > \frac{1}{2} \iff \sigma_t > f_t(\sigma_\iota) \text{ for } t = \iota, \gamma\}$. The (closure of) set R can also be written explicitly as $R = R^1 \cup R^2$ where $R^1 = \{(\sigma_\iota, \sigma_\gamma) : \sigma_\iota \in [\{\sigma_\iota' | f_\iota(\sigma_\iota') = \frac{1}{2}\}, \frac{1}{2}], \sigma_\gamma \in [\frac{1}{2}, min\{f_\iota(\sigma_\iota), 1\}]\}$ and $R^2 = \{(\sigma_\iota, \sigma_\gamma) : \sigma_\iota \in [\frac{1}{2}, \{\sigma_\iota' | f_\iota(\sigma_\iota') = 1\}], \sigma_\gamma \in [f_\iota(\sigma_\iota), 1]\}$ are two non-rectangular components.[40]

To complete the characterization of S, we need to find expressions for U_0, U_γ, and U_ι whereby one type of player or the other faces higher stakes. U_0, where both types face the same stakes, is defined by $(\sigma_\iota, \sigma_\gamma)$ that satisfy

$$\begin{aligned} &|(1-\mu)\pi(\pi\sigma_\gamma + (1-\pi)\sigma_\iota)^{N-1} - \mu(1-\pi)((1-\pi)\sigma_\gamma + \pi\sigma_\iota)^{N-1}| = \\ &|(1-\mu)(1-\pi)(\pi\sigma_\gamma + (1-\pi)\sigma_\iota)^{N-1} - \mu\pi((1-\pi)\sigma_\gamma + \pi\sigma_\iota)^{N-1}| \end{aligned}$$

Inspecting the left and right hand sides, we observe that this holds only when

$$\begin{aligned} -\mu(1-\pi)((1-\pi)\sigma_\gamma + \pi\sigma_\iota)^{N-1} + \pi(1-\mu)(\pi\sigma_\gamma + (1-\pi)\sigma_\iota)^{N-1} &= \\ \mu\pi((1-\pi)\sigma_\gamma + \pi\sigma_\iota)^{N-1} - (1-\pi)(1-\mu)(\pi\sigma_\gamma + (1-\pi)\sigma_\iota)^{N-1} & \quad \iff \end{aligned}$$

$$\sigma_\gamma = \underbrace{\frac{\mu^{1/(N-1)}\pi - (1-\mu)^{1/(N-1)}(1-\pi)}{(1-\mu)^{1/(N-1)}\pi - \mu^{1/(N-1)}(1-\pi)} \cdot \sigma_\iota}_{\equiv g(\sigma_\iota)}$$

where for convenience we have defined $g(\sigma_\iota)$. We note that $g(\sigma_\iota)$ is an upward sloping ray from the origin that falls between $f_\gamma(\sigma_\iota)$ and $f_\iota(\sigma_\iota)$: $f_\gamma(\sigma_\iota) < g(\sigma_\iota) < f_\iota(\sigma_\iota)$ for all $\sigma_\iota > 0$. It is now easy to show that

$$\begin{aligned} U_0 &= \{(\sigma_\iota, \sigma_\gamma) : \sigma_\gamma = g(\sigma_\iota)\}, \\ U_\gamma &= \{(\sigma_\iota, \sigma_\gamma) : \sigma_\gamma > g(\sigma_\iota)\}, \text{ and} \\ U_\iota &= \{(\sigma_\iota, \sigma_\gamma) : \sigma_\gamma < g(\sigma_\iota)\}. \end{aligned}$$

Hence, U_0 is the line defined by $g(\sigma_\iota)$, U_γ is the half-space above it, and U_ι is the

[40] If $f_\iota(\frac{1}{2}) \geq 1$, then the second component is degenerate: $R^2 = \emptyset$.

half-space below it.

