Quantal Response Equilibrium with a Continuum of Types: Characterization and Nonparametric Identification

Evan Friedman[*] Duarte Gonçalves[†]

Abstract

Quantal response equilibrium (QRE), a statistical generalization of Nash equilibrium, is a standard benchmark in the analysis of experimental data. Despite its influence, nonparametric characterizations and tests of QRE are unavailable beyond the case of finite games. We address this gap by completely characterizing the set of QRE in a class of binary-action games with a continuum of types. Our characterization provides sharp predictions in settings such as global games, the volunteer's dilemma, and the compromise game. Further, we leverage our results to develop nonparametric tests of QRE. As an empirical application, we revisit the experimental data from Carrillo and Palfrey (2009) on the compromise game.

Keywords: quantal response; Bayesian games; global games; compromise game; nonparametric analysis.
JEL Classifications: C44, C72, C92.

[*] Department Economics, Paris School of Economics; evan.friedman@psemail.eu.
[†] Department of Economics, University College London; duarte.goncalves@ucl.ac.uk.
We thank Olivier Compte and seminar participants at PSE for helpful comments.
This draft: 1 March 2023.

1. Introduction

Quantal Response Equilibrium (QRE) (McKelvey and Palfrey, 1995), a statistical generalization of Nash equilibrium (NE), has found significant success in explaining experimental data (Goeree, Holt, and Palfrey, 2016). In a QRE, players make probabilistic mistakes in best responding to their beliefs, but their beliefs are correct, taking into account the mistakes of others. This leads to systematic deviations from NE that can capture a range of observed behavioral phenomena, such as turnout in large elections (Levine and Palfrey, 2007) and overbidding in auctions (Goeree et al., 2002). However, while this solution concept has been influential, there are limited results on characterizing and testing QRE beyond the case of finite games.

In this paper, we consider a class of infinite games, those with binary actions and a continuum of types, for which we provide a complete characterization of the set of quantal response equilibria. In this environment, a QRE is a function mapping types to the probability of taking a given action. Under a monotonicity condition on payoffs, our characterization is as follows: any QRE is a continuous, strictly monotone function such that uniform randomization implies indifference between actions. Further, we provide a converse: any such function is a QRE, and thus we fully describe the set of equilibria. Using this result, we characterize QRE in a number of classic applications, including global games. Finally, we leverage our results to develop novel methods for nonparametric QRE analysis, which we apply to the experimental data of Carrillo and Palfrey (2009) on the compromise game.

The games we study are symmetric, binary-action Bayesian games, with any number of players. Before taking an action, each player learns their *type*, which potentially takes a continuum of values. This can either be a parameter of the utility function or a signal about some payoff-relevant state. To derive our main result, we assume only minimal structure on payoffs. Namely, we require that (1) (interim) expected payoffs are monotone in type whenever opponents' choice probabilities are also monotone in type, and that (2) for each of the two actions, there exists a type for which that action yields the higher payoff. While a simple class of infinite games, it is rich enough to include many games of significant theoretical interest.

The primitive of all QRE models is a *quantal response function* — the mapping from expected payoffs to a distribution over actions — and a QRE obtains when all players' behavior is consistent with the quantal response. Following McKelvey and Palfrey (1998), we study *agent QRE in symmetric strategies*, or simply *QRE* for short. In such an equilibrium, each type represents an agent who acts independently, and different players with the same type have the same behavior. In this way, each type has only two actions, allowing us to focus on the role of an infinite type space (as opposed to an infinite strategy space). A QRE is thus a function mapping from types to the probability of taking a given action.

Rather than imposing parametric structure, in our approach, we take inspiration from recent work characterizing the set of QRE based on minimal restrictions on the quantal response function. Goeree, Holt, and Palfrey (2005) define a *regular QRE* as any for which the quantal response function satisfies *monotonicity* and *responsiveness*, as well as other technical axioms. These require that actions with higher expected payoffs are taken more often and that an all-else-equal increase in the payoff to some action means it is played even more often. Friedman and Mauersberger (2022) study *symmetric QRE*, a refinement of regular QRE, whereby the quantal response function also satisfies various symmetries across players and actions.[1] In this paper, we consider both regular and symmetric QRE.

Our main result is that a function from types to choice probabilities is a QRE for some quantal response function satisfying the axioms if and only if the function is continuous, monotone, and uniform randomization implies indifference between actions. Furthermore, symmetric QRE are exactly identified by functions satisfying these conditions together with an additional symmetry condition. The key behind this result is that we impose only weak restrictions on quantal response: while any function satisfying the conditions of our theorem is consistent with some quantal response function, it would not in general be consistent with any given parametric form such as logit. Hence, we characterize the *set* of QRE without reference to the underlying quantal response functions. We show, however, how to recover the quantal response function associated with a given QRE — a by-product of our characterization result.

Our result transforms the problem of finding a fixed point in an infinite-dimensional function space to that of constructing a monotone function with a unique indifferent type who uniformly randomizes. This is still a difficult problem, but we show that it is feasible in the context of several classic applications. A key observation that makes the problem tractable is that expected payoffs often depend only on just a few features of the equilibrium strategy, for example its mean.[2] In such cases, it is often easy to characterize the *set of possible indifferent types* without fully specifying

[1]The three models we've mentioned are nested: logit QRE $\subset$ symmetric QRE $\subset$ regular QRE. Hence, symmetric QRE, by refining regular QRE also implies new bounds on logit QRE.

[2]In our three applications, the payoff-relevant statistics are the mean, the distance of one's type to a particular reference type, and the mean conditional on having a lower type.

the underlying equilibrium, and then, for each indifferent type, construct the set of supporting mappings that satisfy the necessary features.

We then consider three applications: the volunteer's dilemma (Diekmann, 1986) with a continuum of participation costs, global games (e.g. Carlsson and van Damme, 1993; Morris and Shin, 1998) with a continuum signals about the state of the world, and the compromise game (Carrillo and Palfrey, 2009) with a continuum of "strengths." In each case, we apply our result to characterize the set of QRE, deriving a sharper characterization that depends on specific features of the games. The games were chosen to showcase a breadth of different arguments, with the goal of suggesting new applications.

We view our contribution as two-fold. Our first contribution is theoretical as our results expand the universe of games that are amenable to nonparametric QRE analysis. We also derive economic insights, showing in each of our applications the precise sense in which QRE deviates systematically from NE. And while our focus is on characterizing sets of QRE nonparametrically, since the common parametric models are contained within the axiomatic families we study, our result implies bounds on these models as well.

Our second contribution is empirical. We show our characterization results can be used as the basis for nonparametric tests of QRE. Consider the common practice of fitting logit QRE to data. If the parametric model does *not* fit well, it is unclear to what extent this is due to the logit structure or a general limitation of QRE. By characterizing the set of QRE, our results allow us to ask whether the data is consistent with *some* QRE. This is tantamount to testing (1) if choice probabilities are monotone with respect to types, and (2) whether the type uniformly randomizing is indifferent between the two actions — a simple moment condition.

As an empirical application, we revisit the experimental study of Carrillo and Palfrey (2009) on the compromise game, which we re-analyze through the lens of our results. While we cannot reject monotonicity, we find a violation of the second condition: the type uniformly randomizing has a strictly higher expected payoff for one of the two actions.

This paper is organized as follows. In the remainder of this section, we discuss related literature. In Section 2, we introduce the family of games we consider, the definition of equilibrium, and provide general existence and characterization results. Section 3 applies our results to characterize QRE in three applications: the volunteer's dilemma, global games, and the compromise game. Section 4 analyzes experimental data on the compromise game, and Section 5 concludes.

Related literature. In response to concerns that some forms of QRE might lack empirical content (e.g. Haile et al., 2008), Goeree et al. (2005) offer the first systematic consideration of more general, non-parametric forms. They introduce the axiomatic regular QRE and show it is falsifiable, which implies the same for the models it nests, namely logit QRE and, more generally, all "structural" QRE with i.i.d. errors. Goeree et al. (2005) also characterize the *set of* regular QRE in some prominent examples. Only recently, however, Goeree and Louis (2021) introduced M equilibrium, an explicitly set-valued concept that nests a number of existing concepts. In particular, the union of M-equilibrium *choice sets* coincides with the set of regular QRE. Goeree and Louis (2021) show that this set is semi-algebraic, i.e. characterized by a finite number of polynomial (in)equalities, and therefore computable by a finite algorithm. This establishes the main insight that is relevant to our paper: by imposing only weak restrictions on quantal response, the resulting set of QRE is a tractable object. Friedman and Mauersberger (2022) refines regular QRE by augmenting it with various forms of symmetry across players and actions, and show how to analyze the resulting sets of equilibria.[3] The model is similarly tractable as regular QRE and implies much tighter bounds on the models nested within it, such as logit. Whereas all of these previous papers focus on finite games, the main novelty in our paper is to provide characterization results for nonparametric QRE in a class of infinite games — those with a continuum of types. This exercise is perfectly analogous to what has been done for finite games, but requires new methods altogether.

While fitting parametric QRE models — and comparing their fit to other parametric models — is common practice, there is surprisingly little work that develops formal tests. A prominent exception is (Melo et al., 2018),[4] which derives a test for structural QRE in sets of finite games. Their result is very general in that it applies to an arbitrary set of games, but here too the result requires that the games are finite.

Parametric QRE has been successfully applied to infinite games. In terms of games with a continuous type space, we are unaware of any systematic exploration such as we undergo in this paper. However, an excellent experimental study in this area is that of Carrillo and Palfrey (2009) who numerically approximate logit QRE in the compromise game. The class of games we study includes this game, and so we revisit their data in Section 4.

In terms of games with continuous action spaces, Anderson et al. (2002) study logit QRE in a family of "auctionlike" games with "payoff functions that depend on rank, such as whether a

[3]Friedman (2022) provides some comparative static results for regular QRE augmented with *translation invariance.*
[4]See also Aguirregabiria and Magesan (2020).

player's decision is higher or lower than another's." Here, a logit QRE is a choice density that satisfies a differential equation. While there is no closed form expression for this choice density, Anderson et al. (2002) establish existence, uniqueness, and comparative statics. In a similar vein, Anderson et al. (2001) study logit QRE of a continuous minimum-effort coordination game, and Baye and Morgan (2004) study the parametric "Luce" QRE in a continuous Bertrand pricing game. In our paper, we study a non-overlapping family of games, so our results are complementary, but a natural direction for future work is to extend our non-parametric analysis to games with larger action spaces.

QRE, which requires being able to assign a positive probability or density to all strategies, is not well-defined when the strategy space is an infinite-dimensional function space. Hence, we simplify the strategy space by considering interim or agent QRE. An alternative approach would be to impose a priori restrictions on strategies, as in Compte and Postlewaite (2019). After imposing restrictions, possibly allowing for a family of stochastic strategies, Compte and Postlewaite (2019) study NE of the restricted game. Alternatively, one could study QRE of the restricted game.[5] In this way, it would be natural to combine non-parametric QRE methods with strategy restrictions, which could be a powerful approach for complex games with large strategy spaces.[6]

2. The games, equilibrium, and characterization

We introduce the class of games, equilibrium concept, and restrictions on quantal response we consider. We then establish existence and characterize equilibria.

2.1. Binary-action games with a continuum of types

There is a set of players I, which can be either finite (with at least 2 players) or a continuum. They play a symmetric binary-action game in which each player i has the same binary action set, $A := \{0, 1\}$. There is an unknown state of the world $\theta \in \Theta := [0, 1]$, distributed according to density h. Before taking an action, each player i observes their private type $t_i \in T := [0, 1]$, independently drawn conditional on the state according to density f^θ; we require $f := \int_\Theta f^\theta h(\theta) d\theta$ to have full

[5] Carrillo and Palfrey (2009), in their logit QRE analysis of the compromise game, consider two versions. The first is agent QRE, which is a parametric form of the model we study in this paper. The second, which they refer to as "cutpoint QRE", imposes that each player only considers threshold strategies and then studies the QRE of this restricted game. Compte and Postlewaite (2019), in several of their applications, restrict players to choose "target" actions that are implemented with exogenous trembles, but the resulting equilibria cannot be interpreted as QRE of a restricted game.

[6] See also Arad and Rubinstein (2019) for a theory of behavior in complex games.

support on T. Players act independently, conditional on the state of the world and a player's payoff function is given by $u : A^I \times T^I \times \Theta \to \mathbb{R}$, a measurable real-valued mapping depending on the players' action profile, the realized type profile, and the state of the world.

Anticipating the symmetric nature of the solution concept we consider, we focus on symmetric Lebesgue measurable strategies, $\sigma : T \to [0,1]$. Given continuity properties imposed on the payoffs below and the fact that the distribution of types admits a density, we take the strategy space Σ as the set of $L^1(T)$ functions endowed with the L^1-norm $\|\cdot\|_{L^1}$. We will denote the expected payoff to a player with type t_i choosing action $a \in A$ given their opponents all follow strategy σ as $\bar{u}^a_{t_i}(\sigma)$, formally given by

$$\bar{u}^a_{t_i}(\sigma) := \mathbb{E}_{\theta\sim h}[\mathbb{E}_{t_j\sim f^\theta,\forall j\neq i}[\mathbb{E}_{a_j\sim\sigma(t_j)}[u(a,a_{-i},t_i,t_{-i},\theta)]] \mid t_i].$$

Further, define $\Delta\bar{u}_{t_i}(\sigma) := \bar{u}^1_{t_i}(\sigma) - \bar{u}^0_{t_i}(\sigma)$ as the corresponding expected utility difference between taking actions 1 and 0. Because of the symmetric nature of the environment, we henceforth omit player subscripts, referring to arbitrary types as t, t', or similar.

The above formulation is general enough to encompass many types of games. In particular, we note that a player's type t can simply be a parameter of the utility function, or it can be a signal about the unknown state. We consider applications with both interpretations.

We impose the following restrictions on payoffs:

(A1) **Continuity**: For all $a \in A$, $\bar{u}^a_t(\sigma)$ is jointly continuous in (t,σ) with respect to the product topology.

(A2) **Payoff-responsiveness**: If there are $t < t'$ such that $\sigma(t) \le \sigma(t')$, then there exist $\hat{t} \neq \hat{t}'$ satisfying (i) $\sigma(\hat{t}) \le \sigma(\hat{t}')$ and (ii) $\bar{u}^1_{\hat{t}}(\sigma) \ge \bar{u}^1_{\hat{t}'}(\sigma)$ and $\bar{u}^0_{\hat{t}}(\sigma) \le \bar{u}^0_{\hat{t}'}(\sigma)$, with at least one of these inequalities strict.

(A3) **Payoff-monotonicity**: For any strictly decreasing σ, for any $t < t'$, $\bar{u}^1_t(\sigma) \ge \bar{u}^1_{t'}(\sigma)$ and $\bar{u}^0_t(\sigma) \le \bar{u}^0_{t'}(\sigma)$, with at least one of these inequalities strict.

