The Foundations of Empirical Equilibrium*

Rodrigo A. Velez† and Alexander L. Brown‡

Department of Economics, Texas A&M University, College Station, TX 77843

March 15, 2023

Abstract

We study the foundations of *empirical equilibrium*, a refinement of Nash equilibrium that is based on a non-parametric characterization of empirical distributions of behavior in games with observable payoffs (Velez and Brown, 2020b). We show that the refinement can be alternatively defined as those Nash equilibria that are the limits of increasingly sophisticated regular QRE of Goeree et al. (2005). By contrast, some empirical equilibria cannot be approximated by all monotone additive randomly perturbed payoff models, including monotone structural QRE of McKelvey and Palfrey (1995). As a byproduct, we answer a question posed by Goeree et al. (2005) regarding the foundations of QRE models: For a fixed payoff matrix, the empirical content of regular and monotone structural QRE differ in fundamental ways.

JEL classification: C72, D47, D91.

Keywords: equilibrium refinements; behavioral game theory; structural monotone QRE; regular QRE; empirical equilibrium; randomly perturbed payoff models; deterministically perturbed payoff models.

*This paper benefited from comments of Dustin Beckett, Antonio Cabrales, Thomas Palfrey, Brian Rogers, and Anastasia Zervou as well as audiences at the 25+ Years of QRE conference at Caltech, the 2018 Naples Workshop on Equilibrium Analysis, the 2019 World Economic Science Association Meetings, the 6th annual Texas Experimental Association Symposium, and seminars at Boston College, Chapman U., Kellogg (MEDS), NCSU, Ohio State University, TETC18, UCSD, U. Maryland, U. Rochester, UTDallas and U. Virginia. The concept of empirical equilibrium emerged from the need to inform the mechanism design paradigm with the accumulated experience in experimental and empirical data. Partial progress in this endeavor was communicated in a series of unpublished manuscripts that contain parts of the material in the current paper and were never submitted for publication. In order to organize the material for publication we have separated results in four independent papers. We credit Velez and Brown (2020b) with introducing the refinement in the problem of dominant strategy implementation, so far its most prominent application. The current paper studies the foundations of the refinement. It subsumes results circulated in Velez and Brown (2020a, 2019). Two additional papers study the application of the refinement to the problem of Nash implementation (Velez and Brown, 2020a; Brown and Velez, 2020). All errors are our own.

†rvelezca@tamu.edu; https://sites.google.com/site/rodrigoavelezswebpage/home

‡alexbrown@tamu.edu; http://people.tamu.edu/~alexbrown

1 Introduction

This paper studies the foundations of Empirical Equilibrium (EE), a refinement of Nash equilibrium introduced by Velez and Brown (2020b). The refinement informs full implementation theory with the insights accumulated from experiments in finite simultaneous-move games. This paper provides foundations for this refinement and settles its relationship with previous refinements in the literature. As a byproduct, it considerably advances our understanding of perturbed payoff models.

A mechanism designer evaluates economic institutions by analyzing the properties of the outcomes predicted by a non-cooperative solution for the games that ensue under the institutions. The *full* implementation approach to mechanism design has two objectives: (i) to provide agents the oppurtunity to non-cooperatively coordinate on optimal outcomes; and (ii) that whenever a non-cooperative solution is reached, the society obtains optimal outcomes.

When the Nash equilibrium prediction is used as the basis for this analysis, condition (ii) is unattainable by any simultaneous-move mechanism for several different social objectives (Maskin, 1999; Palfrey and Srivastava, 1989; Jackson, 1991). If we allow that some Nash equilibria are "implausible" and should therefore be excluded from this analysis, then this conclusion is overly pessimistic. Indeed, when one simply restricts the prediction to Nash equilibria that involve no weakly dominated behavior, condition (ii) is easily satisfied in complete information environments (Palfrey and Srivastava, 1991; Jackson, 1992). Provided one's definition of "implausible" excludes behavior that is commonly and robustly found in experimental games, this conclusion would appear too optimistic. Weakly dominated behavior can be persistently observed in experiments, even when frequencies of play approximate a Nash equilibrium (see Velez and Brown, 2020b).

Retaining the spirit of this full implementation approach, Velez and Brown (2020b) refine Nash equilibrium based on empirical regularities in experiments, to obtain a more evidence-based approach for design. They note that in experimental games repeated multiple times—where agents have a chance to form rational expectations of the other agents' frequencies of play—Monotone Noisy Best Response Equilibrium (MNBRE) models (including Monotone Structural Quantal Response Equilibrium (MSQRE) of McKelvey and Palfrey, 1995) typically predict final period averages and comparative statics across treatments well (Goeree et al., 2016, 2018b).

Velez and Brown (2020b) reason that if one endorses the hypothesis that a particular theory is well specified for behavior in a game, one immediately produces a refinement of Nash equilibrium. For a Nash equilibrium to be empirically relevant (under the hypothesis that the theory is well specified) it needs to be in the closure of the empirical content of the theory. If instead, a Nash equilibrium exists that cannot be arbitrarily approximated by data generated by the theory, datasets that resemble such Nash equilibrium are necessarily rejections of the theory.

The family of MNBRE models contains a significant larger set of models than McKelvey and Palfrey (1995)'s MSQRE. It includes the Regular QRE (RQRE) model of McKelvey and Palfrey (1996) and Goeree et al. (2005), the control cost model of van Damme (1991), and the Harsanyi (1973)-type exchangeable perturbation additive randomly perturbed payoff models of van Damme (1991). Velez and Brown (2020b) note that all of these models satisfy *weak payoff-monotonicity*, i.e., between two actions available to an agent, say a and b, given what other agents are doing, the agent places greater probability on an action a only if its expected payoff is greater than that of action b.[1] Thus instead of defining a refinement based on a particular parametric theory, they define a refinement based on this property. An EE is a Nash equilibrium that is in the closure of weakly payoff-monotone behavior. It logically follows that should data exist that show evidence of an equilibrium that is not empirical, such data would also falsify all models that satisfy weak payoff-monotonicity.