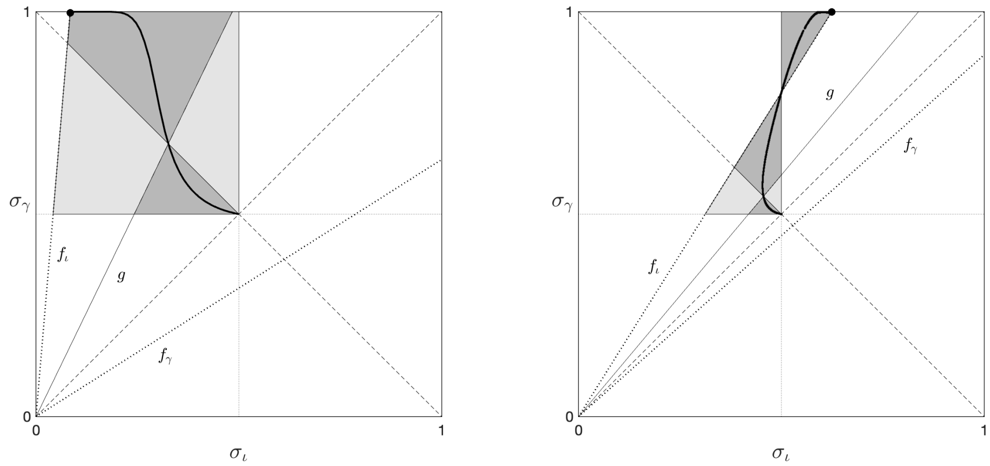

Figure 14: *Jury voting.* Left panel: $\pi = 0.8$, $\mu = 0.7$, and $N = 3$. Right panel: $\pi = 0.8$, $\mu = 0.7$, and $N = 9$.

With E_0, E_γ, and E_ι defined as in (3), but grouping players of the same type, we characterize S as in Corollary 1. Figure 4 plots R, S, and L for two parameterizations of the game (parameters given in the figure). Again, we see that S is given as the union of polytopes, that it significantly refines R, and that it imposes significant bounds on L–including giving some LQRE exactly as the solution to linear systems.

For all parameters, all RQRE and SQRE are such that γ-jurors are more likely to vote for conviction than are ι-jurors. In the left panel, parameters are such that γ-jurors tend to convict whereas ι-jurors tend to acquit, and so voters tend to vote with their signal. More interestingly, in the right panel, parameters are such that there are some QRE for which both types vote to convict more often than not. We find this interesting because, in the semi-symmetric NE, ι-jurors are indifferent and vote to convict at high rates simply to support the equilibrium. For nearby QRE, however, the ι-jurors actually strictly prefer voting to convict, that is, to vote against their signal.

10.2.2 Infinitely repeated prisoner's dilemma

We consider the infinitely repeated prisoner's dilemma (PD) whose stage game is given in Table 6. The stage game is symmetric, with each of two players having the opportunity to cooperate (C) or defect (D). Defection is a dominant strategy in the stage game.

For simplicity, we consider the case that the difference in stage game payoffs between defecting and cooperating is $d > 0$, a constant that does not depend on the opponent's action. The stage game captures all such PD games up to affine transformations of payoffs, which does not affect the set of SQRE. As is standard, we assume super game payoffs are the discounted sum of stage game payoffs with discount factor $\delta \in [0, 1)$.

	D	**C**
D	$0, 0$	$1 + d, -d$
C	$-d, 1 + d$	$1, 1$

Table 6: *Stage game of the infinitely repeated prisoner's dilemma.*

For a similar game, Goeree et al. [2016] (chapter 5) numerically approximate Markov-perfect LQRE whereby behavior is measurable with respect to two states. Closely following their approach, we explicitly find the entire set of Markov-perfect SQRE for the same specification of states.

The first state is the "cooperative" state defined by any history in which no player has defected. The second is the "non-cooperative" state defined by any history in which at least one player has defected previously. Let σ_1 and σ_2 denote the probabilities of defecting in the cooperative and non-cooperative states, respectively. The continuation value in the non-cooperative state is given by $V_2 = 0\sigma_2^2 + (1 + d)\sigma_2(1 - \sigma_2) + (-d)(1 - \sigma_2)\sigma_2 + (1 - \sigma_2)^2 + \delta V_2 = 1 - \sigma_2 + \delta V_2$ and so

$$V_2 = \frac{1 - \sigma_2}{1 - \delta}.$$

The continuation value in the cooperative state is given by $V_1 = 1 - \sigma_1 + \delta(1 - \sigma_1)^2 V_1 + \delta[1 - (1 - \sigma_1)^2]V_2$ and so

$$V_1 = \frac{(1 - \delta)(1 - \sigma_1) + \delta(1 - (1 - \sigma_1)^2)(1 - \sigma_2)}{(1 - \delta)(1 - \delta(1 - \sigma_1)^2)}.$$

In the non-cooperative state, the payoffs from defecting and cooperating are $(1 + d)(1 - \sigma_2) + \delta V_2$ and $-d(\sigma_2) + (1 - \sigma_2) + \delta V_2$, respectively. Hence the difference is $d > 0$, meaning it is always payoff maximizing to defect. This is intuitive because the non-cooperative state is absorbing. Hence, the opponent's future behavior will not be affected by the choice of action, and since defecting is dominant in the stage game, it is optimal here as well.