(A4) **Non-triviality**: For any σ such that $\sigma(t') > \frac{1}{2}$ (resp. $\sigma(t') < \frac{1}{2}$) for all t', then $\Delta\bar{u}_t(\sigma) < 0$ (resp. $\Delta\bar{u}_t(\sigma) > 0$) for some t.

The first assumption (A1) is continuity of expected payoffs, which is relatively innocuous, and guarantees that σ will be continuous in equilibrium. The second assumption (A2) imposes that, if σ is increasing at some point, then there exists two types such that one plays action 1 more often despite facing payoffs that are relatively less favorable. This will be shown to imply that

σ is strictly decreasing in equilibrium. It is still fairly weak in that it allows for the existence of σ such that payoffs are non-monotone in type, but precludes such σ from being part of the equilibrium. This flexibility is important to accommodate some interesting applications, such as global games (Section 3.2). The third assumption (A3) requires that monotone strategies entail payoff-monotonicity in type, which will allow us to characterize the full set of QRE. The last assumption (A4) ensures some degree of strategic substitutability. Without this assumption, games admit equilibria in which there is an action that all types take more often than not. Such equilibria are not un-interesting, but their QRE analysis turns out to be somewhat trivial. Hence, we think of this final assumption as a non-triviality constraint that allows us to focus on the most interesting cases.

2.2. Quantal response equilibrium

We assume that each type's behavior is governed by the same *quantal response function* $Q : \mathbb{R}^2 \to [0,1]$, which maps from expected payoffs to a mixed action that we identify with the probability of choosing action 1.[7] We denote by $\bar{u}_t(\sigma) := (\bar{u}_t^1(\sigma), \bar{u}_t^0(\sigma))$ the vector of expected utilities for a player with type $t \in T$ and strategy $\sigma : T \to [0,1]$. Stated formally after imposing restrictions on Q, a quantal response equilibrium will be defined as a strategy σ such that all types' behavior is consistent with quantal response: $\sigma(t) = Q(\bar{u}_t(\sigma))$ for all $t \in T$.[8]

Without restrictions on Q, this poses virtually no restrictions on observable behavior. Hence, we impose on Q weak restrictions or *axioms* that are found in the literature. The axioms are defined for arbitrary finite numbers of actions, but we present them in a binary-action form. Throughout the paper, we always assume Q satisfies the *regularity* axioms (R1)-(R4) below, which are due to Goeree et al. (2005).

(R1) **Interiority**: $Q(v) \in (0,1)$ for all $v = (v^1, v^0) \in \mathbb{R}^2$.

(R2) **Continuity**: Q is continuous.

(R3) **Responsiveness**: $\frac{\partial Q(v)}{\partial v^1} > 0 > \frac{\partial Q(v)}{\partial v^0}$ for all $v = (v^1, v^0) \in \mathbb{R}^2$.

[7]One can view Q as the representative quantal response for a population of individuals with potentially heterogeneous quantal responses. When quantal responses arise from additive i.i.d. payoff disturbances and the action space is binary, Golman (2011) shows a representative quantal response emerges that is also based on additive i.i.d. payoff disturbances.

[8]Note that this corresponds to *agent QRE* (McKelvey and Palfrey, 1998), which is common in the literature. Introduced to analyze extensive-form games, agent QRE treats the same player at different nodes — or of different types — as separate agents who mix independently. This is often viewed as a simplification as each agent has a smaller strategy space than the player. In this paper, each agent has exactly two actions.

(R4) **Monotonicity**: $v^1 > v^0 \iff Q(v) > 1 - Q(v)$.

We now state the definition of the solution concept formally:

Definition 1. A *quantal response equilibrium* (QRE) is a strategy $\sigma : T \to [0,1]$ such that there is a $Q : \mathbb{R}^2 \to [0,1]$ satisfying (R1)-(R4) for which $\sigma(t) = Q(\bar{u}_t(\sigma))$ for all $t \in T$.

It is important to note that a QRE is a fixed point in a function space. Formally, define $\bar{u}(\sigma) := (\bar{u}_t(\sigma))_{t\in[0,1]}$ as the vectors of expected payoffs faced by all types. Recalling that Σ denotes the space of strategies mapping from T to $[0,1]$, we define the operator $q : \Sigma \to \Sigma$ to be such that, for all $t \in T$, $q(\sigma)(t) := Q(\bar{u}_t(\sigma))$. An equivalent definition of QRE is then as a fixed point of q: $\sigma \in \Sigma$ is a QRE if $\sigma = q(\sigma)$.

In addition to QRE, we consider a refinement that also imposes the *symmetry* axioms (S1)-(S2) below, also introduced in Goeree et al. (2005).

(S1) **Translation invariance**: $Q(v + \gamma e) = Q(v)$ for all $v = (v^1, v^0) \in \mathbb{R}^2$, $\gamma \in \mathbb{R}$ and $e = (1,1)$.

(S2) **Label independence**: For any $v = (v^1, v^0), \tilde{v} = (\tilde{v}^1, \tilde{v}^0) \in \mathbb{R}^2$, if $v^1 = \tilde{v}^0$ and $v^0 = \tilde{v}^1$, then $Q(v) = 1 - Q(\tilde{v})$.

Following Friedman and Mauersberger (2022), whenever Q satisfies (R1)-(R4) and (S1)-(S2), we refer to the resulting model as *symmetric QRE* or *SQRE*.

Definition 2. A *symmetric quantal response equilibrium* (SQRE) is a strategy $\sigma : T \to [0,1]$ such that there is a $Q : \mathbb{R}^2 \to [0,1]$ satisfying (R1)-(R4) and (S1)-(S2) for which $\sigma(t) = Q(\bar{u}_t(\sigma))$ for all $t \in T$.

The axioms (R1)-(R2) impose the key technical conditions — that all actions are played with positive probability and that behavior is continuous in payoffs. (R3)-(R4) are the main behavioral axioms, imposing a weak form of rationality: that higher payoff actions are played more often and that an all-else equal increase in the payoff to some action leads to it being played even more often. (S1) ensures that quantal response is invariant to adding the same constant to both payoffs, and (S2) imposes that only actions' payoffs — and not their labels — matter for quantal response. (S1)-(S2) are not implied by (R1)-(R4). They do hold, however, under the common "structural" approach in which quantal response is induced by additive errors if the errors are exchangeable with respect to actions (weaker than i.i.d.) and invariant to the payoffs themselves. In virtually all applications, (R1)-(R4) are satisfied; and in the large majority of applications, including the common logit QRE, (S1)-(S2) are also satisfied. In this paper, we study both QRE ((R1)-(R4)) and SQRE ((R1)-(R4) and (S1)-(S2)), which allows us to isolate the effects of symmetry.

2.3. Existence and characterization

Because of the infinite nature of the game, a QRE is a function $\sigma \in \Sigma = [0,1]^T$ that is a fixed point of the operator q. Our first step is to show that under general conditions, such a fixed point exists. The crucial step of the proof is to invoke Schauder's fixed-point theorem, a generalization of Brouwer's fixed-point theorem for infinite dimensional spaces.

Lemma 1. *For any Q satisfying (R2)-(R3) and game satisfying (A1)-(A2), q admits a fixed point.*

Proof. See Appendix A. □

We next provide a characterization of QRE. We find that any QRE is continuous, strictly decreasing, interior, and has a unique *indifferent type* that uniformly randomizes. Furthermore, these properties deliver a converse, and so we completely characterize the set of QRE.

Theorem 1. *Assume (A1)-(A4). A strategy $\sigma : T \to (0,1)$ is a QRE if and only if (i) σ is continuous and strictly decreasing, (ii) $\sigma(t) \in (0,1)\, \forall t \in T$, and (iii) there exists a unique type $\tilde{t} \in (0,1)$ such that $\sigma(\tilde{t}) = \frac{1}{2}$ and $\Delta \bar{u}_{\tilde{t}}(\sigma) = 0$.*

Proof. Only if: That σ is continuous and strictly decreasing follows from the proof of Lemma 1. Note that if payoffs satisfy (A1), since the relevant domain for Q is bounded set $\mathscr{U} \subset \mathbb{R}^2$ as defined in the proof of Lemma 1, from (R1) it will also be the case that $\sigma(t) \in (0,1)$ for all t. Finally, we note that for any fixed point $\sigma = q(\sigma)$ there is a unique $\tilde{t}$ such that $\sigma(\tilde{t}) = 1/2$. That there is at most one follows from the fact that σ must be continuous and strictly decreasing. Suppose now that there is no such type and instead $\sigma(t^{'}) > 1/2\ \forall t^{'} \in T$. Then by (A4), there exists t such that $\Delta \bar{u}_t(\sigma) < 0 \Longrightarrow \bar{u}_t^1(\sigma) < \bar{u}_t^0(\sigma)$. From (R4) we then get that $Q(\bar{u}_t(\sigma)) = \sigma(t) < 1/2$, a contradiction. A symmetric contradiction is obtained when assuming that $\sigma(t^{'}) < 1/2\ \forall t^{'} \in T$. Further, by (R4), $\sigma(\tilde{t}) = 1/2 \Longrightarrow \Delta \bar{u}_{\tilde{t}}(\sigma) = 0$.

If: From (A3), as σ is strictly decreasing, $\Delta \bar{u}_t(\sigma)$ is strictly decreasing in t. Let $\delta : T \to \mathbb{R}$ be given by $\delta(t) := \Delta \bar{u}_t(\sigma)$ and define $\tilde{Q} : [\delta(1), \delta(0)] \to [0,1]$ by $\tilde{Q}(d) = \sigma(\delta^{-1}(d))$, which is well-defined since δ is strictly decreasing. Extend this to the whole real line in any arbitrary way such that $\tilde{Q} : \mathbb{R} \to [0,1]$ is continuous, strictly increasing, and interior. Finally, extend this to a quantal response function Q defined over $\mathbb{R}^2$ as $Q : \mathbb{R}^2 \to (0,1)$ where $Q(v^1, v^0) = \tilde{Q}(v^1 - v^0)$. By construction, Q satisfies (R1)-(R4) and $Q(\bar{u}_t(\sigma)) = \tilde{Q}(\Delta\bar{u}_t(\sigma)) = \sigma(t)\ \forall t$. □

Intuitively, we find that an SQRE is a QRE with an additional symmetry condition across types. For ease of reference, we define *symmetry* formally as a condition on σ (and the expected payoffs induced by σ and u).

Definition 3. A strategy σ is *symmetric* if $\sigma(t) = 1 - \sigma(t') \iff \Delta\bar{u}_t(\sigma) = -\Delta\bar{u}_{t'}(\sigma)$.

The next result delivers a characterization of SQRE. It is the same as Theorem 1 but includes the above symmetry condition.

Theorem 2. *Assume (A1)-(A4). A strategy $\sigma : T \to (0,1)$ is an SQRE if and only if (i) σ is continuous and strictly decreasing, (ii) $\sigma(t) \in (0,1)\,\forall t \in T$, (iii) there exists a unique type $\tilde{t} \in (0,1)$ such that $\sigma(\tilde{t}) = \frac{1}{2}$ and $\Delta\bar{u}_{\tilde{t}}(\sigma) = 0$, and (iv) σ is symmetric.*

Proof. Only if: For any QRE $\sigma = q(\sigma)$, properties (i)-(iii) follow from Theorem 1; we now show (iv), the symmetry of σ. Let $\tilde{Q}(v^1 - v^0) := Q((v^1, v^0))$. First note that if Q satisfies (S1), then, for any two payoff functions u, u' such that $\Delta\bar{u}_t(\sigma) = \Delta\bar{u}'_t(\sigma)$ for all σ, we have $\tilde{Q}(\Delta\bar{u}_t(\sigma)) = Q(\bar{u}_t(\sigma)) = Q(\bar{u}'_t(\sigma))$ for all σ. If Q also satisfies (S2), we have that $\tilde{Q}(\Delta\bar{u}_t(\sigma)) = Q((\bar{u}^1_t(\sigma), \bar{u}^0_t(\sigma))) = 1 - Q((\bar{u}^0_t(\sigma), \bar{u}^1_t(\sigma))) = 1 - \tilde{Q}(-\Delta\bar{u}_t(\sigma))$, and so at any SQRE, $\sigma(t) = 1 - \sigma(t') \iff \tilde{Q}(-\Delta\bar{u}_t(\sigma)) = \tilde{Q}(\Delta\bar{u}_{t'}(\sigma)) \iff -\Delta\bar{u}_t(\sigma) = \Delta\bar{u}_{t'}(\sigma)$, where the last equivalence follows from (R3), proving σ is symmetric.

If: Construct $Q : \mathbb{R}^2 \to (0,1)$ where $Q(v^1, v^0) = \tilde{Q}(v^1 - v^0)$ exactly as in the "if" direction of the proof of Theorem 1, except also require that $\tilde{Q}(d) = 1 - \tilde{Q}(-d)$, which is possible since σ is symmetric. That Q satisfies (S1) and (S2) (as well as (R1)-(R4)) and $Q(\bar{u}_t(\sigma)) = \tilde{Q}(\Delta\bar{u}_t(\sigma)) = \sigma(t)\ \forall t$ follows from the construction. □

Due to these results, one can now characterize QRE without the need to solve for high-dimensional fixed points. One simply needs to construct the equilibrium in a way that preserves the necessary features. Furthermore, as we discuss in Section 4, Theorems 1 and 2 pave the way for a general methodology to nonparametrically test the ability of QRE and SQRE to rationalize data.

Our last general result pertains to the quantal response function Q. While our characterization of QRE makes no reference to the underlying quantal response function $Q : \mathbb{R}^2 \to [0,1]$, we show how to partially recover Q from equilibrium play. To this end, let $V : \Sigma \rightrightarrows \mathbb{R}^2$ be such that $V(\sigma) :=$

$\{(v^1, v^0) \in \mathbb{R}^2 \mid \exists t \in T : \Delta\bar{u}_t(\sigma) = v^1 - v^0\}$. That is, if σ is an equilibrium, $V(\sigma)$ delivers the set of expected payoff vectors that arise in equilibrium as well as all translations of such vectors. For all $v \in V(\sigma)$, we also define $t(v) \in T$ as the unique type satisfying $\Delta\bar{u}_{t(v)}(\sigma) = v^1 - v^0$.

For any QRE σ, our next result provides a construction for $Q|_{V(\sigma)}$ that is consistent with σ. If σ is also an SQRE, the same construction provides the *unique* such $Q|_{V(\sigma)}$.

Corollary 1. *Assume (A1)-(A4).*

(1) If σ is a QRE, then there is a $Q : \mathbb{R}^2 \to [0,1]$ (satisfying (R1)-(R4) and (S1)) with $Q|_{V(\sigma)}(v) = \sigma(t(v))\ \forall v \in V(\sigma)$ such that $\sigma(t) = Q(\bar{u}_t(\sigma))\ \forall t \in T$.