Mechanism design built upon EE allows the researcher to inform the design with the accumulated evidence from experiments. MNBRE models capture regularities in games and provide a rationale why weakly payoff monotone behavior can persist, a point observed since the introduction of MSQRE (see McKelvey and Palfrey, 1995). Thus, a design based on EE is balanced between our aforementioned optimistic and pessimistic views. The approach optimistically recognizes that not all Nash equilibria are plausible. For an equilibrium to concern the mechanism designer, it must be in the proximity of behavior that can be fit by some weakly payoff-monotone distribution. To be pessimistically cautious, the approach accounts for all equilibria that are in the proximity of weakly payoff-monotone behavior.[2] Thus, the designer robustly accounts for all behavior rationalized by MNBRE models.

This paper provides a foundation for Velez and Brown (2020b)'s choice of weak payoff-monotonicity as the basis of plausibility in games and the subsequent definition of EE.

First, we show that EE is on a sweet spot of tractability. On the one hand, since it is defined based on a general class of theories, it is easier to prove that a given equilibrium is empirical by approximating it by weakly payoff-monotone behavior instead of constructing approximation with a particular parametric theory. On the other hand, our first result,

[1] This property is universally assumed in the related random choice models where payoffs are not observable (c.f., Fudenberg et al., 2015).

[2] A predecessor of this approach is Tumennasan (2013), who proposes to complete information implementation in the limits of Logistic QRE. This author makes an additional requirement of convergence that restricts the admissible mechanisms. See (Velez and Brown, 2020b) for details.

Theorem 1, reveals that EE can be defined by requiring proximity to behavior satisfying two more stringent conditions, *interiority* and *payoff-monotonicity* (i.e., behavior is ordinally equivalent to expected payoffs). Since analyzing behavior that satisfies these properties rules out corner cases, this result significantly reduces the computation of EE in applications (Sec. 4.1).

Defining EE by means of weak payoff-monotonicity provides a cautious and tractable base of design.[3] Unfortunately, as a non-parametric theory, this property by itself does not allow us to articulate the idea of increasing sophistication. Consider a two-person game with a Nash equilibrium in which one agent, Ann, uniformly randomizes on her action set. Suppose that one can approximate this equilibrium with a sequence of Logistic QRE behavior in which the parameter of sophistication for Ann converges to zero (uniform random play operator), and the parameter of sophistication of the other agent converges to infinity (best response operator).[4] Since each Logistic QRE is weakly payoff-monotone, this equilibrium is empirical. However, it is not clear that one can interpret this equilibrium, and EE in general, as selecting the limits of increasingly sophisticated, albeit noisy, behavior.

It turns out that this limitation of EE is only superficial. We say that a sequence of RQRE models is utility maximizing in the limit if any corresponding convergent sequence of RQRE equilibria needs to converge necessarily to a Nash equilibrium. It turns out that each EE is the limit of behavior in a sequence of RQRE models that are utility maximizers in the limit (Theorem 2). This result provides a clear connection of the EE refinement with the practice in experimental economics: A Nash equilibrium is empirical if and only if in any neighborhood of it there is a data set consistent with the RQRE model. Moreover, as the neighborhood is selected smaller, one can always select a near-best-response RQRE that generates the data.

Given the popularity of the Logistic QRE and, in general, of MSQRE, we determine if EE can be defined by means of these theories (the Logistic QRE is a particular form of MSQRE). We show a general result stating that, when at least an agent has at least three actions available, each family of monotone additive randomly perturbed payoff model, as defined by Govindan et al. (2003), generates a strict selection form EE (Theorem 3). These models include both MSQRE and the exchangeable perturbation additive randomly perturbed payoff models of van Damme (1991). This result explains to some extent the difficulty that one has in characterizing equilibria that are approachable by these models.

[3]For instance, it is possible to characterize all social choice functions that fully and robustly implement themselves in EE in any finite private values model (Velez and Brown, 2020b); and the characterization of empirical equilibria in winner-bid and loser-bid auctions in a partnership dissolution with complete information (Velez and Brown, 2020a).

[4]Such a sequence can be easily constructed in a symmetric matching pennies game.

The selection depends on the additive form in which random shocks affect payoffs.

Besides clarifying the relationship of EE and refinements based on monotone additive randomly perturbed payoff models, Theorem 3 significantly advances our understanding of these models in both the strategic and random choice domains. Indeed, this theorem resolves the long standing open question posed by Goeree et al. (2005) of whether the empirical content of RQRE and MSQRE coincide for a fixed payoff matrix. When one agent has at least three actions available, the answer is decidedly negative. Combined with our results on approximation by deterministic perturbed payoff models (Sec. 4.4), this theorem establishes a fundamental difference of these models in both strategic and random choice domains (see Sec. 5).

Our final results show that EE can be equivalently defined by approximation by means of behavior in equilibrium games with deterministically perturbed payoff models, also known as control cost models (van Damme, 1991). With the usual notation an agent has payoffs

$$\mathrm{E}_\sigma u_i - \sum_{a_i \in A_i} c_i(\sigma_i(a_i)),$$

where c_i is a smooth and convex decreasing function that can be interpreted as a deterministic perturbation. Theorem 4 states that one can approximate each EE by behavior in deterministically perturbed models with vanishing perturbations. Moreover, our proof of this result shows that perturbation functions can always be constructed as splines whose number of segments is essentially bounded by the number of actions available to the agent. Deterministically perturbed models have been central in the study of probabilistic choice and have been used as a basis of perturbed behavior in the theory of learning in games (see Fudenberg et al., 2015, and references therein). Thus, Theorem 4 provides not only a low dimensional base for EE, but also a clear connection with one of the leading random choice and perturbed games theories. Since deterministically perturbed models are RQRE models, Theorem 2 is a corollary of Theorem 4.