In the cooperative state, the payoffs from defecting and cooperating are $(1 + d)(1 -$

$\sigma_1)+\delta V_2$ and $-d(\sigma_1)+(1-\sigma_1)+\delta(V_2\sigma_1+V_1(1-\sigma_1))$, respectively. Hence, the difference is $d-\delta(1-\sigma_1)(V_1-V_2)$. Substituting expressions for V_1 and V_2, this difference becomes:

$$d - \frac{\delta(1-\sigma_1)(\sigma_2-\sigma_1)}{1-\delta(1-\sigma_1)^2}.$$

Hence, it may be optimal to cooperate, depending on the discount factor and cooperation rates in both phases. Inspecting the above expression, defection in the cooperation phase is optimal if and only if

$$\sigma_2 \leq \underbrace{\frac{d(1-\delta(1-\sigma_1)^2)}{\delta(1-\sigma_1)} + \sigma_1}_{=f(\sigma_1)},$$

where we have defined $f(\sigma_1)$ for convenience.

Using the above inequality and the fact that it is always optimal to defect in the non-cooperative state, R is found by collecting regions of the unit square that are consistent with best responding more often than not: $R = \{(\sigma_1,\sigma_2)|\sigma_1 > \frac{1}{2} \iff \sigma_2 < f(\sigma_1),\ \sigma_2 > \frac{1}{2}\}$. The (closure of) set R can also be written explicitly as $R = R^1 \cup R^2$ where $R^1 = \{(\sigma_1,\sigma_2) : \sigma_1 \in [0, min\{\frac{1}{2}, \{\sigma_1^{'}|f(\sigma_1^{'}) = 1\}\}], \sigma_2 \in [max\{\frac{1}{2}, f(\sigma_1)\}, 1]\}$ and $R^2 = \{(\sigma_1,\sigma_2) : \sigma_1 \in [\frac{1}{2},1], \sigma_2 \in [\frac{1}{2}, min\{f(\sigma_1),1\}]\}$ are two non-rectangular components.[41]

To complete the characterization of S, we need to find expressions for U_0, U_1, and U_2 whereby players in one state or the other face higher stakes. U_0, whereby players in both states face the same stakes, is defined by (σ_1,σ_2) that satisfy

$$|d - \frac{\delta(1-\sigma_1)(\sigma_2-\sigma_1)}{1-\delta(1-\sigma_1)^2}| = |d|.$$

Since $\frac{\delta(1-\sigma_1)(\sigma_2-\sigma_1)}{1-\delta(1-\sigma_1)^2} \geq 0$, this holds if

$$\frac{\delta(1-\sigma_1)(\sigma_2-\sigma_1)}{1-\delta(1-\sigma_1)^2} - d = d \iff$$

$$\sigma_2 = \underbrace{\frac{2d(1-\delta(1-\sigma_1)^2)}{\delta(1-\sigma_1)} + \sigma_1}_{=g(\sigma_1)},$$

[41] If $f(0) > 1 \iff \delta < \frac{d}{d+1}$, then the first component is degenerate: $R^1 = \emptyset$.

where we have defined $g(\sigma_1)$ for convenience. Noting also that $\frac{\delta(1-\sigma_1)(\sigma_2-\sigma_1)}{1-\delta(1-\sigma_1)^2} = 0$ whenever $\sigma_2 = \sigma_1$ or $\sigma_1 = 1$, we have that

$$U_0 = \{(\sigma_1, \sigma_2) : \sigma_2 = g(\sigma_1) \text{ or } \sigma_2 = \sigma_1 \text{ or } \sigma_1 = 1\}.$$

It is not hard to show that

$$U_1 = \{(\sigma_1, \sigma_2) : \sigma_1 \in [0, 1), \sigma_2 \in [0, \sigma_1) \cup (g(\sigma_1), 1)\}, \text{ and}$$
$$U_2 = \{(\sigma_1, \sigma_2) : \sigma_1 \in [0, 1), \sigma_2 \in (\sigma_1, g(\sigma_1))\}.$$

Hence, U_0 is the union of the line defined by $g(\sigma_1)$, the upward-slowing diagonal, and the vertical line such that $\sigma_1 = 1$; U_2 is the region between the upward-sloping diagonal and $g(\sigma_1)$; and U_1 is what's left over.

With E_0, E_1, and E_2 defined as in (3), but grouping players by state, we characterize S as in Corollary 1. Figure 15 plots R, S, and L for $d = 0.2$ and two different discount factors, $\delta = 0.6$ on the left and $\delta = 0.9$ on the right. We see that S is given as the union of several regions (in this case with non-linear boundaries), that it significantly refines R, and that it imposes significant bounds on L–including giving some LQRE exactly as the solution to polynomial systems that can be found in closed form.