(2) If σ is an SQRE, then all $Q : \mathbb{R}^2 \to [0,1]$ such that $\sigma(t) = Q(\bar{u}_t(\sigma))\ \forall t \in T$ satisfy $Q|_{V(\sigma)}(v) = \sigma(t(v))\ \forall v \in V(\sigma)$.

Proof. (1): This follows exactly from the construction in the "if" direction of the proof of Theorem 1. *(2)*: This is the construction in the "if" direction of the proof of Theorem 2. That this is unique follows from the fact that σ uniquely identifies Q restricted to the payoff vectors observed in equilibrium, which, by (S1), extends uniquely to $Q|_{V(\sigma)}$. □

To summarize, our results transform the problem of finding a fixed point in a function space to that of constructing a monotonic function with very specific features. However, it is not obvious, a priori, that this new problem is any easier. One must construct σ such that there is a unique indifferent type $\tilde{t}$ who uniformly mixes while ensuring that σ satisfies other properties.

As we show in our next section, this problem is further simplified by the fact that, in applications, expected payoffs often do not depend on σ in its entirety, but rather on just a few of its properties. For example, it could be that the payoffs to type t' depend on σ only through the mean $\mathbb{E}[\sigma(t)]$, the mean among lower types $\mathbb{E}[\sigma(t)|t \le t']$, or the distance to some special reference type $t^*(\sigma)$. In such cases, the problem becomes particularly tractable as one may be able to characterize the *set of indifferent types* without being precise about the supporting strategies, and only then construct the set of strategies (satisfying the relevant conditions) that can support each indifferent type.

3. Applications

We consider three games: the volunteer's dilemma, the compromise game, and a global game. The games were chosen to showcase a broad range of possible applications. While we invoke

Theorems 1 and 2 in all applications, the arguments are unique in each case.

3.1. The volunteer's dilemma

Two players simultaneously decide whether to *volunteer* to perform a task or to *abstain*. If at least one player volunteers, both receive $B \in (1,2)$. However, volunteering is costly. A player's private cost c is distributed as $c \sim U[0,1]$, i.i.d. across players. Let $\sigma(c)$ denote the probability that a player with cost c *volunteers*.[9]

Volunteering ensures that the benefit is received and the cost is paid, so the value for type c of volunteering is $B-c$. By abstaining, type c forgoes the cost, but only benefits if the other player volunteers; hence, this yields an expected payoff of $B\int_0^1 \sigma(c^{'})dc^{'}$. The difference is

$$\Delta \bar{u}_c(\sigma) = B\int_0^1 [1-\sigma(c^{'})]dc^{'} - c.$$

Hence, the payoff difference depends on σ only through the mean $\mathbb{E}(\sigma) = \int_0^1 \sigma(c^{'})dc^{'}$ and is additively separable in type c. These features make the analysis particularly simple.

As a benchmark, consider first the (essentially) unique NE, which is in symmetric threshold strategies: $\sigma^{NE}(c) = \mathbf{1}\{c < \frac{B}{B+1}\}$.[10] Hence, low-cost types volunteer, high-cost types abstain, and there is a unique indifferent type $\tilde{c}^{NE} = \frac{B}{B+1}$ that can mix arbitrarily. Intuitively, by injecting noise as in QRE, this step function will be smoothed out, and the flexibility in the admissible noise structures lead to a range of possible indifferent types. Letting $\tilde{R}$ denote the set of indifferent types for QRE, we obtain the following characterization:

Proposition 1. *$\sigma : [0,1] \to (0,1)$ is a QRE if and only if the indifferent type is $\tilde{c} \in \tilde{R} = (\frac{B}{B+2}, \frac{2B}{B+2})$, σ is continuous and strictly decreasing, and $\mathbb{E}(\sigma(c)|c \in [0,\tilde{c}]) = -\mathbb{E}(\sigma(c)|c \in [\tilde{c},1])(\frac{1-\tilde{c}}{\tilde{c}}) + \frac{B-\tilde{c}}{B\tilde{c}}$.*

Proof. Type $\tilde{c}$ is indifferent if and only if $B\int_0^1[1-\sigma(c^{'})]dc^{'} - \tilde{c} = 0$. Letting $\tilde{\sigma}_L = \frac{1}{\tilde{c}}\int_0^{\tilde{c}} \sigma(c^{'})dc^{'}$ and $\tilde{\sigma}_H = \frac{1}{1-\tilde{c}}\int_{\tilde{c}}^1 \sigma(c^{'})dc^{'}$ denote the average actions of types lower and higher than $\tilde{c}$, respectively, we re-write the indifference condition:

$$B\int_0^1 [1-\sigma(c^{'})]dc^{'} - \tilde{c} = 0 \iff$$
$$B[(1-\tilde{\sigma}_L)\tilde{c} + (1-\tilde{\sigma}_H)(1-\tilde{c})] - \tilde{c} = 0 \iff$$
$$(1-\tilde{\sigma}_L)\tilde{c} + (1-\tilde{\sigma}_H)(1-\tilde{c}) = \frac{\tilde{c}}{B}.$$

[9] In the version introduced by Diekmann (1986), there are N players, each of whom has the same cost.
[10] If $c = \frac{B}{B+1}$, the player is indifferent and may mix arbitrarily.

By Theorem 1, in any QRE, σ must be strictly decreasing and $\sigma(\tilde{c}) = \frac{1}{2}$, and therefore it must be that $1-\tilde{\sigma}_L < \frac{1}{2}$ and $1-\tilde{\sigma}_H > \frac{1}{2}$. Hence the supremum of the left-hand-side (LHS) of the indifference condition is $\frac{1}{2}\tilde{c}+(1-\tilde{c}) = 1-\frac{1}{2}\tilde{c}$, and so type $\tilde{c}$ cannot be indifferent if $1-\frac{1}{2}\tilde{c} < \frac{\tilde{c}}{B} \iff \tilde{c} > \frac{2B}{2+B}$. Similarly, the infimum of the LHS is $\frac{1}{2}(1-\tilde{c}) = \frac{1}{2}-\frac{1}{2}\tilde{c}$, and so type $\tilde{c}$ cannot be indifferent if $\frac{1}{2}-\frac{1}{2}\tilde{c} > \frac{\tilde{c}}{B} \iff \tilde{c} < \frac{B}{2+B}$. Conversely, solving the indifference condition yields $\tilde{\sigma}_L = -\tilde{\sigma}_H(\frac{1-\tilde{c}}{\tilde{c}})+\frac{B-\tilde{c}}{B\tilde{c}}$, which gives some $\tilde{\sigma}_L \in (\frac{1}{2},1)$ for any $\tilde{\sigma}_H \in (0,\frac{1}{2})$ if $\tilde{c} \in (\frac{B}{2+B},\frac{2B}{2+B})$, and hence any such $\tilde{c}$ can be supported as the indifferent type for some strictly decreasing σ. The result now follows directly from Theorem 1. □

QRE is very flexible relative to NE. Much of this flexibility comes from the fact that, while types must tend to take the action that yields the higher payoff (and uniformly mix when indifferent), they may still be biased in favor of a particular action. For example, if $\sigma(c) > 1-\sigma(c')$ for some $c < \tilde{c} < c'$ and $|\Delta\bar{u}_c(\sigma)| \leq |\Delta\bar{u}_{c'}(\sigma)|$, then there is a (local) bias in favor of volunteering. SQRE, by imposing symmetry, rules out precisely these biases. In the sequel, we define $\tilde{S}$ as the set of indifferent types for SQRE.

Proposition 2. *$\sigma : [0,1] \to (0,1)$ is an SQRE if and only if the indifferent type is $\tilde{c} \in \tilde{S} = (\frac{B}{B+1},\frac{1}{2}B)$, σ is continuous and strictly decreasing, $\mathbb{E}(\sigma(c)|c \in [0,2\tilde{c}-1]) = \frac{B\tilde{c}-\tilde{c}}{2B\tilde{c}-B}$, and σ is symmetric: $\sigma(\tilde{c}-\epsilon) = 1-\sigma(\tilde{c}+\epsilon)$ for all $\epsilon \in [0,1-\tilde{c}]$.*

Proof. Borrowing from the proof of Proposition 1, indifference requires

$$(1-\tilde{\sigma}_L)\tilde{c}+(1-\tilde{\sigma}_H)(1-\tilde{c}) = \frac{\tilde{c}}{B}.$$

In any SQRE, this cannot be satisfied if $\tilde{c} \leq \frac{1}{2}$. To see this, SQRE requires, by Theorem 2, that $\sigma(c) = 1-\sigma(c')$ if and only if $\Delta\bar{u}_c(\sigma) = -\Delta\bar{u}_{c'}(\sigma)$, which in this case means that $\sigma(\tilde{c}-\epsilon) = 1-\sigma(\tilde{c}+\epsilon)$ for all $\epsilon \in [0, min\{\tilde{c},1-\tilde{c}\}]$. Hence, symmetry and $\tilde{c} \leq \frac{1}{2}$ implies $(1-\tilde{\sigma}_H) \geq 1-(1-\tilde{\sigma}_L) \iff (1-\tilde{\sigma}_H)+(1-\tilde{\sigma}_L) \geq 1$. This implies that the LHS of the indifference condition is greater than a linear combination of $\tilde{c}$ and $1-\tilde{c}$ and thus is also greater than or equal to $min\{\tilde{c},(1-\tilde{c})\} = \tilde{c}$, which is strictly greater than $\frac{\tilde{c}}{B}$, meaning we cannot have indifference.

Hence, suppose $\tilde{c} > \frac{1}{2}$. We partition types into three intervals: $[0,2\tilde{c}-1]$, $[2\tilde{c}-1,\tilde{c}]$, and $[\tilde{c},1]$. Let the average probability of abstaining in each of the three intervals be $1-\sigma_1 = \frac{1}{2\tilde{c}-1}\int_0^{2\tilde{c}-1}\sigma(c')dc'$, $1-\sigma_2 = \frac{1}{1-\tilde{c}}\int_{2\tilde{c}-1}^{\tilde{c}}\sigma(c')dc'$, and $1-\sigma_3 = \frac{1}{1-\tilde{c}}\int_{\tilde{c}}^{1}\sigma(c')dc'$. By symmetry, since the length of the second two intervals is the same, we have that $1-\sigma_3 = \sigma_2$. The indifference condition can thus be re-written

$$\begin{aligned}(1-\sigma_1)(2\tilde{c}-1)+(1-\sigma_2)(1-\tilde{c})+(1-\sigma_3)(1-\tilde{c}) &= \frac{\tilde{c}}{B} \iff\\ (1-\sigma_1)(2\tilde{c}-1)+(1-\sigma_2)(1-\tilde{c})+\sigma_2(1-\tilde{c}) &= \frac{\tilde{c}}{B} \iff\\ (1-\sigma_1)(2\tilde{c}-1)+(1-\tilde{c}) &= \frac{\tilde{c}}{B} \iff\\ -\sigma_1(2\tilde{c}-1)+\tilde{c} &= \frac{\tilde{c}}{B}.\end{aligned}$$

Since σ_2 drops out, the only restriction is that σ_1 satisfies the above such that σ can be strictly decreasing, i.e. that $\sigma_1 \in (\frac{1}{2}, 1)$ and $\sigma_1 > \sigma_2 \in (\frac{1}{2}, 1)$. Solving gives $\sigma_1 = \frac{B\tilde{c}-\tilde{c}}{2B\tilde{c}-B}$, which is between $\frac{1}{2}$ and 1 if and only if $\tilde{c} \in (\frac{B}{B+1}, \frac{1}{2}B)$. Hence, any $\tilde{c} \in (\frac{B}{B+1}, \frac{1}{2}B)$ can be supported as the indifferent type for a strictly decreasing σ satisfying the relevant conditions. The result now follows directly from Theorem 2. □

The left panel of Figure 1 draws representative QRE (in red) and SQRE (in blue) for the case that $B = 1.5$. The thick horizontal lines at $\frac{1}{2}$ represent the sets of possible indifferent types $\tilde{R} = (\frac{B}{B+2}, \frac{2B}{B+2})$ and $\tilde{S} = (\frac{B}{B+1}, \frac{1}{2}B)$. The thin horizontal lines give average strategies within the intervals highlighted in Propositions 1 and 2. While all SQRE satisfy symmetry, this particular QRE is drawn with a bias in favor of volunteering: there exists $c < \tilde{c} < c^{'}$ such that $\sigma(c) > 1-\sigma(c^{'})$ and $|\Delta\bar{u}_c(\sigma)| \leq |\Delta\bar{u}_{c^{'}}(\sigma)|$.

We find that, while the flexibility in both QRE models gives rise to a range of possible behaviors, SQRE gives much more precise predictions. For instance, consider the measures of $\tilde{R}$ and $\tilde{S}$, which are $|\tilde{R}| = \frac{B}{B+2}$ and $|\tilde{S}| = \frac{B^2-B}{2B+2}$, respectively. Plotting these measures in the right panel of Figure 1 as a function of B, we see that they are larger under QRE than SQRE for all values of B.

There are also important qualitative differences between QRE and SQRE. We find that $\tilde{c}^{NE} = \frac{B}{B+1}$, the indifferent type under NE, is always in the interior of $\tilde{R}$, whereas $\tilde{c}^{NE}$ is precisely the infimum of $\tilde{S}$. This gives a sense in which SQRE leads to more systematic deviations from NE. Furthermore, since the indifference condition yields $\mathbb{E}(\sigma) = 1 - \frac{\tilde{c}}{B}$, one can calculate the ex-ante probability of volunteering as a function of $\tilde{c}$. This allows us to obtain, as a corollary, that SQRE always yields a lower probability of volunteering than in NE, whereas QRE can yield a lower or higher value depending on the direction of bias.