The remainder of the paper is organized as follows. Sec. 2 introduces the model and definitions. Sec. 3 presents a series of examples that allow the reader to familiarize with EE and show that this refinement is independent from the most prominent tremble-based refinements previously defined in the literature. Sec. 4 presents our results. Sec. 5 discusses and concludes.

2 Model

We study the plausibility of Nash equilibria in a finite normal-form game $\Gamma(u) := (N, A, u)$ where $N := \{1, ..., n\}$ is a set of agents; $(A_i)_{i\in N}$ are the corresponding action spaces and $A := A_1 \times \cdots \times A_n$ the set of action profiles; and $u := (u_i)_{i\in N}$ is the profile of expected utility indices, i.e., functions $u_i : A \to \mathbb{R}$. To avoid trivialities we assume that each agent has at least two actions available, i.e., for each $i \in N$, $|A_i| \geq 2$. Let $\mathcal{U}$ be the set of all utility profiles. Our interpretation of the game is standard. Agents simultaneously choose an action. Given that action profile $a := (a_i)_{i\in N}$ is chosen, agent i's payoff is $u_i(a)$. Our analysis will not involve comparisons of behavior across games with different agent sets or action spaces. Thus, N and A are fixed throughout.

A strategy for agent i is a probability distribution on A_i, denoted generically by $\sigma_i \in \Delta(A_i)$. A pure strategy places probability one on a given action. We identify pure strategies with the actions themselves. A strategy is interior if it places positive probability on each possible action. A profile of strategies is denoted by $\sigma := (\sigma_i)_{i\in N} \in \Sigma(A) := \Delta(A_1) \times \cdots \times \Delta(A_n)$. Given $S \subseteq N$, we denote a subprofile of strategies for these agents by σ_S. When $S = N\setminus\{i\}$, we simply write $\sigma_{-i} \in \Sigma(A)_{-i} := \times_{j\in N\setminus\{i\}}\Delta(A_j)$. Consistently, we concatenate partial strategy profiles as in (σ_{-i}, μ_i). We consistently use this convention when operating with vectors, as with action profiles.

Agent i's expected utility given strategy profile σ is

$$\mathrm{E}_\sigma u_i = \sum_{a\in A} u_i(a)\sigma(a),$$

where $\sigma(a) = \sigma_1(a_1)\ldots\sigma_n(a_n)$. Following our convention of identifying pure strategies with actions, we write $\mathrm{E}_{(\sigma_{-i},a_i)}u_i$ for the utility that agent i gets from playing actions a_i when the other agents play σ_{-i}. We say that an action $a_i \in A_i$ is weakly dominated by action $\hat{a}_i \in A_i$ if for each $a_{-i} \in A_{-i}$, $u_i(a_{-i}, \hat{a}_i) \geq u_i(a)$ with strict inequality for at least an element of A_{-i}. We say that $a_i \in A_i$ is a weakly dominated action if there is another action that weakly dominates it.

The following are the basic prediction for game $\Gamma(u)$ and three of its most prominent refinements.

1. (Nash, 1951) A *Nash equilibrium of* $\Gamma(u)$ is a profile of strategies σ, such that for each $i \in N$ and each $\mu_i \in \Delta(A_i)$, $\mathrm{E}_\sigma u_i \geq \mathrm{E}_{(\sigma_{-i},\mu_i)}u_i$. We denote this set by $\mathrm{N}(u)$.

2. An *undominated Nash equilibrium of* $\Gamma(u)$ is a Nash equilibrium of Γ in which no agent plays with positive probability a weakly dominated action. We denote this set

by $\mathrm{U}(u)$.

3. (Selten, 1975) A *perfect equilibrium* of Γ is a profile of strategies σ that is the limit of a sequence of interior strategy profiles $\{\sigma^\lambda\}_{\lambda\in\mathbb{N}}$ such that for each $\lambda \in \mathbb{N}$ and each $i \in N$, σ_i^λ places probability greater than $1/\lambda$ on a given action only if it is a best response to σ_{-i}^λ. We denote this set by $\mathrm{T}(u)$.[5]

4. (Myerson, 1978) A *proper equilibrium* of $\Gamma(u)$ is a profile of strategies σ that is the limit of a sequence of interior strategy profiles $\{\sigma^\lambda\}_{\lambda\in\mathbb{N}}$ such that for each $\lambda \in \mathbb{N}$, each $i \in N$, and each pair of actions $\{a_i, \hat{a}_i\} \subseteq A_i$, if $\mathrm{E}_{(\sigma_{-i}^\lambda, a_i)} u_i > \mathrm{E}_{(\sigma_{-i}^\lambda, \hat{a}_i)} u_i$, then $\sigma_i^\lambda(\hat{a}_i) \leq (1/\lambda)\sigma_i^\lambda(a_i)$. We denote this set by $\mathrm{P}(u)$.

Our main objective is to study Empirical Equilibrium, a refinement of Nash equilibrium that is based on an empirical characterization of behavior. That is, we envision that the researcher samples empirical distributions of behavior in a finite set of normal-form games. Based on the analysis of the data, the researcher constructs a refutable theory that explains this behavior. Then, uses this theory to determine the plausibility of Nash equilibria in all normal-form games. If a Nash equilibrium is not in the closure of the empirical content of the researchers theory, the researcher would be able to reject the specification of the theory were they to observe this equilibrium. Thus, under the hypothesis that the researcher's theory is well specified, each Nash equilibrium that does not belong to the closure of the empirical content of the researcher's theory is implausible. Empirical equilibrium is the refinement so defined when the researcher endorses the non-parametric theory that each agent chooses actions with higher probability only when they are better for her given what the other agents are doing.