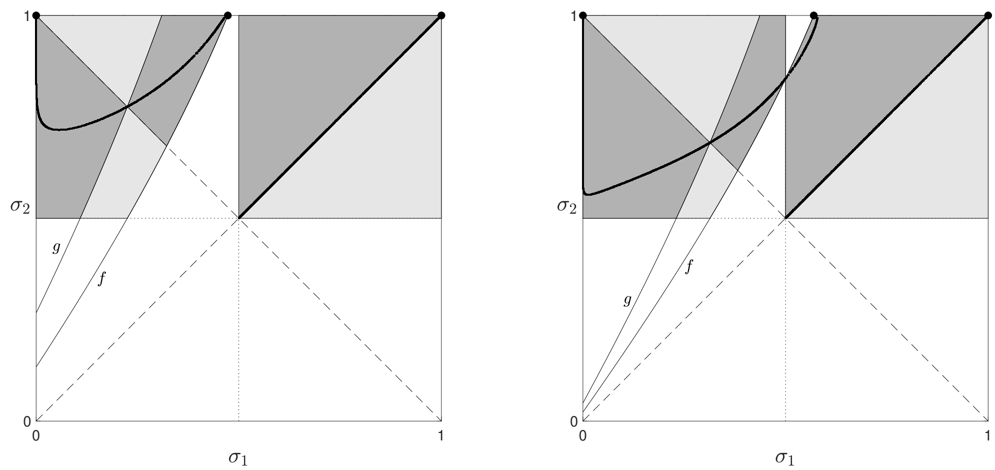

Figure 15: *Infinitely repeated prisoner's dilemma.* Left panel: $d = 0.2$ and $\delta = 0.6$. Right panel: $d = 0.2$ and $\delta = 0.9$.

As long as the discount factor is not too low, i.e. if $\delta > \frac{d}{d+1}$, as is the case in both panels of Figure 15, then there are three Markov-perfect NE of this game, all of which are limit points of Markov-perfect LQRE. These equilibria are $(\sigma_1^{*D}, \sigma_2^{*D}) = (1, 1)$ in which players defect in all states, $(\sigma_1^{*G}, \sigma_2^{*G}) = (0, 1)$ in which there is cooperation on-path supported by the "grim trigger" strategy of always defecting following any defection, and $(\sigma_1^{*M}, \sigma_2^{*M}) = (1 - \sqrt{1 - \frac{\delta(d+1)-d}{\delta(d+1)}}, 1)$ whereby players mix until any defection is realized which results in the grim trigger. If the discount factor is too low, i.e. if $\delta < \frac{d}{d+1}$, then the only Markov-perfect NE is $(\sigma_1^{*D}, \sigma_2^{*D}) = (1, 1)$ in which players defect in all states.[42]

For intermediate discount factors, i.e. $\delta \in (\frac{d}{d+1}, \frac{4d}{d+1})$, as in the left hand panel of Figure 15, only the NE $(\sigma_1^{*D}, \sigma_2^{*D}) = (1, 1)$ is path-connected via R to the centroid. Hence, this equilibrium is uniquely selected by LQRE and, indeed, by RQRE as well, making this a very robust prediction. For $\delta > \frac{4d}{d+1}$, as in the right panel, all three NE are path-connected to the centroid via the smaller set S. Hence, when the discount factor is sufficiently high, a lot of structure is needed to make a unique selection, which accords with "folk theorem"-type intuition that many types of behavior can be supported if players are sufficiently patient. Interestingly, for all parameters, LQRE always selects the non-cooperative NE $(\sigma_1^{*D}, \sigma_2^{*D}) = (1, 1)$.[43]

Finally, for all parameter values, there are RQRE in which players defect *more* often in the cooperative state. Such equilibria are clearly implausible and ruled out by symmetry as no SQRE involves this feature.

10.3 Extensions

10.3.1 Sets of games

It is well known that QRE makes comparative static predictions with respect to changes in a game's payoff parameters, assuming that the quantal response function is held fixed. For example, Goeree et al. [2003] show that RQRE predicts the so-called "own-payoff effect" (Ochs [1995]). Put differently, QRE models impose restrictions on the data from

[42] This corresponds to the case that the first component of R is degenerate, i.e. $R^1 = \emptyset$; see Footnote 41.