Corollary 2. *The set of attainable values of $\mathbb{E}(\sigma)$ is $(\frac{B}{B+2}, \frac{B+1}{B+2})$ in QRE, $(\frac{1}{2}, \frac{B}{B+1})$ in SQRE, and $\{\frac{B}{B+1}\}$ in NE.*

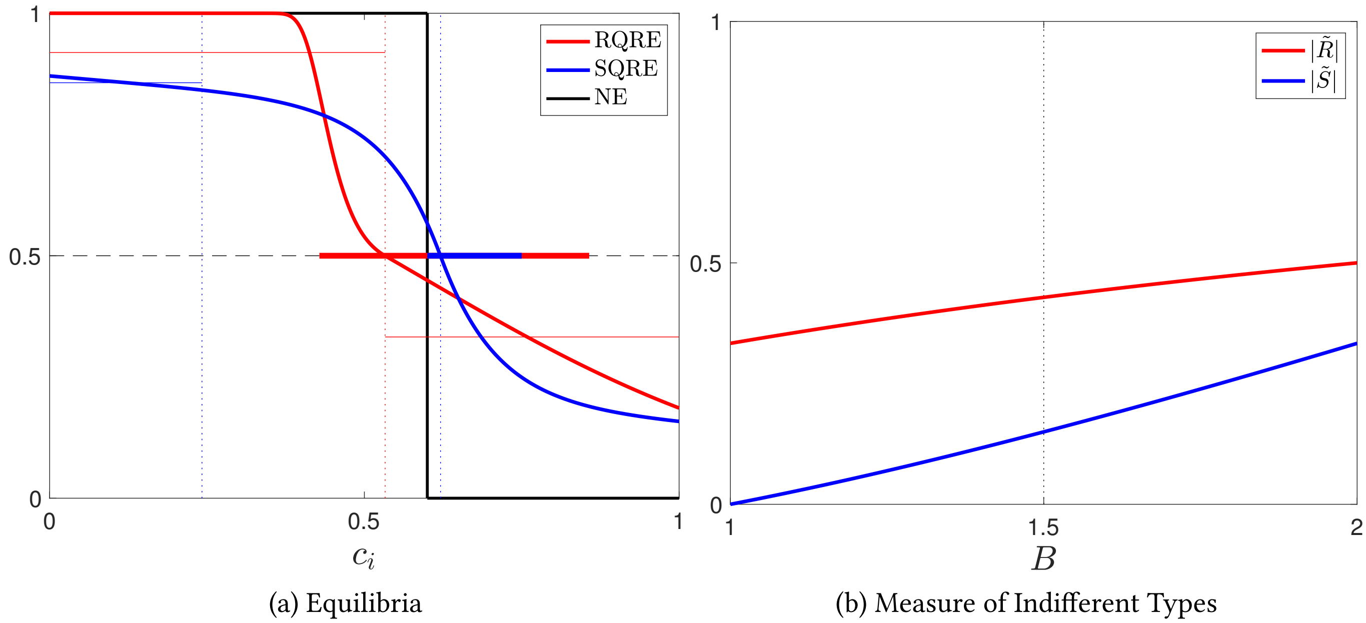

Figure 1: QRE and SQRE in the Volunteer's dilemma

Notes: Panel 1a depicts the NE (black) as well as illustrative QRE (red) and SQRE (blue), for $B = 1.5$. Panel 1b exhibits the measure of the sets of possible indifferent types for QRE (red) and SQRE (blue) for $B \in (0, 2)$.

We conclude this example with two remarks. First, as already mentioned, this is one of the simplest possible examples because payoffs depend on σ only through its expectation $\mathbb{E}(\sigma)$, and payoffs are additively separable in type c. Because of this, it is very easy to characterize the set of indifferent types, and symmetry takes a particularly simple form. We will see in the subsequent sections, however, that even without these features, we may still obtain precise characterization results. Second, this example illustrates the role of the non-triviality constraint (A4), which here would be violated if $B > 2$. In such a case, there would exist QRE and SQRE in which all types volunteer more often than not. Moreoever, all such QRE and SQRE would coincide as it would be that $\sigma(c) > \frac{1}{2}$ and $\Delta \bar{u}_c(\sigma) > 0$ for all c, meaning no two types face the same absolute payoff difference. While it would not be hard to analyze this case, we keep the discussion focused, and highlight the role of symmetry, by ruling out such cases a priori.

3.2. Global games

A continuum of players decide whether to *attack* a regime (e.g. a currency peg) or to *abstain*. The attack is successful if and only if (strictly) more than $\frac{1}{2}$ of the mass of players attack. The value of a successful attack is $\theta - k$, where $\theta \sim U[0, 1]$ is the state of the world and $k \in (0, 1)$ is the cost of attacking. If an attack fails, attackers receive $\theta - k - c$ where $c \in (0,1)$ is the penalty for failure. We assume that $c + k \in (0,1)$. Abstaining yields the safe payoff of 0.

Prior to taking an action, each player privately observes their type x, which is a signal about θ. Let $\epsilon > 0$ be a parameter representing the imprecision of this signal, and let it be "small", satisfying the technical condition $0 \le \epsilon < \min\{k, 1-k-c\}$. For each $\theta \in [\epsilon, 1-\epsilon]$, signals are distributed as $x \sim U[\theta-\epsilon, \theta+\epsilon]$, i.i.d. across players (conditional on the state); and thus $\mathbb{E}(\theta|x) = \frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon} \theta' d\theta' = x$. For $\theta \notin [\epsilon, 1-\epsilon]$, the form of signals is slightly different,[11] but it is still the case that $\mathbb{E}(\theta|x) = x$. Let $\sigma(x)$ denote the probability that a player with signal x chooses to *attack*.

Remark 1. By setting σ as the probability of attacking (as opposed to the probability of abstaining), a QRE will now be an *increasing* function. This choice is arbitrary, but more consistent with convention for these games.

This particular variant of global game deviates from that studied in Morris and Shin (1998),[12] but it still has the classic features with respect to multiplicity of equilibria. To be precise, if $\epsilon = 0$, and θ is publically observed, there are multiple NE. If $\theta < k := \underline{\theta}$, then it is strictly dominant to abstain and so all abstain in equilibrium. If $\theta > k + c := \bar{\theta}$, then it is strictly dominant to attack and so all attack in equilibrium. For each $\theta \in [\underline{\theta}, \bar{\theta}]$, players can coordinate on either attacking or abstaining with probability one. However, if $\epsilon > 0$, and each player privately has a slightly different assessment of θ, then the Bayesian NE is unique: $\sigma^{NE}(x) = \mathbf{1}\{x > \frac{2k+c}{2}\}$.[13]

What about QRE? It is clear that, if $\epsilon = 0$, there are multiple QRE. For any given $\boldsymbol{Q}$, if $\theta \le \underline{\theta}$, then players abstain strictly more often than not in equilibrium. If $\theta \ge \bar{\theta}$, players attack strictly more often than not in equilibrium. For all $\theta \in (\underline{\theta}, \bar{\theta})$, however, players may coordinate on tending to attack or tending to abstain.[14] We also note that, for $\epsilon = 0$, conditions (A1)-(A4) are violated and so σ need not be continuous nor increasing. If $\epsilon > 0$, however, (A1)-(A4) are satisfied, and we will show that the QRE is unique for any given quantal response function. We assume $\epsilon > 0$ from now on.

[11] If $\theta < \epsilon$, then $x \sim U[0, 2\theta]$. If $\theta > 1-\epsilon$, $x \sim U[2\theta - 1, 1]$.

[12] In Morris and Shin (1998), θ represents the strength of the regime — the required mass of attacking players for regime change. In our version, it represents the direct value of attacking, regardless of whether or not the attack succeeds.

[13] If $x = \frac{2k+c}{2}$, they are indifferent and may mix arbitrarily.

[14] Unlike in NE for which players can coordinate either way when $\theta \in \{\underline{\theta}, \bar{\theta}\}$, here they must tend in one particular direction. Also note that, for $\theta \in (\underline{\theta}, \bar{\theta})$, the minimum attack probability required to support "tending to attack" (and similarly for abstaining) depends on θ.

We define some objects that will be used in our analysis and make a few useful observations. We define the *subjective failure probability* $P(x,\sigma,\epsilon) = \frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon} \mathbf{1}\left\{\frac{1}{2\epsilon}\int_{\theta'-\epsilon}^{\theta'+\epsilon}\sigma(x')dx' \leq \frac{1}{2}\right\}d\theta'$ as the subjective probability the attack will fail for a player who receives signal x, given the strategy σ and parameter ϵ. To understand the expression, a player who receives signal x forms the posterior that $\theta \sim U[x-\epsilon, x+\epsilon]$. Conditional on each value θ' in the support, the attack will fail if and only if $\frac{1}{2\epsilon}\int_{\theta'-\epsilon}^{\theta'+\epsilon}\sigma(x')dx' \leq \frac{1}{2}$. Hence, $P(x,\sigma,\epsilon)$ gives the subjective probability that some θ' satisfying $\frac{1}{2\epsilon}\int_{\theta'-\epsilon}^{\theta'+\epsilon}\sigma(x')dx' \leq \frac{1}{2}$ realizes. Using this expression, expected payoff differences are given by

$$\begin{aligned}\Delta\bar{u}_x(\sigma) &= [\mathbb{E}(\theta|x) - k - cP(x,\sigma,\epsilon)] - 0 \\ &= x - k - cP(x,\sigma,\epsilon).\end{aligned}$$

From Theorem 1, any QRE will be a strictly increasing function. There will thus be a *threshold state* $\theta^*(\sigma,\epsilon) = \{\theta' | \frac{1}{2\epsilon}\int_{\theta'-\epsilon}^{\theta'+\epsilon}\sigma(x')dx' = \frac{1}{2}\}$ (that does not depend on x) such that attacks will fail for $\theta \leq \theta^*(\sigma,\epsilon)$ and attacks will succeed for $\theta > \theta^*(\sigma,\epsilon)$. Hence, in any QRE

$$\begin{aligned}P(x,\sigma,\epsilon) &= \frac{1}{2\epsilon}\int_{x-\epsilon}^{x+\epsilon}\mathbf{1}\{\theta' \leq \theta^*(\sigma,\epsilon)\}d\theta' \\ &= \begin{cases} 1 & \iff x \leq \theta^*(\sigma,\epsilon) - \epsilon \\ \frac{\theta^*(\sigma,\epsilon) - x + \epsilon}{2\epsilon} & \iff x \in (\theta^*(\sigma,\epsilon)-\epsilon, \theta^*(\sigma,\epsilon)+\epsilon) \\ 0 & \iff x \geq \theta^*(\sigma,\epsilon) + \epsilon. \end{cases}\end{aligned} \tag{1}$$

In any QRE, the payoff difference $\Delta\bar{u}_x(\sigma)$ thus depends only on x and the threshold state $\theta^*(\sigma,\epsilon)$. An implication is that any given indifferent type $\tilde{x}$ pins down both $P(\tilde{x},\sigma,\epsilon)$ and $\theta^*(\sigma,\epsilon)$, and thus the payoffs to all other types. It is this limited dependence of payoffs on σ that makes the problem tractable. In particular, we make use of the following lemma, which follows from the expressions for $\Delta\bar{u}_x$ and $P(x,\sigma,\epsilon)$.

Lemma 2. *In any QRE σ,*

(1) $\tilde{x} \in (k, k+c)$,

(2) $P(\tilde{x},\sigma,\epsilon) = \frac{\tilde{x}-k}{c}$,

(3) $\theta^*(\sigma,\epsilon) = \tilde{x}(\frac{2\epsilon+c}{c}) - \frac{2k\epsilon+c\epsilon}{c} \in (\tilde{x}-\epsilon, \tilde{x}+\epsilon)$*, and*

(4) for $x' > x$, $\Delta\bar{u}_{x'} - \Delta\bar{u}_x = (x'-x) + c\frac{|[\theta^*(\sigma,\epsilon)-\epsilon, \theta^*(\sigma,\epsilon)+\epsilon] \cap [x,x']|}{2\epsilon} > 0$.

Proof. First, note that the indifference condition $\Delta\bar{u}_{\tilde{x}}(\sigma) = 0 = \tilde{x} - k - cP(\tilde{x},\sigma,\epsilon)$ yields $P(\tilde{x},\sigma,\epsilon) = \frac{\tilde{x}-k}{c}$, which is a valid probability for any $\tilde{x} \in [k, k+c]$.

It must be that $\tilde{x} \in [k, k+c]$ because abstaining is strictly dominant for $\tilde{x} < k$ and attacking is strictly dominant for $\tilde{x} > k+c$. Suppose $\tilde{x} = k$. In such a case, it must be that $\sigma(k) = \frac{1}{2}$ and $P(k,\sigma,\epsilon) = 0$ (from the expression $P(\tilde{x},\sigma,\epsilon) = \frac{\tilde{x}-k}{c}$, derived above). However, since Theorem 1 implies that $\sigma(x) < \frac{1}{2}$ for $x < k$ and that σ is continuous, it must be that $P(k,\sigma,\epsilon) > 0$ by definition of $P(\tilde{x},\sigma,\epsilon)$, a contradiction. A symmetric argument shows that it also cannot be that $\tilde{x} = k+c$. Hence, $\tilde{x} \in (k, k+c)$ and $P(\tilde{x},\sigma,\epsilon) = \frac{\tilde{x}-k}{c}$, which shows *(1)* and *(2)*.

Since $\tilde{x} \in (k, k+c)$, $P(\tilde{x},\sigma,\epsilon) = \frac{\tilde{x}-k}{c} \in (0,1)$, which implies $\tilde{x} \in (\theta^*(\sigma,\epsilon) - \epsilon, \theta^*(\sigma,\epsilon) + \epsilon)$, or equivalently $\theta^*(\sigma,\epsilon) \in (\tilde{x} - \epsilon, \tilde{x} + \epsilon)$, from (1). (1) also implies that $P(\tilde{x},\sigma,\epsilon) = \frac{\theta^*(\sigma,\epsilon) - \tilde{x} + \epsilon}{2\epsilon}$. Setting this equal to $P(\tilde{x},\sigma,\epsilon) = \frac{\tilde{x}-k}{c}$ and solving yields $\theta^*(\sigma,\epsilon) = \tilde{x}(\frac{2\epsilon+c}{c}) - \frac{2k\epsilon + c\epsilon}{c}$, which shows *(3)*.

For $x' > x$, $\Delta\bar{u}_{x'}(\sigma) - \Delta\bar{u}_x(\sigma) = [x' - k - cP(x',\sigma,\epsilon)] - [x - k - cP(x,\sigma,\epsilon)] = (x' - x) + c(P(x,\sigma,\epsilon) - P(x',\sigma,\epsilon))$. That $P(x,\sigma,\epsilon) - P(x',\sigma,\epsilon) = \frac{|[\theta^*(\sigma,\epsilon)-\epsilon, \theta^*(\sigma,\epsilon)+\epsilon] \cap [x,x']|}{2\epsilon}$ follows directly from (1), which shows *(4)*. □

Part *(1)* bounds the set of indifferent types. Parts *(2)* and *(3)* give, respectively, expressions for the indifferent type's subjective failure probability and the threshold state, as a function of the indifferent type. Part *(4)* makes clear that, as the signal varies between x and $x' > x$, the *difference in differences* of payoffs can be decomposed into a component related to changes in the direct value of attacking that is proportional to $|x' - x|$ and a component related to changes in the subjective failure probability. This latter component is proportional to how much the interval of change $[x,x']$ overlaps with the ϵ-neighborhood of the threshold state $[\theta^*(\sigma,\epsilon) - \epsilon, \theta^*(\sigma,\epsilon) + \epsilon]$.