Definition 1 (Velez and Brown, 2020b)**.** $\sigma \in \Sigma(A)$ is *weakly payoff-monotone for* u if for each $i \in N$ and each pair of actions $\{a_i, \hat{a}_i\} \subseteq A_i$ such that $\sigma_i(a_i) > \sigma_i(\hat{a}_i)$, we have that $\mathrm{E}_{(\sigma_{-i}, a_i)} u_i > \mathrm{E}_{(\sigma_{-i}, \hat{a}_i)} u_i$.

Intuitively, a profile of strategies is weakly payoff-monotone for a game if differences in behavior reveal differences in expected payoffs.

Definition 2 (Velez and Brown, 2020b)**.** An *empirical equilibrium of* $\Gamma(u)$ is a Nash equilibrium of $\Gamma(u)$ that is the limit of a sequence of weakly payoff-monotone strategies for u. We denote this set by $\mathrm{EE}(u)$.

[5] Our definition of perfect equilibrium corresponds to Myerson (1978)'s characterization of Selten (1975)'s perfect equilibrium.

(a) u^1

		Player 2 b_1	b_2
Player 1	a_1	$1,1$	$0,0$
	a_2	$0,0$	$0,0$

(b) v^1

		Player 2 b_1	b_2
Player 1	a_1	$2,2$	$2,1$
	a_2	$2,3$	$0,0$

Table 1: In two games shown $N = \{1, 2\}$, $A_1 = \{a_1, a_2\}$, $A_2 = \{b_1, b_2\}$, and payoffs are shown in the corresponding table; (a) a game in which the set of empirical equilibria is a proper subset of the set of Nash equilibria; (b) a game in which there are empirical equilibria in which player 1 chooses a weakly dominated strategy with positive probability.

It is well known that McKelvey and Palfrey (1995)'s Logistic Quantal Response Equilibria are weakly payoff-monotone distributions. For a sequence of sophistication parameters converging to utility maximizing behavior, a corresponding sequence of Logistic QRE contains a subsequence converging to a Nash equilibrium (McKelvey and Palfrey, 1995). This equilibrium is then an EE. Thus empirical equilibria exist for each finite game.

The following property of a strategy profile, which implies weak payoff monotonicity, will allow us to provide an alternative useful characterization of EE.

Definition 3. $\sigma \in \Sigma(A)$ is *payoff-monotone for u* if for each $i \in N$ and each pair of actions $\{a_i, \hat{a}_i\} \subseteq A_i$, $\sigma_i(a_i) \geq \sigma_i(\hat{a}_i)$ if and only if $\mathrm{E}_{(\sigma_{-i},a_i)} u_i \geq \mathrm{E}_{(\sigma_{-i},\hat{a}_i)} u_i$.

3 Empirical equilibrium and tremble based refinements

In this section we study the relationship between empirical EE and undominated, perfect, and proper equilibria. We do so by analyzing a series of examples showing that these equilibrium concepts are independent. The main purpose of this discussion is to provide the reader with clear intuition about EE by contrasting it with these more familiar refinements.[6]

Example 1. Consider game $\Gamma(u^1)$ in Table 1 (a). This game was proposed by Myerson (1978) to illustrate that some Nash equilibria are intuitively implausible. There are two Nash equilibria in $\Gamma(u^1)$, (a_1, b_1) and (a_2, b_2). Only (a_1, b_1) is an EE in this game. Indeed, for each distribution of actions of player 2, player 1's utility from playing a_1 is greater than or equal to the utility from playing a_2; thus, in a profile of weakly payoff-monotone distributions of play, agent 1 will always play a_1 with probability at least 1/2 (Fig. 1

[6] In Sec. 4 we discuss two refinements introduced by van Damme (1991) and McKelvey and Palfrey (1995) that are subrefinements of EE. These studies provide examples showing that these refinements may not be contained in the set of undominated equilibria. Thus, one can conclude from the examples in van Damme (1991) and McKelvey and Palfrey (1995) that empirical equilibria may involve weakly dominated actions are played with positive probability.

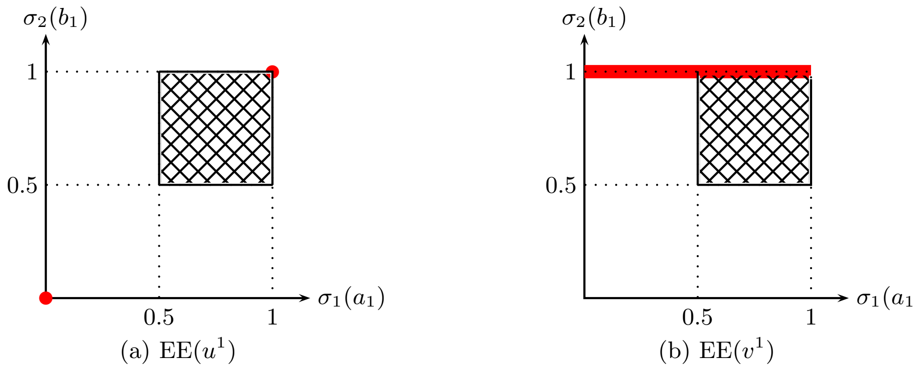

Figure 1: (a) Weakly payoff-monotone distributions (shaded area) and Nash equilibria of $\Gamma(u^1)$; $\sigma_1(a_1)$ is the probability with which agent 1 plays a_1. Equilibrium (a_1, b_1) can be approximated by weakly payoff-monotone behavior. Thus, it is an EE of $\Gamma(u^1)$. Equilibrium (a_2, b_2) cannot be approximated by weakly payoff-monotone behavior. Thus, it is not an empirical equilibrium of $\Gamma(u^1)$. (b) Weakly payoff-monotone distributions and Nash equilibria of $\Gamma(v^1)$; each Nash equilibrium in which agent 1 plays a_1 with probability at least 1/2 is an empirical equilibrium.