[43] *Proof.* The result is proven if we can construct a path of LQRE indexed by $\lambda \in [0, \infty)$ that goes from $(\sigma_1, \sigma_2) = (\frac{1}{2}, \frac{1}{2})$ to $(\sigma_1, \sigma_2) = (1, 1)$. To this end, for any $\sigma' \in [\frac{1}{2}, 1)$, define $\bar{\lambda}(\sigma')$ to be the unique value of λ such that

$$\frac{e^{\lambda d}}{e^{\lambda d} + 1} = \sigma'.$$

The function $\bar{\lambda} : [\frac{1}{2}, 1) \to [0, \infty)$ is one-to-one and strictly increasing. If $\sigma_1 = \sigma_2$, then the expected payoff difference between defecting and cooperating in the cooperative phase is $d - \frac{\delta(1-\sigma_1)(\sigma_2-\sigma_1)}{1-\delta(1-\sigma_1)^2} = d > 0$, which equals the expected payoff difference in the non-cooperative phase. Hence, for any $\sigma' \in [\frac{1}{2}, 1)$, $(\sigma_1, \sigma_2) = (\sigma', \sigma')$ is an LQRE associated with $\lambda = \bar{\lambda}(\sigma')$, which defines our path and completes the proof. $\square$

sets of games that are stronger than the product of individual-game restrictions.

In this section, we show that Theorem 1 generalizes to *sets of* games in the natural way. To this end, we consider $\{\Gamma^{N_1\times 2}, \Gamma^{N_2\times 2}, ..., \Gamma^{N_M\times 2}\}$, a set of M binary-action games where the mth game $\Gamma^{N_m\times 2} = \{\mathcal{N}_m, S^m, u^m\}$ is defined by the set players $\mathcal{N}_m = \{1, ..., N_m\}$, action space S^m such that $|S_i^m| = 2$ for all $i \in \mathcal{N}_m$, and payoff function u^m. We also define for $\Gamma^{N_m\times 2}$ the vectors of expected payoffs $\bar{u}^m = (\bar{u}_1^m, \bar{u}_2^m, ..., \bar{u}_{N_m}^m)$ and mixed action space Δ^m with $\sigma^m = (\sigma_1^m, ..., \sigma_{N_m}^m) \in \Delta^m$ an arbitrary element.

We are interested in representing the following set, which is the set of multi-game action profiles that are consistent with SQRE for a *single* quantal response function held fixed across games:

$$\mathcal{S} = \{\times_m \sigma^m \in \times_m \Delta^m : Q \circ \bar{u}^m(\sigma^m) = \sigma^m \ \forall m \text{ for some } Q \text{ satisfying (R1)-(R4), (S1)-(S3)}\}.$$

To represent $\mathcal{S}$, we introduce objects analogous to those used in Theorem 1. First, using R^m to denote the set of RQRE in game m, the product of sets of RQRE from all games is given by

$$\mathcal{R} = \times_m R^m.$$

We emphasize that, in the definition of $\mathcal{R}$, unlike in the definition of $\mathcal{S}$, there is no restriction that the underlying quantal response functions be the same across games.

Next, we define an analogue of X that gives the set of multi-game action profiles such that players both *within* and *across* all games have the same ranking by extremeness and stakes.

$$\begin{aligned}\mathcal{X} =& \{\times_m \sigma^m \in \times_m \Delta^m : \delta(\bar{u}_i^{m'}) > (=)\ \delta(\bar{u}_j^{m''}) \iff \\ & \qquad max(\sigma_{i1}^{m'}, 1 - \sigma_{i1}^{m'}) > (=)\ max(\sigma_{j1}^{m''}, 1 - \sigma_{j1}^{m''})\}\end{aligned}$$

The representation theorem closely parallels Theorem 1 with $\mathcal{S}$ as the intersection of $\mathcal{R}$ and $\mathcal{X}$. The proof is nearly identical to that of Theorem 1 and thus omitted.

Theorem 4. *Fix* $\{\Gamma^{N_1\times 2}, \Gamma^{N_2\times 2}, ..., \Gamma^{N_M\times 2}\}$. $\mathcal{S} = \mathcal{X} \cap \mathcal{R}$.

Hence, a multi-game action profile is consistent with SQRE for a single quantal response function if and only if two conditions are satisfied. First, for each game *individually*, it must be that the associated action profile is an RQRE. Second, pooling *across games*, it must be that the ranking of players by extremeness is the same as their ranking by stakes.

		A				**B**	
		L	**R**			**L**	**R**
U		4,0	0,1	**U**		9,0	0,1
D		0,1	1,0	**D**		0,1	1,0

Table 7: *An example set of matching pennies games.*

To demonstrate Theorem 4, we consider the example games A and B of Table 7, which are matching pennies games that are the same up to a single payoff parameter for player 1.[44] The proposition below applies Theorem 4 to characterize $\mathcal{S}$.