Our first result establishes uniqueness. It is proved through repeated application of Lemma 2, but the intuition is similar to that of the classic result of Morris and Shin (1998): without a publically observed state, coordination is impossible.

Proposition 3. *For any Q satisfying (R1)-(R4), the QRE is unique.*

Proof. See Appendix A. □

With uniqueness established, we first characterize QRE.

Proposition 4. $\sigma : [0,1] \to (0,1)$ *is a QRE if and only if the indifferent type is* $\tilde{x} \in \tilde{R} = (k, k+c)$, σ *is continuous and strictly increasing, and the indifferent type's subjective failure probability is* $P(\tilde{x},\sigma,\epsilon) = \frac{\tilde{x}-k}{c}$.

Proof. Part *(1)* of Lemma 2 shows that $\tilde{R} \subset (k, k+c)$. Conversely, any $\tilde{x} \in (k, k+c)$ can be supported for some strictly increasing σ, as we show below, and thus $\tilde{R} = (k, k+c)$.

For any $\tilde{x} \in (k, k+c)$, the corresponding threshold state $\theta^*(\sigma,\epsilon) \in (\tilde{x}-\epsilon, \tilde{x}+\epsilon)$ required to support it as the indifferent type is unique from part *(3)* of Lemma 2. Conversely, any σ that is strictly increasing with $\sigma(\tilde{x}) = \frac{1}{2}$ such that $\frac{1}{2\epsilon}\int_{\theta^*(\sigma,\epsilon)-\epsilon}^{\theta^*(\sigma,\epsilon)+\epsilon} \sigma(x^{'})dx^{'} = \frac{1}{2}$ supports $\tilde{x}$ as the indifferent type. In what follows, fixing some $\tilde{x} \in (k, k+c)$ and letting $\tilde{\theta} \in (\tilde{x}-\epsilon, \tilde{x}+\epsilon)$ be *arbitrary*, we construct a function σ that is strictly increasing with $\sigma(\tilde{x}) = \frac{1}{2}$ such that $\frac{1}{2\epsilon}\int_{\tilde{\theta}-\epsilon}^{\tilde{\theta}+\epsilon} \sigma(x^{'})dx^{'} = \frac{1}{2}$. Since this can done for arbitrary $\tilde{\theta} \in (\tilde{x}-\epsilon, \tilde{x}+\epsilon)$, it can be done for $\theta^*(\sigma,\epsilon)$, and so we are done.

Consider the family of step functions $\hat{\sigma}(x) = \begin{cases} \tilde{\sigma}_L & x \le \tilde{x} \\ \tilde{\sigma}_H & x > \tilde{x} \end{cases}$, where $\tilde{\sigma}_L \in (0, \frac{1}{2})$ and $\tilde{\sigma}_H \in (\frac{1}{2}, 1)$ are constants. This family is neither continuous nor strictly increasing, but we can choose $\tilde{\sigma}_L$ and $\tilde{\sigma}_H$ such that $\frac{1}{2\epsilon}\int_{\tilde{\theta}-\epsilon}^{\tilde{\theta}+\epsilon} \hat{\sigma}(x^{'})dx^{'} = \frac{1}{2}$. Furthermore, $\hat{\sigma}$ can be approximated arbitrarily well by a continuous, strictly increasing σ satisfying $\sigma(\tilde{x}) = \frac{1}{2}$, and so this is enough for our purposes. Consider the case that $\tilde{\theta} < \tilde{x}$ (the case that $\tilde{\theta} \ge \tilde{x}$ is similar). In this case, $\frac{1}{2\epsilon}\int_{\tilde{\theta}-\epsilon}^{\tilde{\theta}+\epsilon} \hat{\sigma}(x^{'})dx^{'} = \frac{1}{2\epsilon}(\tilde{\sigma}_L(\epsilon+\tilde{x}-\tilde{\theta}) + \tilde{\sigma}_H(\tilde{\theta}+\epsilon-\tilde{x}))$. Constants $\tilde{\sigma}_L$ and $\tilde{\sigma}_H$ can always be chosen so that this expression equals $\frac{1}{2}$. This is because it equals $\tilde{\sigma}_L\alpha + \tilde{\sigma}_H(1-\alpha)$ for some $\alpha \in (\frac{1}{2}, 1)$, i.e. a linear combination of $\tilde{\sigma}_L$ and $\tilde{\sigma}_H$, and $\tilde{\sigma}_L$ can be chosen arbitrarily within $(0, \frac{1}{2})$ and $\tilde{\sigma}_H$ can be chosen arbitrarily within $(\frac{1}{2}, 1)$. The result now follows directly from Theorem 1. □

For QRE, the set of possible indifferent types is very large; it is the entire set of types for which neither action is dominant. Furthermore, since the threshold state $\theta^*(\sigma,\epsilon)$ depends only on $\sigma(x^{'})$ for $x^{'} \in (\theta^*(\sigma,\epsilon)-\epsilon, \theta^*(\sigma,\epsilon)+\epsilon)$, and the expected payoffs to each type x depend only on its distance to $\theta^*(\sigma,\epsilon)$, there are many σ that are consistent with any given indifferent type. By contrast, all SQRE are associated with a unique indifferent type, and σ must be symmetric about that indifferent type.

Proposition 5. *$\sigma : [0,1] \to (0,1)$ is an SQRE if and only if the indifferent type is $\tilde{x} \in \tilde{S} = \{\frac{2k+c}{2}\}$, σ is continuous and strictly increasing, and σ is symmetric: $\sigma(\tilde{x}-\delta) = 1-\sigma(\tilde{x}+\delta)$ for all $\delta \in [0, min\{\tilde{x}, 1-\tilde{x}\}]$.*

Proof. We show that, in an SQRE, it is impossible to have $P(\tilde{x},\sigma,\epsilon) > \frac{1}{2}$, and thus it must be that $P(\tilde{x},\sigma,\epsilon) \le \frac{1}{2}$. To this end, suppose $P(\tilde{x},\sigma,\epsilon) > \frac{1}{2}$. First, since $\theta^*(\sigma,\epsilon) \in (\tilde{x}-\epsilon, \tilde{x}+\epsilon) \iff \tilde{x} \in (\theta^*(\sigma,\epsilon)-\epsilon, \theta^*(\sigma,\epsilon)+\epsilon)$ by part *(3)* of Lemma 2, it must be that $P(\tilde{x},\sigma,\epsilon) = \frac{\theta^*(\sigma,\epsilon)-\tilde{x}+\epsilon}{2\epsilon}$ from (1). $P(\tilde{x},\sigma,\epsilon) > \frac{1}{2}$ thus implies that $\tilde{x} \in (\theta^*(\sigma,\epsilon)-\epsilon, \theta^*(\sigma,\epsilon))$. It is immediate from part *(4)* of Lemma 2 that $|\Delta\bar{u}_{\tilde{x}+\delta}(\sigma) - \Delta\bar{u}_{\tilde{x}}(\sigma)| \ge |\Delta\bar{u}_{\tilde{x}-\delta}(\sigma) - \Delta\bar{u}_{\tilde{x}}(\sigma)|$ for all $\delta \in [0,\epsilon]$. But, by symmetry, this implies that $\sigma(\tilde{x}+\delta) \geq \sigma(\tilde{x}-\delta)$ for all $\delta \in [0,\epsilon]$, which in turn, implies by definition of $\theta^*(\sigma,\epsilon)$ that $\theta^*(\sigma,\epsilon) \leq \tilde{x}$, a contradiction. Hence, it cannot be that $P(\tilde{x},\sigma,\epsilon) > \frac{1}{2}$.

A symmetric argument shows that it cannot be that $P(\tilde{x},\sigma,\epsilon) < \frac{1}{2}$. Hence, it must be that $P(\tilde{x},\sigma,\epsilon) = \frac{1}{2}$ and thus that $\tilde{x} = \theta^*(\sigma,\epsilon)$ from $P(\tilde{x},\sigma,\epsilon) = \frac{\theta^*(\sigma,\epsilon)-\tilde{x}+\epsilon}{2\epsilon}$ and $\tilde{x} = \frac{2k+c}{2}$ from part *(3)* of Lemma 2. In this case, part *(4)* of Lemma 2 implies that $|\Delta\bar{u}_{\tilde{x}+\delta}(\sigma) - \Delta\bar{u}_{\tilde{x}}(\sigma)| = |\Delta\bar{u}_{\tilde{x}-\delta}(\sigma) - \Delta\bar{u}_{\tilde{x}}(\sigma)|$ for all $\delta \in [0, min\{\tilde{x}, 1-\tilde{x}\}]$ and thus symmetry requires $\sigma(\tilde{x}-\delta) = 1-\sigma(\tilde{x}+\delta)$ for all $\delta \in [0, min\{\tilde{x}, 1-\tilde{x}\}]$. The result now follows directly from Theorem 2. □

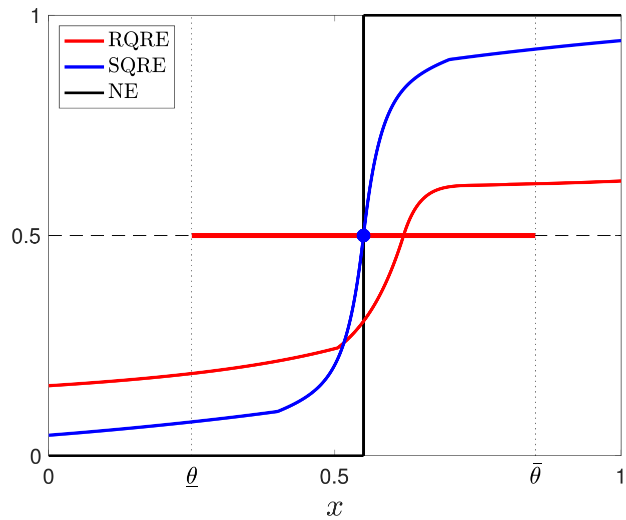

Figure 2: QRE and SQRE in the Global game

Notes: The figure depicts the NE (black) as well as illustrative QRE (red) and SQRE (blue), for $(k,c,\epsilon) = (0.25, 0.60, 0.15)$.

Figure 2 draws representative QRE (in red) and SQRE (in blue) for parameters $(k,c,\epsilon) = (0.25, 0.60, 0.15)$. QRE is consistent with a range of possible indifferent types, drawn as a horizontal red line. This particular QRE features a bias in favor abstaining: there exists $x < \tilde{x} < x'$ such that $\sigma(x) < 1-\sigma(x')$ and $|\Delta\bar{u}_x(\sigma)| \geq |\Delta\bar{u}_{x'}(\sigma)|$. By contrast, the SQRE is consistent with only one indifferent type, $\tilde{x} = \frac{2k+c}{2} = 0.55$, which coincides with that under NE. We also see that, like the NE, the SQRE is symmetric about this indifferent type.

3.3. Compromise Game

Following Carrillo and Palfrey (2009), two players simultaneously decide whether to fight or flee. Each player has a private type representing their strength $s \sim U[0,1]$, i.i.d. across players. If at least one player fights, the stronger player receives the high payoff of 1 and the weaker player receives the low payoff of 0.[15] If both players flee, each receives the compromise payoff $M \in (0, \frac{1}{2}]$. Let $\sigma(s)$ denote the probability that a player with strength s chooses to*flee*.

For type s, fleeing yields a payoff of M if the other player flees, 1 if the opponent fights *and* is of a lower type, and 0 otherwise. This gives an expected payoff of[16]

$$M \int_0^1 \sigma(s')ds' + s \cdot \frac{1}{s} \int_0^s (1 - \sigma(s'))ds'.$$

Fighting yields type s a payoff of 1 if the other player is of a lower type, resulting in an expected payoff of $\int_0^1 \mathbf{1}\{s' < s\}ds' = s$. The difference in expected payoffs is thus

$$\Delta_s \bar{u}(\sigma) = M\bar{\sigma} - s\sigma_L(s),$$

where, for convenience, we have defined $\bar{\sigma} = \int_0^1 \sigma(s')ds'$ and $\sigma_L(s) = \frac{1}{s}\int_0^s \sigma(s')ds'$ as the expected fleeing probabilities among all types and among types lower than s, respectively. Hence, the payoff difference for type s depends on σ only through the mean $\bar{\sigma}$ *and* the mean among lower types $\sigma_L(s)$. Unlike the example of Section 3.1, σ and s enter payoffs non-separably. This complicates the analysis, but we are still able to derive a detailed characterization. For some results that follow, it is convenient to also define the expected fleeing probability among types *higher* than s as $\sigma_H(s) = \frac{1}{1-s}\int_s^1 \sigma(s')ds'$. This allows us to rewrite $\bar{\sigma}$ as the weighted average $\bar{\sigma} = s\sigma_L(s) + (1-s)\sigma_H(s)$ for any s.

In this game, the unique NE is for both players to fight no matter the strength: $\sigma^{NE}(s) = 0$ for all s. One intuition is as follows. Clearly, there can be no NE in which higher types flee and lower types fight. But there can also not be a NE in which higher types fight when lower types flee. This is because a player is only pivotal when the other player flees; and so, conditioning on the pivotal event, the flee-ers with the highest types would deviate. In QRE, a trivial implication is that all types will flee with some probability. A more interesting implication is that some low types will actually prefer to flee, as fleeing is a best response to a higher type that also flees.

Proposition 6. *$\sigma : [0,1] \to (0,1)$ is a QRE if and only if the indifferent type is $\tilde{s} \in \tilde{R} = (0, M)$, σ is continuous and strictly decreasing, and $\mathbb{E}(\sigma(s)|s \in [\tilde{s}, 1]) = \mathbb{E}(\sigma(s)|s \in [0, \tilde{s}])(\tilde{s} - M\tilde{s})/(M - M\tilde{s})$.*

Proof. Letting $\tilde{s}$ denote the indifferent type, and setting $\tilde{\sigma}_L = \sigma_L(\tilde{s})$ and $\tilde{\sigma}_H = \sigma_H(\tilde{s})$, we write the

[15]In the measure zero event both players fight and have the same type, suppose they each receive 0.