(a)); thus, (a_2, b_2) cannot be approximated by weakly payoff-monotone behavior. If this game is played and agents behavior is weakly payoff-monotone and approximates a Nash equilibrium, it is necessarily (a_1, b_1). $\square$

Each refinement that rules out weakly dominated behavior coincides with EE in game $\Gamma(u^1)$. Undominated equilibria and empirical equilibria are independent, however.

Example 2. Consider game $\Gamma(v^1)$ in Table 1 (b). Player 2 has a strictly dominant strategy in this game. Thus, in each Nash equilibrium σ of $\Gamma(v^1)$, $\sigma_2(b_1) = 1$. Agent 1 is indifferent between both actions if agent 2 plays b_1. Thus, the set of Nash equilibria of this game is the distributions in which agent 1 randomizes between both actions and agent 2 plays b_1. Now, let σ be a weakly payoff-monotone distribution for $\Gamma(v^1)$. Since b_1 strictly dominates b_2, $\sigma_2(b_1) \geq \sigma_2(b_2)$. If $\sigma_2(b_2) > 0$, $E_{(\sigma_2,a_1)}v_1^1 > E_{(\sigma_2,a_2)}v_1^1$. Thus, it must be the case that $\sigma_1(a_1) \geq \sigma_1(a_2)$. If $\sigma_2(b_2) = 0$, $E_{(\sigma_2,a_1)}v_1^1 = E_{(\sigma_2,a_1)}v_1^1$. Thus, $\sigma_1(a_1) = \sigma_1(a_2)$. Thus, the set of weakly payoff distributions for $\Gamma(v^1)$ are those at which $\sigma_1(a_1) \geq 1/2$ and $\sigma_2(b_1) \geq 1/2$, except those at which $\sigma_2(b_1) = 1$ and $\sigma_1(a_1) < 1/2$ (Fig. 1 (b)). The set of empirical equilibria of $\Gamma(v^1)$ are the Nash equilibria in which agent 1 plays a_1 with probability at least 1/2. Since a_2 is weakly dominated by a_1 for player 1, almost all of these empirical equilibria involve one player playing a weakly dominated action with positive probability. $\square$

Empirical equilibrium does a subtle selection from the Nash equilibrium set. It determines the plausibility of a strategy based on its relative merits with respect to the alternative actions that the agent may choose. The following example drives this point home. It il-

		Player 2		
		b_1	b_2	b_3
	a_1	$1, 1$	$0, 0$	$-7 - c_1, -7 - c_2$
Player 1	a_2	$0, 0$	$0, 0$	$-7, -7$
	a_3	$-7 - c_1, -7 - c_2$	$-7, -7$	$-7, -7$

Table 2: $N = \{1, 2\}$, $A_1 = \{a_1, a_2, a_3\}$, $A_2 = \{b_1, b_2, b_3\}$, and payoffs u^c given in the table, where $c := (c_1, c_2)$, $c_1 > 0$, and $c_2 > 0$.

lustrates it for a parametric family of games. This family is a generalization of a game proposed by Myerson (1978) to show that it is possible to introduce weakly dominated actions in $\Gamma(u^1)$, and considerably change its set of trembling hand perfect equilibria.

Example 3. Consider game $\Gamma(u^c)$ for some $c := (c_1, c_2)$, $c_1 > 0$, and $c_2 > 0$ (Table 2). Standard arguments show that for each $c > 0$,

$$\begin{aligned}
&\mathrm{N}(u^c) = \{(a_1, b_1), (a_2, b_2), (a_3, b_3)\},\\
&\mathrm{T}(u^c) = \mathrm{U}(u^c) = \{(a_1, b_1), (a_2, b_2)\},\\
&\mathrm{P}(u^c) = \{(a_1, b_1)\}.
\end{aligned}$$

In contrast to these refinements, the empirical equilibrium set of $\Gamma(u^c)$ depends on c. First, note that for no $c > 0$, (a_3, b_3) is an empirical equilibrium of $\Gamma(u^c)$. This is so because a_2 weakly dominates a_3 for player 1. Thus, in any weakly payoff-monotone distribution for $\Gamma(u^c)$, player 1 plays a_2 with a probability that is at least the probability with which she plays action a_3. Thus, no sequence of weakly payoff-monotone distributions for $\Gamma(u^c)$ converges to (a_3, b_3). On the other hand for each $c > 0$, $(a_1, b_1) \in \mathrm{EE}(u^c)$. Indeed, (a_1, b_1) is a strict Nash equilibrium, i.e., a profile of actions that constitute unique mutual best responses. These type of equilibria are weakly payoff-monotone distributions themselves, so a constant sequence sustains them as EE.

Let us now examine the plausibility of (a_2, b_2) in $\Gamma(u^c)$. Think of the payoffs in the game as dollar amounts. Consider first a small c, say $c_1 \approx c_2 \approx 0.01$. Let σ be an empirical distribution of play that approximates (a_2, b_2). In such a situation, $E_{(\sigma_2, a_1)} u_1^c \approx 0 > E_{(\sigma_2, a_3)} u_1^c \approx -7$ and $E_{(\sigma_2, b_1)} u_2^c \approx 0 > E_{(\sigma_2, b_3)} u_2^c \approx -7$. Thus, if expected utility guides the choices of the players, one can expect that player 1 will play a_1 at least as often as a_3, and player 2 will play b_1 at least as often as b_3. If this is so, action a_1 will have a greater utility than action a_2 for player 1, and action b_1 will have a greater utility than action b_2 for player 2. Thus, if expected utility guides the choices of the agents, σ will not be close to (a_2, b_2). Thus, a plausible empirical distribution, i.e., one that is informed by expected utility for this game, will never be close to (a_2, b_2).