Proposition 4. *In games A and B of Figure 7, $(q^A, p^A) \times (q^B, p^B) \in \mathcal{S}$ if and only if $(q^A, p^A) \in S^A$, $(q^B, p^B) \in S^B$, $p^B > p^A$, and $\frac{1}{2}q^A < q^B < q^A$.*

Proof. Suppose that $(q^A, p^A) \times (q^B, p^B) \in \mathcal{S}$. Goeree et al. [2005] show that, since behavior in each game is consistent with RQRE, $q^A > \frac{1}{5}$, $p^A > \frac{1}{2}$, $q^B > \frac{1}{10}$ and $p^B > \frac{1}{2}$. They also show that, since each player has the same regular quantal response function across games, $p^B > p^A$ and $q^B < q^A$. It must also be that $(q^A, p^A) \in S^A$ and $(q^B, p^B) \in S^B$ trivially by definition of $\mathcal{S}$. Suppose for contradiction that $\frac{1}{2}q^A \geq q^B$. In this case, we have that $\delta(\bar{u}_1^A(q^A)) = |5q^A - 1| \geq |10q^B - 1| = \delta(\bar{u}_1^B(q^B))$ (since $q^A > \frac{1}{5}$ and $q^B > \frac{1}{10}$) and $max(p^A, 1 - p^A) < max(p^B, 1 - p^B)$ (since $p^B > p^A$, $p^A > \frac{1}{2}$, and $p^B > \frac{1}{2}$), but this is inconsistent with player 1 using a single quantal response function in both games that satisfies the SQRE axioms, a contradiction. Hence, it must be that $\frac{1}{2}q^A < q^B$, and thus we have that $(q^A, p^A) \in S^A$, $(q^B, p^B) \in S^B$, $p^B > p^A$, and $\frac{1}{2}q^A < q^B < q^A$.

Conversely, suppose that $(q^A, p^A) \in S^A$, $(q^B, p^B) \in S^B$, $p^B > p^A$, and $\frac{1}{2}q^A < q^B < q^A$. Since $S^A \subset R^A$ and $S^B \subset R^B$, we have that $(q^A, p^A) \times (q^B, p^B) \in \mathcal{R}$. All that remains is to show that $(q^A, p^A) \times (q^B, p^B) \in \mathcal{X}$, i.e. that players both within and across games have the same ranking by stakes and extremeness, which would imply that $(q^A, p^A) \times (q^B, p^B) \in \mathcal{S}$ by Theorem 4. Since $(q^A, p^A) \in S^A$ and $(q^B, p^B) \in S^B$, players have the same ranking by stakes as extremeness *within-game* (Corollary 1), so we need only check that players have the same ranking by stakes as extremeness *across games*. Specifically, we need the

[44] Such games were played in the lab in Ochs [1995] and McKelvey et al. [2000] and are the canonical games for illustrating the "own-payoff effect".

following equivalences to hold:

$$
\begin{aligned}
max(p^A, 1-p^A) > (=)\; max(p^B, 1-p^B) &\iff \delta(\bar{u}_1^A(q^A)) > (=)\; \delta(\bar{u}_1^B(q^B)),\\
max(q^A, 1-q^A) > (=)\; max(q^B, 1-q^B) &\iff \delta(\bar{u}_2^A(p^A)) > (=)\; \delta(\bar{u}_2^B(p^B)),\\
max(p^A, 1-p^A) > (=)\; max(q^B, 1-q^B) &\iff \delta(\bar{u}_1^A(q^A)) > (=)\; \delta(\bar{u}_2^B(p^B)), \text{ and}\\
max(q^A, 1-q^A) > (=)\; max(p^B, 1-p^B) &\iff \delta(\bar{u}_2^A(p^A)) > (=)\; \delta(\bar{u}_1^B(q^B)).
\end{aligned}
$$

It easy to check that all of these equivalences do hold. For instance, the first equivalence holds because $max(p^A, 1-p^A) < max(p^B, 1-p^B)$ (since $p^B > p^A$, $p^A > \frac{1}{2}$, and $p^B > \frac{1}{2}$), and $\delta(\bar{u}_1^A(q^A)) < \delta(\bar{u}_1^B(q^B))$ (since $\frac{1}{2}q^A < q^B$, $q^A > \frac{1}{5}$, and $q^B > \frac{1}{10}$, which together imply that $\delta(\bar{u}_1^A(q^A)) = |5q^A - 1| < |10q^B - 1| = \delta(\bar{u}_1^B(q^B))$). The remaining three equivalences can be shown to hold in a similar fashion. $\square$