[16]To understand the second term, $\int_0^1 \mathbf{1}\{s' < s\}ds' = s$ is the probability the other player is of a lower type, and $\frac{1}{s}\int_0^s(1 - \sigma(s'))ds'$ is the expected probability of fighting *conditional* on being a lower type.

indifference condition as

$$M(\tilde{s}\tilde{\sigma}_L + (1-\tilde{s})\tilde{\sigma}_H) - \tilde{s}\tilde{\sigma}_L = 0.$$

Solving for $\tilde{s}$ yields

$$\tilde{s} = \frac{\tilde{\sigma}_H M}{\tilde{\sigma}_H M + \tilde{\sigma}_L(1-M)}.$$

In any QRE, it must be that $\sigma(\tilde{s}) = \frac{1}{2}$ and σ is strictly decreasing by Theorem 1. The only restriction this imposes is $0 < \tilde{\sigma}_H < \frac{1}{2} < \tilde{\sigma}_L < 1$. Hence, it is easy to show that $\tilde{s}$ can be supported as the indifferent type if and only if $\tilde{s} \in (0, M)$. The result now follows directly from Theorem 1 $\square$

The top left panel of Figure 3 shows a QRE (red). We also depict the indifference condition graphically. The area of the pink rectangle is $\tilde{s}\tilde{\sigma}_L$, the area of the green rectangle is $M\bar{\sigma}$, and these areas must be equal to support $\tilde{s}$ as the indifferent type. Though it requires an argument that we develop later on, we claim that this is an example in which types are biased in favor of fighting in the sense that $\sigma(s) < 1 - \sigma(s^{'})$ whenever $s < \tilde{s} < s^{'}$ and $|\Delta \bar{u}_s(\sigma)| = |\Delta \bar{u}_{s'}(\sigma)|$.

In the next result, we show that imposing symmetry does not reduce the set of possible indifferent types, but it does impose significant structure on σ. Our characterization of symmetry is constructive. We show that, under symmetry, $\sigma|_{[0,\tilde{s}]}$ (σ restricted to types below $\tilde{s}$) completely pins down $\sigma|_{[\tilde{s},\{s'|\sigma(s')=1-\sigma(0)\}]}$ (σ restricted to types above $\tilde{s}$ whose behavior is not more extreme than that of the types below). Moreover, for any given $\sigma|_{[0,\tilde{s}]}$, symmetry is satisfied if and only if $\sigma|_{[\tilde{s},\{s'|\sigma(s')=1-\sigma(0)\}]}$ is the limit of a sequence of functions that depends only on $\sigma|_{[0,\tilde{s}]}$. Because of the construtive nature of the characterization, we give the formal statement of symmetry in Appendix B.

Proposition 7. *$\sigma : [0,1] \to (0,1)$ is an SQRE if and only if the indifferent type is $\tilde{s} \in \tilde{S} = (0, M)$, σ is continuous and strictly decreasing, $\mathbb{E}(\sigma(s)|s \in [\tilde{s}, 1]) = \mathbb{E}(\sigma(s)|s \in [0, \tilde{s}])(\tilde{s} - M\tilde{s})/(M - M\tilde{s})$, and σ is symmetric (see Appendix B for the formal statement of symmetry).*

Proof. Solving the indifference condition (borrowing notation from the proof of Proposition 6) yields

$$\tilde{s} = \frac{\tilde{\sigma}_H M}{\tilde{\sigma}_H M + \tilde{\sigma}_L(1-M)}.$$

We already know that $\tilde{S} \subset \tilde{R} = (0, M)$. Conversely, we show that any $\tilde{s} \in (0, M)$ can be achieved through a symmetric, strictly decreasing σ, and thus $\tilde{S} = (0, M)$. To see this, consider the family of step functions $\hat{\sigma}(s) = \begin{cases} \sigma_L^{'} & s \leq \tilde{s} \\ 1-\sigma_L^{'} & s > \tilde{s} \end{cases}$, where $\sigma_L^{'} \in (\frac{1}{2}, 1)$. This family is neither continuous nor strictly decreasing, but it does satisfy the implication of symmetry that $\Delta\bar{u}_s(\hat{\sigma}) = \Delta\bar{u}_{s'}(\hat{\sigma}) \implies \hat{\sigma}(s) = 1 - \hat{\sigma}(s^{'})$. Furthermore, it can be approximated arbitrarily well by a strictly decreasing, symmetric σ. Hence, any $\tilde{s}$ that can be attained via this family of functions can also be attained via σ satisfying the necessary conditions. Substituting $\hat{\sigma}$ into the indifference condition yields

$$\tilde{s} = \frac{(1-\sigma_L^{'})M}{(1-\sigma_L^{'})M + \sigma_L^{'}(1-M)},$$

which attains all $\tilde{s} \in (0, M)$ for some $\sigma_L^{'} \in (\frac{1}{2}, 1)$, and so $\tilde{S} = (0, M)$.

That strictly decreasing σ and $\mathbb{E}(\sigma(s)|s \in [\tilde{s}, 1]) = \mathbb{E}(\sigma(s)|s \in [0, \tilde{s}])(\tilde{s} - M\tilde{s})/(M - M\tilde{s})$ are necessary is immediate from Proposition 6. If, in addition, symmetry is satisfied, the result follows directly from Theorem 2. □

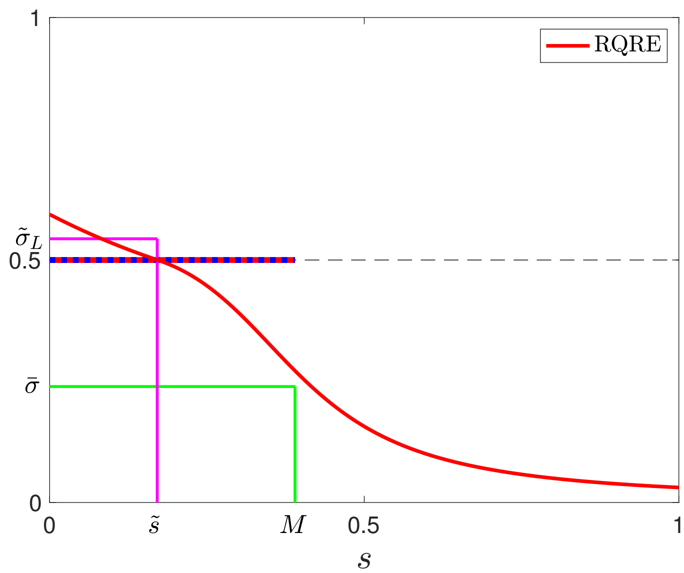

(a) Compromise Payoff M = .50

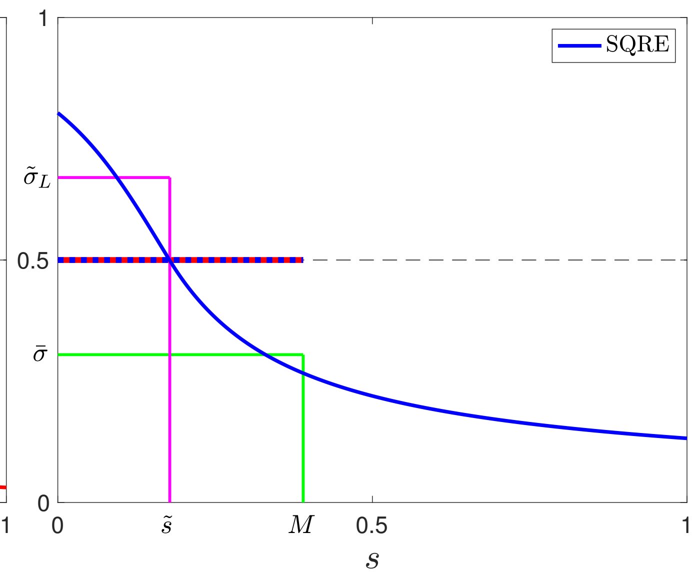

(b) Compromise Payoff M = .39

Figure 3: QRE and SQRE in the Compromise game

Notes: Panel 3a depicts a QRE (red) that is biased in favor of fighting, and panel 3b gives an SQRE (blue); in both cases $M = 0.39$.

The top right panel of Figure 3 shows an SQRE (blue), which unlike the QRE in the left panel, satisfies symmetry. While the full statement of symmetry is given in Appendix B, the next result gives two necessary conditions for symmetry that may be useful in applications for identifying symmetry violations.

Corollary 3. *In any SQRE, (1)* $\sigma(\tilde{s}-\epsilon) > 1-\sigma(\tilde{s}+\epsilon)$ *for any* $\epsilon \in [0,\tilde{s}]$*, and* (2) *if* $M < \frac{1}{2}$*, there exists* $\underline{S} \in (\tilde{s},1)$ *such that* $\sigma(s) > 1-\sigma(0)$ *for all* $s > \underline{S}$*.*

Proof. See Appendix A. □

Violations of these two conditions point to qualitatively different biases that may be part of QRE, but not SQRE. For instance, if *(1)* is violated, some types are biased in favor of fighting. If *(2)* is violated, some types are biased in favor of fleeing. An example of the former is given in the top left panel of Figure 3.

4. Empirical analysis of the compromise game

In this section, we leverage our analysis of the compromise game to nonparametrically test whether QRE is able rationalize the experimental data in Carrillo and Palfrey (2009).

The experiment has two variants corresponding to two values for the compromise payoff, $M \in \{.39, .50\}$, with types drawn uniformly at random from $[0,1]$, in increments of .01. The 56 recruited subjects were students at Princeton University and played for 20 incentivized rounds with random rematching and randomly redrawn types.[17]

Carrillo and Palfrey (2009) discuss the support for quantal response equilibrium based on comparing goodness-of-fit to other models. The authors first observe that, contra Nash equilibrium, compromise rates are strictly positive, decreasing in strength s, and increasing in the compromise payoff M — features consistent with QRE. They then fit different parametric models to the data — variations of logit QRE, Poisson-based cognitive hierarchy (Camerer et al., 2004), and cursed equilibrium (Eyster and Rabin, 2005), among others. They find that, while the QRE models provide a fairly good fit, it fails to capture the tendency of subjects to "fight with probability close to one when their strength is sufficiently high and with probability close to zero when their strength is sufficiently low." In order to capture this feature, they augment QRE with a cursedness parameter (α-QRE) and find a statistically better fit (rejecting $\alpha = 0$).

Importantly, the conclusions drawn in Carrillo and Palfrey (2009) regarding QRE are entirely

[17] Their experiment also includes two variants where choices are sequential; these were explained and played only after the simultaneous choice rounds were. We focus on the data with simultaneous choices as it matches our theoretical application from Section 3.3.

based on the logit functional form. To the extent these QRE models do not fully explain the data, a natural question then is whether QRE with a more general error structure can, in which case one need not posit additional behavioral parameters.

To answer this question, we will test whether *any* QRE is able to rationalize the data. We start by highlighting that Theorems 1 and 2 suggest a general methodology to nonparametrically test for the adequacy of QRE and SQRE to rationalize data. Specifically, by observing types and actions, one could in principle nonparametrically estimate σ and test whether (i) σ is decreasing and (ii) whether there exists a type $\tilde{t}$ such that $\sigma(\tilde{t}) = 1/2$ and $\Delta\bar{u}_{\tilde{t}}(\sigma) = 0$; continuity of σ, *strict* monotonicity, and interiority are not falsifiable. In specific applications, including the compromise game, we can relate the indifference condition $\Delta\bar{u}_{\tilde{t}}(\sigma) = 0$ to specific parameters and moment conditions based on σ, which further simplifies the analysis.

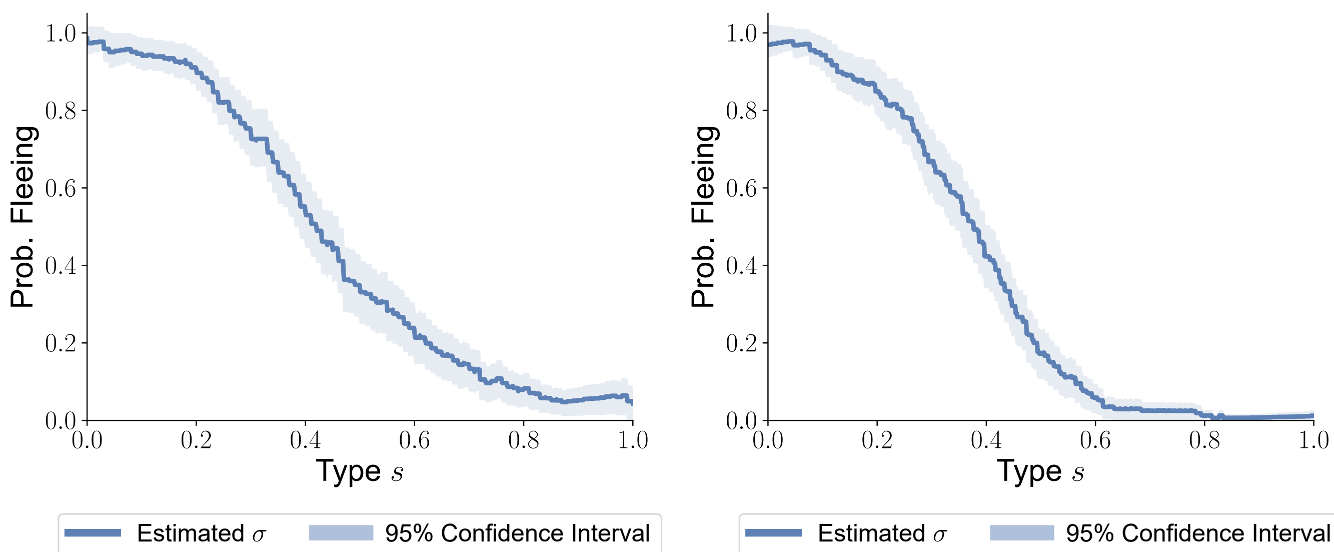

(a) Compromise Payoff M = .50

(b) Compromise Payoff M = .39

Figure 4: Estimated Choice Probability

Notes: The figure displays the estimated choice probability per type (line) and its 95% bootstrap confidence interval (shaded region) for the compromise game (Section 3.3). Point estimates were based on nonparametric kernel estimation using a uniform kernel with a bandwidth $h = n^{-1/5}\hat{\rho}$, where n refers to the number of observations and $\hat{\rho}$ the estimated standard deviation in choices. The data is from Carrillo and Palfrey (2009); we focus on simultaneous choice treatments.

Proposition 6 states three empirically testable conditions: (i) σ ought to be decreasing, (ii) the type uniformly randomizing is greater than 0 and lower than the compromise payoff M, and (iii) that $\beta = 0$, where $\beta := \mathbb{E}[\sigma(s) \mid s \in [\tilde{s}, 1]](1-\tilde{s}) - \mathbb{E}[\sigma(s) \mid s \in [0, \tilde{s}]]\tilde{s}\frac{1-M}{M}$. Figure 4 exhibits the estimated mapping from types to choice probabilities based on nonparametric kernel density estimation with a uniform kernel.[18] Although we are underpowered for an adequate test of monotonicity, both panels 4a and 4b depict decreasing estimated mappings. Furthermore, in both cases the indifferent type $\tilde{s}$ is lower than the compromise M as required by Proposition 6.

Finally, we test (iii) by nonparametrically estimating β and deriving confidence intervals based on bootstrapping. The resulting confidence intervals are shown in Table 1, which decidedly reject QRE — and therefore SQRE as well. In particular, the negative value of the statistic indicates the type who is independently randomizing would be better off by fighting. In other words, there seems to be a bias toward fleeing, which suggests subjects may be underestimating the frequencies with which lower types flee and/or higher types fight.