Now, consider a large c, say $c_1 \approx c_2 \approx 100{,}000$. Again, if σ is an empirical distribution of play that approximates (a_2, b_2) and is guided by expected utility, player 1 will be playing a_1 at least as often as a_3, and player 2 will be playing b_1 at least as often as b_3. In contrast with our previous case, it does not follow that necessarily action a_1 will have a greater utility than action a_2 for player 1, and action b_1 will have a greater utility than action b_2 for player 2. This will only happen if player 1 is playing a_1 at least one hundred thousand times as often as a_3, and player 2 is playing b_1 at least one hundred thousand times as often as b_3. Thus, it is possible that σ is informed by expected utility, i.e., $\sigma_1(a_1) > \sigma_1(a_3)$ and $\sigma_2(b_1) > \sigma_2(b_3)$, and at the same time $E_{(\sigma_2,a_2)}u_1^c > E_{(\sigma_2,a_1)}u_1^c$, $E_{(\sigma_1,b_2)}u_2^c > E_{(\sigma_1,b_1)}u_2^c$, $\sigma_1(a_2) \approx 1$, and $\sigma_2(b_2) \approx 1$. Essentially, since the possible loss for player 1 from playing a_1 is about 100,000.00, player 1 can be scared away from playing a_1 if player 2 is playing b_3 more than once each 100,000 times she plays b_1. This is still compatible with b_3 being the worst alternative given what the other agent is doing.

These arguments can be easily formalized to show that

$$\mathrm{EE}(u^c) = \begin{cases} \{(a_1, b_1)\} & \text{if } \min\{c_1, c_2\} \leq 1, \\ \{(a_1, b_1), (a_2, b_2)\} & \text{Otherwise.} \end{cases}$$

One cannot expect that if one brings these games to a laboratory setting or has the opportunity to collect field data on them, the threshold $\min\{c_1, c_2\} = 1$ will be a good predictor of a structural change in the behavior of the agents. However, it is reasonable that behavior in this game will depend on the size of c, as empirical equilibrium predicts, i.e., equilibrium (a_2, b_2) will be relevant only for high values of c. Undominated equilibria, perfect equilibria, and proper equilibria all miss this point. Undominated equilibrium and perfect equilibrium miss that when c is too low, actions a_2 and b_2 are de facto "weakly dominated" when they are played with almost certainty. That is, if they were going to be played with probability close to one, actions a_1 and b_1, would be preferred for the respective players. Thus, we can rule this equilibrium out by means of the following observation. It is not reasonable that we will observe a distribution of play in which an agent is not playing her unique maximizer of utility with high probability, say more than random play.

Finally, proper equilibrium dismisses (a_2, b_2) independently of c. Think of our example with high c.[7] For (a_2, b_2) to be a proper equilibrium of $\Gamma(u^c)$, for large λ there must be a distribution of play σ^λ satisfying two conditions: (i) σ^λ is close to (a_2, b_2), and thus $E_{(\sigma_2^\lambda,a_1)}u_1^c \approx 0 > E_{(\sigma_2^\lambda,a_3)}u_1^c \approx -7$ and $E_{(\sigma_2^\lambda,b_1)}u_2^c \approx 0 > E_{(\sigma_2^\lambda,b_3)}u_2^c \approx -7$; and (ii)

[7] Not every Proper equilibrium is empirical. Consider for instance the null game $u = 0$. Each profile of strategies is a Proper equilibrium of this game. By contrast, only uniform random play is an Empirical Equilibrium. We thank Yuval Heller for suggesting this simple example.

$\sigma_1^\lambda(a_1) > \lambda\sigma_1^\lambda(a_3)$ and $\sigma_2^\lambda(b_1) > \lambda\sigma_2^\lambda(b_3)$. For distributions where $\lambda \geq 100{,}000$, a_2 and b_2 are not maximizing choices for players 1 and 2, respectively, meaning (a_2, b_2) cannot be a proper equilibrium. Thus, the reason why proper equilibrium dismisses (a_2, b_2) for high c is that it uses the same parameter for proximity to (a_2, b_2) and for the agents' reactivity to differences in expected utility. This allows us to draw a stark difference of this refinement and empirical equilibrium. Proper equilibrium is a decision-theoretical, thought experiment in which a utility maximizing agent considers the possibility that another utility maximizing agent makes a mistake. Confronted with this thought, a utility maximizing agent will determine a Nash equilibrium as implausible because it is impossible that agents who are infinitely reactive to expected utility make self-sustaining mistakes arbitrarily close to the equilibrium. By contrast, empirical equilibrium is an exercise performed by an observer based on weak payoff monotonicity, a testable property of behavior. The observer knows that if this property is satisfied by empirical frequencies, only empirical equilibria can be approximated by data. $\square$

4 Results

4.1 Approachability by payoff-monotone behavior

Empirical equilibrium can be equivalently defined by proximity of interior payoff-monotone behavior.

Theorem 1. $\sigma \in \mathrm{EE}(u)$ if and only if $\sigma \in \mathrm{N}(u)$ and there is a convergent sequence of interior payoff-monotone distributions for u whose limit is σ.

The characterization of the set of empirical equilibria by means of approximation of interior payoff-monotone distributions simplifies its computation. In applications, the computation of this set usually requires two steps. First, one needs to identify conditions satisfied by each possible empirical equilibrium. Then, one needs to show that each equilibrium satisfying these properties is an empirical equilibrium. It is in the first step of this process that Theorem 1 proves to be essential. One can identify all the implications of proximity to weakly payoff-monotone behavior by assuming only proximity to interior payoff-monotone behavior. The gain can be considerable in applications, for this avoids the analysis of corner cases (e.g. Velez and Brown, 2020a).

It is worth noting that computing the set of empirical equilibria of a game usually also involves a non-trivial second step in which one proves that the candidates one identified as possible empirical equilibria are actually so. At this point it is more convenient to construct a sequence of weakly payoff-monotone behavior that converges to the candidate, that is actually neither interior, nor payoff-monotone (e.g. Velez and Brown, 2020a). In this sense, if one were to equivalently define empirical equilibrium based on the empirical content of payoff monotonicity, Theorem 1 would still be a valuable tool in its characterization in applications.