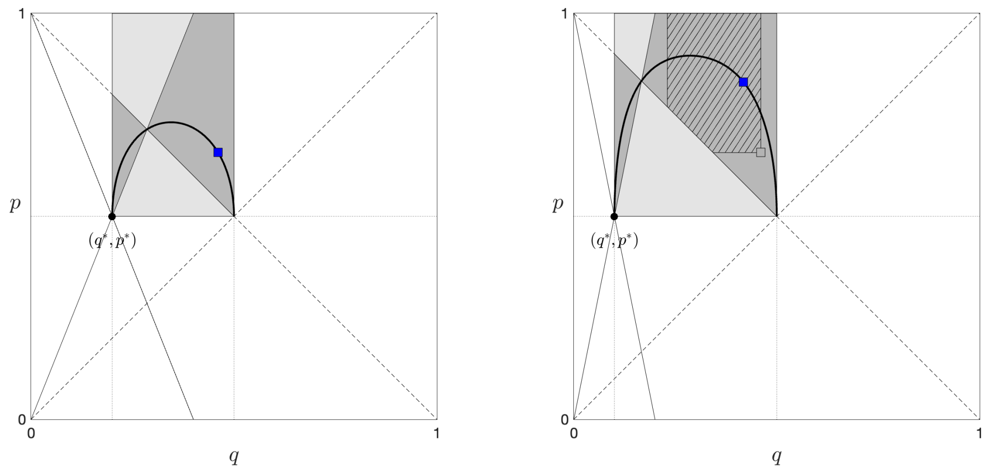

Figure 16: *SQRE in sets of games.*

Hence, the data from both games jointly is consistent with SQRE using a single quantal response function if and only if it is consistent with SQRE game-by-game, and, in addition, satisfies some additional across-game restrictions (i.e. comparative statics).

We illustrate the proposition in Figure 16. In the left panel, we plot the set of SQRE for game A, and in the right panel, we plot the set of SQRE for game B. We also plot the sets of LQRE as black curves, with the LQRE associated with $\lambda = 0.5$ highlighted as blue squares. If we consider the blue square in the left panel as an SQRE in game A (not necessarily an LQRE), then a point in the right panel would be an SQRE in game B consistent with the same quantal response function if and only if it is in the crosshatched region.

10.3.2 Alternatives to translation invariance

While the large majority of applications of QRE involve quantal response functions that satisfy *translation invariance*, some non-translation invariant models have been introduced. In particular, there is the generalized Luce model for strictly positive payoffs:

$$Q_{ij}(\bar{u}_i; \mu) = \frac{\bar{u}_{ij}^{(1/\mu)}}{\sum_{k=1}^{J(i)} \bar{u}_{ik}^{(1/\mu)}}, \tag{6}$$

where the parameter $\mu \in (0, \infty)$ governs the degree of noisy behavior and is analogous to LQRE λ. After LQRE, this is the most commonly used parametric form of QRE, and it has been shown to outperform LQRE in some situations.[45]

The generalized Luce model is *scale invariant* in the sense that scaling payoffs by some positive factor does not affect stochastic choice. Formally, we state this as an axiom:

(S2$'$) **Scale invariance**: $Q_i(\beta\bar{u}_i) = Q_i(\bar{u}_i)$ for any $\bar{u}_i \in \mathbb{R}_{++}^{J(i)}$ and $\beta > 0$.

Luce QRE satisfies the regularity axioms (R1)-(R4) as well as *player symmetry* (S1), *scale invariance* (S2$'$), and *label independence* (S3), but it does not satisfy *translation invariance* (S2). In fact, Friedman [2022] shows that no regular QRE model can satisfy both *scale invariance* and *translation invariance*; and so we consider a version of SQRE that imposes *scale invariance* instead of *translation invariance.*

Definition 6. The set of scale invariant-SQRE is defined as

$$S' = \{\sigma \in \Delta : Q \circ \bar{u}(\sigma) = \sigma \text{ for some } Q \text{ satisfying (R1)-(R4), (S1), (S2}'\text{), and (S3)}\},$$

and we say that if $\sigma \in S'$, then σ is a scale invariant-SQRE.

To derive a representation theorem analogous to Theorem 1, we define a player order based on *relative stakes*. We say that a player *faces higher relative stakes than* another if the *ratio* of expected payoffs between her actions is higher. Formally, defining player i's relative stakes in taking an action as $\rho(\bar{u}_i) \equiv max\{\bar{u}_{ij}/\bar{u}_{ik}, \bar{u}_{ik}/\bar{u}_{ij}\}$, we have the following order:

[45] See Goeree et al. [2005] for a discussion.

Definition 7. Suppose players i and j have two actions each. Player i *faces higher relative stakes than* player j if $\rho(\bar{u}_i) \geq \rho(\bar{u}_j)$ (and *faces strictly higher relative stakes* if $\rho(\bar{u}_i) > \rho(\bar{u}_j)$.