Carrillo and Palfrey (2009) find that logit QRE does not fully explain their data, which leads them to consider other behavioral models. We view our result as strengthening this conclusion: since a general form of QRE is rejected, one must step outside of the QRE family in order to fully explain the data.

Compromise Payoff	Confidence Interval 95%	Confidence Interval 99%
(1)	(2)	(3)
50	(-0.311 , -0.189)	(-0.330 , -0.170)
39	(-0.480 , -0.315)	(-0.500 , -0.291)

Table 1: Testing QRE

Notes: This table tests the adequacy of QRE in rationalizing the data for the compromise game (Section 3.3) from Carrillo and Palfrey (2009); we focus on simultaneous choice treatments. It provides bootstrapped confidence intervals for $\beta = \mathbb{E}[\sigma(s) \mid s \in [\tilde{s}, 1]](1-\tilde{s}) - \mathbb{E}[\sigma(s) \mid s \in [0, \tilde{s}]]\tilde{s}\frac{1-M}{M}$. By Proposition 6, this quantity should equal zero.

5. Conclusion

Quantal response equilibrium (QRE) explains many of the well-known deviations from Nash equilibrium (NE) observed in the lab. It should also be regarded as an important theoretical benchmark in that it deviates in a minimal way from NE. Nevertheless, its influence in theoretical applications is limited, perhaps due to concerns over its tractability. Recent work, focusing on finite games,

[18]The bandwidth chosen was $h = n^{-1/5}\hat{\rho}$, where n refers to the number of observations and $\hat{\rho}$ the estimated standard deviation in choices, based on Silverman's rule-of-thumb.

has made great strides by analyzing more general non-parametric forms. This has opened up the potential for richer applications and new ways of organizing experimental data.

In this paper, we provide analogous results for a common class of infinite games, those with binary actions and a continuum of types. Specifically, under a weak monotonicity condition on payoffs, we show that the full set of QRE is characterized by three simple conditions on choice probabilities: continuity, monotonicity, and uniform mixing of indifferent types. Further, we show how to recover the quantal response function from observable choices and types. We then apply our results to characterize QRE in a number of classic games and obtain sharp predictions. We conclude by illustrating the usefulness of our characterization in developing nonparametric tests of QRE. We believe that these results will inform both theoretical and empirical research, reducing reliance on parametric assumptions.

6. References

Aguirregabiria, V. and A. Magesan (2020). Identification and estimation of dynamic discrete games when players' beliefs are not in equilibrium. *Review of Economic Studies.*

Anderson, S., J. Goeree, and C. Holt (2001). Minimum-effort coordination games: Stochastic potential and logit equilibrium. *Games and Economic Behavior*.

Anderson, S., J. Goeree, and C. Holt (2002). The logit equilibrium: A perspective on intuitive behavioral anomalies. *Southern Economic Journal.*

Arad, A. and A. Rubinstein (2019). Multidimensional reasoning in games: Framework, equilibrium, and applications. *American Economic Journal: Microeconomics.*

Baye, M. and J. Morgan (2004). Price dispersion in the lab and on the internet: Theory and evidence. *The RAND Journal of Economics.*

Camerer, C., T.-H. Ho, and J.-K. Chong (2004). Behavioral game theory: thinking, learning, and teaching. *Advances in Understanding Strategic Behavior, Essays in Honor of Werner Guth.*

Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica.*

Carrillo, J. and T. Palfrey (2009). The compromise game: Two-sided adverse selection in the laboratory. *American Economic Journal: Microeconomics.*

Compte, O. and A. Postlewaite (2019). Ignorance and uncertainty. *Econometric Society Monographs.*

Diekmann, A. (1986). 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*.

Eyster, E. and M. Rabin (2005). Cursed equilibrium. *Econometrica*.

Friedman, E. (2022). Stochastic equilibria: Noise in actions or beliefs? *American Economic Journal: Microeconomics*.

Friedman, E. and F. Mauersberger (2022). Quantal response equilibrium with symmetry: Representation and applications. *Working paper*.

Goeree, J., C. Holt, and T. Palfrey (2005). Regular quantal response equilibrium. *Experimental Economics*.

Goeree, J., C. Holt, and T. Palfrey (2016). Quantal response equilibrium. *Princeton University Press*.

Goeree, J. and P. Louis (2021). M equilibrium: A theory of beliefs and choices in games. *American Economic Review*.

Goeree, J. K., C. A. Holt, and T. Palfrey (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*.

Golman, R. (2011). Quantal response equilibria with heterogeneous agents. *Journal of Economic Theory*.

Haile, P., A. Hortacsu, and G. Kosenok (2008). On the empirical content of quantal response equilibrium. *American Economic Review*.

Levine, D. and T. Palfrey (2007). The paradox of voter participation? a laboratory study. *American Political Science Review*.

McKelvey, R. and T. Palfrey (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*.

McKelvey, R. and T. Palfrey (1998). Quantal response equilibria for extensive form games. *Experimental Economics*.

Melo, E., K. Pogorelskiy, and M. Shum (2018). Testing the quantal response hypothesis. *International Economic Review*.

Morris, S. and H. S. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*.

Appendix A. Omitted Proofs

A.1. Proof of Lemma 1

Proof. Take any $\sigma \in \Sigma$ satisfying $\sigma = q(\sigma)$. We first show that any such σ must be continuous. Fix any $t \in T$ and take an arbitrary sequence $\{t_n\}_{n\in\mathbb{N}} \subseteq T$ converging to t. Then, by (A1), $\bar{u}_{t_n}(\sigma) \to \bar{u}_t(\sigma)$ which, by (R2) implies $Q(\bar{u}_{t_n}(\sigma)) \to Q(\bar{u}_t(\sigma))$ and thus $\sigma(t_n) = q(\sigma)(t_n) = Q(\bar{u}_{t_n}(\sigma)) \to Q(\bar{u}_t(\sigma)) = \sigma(t)$. We now prove that any such σ must be strictly decreasing. Suppose for the purpose of contradiction that there are $t < t'$ such that $\sigma(t) \leq \sigma(t')$. Then, by (A2), there are $\hat{t}, \hat{t}'$ such that $\sigma(\hat{t}) \leq \sigma(\hat{t}')$ and $\bar{u}^1_{\hat{t}}(\sigma) \geq \bar{u}^1_{\hat{t}'}(\sigma)$ and $\bar{u}^0_{\hat{t}}(\sigma) \leq \bar{u}^0_{\hat{t}'}(\sigma)$ with at least one of the payoff inequalities being strict. From (R3), $\sigma(\hat{t}) = q(\sigma)(\hat{t}) = Q(\bar{u}_{\hat{t}}(\sigma)) > Q(\bar{u}_{\hat{t}'}(\sigma)) = q(\sigma)(\hat{t}') = \sigma(\hat{t}')$, a contradiction.

We then note that the image of q, $q(\Sigma)$, is a subset of the space of continuous nonincreasing functions mapping from compact set T to $[0,1]$, denoted $\mathscr{S}$, which is compact with respect to the bounded variation norm, $\|\cdot\|_{BV}$.[19]

We now want to show that q (restricted to $\mathscr{S}$, the relevant domain) is continuous with respect to $\|\cdot\|_{BV}$. Let $\|\bar{u}\|_\infty := \max_{t\in T, \sigma\in\mathscr{S}} \|\bar{u}_t(\sigma)\|_\infty$, which is well-defined by Weierstrass extremum theorem and (A1). Let $\mathscr{U} := [-\|\bar{u}\|_\infty, \|\bar{u}\|_\infty]^2 \subset \mathbb{R}^2$, which is a compact superset of the domain of expected payoffs when restricting opponents' (symmetric) strategies to $\mathscr{S}$. Restricting Q to $\mathscr{U}$ renders it uniformly continuous by Heine-Cantor's theorem. Take any sequence $\{\sigma_n\}_n \subseteq \mathscr{S}$ such that $\|\sigma_n - \sigma\|_{BV} \to 0$. Since $\|\sigma_n - \sigma\|_{BV} \to 0 \Longrightarrow \|\sigma_n - \sigma\|_{L^1} \to 0$, then by Berge's theorem of the maximum $\max_{t\in T} \|\bar{u}_t(\sigma_n) - \bar{u}_t(\sigma)\|_\infty$ is continuous and converges to zero. This implies that for every $\epsilon > 0$ there exists $N < \infty$ such that for all $n > N$, and all $t \in T$, $\|\bar{u}_t(\sigma_n) - \bar{u}_t(\sigma)\|_\infty < \epsilon$. Combining this with uniform continuity of Q when restricted to the relevant domain, we obtain that for every $\epsilon > 0$ there exists $N < \infty$ such that for all $n > N$, and all $t \in T$, $|q(\sigma_n)(t) - q(\sigma)(t)| = |Q(\bar{u}_t(\sigma_n)) - Q(\bar{u}_t(\sigma))| < \epsilon$, and therefore $q(\sigma_n)$ converges uniformly to $q(\sigma)$, thereby converging also with respect to $\|\cdot\|_{BV}$.

Finally, we observe that $q(\Sigma) \subseteq \mathscr{S}$ and thus $q(\mathscr{S}) \subseteq \mathscr{S}$, and that $\mathscr{S}$ is in turn a subset of the space of functions with bounded variation defined on T, $BV(T) := \{f \in [0,1]^T \mid f \text{ is continuous and } V(f) < \infty\}$, itself a Banach space with respect to $\|\cdot\|_{BV}$. Since $\mathscr{S}$ is compact with respect to $\|\cdot\|_{BV}$ and convex, by Schauder's fixed-point theorem, a fixed point $\sigma = q(\sigma)$

[19] That is, $\|\sigma\|_{BV} := |\sigma(0)| + V(\sigma)$, where $V(\sigma)$ denotes the total variation of $\sigma \in \mathscr{S}$.

exists. □

A.2. Proof of Corollary 3

Proof. *(1)*: Referring to objects defined in Lemma 3, this follows from the fact that $\epsilon_t > s_{t-1} - s_t$ for all p_T and t. To see this,

$$\epsilon_t = \frac{D_{t-1}}{\frac{1}{2}L_{t-1} + \frac{1}{2}L_t} = \frac{\int_{s_t}^{s_{t-1}} \sigma(s^{'})ds^{'}}{\frac{1}{2}L_{t-1} + \frac{1}{2}L_t} = \frac{2\int_{s_t}^{s_{t-1}} \sigma(s^{'})ds^{'}}{L_{t-1} + L_t} > s_{t-1} - s_t.$$

where the inequality follows because $L_{t-1} + L_t < 1$ and $\sigma(s^{'}) > \frac{1}{2}$ for all $s^{'} \in [0, \tilde{s})$, and thus $\frac{2\int_{s_t}^{s_{t-1}} \sigma(s^{'})ds^{'}}{L_{t-1}+L_t} > 2\int_{s_t}^{s_{t-1}} \sigma(s^{'})ds^{'} > 2\int_{s_t}^{s_{t-1}} \frac{1}{2}ds^{'} = s_{t-1} - s_t$.

(2): This follows from the fact that $\Delta\bar{u}_0(\sigma) = M\bar{\sigma} - 0\sigma_L(0) = M\bar{\sigma}$ and $\Delta\bar{u}_1(\sigma) = M\bar{\sigma} - 1\sigma_L(1) = M\bar{\sigma} - \bar{\sigma} = (M-1)\bar{\sigma}$. Since $M < \frac{1}{2}$, $|\Delta\bar{u}_1(\sigma)| > |\Delta\bar{u}_0(\sigma)|$ and so symmetry requires $\sigma(1) > 1 - \sigma(0)$. Because σ is strictly decreasing and continuous, there exists $\underline{S} \in (\tilde{s}, 1)$ such that $\sigma(s) > 1 - \sigma(0)$ for all $s > \underline{S}$. □

A.3. Proof of Proposition 3

Proof. Consider $\sigma \neq \sigma^{'}$ as two candidate QRE for the same Q satisfying (R1)-(R4). We consider three different cases, and for each derive a contradiction. In what follows, we let $\tilde{x}$ and $\tilde{x}^{'}$ be the indifferent types under σ and $\sigma^{'}$, respectively. Further, let $\theta^* = \theta^*(\sigma, \epsilon)$ and $\theta^{*'} = \theta^*(\sigma^{'}, \epsilon)$ be the corresponding threshold states.

Case 1. Suppose $\tilde{x} = \tilde{x}^{'}$. The indifference type pins down all payoffs, and thus, it must be that $\sigma = \sigma^{'}$, a contradiction.

Case 2. Suppose $\tilde{x} < \tilde{x}^{'}$ and $\sigma(x) = \sigma^{'}(x) > \frac{1}{2}$ at some $x \in (\tilde{x}^{'}, \theta^{*'} + \epsilon)$.[20] Since $\tilde{x} < \tilde{x}^{'}$, we have that $\theta^* < \theta^{*'}$ from part *(3)* of Lemma 2 and thus also that $P(x, \sigma, \epsilon) < P(x, \sigma^{'}, \epsilon)$ from (1) since $x \in (\theta^{*'} - \epsilon, \theta^{*'} + \epsilon)$. Hence, it must be that $\Delta\bar{u}_x(\sigma) > \Delta\bar{u}_x(\sigma^{'})$ and thus $\sigma(x) > \sigma^{'}(x)$, a contradiction.

Case 3. Suppose $\tilde{x} < \tilde{x}^{'}$ and $\sigma(x) \geq \sigma^{'}(x)$ for all $x \in [\theta^* - \epsilon, \theta^{*'} + \epsilon]$ and strictly so for $x \in (\theta^* - \epsilon, \theta^{*'} + \epsilon)$.[21] Since $\tilde{x} < \tilde{x}^{'}$, we have that $P(\tilde{x}, \sigma, \epsilon) < P(\tilde{x}^{'}, \sigma^{'}, \epsilon)$ and $\theta^* - \tilde{x} < \theta^{*'} - \tilde{x}^{'}$ by parts *(2)-(3)* of Lemma 2.[22] But then, by part *(4)* of Lemma 2, we have that $\Delta\bar{u}_{\tilde{x}^{'}+\delta}(\sigma^{'}) \geq \Delta\bar{u}_{\tilde{x}+\delta}(\sigma)$ for all $\delta \in [0, \epsilon + \theta^{*'} - \tilde{x}^{'}]$ and strictly so for all $\delta \in (\epsilon + \theta^* - \tilde{x}, \epsilon + \theta^{*'} - \tilde{x}^{'})$. Hence, it must be that $\sigma^{'}(\tilde{x}^{'} + \delta) \geq \sigma(\tilde{x} + \delta)$ for

[20] It cannot be that, for $x \leq \tilde{x}^{'}$, $\sigma(x) = \sigma^{'}(x) > \frac{1}{2}$ since $\sigma^{'}(x) \leq \frac{1}{2}$ for all such x. The case that $\sigma(x) = \sigma^{'}(x) < \frac{1}{2}$ for some $x \in (\theta^* - \epsilon, \tilde{x})$ is symmetric.