Theorem 1 indicates a form of stability of empirical equilibria. Think for instance of an equilibrium in a game that is itself a non-interior weakly payoff-monotone distribution, e.g., a Nash equilibrium in which each agent plays her unique best response.[8] One can conclude that the equilibrium is an empirical equilibrium by taking the respective constant sequence. Theorem 1 implies that this is not the only sequence that will sustain the argument. One will always be able to find a sequence of interior payoff-monotone distributions that converges to the equilibrium. Equivalently, empirical equilibria are never isolated points in the closure of payoff-monotone behavior.[9]

4.2 Approachability by regular QRE

In this section we study the foundation of empirical equilibrium by means of *regular Quantal Response Equilibrium* models (McKelvey and Palfrey, 1996; Goeree et al., 2005). This theory assumes that agents are noisy best responders whose frequencies of play depend solely on the vector of expected utility. Formally, the model is parameterized by a *quantal response function* (QRF), i.e., for agent i a function $p_i : \mathbb{R}^{A_i} \to \Delta(A_i)$. For each $a_i \in A_i$ and each $x \in \mathbb{R}^{A_i}$, $p_{ia_i}(x)$ denotes the value assigned to a_i by $p_i(x)$. A QRF p_i is *regular* if it satisfies the following four properties (Goeree et al., 2005):

(R1) *Interiority*: $p_i > 0$.

(R2) *Continuity*: p_i is a continuous function.

(R3) *Responsiveness*: for $x \in \mathbb{R}^{A_i}$, $\eta > 0$, and $a_i \in A_i$, $p_{ia_i}(x + \eta 1_{a_i}) > p_{ia_i}(x)$.[10]

(R4) *Monotonicity*: for $x \in \mathbb{R}^{A_i}$ and $\{a_i, \hat{a}_i\} \subseteq A_i$ such that $x_{a_i} > x_{\hat{a}_i}$, $p_{ia_i}(x) > p_{i\hat{a}_i}(x)$.

A *quantal response equilibrium* (QRE) of $\Gamma(u)$ with respect to a profile of QRFs, $p := (p_i)_{i \in N}$, is a fixed point of the composition of p and the expected payoff operator in $\Gamma(u)$ (Goeree et al., 2005), i.e., a profile of distributions $(\sigma_i)_{i \in N}$ such that for each $i \in N$, $\sigma_i = p_i((\mathrm{E}_{(\sigma_{-i}, a_i)} u_i)_{a_i \in A_i})$. When p is regular, we refer to it as an RQRF and to the corresponding QRE also as regular or RQRE.

[8] These equilibria are usually referred to as *strict* (c.f., Harsanyi, 1973).

[9] Compare for instance with the Nash equilibria that re also behavior strategies in an M-equilibrium as defined by Goeree and Louis (2021). These equilibria can be isolated M-equilibrium points that are outside of the closure of weakly payoff-monotone behavior. For instance, equilibrium (a_2, b_2) in game u_1 in Example 1, is an M-equilibrium strategy. See Sec. 5 for a discussion.

[10] 1_{a_i} denotes the vector in $\mathbb{R}^{A_i}$ that has 1 in component a_i and 0 otherwise.

Because RQRFs are interior, monotone, and continuous, each RQRE of $\Gamma(u)$ is an interior payoff-monotone distribution for $\Gamma(u)$. The converse statement is also true. This implies the empirical content of RQRE is completely characterized by interiority and payoff-monotonicity.

Lemma 1. Let $\sigma \in \Sigma(A)$. Then, there is a QRF p such that σ is a RQRE with respect to p for u if and only if σ is interior and payoff-monotone for u.

Lemma 1 is a consequence of Remark 1 in Sec. 4.4.[11] By Theorem 1 we know that the closure of weakly payoff-monotone and the closure of interior payoff-monotone behavior coincide. Thus, the closure of the empirical content of RQRE is the closure of weakly payoff-monotone behavior.

RQRE theory opens the possibility that we provide a foundation of empirical equilibrium in terms of approximation by behavior associated with agents who are best responders in the limit. To make this precise we first need to identify the conditions under which this is so for a sequence of QRFs.

Definition 4. A sequence of QRFs, $\{p^\lambda\}_{\lambda\in\mathbb{N}}$ is *utility maximizing in the limit* if for each $u \in \mathcal{U}$ and each convergent sequence of QREs of $\Gamma(u)$ corresponding to a subsequence of $\{p^\lambda\}_{\lambda\in\mathbb{N}}$, its limit is a Nash equilibrium of $\Gamma(u)$.

We can then define a refinement of Nash equilibrium in the same spirit as EE, but taking as basis for plausibility of behavior RQRE for increasingly sophisticated RQRFs.

Definition 5. $\sigma \in \Sigma(A)$ is *approachable by RQRE that are utility maximizing in the limit in* $\Gamma(u)$ if there is a sequence of RQRF profiles, $\{p^\lambda\}_{\lambda\in\mathbb{N}}$, which is utility maximizing in the limit, and a corresponding convergent sequence of QREs for $\Gamma(u)$, whose limit is σ. We denote this set by $\mathrm{R}(u)$.

Clearly, for each $u \in \mathcal{U}$, $\mathrm{R}(u) \subseteq \mathrm{N}(u)$. The interpretation of this refinement is similar to that of EE. It differs in that it explicitly models approximation of Nash equilibria by behavior of agents who are infinitely sophisticated in the limit. One can argue that empirical equilibrium is a more cautious refinement, for it is based only on observables. Indeed, QRFs are not observable, and some of their properties, as continuity, are not refutable with finite data. It turns out that these notions of plausibility of behavior in games coincide, however.

Theorem 2. For each $u \in \mathcal{U}$, $\mathrm{EE}(u) = \mathrm{R}(u)$.

[11] Lemma 1 is also a consequence of a step in the proof of the main result in Goeree et al. (2018a).