We define by X' the set of action profiles such that the ranking of players by extremeness is the same as the ranking of players by relative stakes.

$$X' = \{\sigma \in \Delta : \rho(\bar{u}_i) > (=)\ \rho(\bar{u}_j) \iff max(\sigma_{i1}, 1-\sigma_{i1}) > (=)\ max(\sigma_{j1}, 1-\sigma_{j1})\}.$$

The set of scale invariant-SQRE is precisely the set X' *intersected with* the set of RQRE. The proof is almost identical to that of Theorem 1, except with ρ replacing δ, and hence omitted.

Theorem 5. *Fix* $\Gamma^{N\times 2}$. $S' = X' \cap R$.

Hence, the set of invariant-SQRE can be found in similar manner to the set of SQRE, by first characterizing the set of RQRE and then intersecting it with another set based on rankings of players by stakes. The only difference is that we must consider *relative* as opposed to *absolute* stakes.

To illustrate the theorem, we characterize all scale-invariant SQRE for the matching pennies game of Table 8. The idea is to characterize analogues of the U-sets (2) for relative stakes and invoke a relative version of Corollary 1. The boundaries of the U-sets are now non-linear, but can still be found in closed form.[46] The left panel of Figure 17 plots S' and the set of Luce QRE which is contained within it. For comparison, the right panel plots S and the set of LQRE.

	L	**R**
U	14,5	5,30
D	5,80	6,5

Table 8: *An example matching pennies game.*

Scale invariant-SQRE gives many of the same lessons as SQRE. The characterization of S' leads to a significant refinement of R and places significant bounds on the set of

[46]For example, in a matching pennies game, consider the ratios $\bar{u}_U(q)/\bar{u}_D(q)$ and $\bar{u}_L(p)/\bar{u}_R(p)$. The first is increasing in q and equals 1 when $q = q^*$. The second is decreasing in p and equals 1 when $p = p^*$. Hence, to find the component of U_0 that is northeast of (q^*, p^*), set $\bar{u}_U(q)/\bar{u}_D(q) = \bar{u}_R(p)/\bar{u}_L(p)$ and solve for p as a function of q. Finding the northwest, southwest, and southeast components of U_0 is similar. The four components of U_0 separate the unit square into four regions. As with the translation invariant model, the regions west and east of (q^*, p^*) give U_1, and the regions north and south of (q^*, p^*) give U_2.

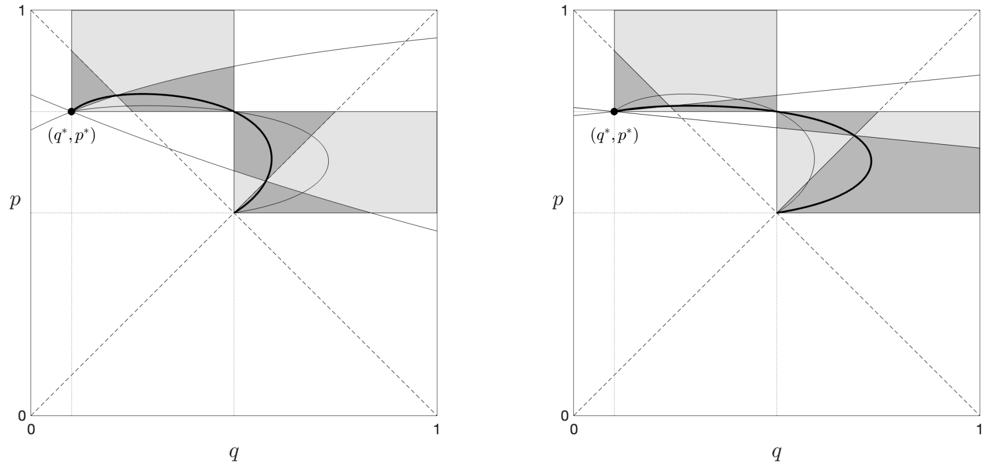

Figure 17: *Scale invariant-SQRE.* For the example game of Table 8, the left panel gives the set of scale invariant-SQRE (dark gray), the set of Luce QRE (thick black curve), and the set of LQRE (thin black curve). The right panel gives the set of translation invariant-SQRE (dark gray), the set of LQRE (thick black curve), and the set of Luce QRE (thin black curve).

Luce QRE. In particular, it allows for us to solve for some special Luce QRE in closed form as the intersection of curves, which, in this example, result in quadratic equations.

Comparing the left and right panels of Figure 17, we see that translation and scale invariance lead to different sets of SQRE. Thinking back to the example of Section 1.7, this is because *scale invariance* completes the "responsiveness order" in a different way than does *translation invariance*.[47] Finally, we note that a corollary of our results is that the set of Luce QRE is distinct from the set of LQRE.

[47]For instance, in the example of Section 1.7, while the data is not consistent with SQRE, it *is* consistent with scale invariant-SQRE since $p > q$ and $3/1 > 5/2$.