[21] It must be that $\sigma(x) = \sigma^{'}(x)$ for $x \geq \theta^{*'} + \epsilon$ as, for all such x, $P(x, \sigma, \epsilon) = P(x, \sigma^{'}, \epsilon) = 0$ and thus $\Delta\bar{u}_x(\sigma) = \Delta\bar{u}_x(\sigma^{'})$.

[22] By part *(3)* of the lemma, it is easy to show that $\theta^*(\sigma, \theta) - \tilde{x}$ is increasing in $\tilde{x}$.

all $\delta \in [0, \epsilon + \theta^{*'} - \tilde{x}']$ and strictly so for all $\delta \in (\epsilon + \theta^* - \tilde{x}, \epsilon + \theta^{*'} - \tilde{x}')$. In other words, σ' must be steeper than σ to the right of their respective indifference points. A symmetric argument shows that σ must be steeper than σ' to the left of their respective indifference points (i.e. when both curves fall below zero). But in this case, it must be that $P(\tilde{x}, \sigma, \epsilon) > P(\tilde{x}', \sigma', \epsilon)$ by definition of $P(\cdot)$, a contradiction. □

Appendix B. Symmetry in the Compromise Game

The following lemma characterizes symmetry in the compromise game. We conclude this section by providing some intution and discussion of the result.

Lemma 3. *Symmetry is satisfied in the Compromise game if and only if* $\sigma|_{[\tilde{s}, \{s'|\sigma(s')=1-\sigma(0)\}]}$ *is the uniform limit of the function sequence* $\{\sigma^{p_T}\}_T$ *defined by the following procedure (which depends on* $\sigma|_{[0,\tilde{s}]}$*, but not* $\sigma|_{(\tilde{s},1]}$*):*

1. Let $p = \{[h_0, h_1], [h_1, h_2], ..., [h_{T-1}, h_T]\}$ be a partition of $[\frac{1}{2}, \sigma(0)]$ into T intervals such that $h_0 = \frac{1}{2}$, $h_{t'} < h_{t''}$ for $t'' > t'$, and $h_T = \sigma(0)$. Let $\delta_t = h_{t+1} - h_t$ be the length of the tth interval.

2. Let $s_t = \sigma^{-1}(h_t)$ be the type associated with h_t, noting that $s_0 = \tilde{s}$ and $s_T = 0$. Let $D_t = \Delta\bar{u}_{s_{t+1}}(\sigma) - \Delta\bar{u}_{s_t}(\sigma) = \int_{s_{t+1}}^{s_t} \sigma(s')ds'$ for $t = 0, 1..., T-1$.[23]

3. Let $L_0 = \frac{1}{2}$ and $L_t = L_{t-1} - \delta_{t-1}$ for $t = 1, 2, ..., T$.

4. Let $\epsilon_t = \frac{D_{t-1}}{\frac{1}{2}L_{t-1} + \frac{1}{2}L_t}$ for $t = 1, 2, ..., T$.

5. Let $(x_0, y_0) = (\tilde{s}, \frac{1}{2})$ and $(x_t, y_t) = (x_{t-1} + \epsilon_t, y_{t-1} - \delta_{t-1})$ for $t = 1, 2, ..., T$.

6. Let $\sigma^p : [\tilde{s}, \{s'|\sigma(s') = 1 - \sigma(0)\}] \to [0, \frac{1}{2}]$ be the piece-wise linear function with vertices defined by $(x_t, y_t)_{t=1,2,...,T}$

7. Let p_T be a partition associated with T intervals, and consider a sequence of partitions $\{p_T\}_T$ as $T \to \infty$ such that the mesh (maximum interval length) goes to 0. Let $\{\sigma^{p_T}\}_T$ be the associated sequence of functions.

Proof. After setting $\sigma|_{[0,\tilde{s}]}$ to be continuous and strictly decreasing with $\sigma(\tilde{s}) = \frac{1}{2}$, $\sigma|_{[\tilde{s}, \{s'|\sigma(s')=1-\sigma(0)\}]}$ is completely pinned down by symmetry. It is the unique function such that, for any $s < s' \in [0, \tilde{s}]$, there exist unique $s'' < s''' \in [0, \{s'|\sigma(s') = 1 - \sigma(0)\}]$ satisfying $\sigma(s) = 1 - \sigma(s''')$,

[23]Note that D_t depends only on $\sigma|_{[0,\tilde{s}]}$ and is strictly greater than zero.

$\sigma(s^{'}) = 1-\sigma(s^{''})$, and $|\Delta_s \bar{u}(\sigma) - \Delta_{s^{'}} \bar{u}(\sigma)| = |\Delta_{s^{''}} \bar{u}(\sigma) - \Delta_{s^{'''}} \bar{u}(\sigma)|$. We show that the uniform limit of $\{\sigma^{p_T}\}_T$ satisfies this property, and therefore, is the function that we seek.

To this end, choose arbitrary $s < s^{'} \in [0, \tilde{s}]$. Since each σ^{p_T} is strictly decreasing and $\sigma^{p_T}(\{s^{'} | \sigma(s^{'}) = 1-\sigma(0)\}) = 1-\sigma(0)$ by construction, for all p_T, there exists unique $s^{''} < s^{'''} \in [0, \{s^{'}_i | \sigma(s^{'}_i) = 1-\sigma(0)\}]$ satisfying $\sigma(s) = 1-\sigma(s^{'''})$ and $\sigma(s^{'}) = 1-\sigma(s^{''})$. What remains is to show is that, for all $\epsilon > 0$, there is sufficiently large T (the mesh of p_T is being sufficiently small) such that $|\Delta_s \bar{u}(\sigma) - \Delta_{s^{'}} \bar{u}(\sigma)| - |\Delta_{s^{''}} \bar{u}(\sigma^{p_T}) - \Delta_{s^{'''}} \bar{u}(\sigma^{p_T})| < \epsilon$. Since $[0, \tilde{s}]$ is compact, the same ϵ works for *all* s and $s^{'}$, and thus this is enough to establish uniform convergence.

For fixed p_T, suppose $s \in [s_{t+1}, s_t]$ and $s^{'} \in [s_{k+1}, s_k]$ for $k < t$. Then, necessarily, $s^{''} \in [x_k, x_{k+1}]$ and $s^{'''} \in [x_t, x_{t+1}]$. Recalling that $\Delta_s \bar{u}(\sigma) - \Delta_{s^{'}} \bar{u}(\sigma) = \int_s^{s^{'}} \sigma(z)dz$, this implies that $\Delta_s \bar{u}(\sigma) - \Delta_{s^{'}} \bar{u}(\sigma) \in [\int_{s_t}^{s_{k+1}} \sigma(z)dz, \int_{s_{t+1}}^{s_k} \sigma(z)dz]$.
A similar argument gives $\Delta_{s^{''}} \bar{u}(\sigma^{p_T}) - \Delta_{s^{'''}} \bar{u}(\sigma^{p_T}) \in [\int_{x_{k+1}}^{x_t} \sigma^{p_T}(z)dz, \int_{x_k}^{x_{t+1}} \sigma^{p_T}(z)dz]$. It is easy to show that $\int_{x_h}^{x_{h+1}} \sigma^{p_T}(z)dz = \epsilon_{h+1} L_{h+1} + \frac{1}{2}\epsilon_{h+1}\delta_h = \int_{s_{h+1}}^{s_h} \sigma(z)dz$ by substituting $\epsilon_{h+1} = \frac{D_h}{\frac{1}{2}L_h + \frac{1}{2}L_{h+1}} = \frac{\int_{s_{h+1}}^{s_h} \sigma(z)dz}{\frac{1}{2}L_h + \frac{1}{2}L_{h+1}}$ and $\delta_h = L_h - L_{h+1}$. Hence, $\int_{x_{k+1}}^{x_t} \sigma^{p_T}(z)dz = \sum_{h=k+1}^{t-1} \epsilon_{h+1} L_{h+1} + \frac{1}{2}\epsilon_{h+1}\delta_h = \sum_{h=k+1}^{t-1} \int_{s_{h+1}}^{s_h} \sigma(z)dz = \int_{s_t}^{s_{k+1}} \sigma(z)dz$. Similarly, $\int_{x_k}^{x_{t+1}} \sigma^{p_T}(z)dz = \int_{s_{t+1}}^{s_k} \sigma(z)dz$, and thus we have that $\Delta_{s^{''}} \bar{u}(\sigma^{p_T}) - \Delta_{s^{'''}} \bar{u}(\sigma^{p_T}) \in [\int_{s_t}^{s_{k+1}} \sigma(z)dz, \int_{s_{t+1}}^{s_k} \sigma(z)dz]$. This implies that $|\Delta_s \bar{u}(\sigma) - \Delta_{s^{'}} \bar{u}(\sigma)| - |\Delta_{s^{''}} \bar{u}(\sigma^{p_T}) - \Delta_{s^{'''}} \bar{u}(\sigma^{p_T})| \leq \int_{s_{t+1}}^{s_k} \sigma(z)dz - \int_{s_t}^{s_{k+1}} \sigma(z)dz = \int_{s_{t+1}}^{s_t} \sigma(z)dz + \int_{s_{k+1}}^{s_k} \sigma(z)dz$. But since $s_t - s_{t+1} < \frac{\epsilon}{2}$ and $s_k - s_{k+1} < \frac{\epsilon}{2}$ for large T and $\sigma(s) < 1$ for all s, $|\Delta_s \bar{u}(\sigma) - \Delta_{s^{'}} \bar{u}(\sigma)| - |\Delta_{s^{''}} \bar{u}(\sigma^{p_T}) - \Delta_{s^{'''}} \bar{u}(\sigma^{p_T})| < \epsilon$, and we are done. □

To see how symmetry is established as the limit of σ^{P_T}, refer to the bottom panel of the figure, which illustrates the construction for $T = 3$. Here, $\sigma|_{[0,\tilde{s}]}$ is fixed whereas $\sigma|_{[\tilde{s}, \{s^{'} | \sigma(s^{'}) = 1-\sigma(0)\}]}$ is approximated through σ^{P_T}. Consider the types $0 = s_3 < s_2 < s_1 < s_0 \equiv \tilde{s}$. The *difference in* expected utility *differences* in going from s_0 to s_1 is $\Delta\bar{u}_{s_1}(\sigma) - \Delta\bar{u}_{s_0}(\sigma) = \{M\bar{\sigma} - s_1\sigma_L(s_1)\} - \{M\bar{\sigma} - s_0\sigma_L(s_0)\} = s_0\sigma_L(s_0) - s_1\sigma_L(s_1) > 0$, which depends only on $\sigma|_{[0,\tilde{s}]}$. The corresponding difference in fleeing probabilities is $\delta_0 \equiv \sigma(s_1) - \sigma(s_0) > 0$. Using these terms, we can linearly approximate σ starting from s_0 and moving a little bit to the right. That is, we build σ^{P_T} by linearly interpolating from $(x_0, y_0) = (s_0, \sigma(s_0))$ to $(x_1, y_1) = (s_0 + \epsilon_1, \sigma(s_0) - \delta_0)$ where ϵ_1 is chosen so that $\Delta\bar{u}_{s_0}(\sigma^{P_T}) - \Delta\bar{u}_{s_0+\epsilon_1}(\sigma^{P_T}) = \Delta\bar{u}_{s_1}(\sigma) - \Delta\bar{u}_{s_0}(\sigma)$. This ensures that type s_1 is *exactly* symmetric to type $s_0 + \epsilon_1$ under the approximation, i.e. that $|\Delta\bar{u}_{s_1}(\sigma)| = |\Delta\bar{u}_{s_0+\epsilon_1}(\sigma^{P_T})|$ and $\sigma(s_1) = 1 - \sigma^{P_T}(s_0 + \epsilon_1)$. Furthermore, there is *approximate* symmetry for all types $s \in [s_1, s_0 + \epsilon_1]$. The rest of σ^{P_T} is built in a similar fashion, using terms $\Delta\bar{u}_{s_{t+1}}(\sigma) - \Delta\bar{u}_{s_t}(\sigma)$ and $\delta_t \equiv \sigma(s_{t+1}) - \sigma(s_t)$ to successively find the next piece-wise linear segment of σ^{P_T} until σ^{P_T} is defined for all $s \in [\tilde{s}, \{s^{'}|\sigma(s^{'}) = 1 - \sigma(0)\}]$.

It is important to note that, even though the limit of $\{\sigma^{p_T}\}_T$ is well-defined for all $\sigma|_{[0,\tilde{s}]}$ and that $\sigma|_{[\tilde{s},\{s'|\sigma(s')=1-\sigma(0)\}]}$ must be the limit of $\{\sigma^{p_T}\}_T$ in any SQRE, not all $\sigma|_{[0,\tilde{s}]}$ are consistent with SQRE. This is because $\sigma|_{[0,\tilde{s}]}$ pins down not only the necessary $\sigma|_{[\tilde{s},\{s'|\sigma(s')=1-\sigma(0)\}]}$, but also that the mean $\bar{\sigma}$ must be $\frac{\tilde{s}\tilde{\sigma}_L}{M}$ from the indifference condition. Hence, it must be that $\sigma|_{[0,\tilde{s}]}$ and the necessary $\sigma|_{[\tilde{s},\{s'|\sigma(s')=1-\sigma(0)\}]}$ leave enough "slack" for $\sigma|_{[\{s'|\sigma(s')=1-\sigma(0)\},1]}$ to be drawn so that $\bar{\sigma} = \frac{\tilde{s}\tilde{\sigma}_L}{M}$.

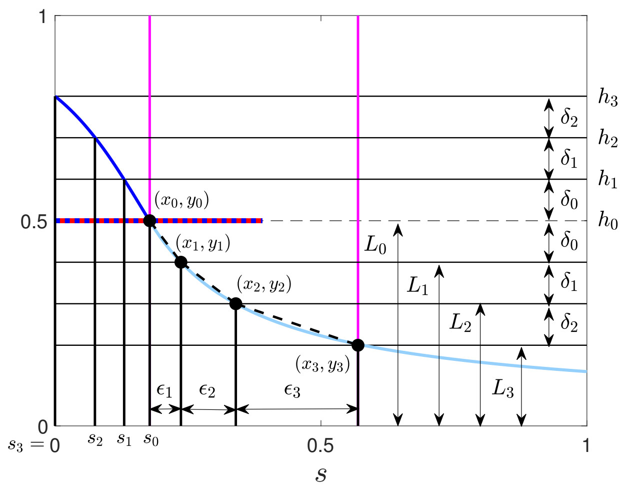

Figure 5: Symmetry in the Compromise game

Notes: The figure shows the construction of $\{\sigma^{p_T}\}_T$ used in Lemma 3 for $T = 3$.