Theorem 2 allows us to alternatively describe EE in a way that speaks closely to the practice in experimental economics. If one expects that as players gain experience their behavior will be fit by increasingly sophisticated RQRE, the only Nash equilibria that can be approximated by data are the EE.

4.3 Approachability by additive randomly perturbed payoffs models

Two predecessors of EE are the Nash equilibria that are limits of behavior in Harsanyi (1973)'s randomly perturbed payoff models for exchangeable perturbations (van Damme, 1991) and Logistic QRE approachable equilibria (McKelvey and Palfrey, 1996). These are subrefinements of EE. In what follows we show that these belong to a general family of refinements of Nash equilibrium that are strict subrefinements of EE. As a by product, which we discuss in Sect. 5, we advance our understanding of the empirical content of monotone randomly perturbed payoff models, a topic that has received attention due to the popularity of these models for the analysis of data from economics experiments (Goeree et al., 2005; Haile et al., 2008; Golman, 2011), and the prominence of random utility models in random choice environments.

We follow Govindan et al. (2003)'s construction of additive randomly perturbed payoff models. Given $\Gamma(u) = (N, A, u)$ and a vector of independent Borel probability measures on $\mathbb{R}^A$, $(\mu_i)_{i\in N}$, let $(\Gamma(u), \mu)$ be the incomplete information game in which $\mu := \mu_1 \times \cdots \times \mu_n$ is a common prior on payoff types. Given type $\eta_i \in \mathbb{R}^A$ for agent i, her expected utility index is $a \in A \mapsto u_i(a) + \eta_i(a)$. Whenever convenient, given an agent, say i, whose action set is $A_i := \{a_1, ..., a_K\}$, we write a vector $x \in \mathbb{R}^A$ as $(x_{a_l})_{a_l \in A_i}$ where $x_{a_l} := (x_{(a_{-i}, a_l)})_{a_{-i}} \in A_{-i}$. The interpretation of these perturbations is that either the observer who models the strategic situation by means of game $\Gamma(u)$ does not observe the real payoffs perceived by the agents (Harsanyi, 1973), or that the agent fails to perfectly recognize the difference of payoffs between the actions and correctly maximize (McKelvey and Palfrey, 1995). We require throughout that:

Definition 6 (Govindan et al., 2003)**.** For each $i \in N$, μ_i is *purifying*, i.e., for each pair of different actions $\{a_k, a_l\} \subseteq A_i$, and each $\sigma_{-i} \in \Sigma(A)_{-i}$, μ_i assigns probability zero to the event $\eta_i(a_k, \cdot) - \eta_i(a_l, \cdot) \in \mathbb{R}^{A_{-i}}$ lies on any single prespecified hyperplane in $\mathbb{R}^{A_{-i}}$ with normal σ_{-i}. Let $\mathcal{B}_i$ be the space purifying Borel probability measures on $\mathbb{R}^A$ and $\mathcal{B} = \mathcal{B}_1 \times \cdots \times \mathcal{B}_n$.

The main two models that fit into our framework are Harsanyi (1973)'s, in which each μ_i is absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^A$; and McKelvey and Palfrey (1995)'s Structural QRE, in which perturbations are perfectly correlated across action profiles with the same action for an agent, so they can be identified with measures on $\mathbb{R}^{A_i}$, that are further assumed to be absolutely continuous with respect to the Lebesgue measure on $\mathbb{R}^{A_i}$.

Consider $i \in N$. The assumption of purifying perturbations implies that for each $\sigma_{-i} \in \Sigma(A)_{-i}$ and for μ_i almost every realization of the perturbation η_i, agent i with type η_i has a unique best response to σ_{-i}. Thus, given μ_i and σ_{-i}, the probability with which agent i is observed playing a given action is uniquely defined by utility maximization. Let

$$B_i^{u_i,\mu_i}(\sigma_{-i}) := (B_{ia_i}^{u_i,\mu_i}(\sigma_{-i}))_{a_i \in A_i} \in \Delta(A_i),$$

be this distribution. It is easy to see that this function is continuous (Govindan et al., 2003), and that the fixed points of $\sigma \in \Sigma(A) \mapsto (B_i^{u_i,\mu_i}(\sigma_{-i}))_{i \in N} \in \Sigma(A)$, which exist by Brouwer's fixed point theorem, are the set of observable strategies in Bayesian Nash equilibria of (Γ, μ). We denote the set of profiles $\sigma \in \Sigma(A)$ that are induced by some Bayesian Nash equilibrium of $(\Gamma(u), \mu)$ by $\mathrm{BNE}(u, \mu)$).

In McKelvey and Palfrey (1995)'s model the best response operator can be further characterized by the function $x \in \mathbb{R}^{A_i} \mapsto Q_i^{\mu_i}(x)$, where

$$\sigma_{-i} \in \Sigma(A)_{-i} \mapsto B_i^{u_i,\mu_i}(\sigma_{-i}) = Q_i^{\mu_i}(\mathrm{E}_{(\sigma_{-i},\cdot)} u_i) \in \Delta(A_i).$$

The function $Q_i^{\mu_i}$ which only depends on N, A, and μ_i, and does not depend on u, is referred to as a *Structural Quantal Response Function* (SQRF) (McKelvey and Palfrey, 1995; Goeree et al., 2005). These functions satisfy R1-R3 in Sec. 4.2, but may violate R4 (Goeree et al., 2005). The SQRF most commonly used in empirical analysis of experimental data is the Logistic form, l^λ, which is associated with the double-exponential i.i.d perturbation (Goeree et al., 2018a), and assigns to each $a_i \in A_i$ and each $x \in \mathbb{R}^{A_i}$ the value,

$$l_{ia_i}^{\lambda}(x) := \frac{e^{\lambda x_{a_i}}}{\sum_{\hat{a}_i \in A_i} e^{\lambda x_{\hat{a}_i}}}. \tag{1}$$

Additive linear randomly perturbed payoff models also allow us to articulate the idea that agents' behavior approximates that of utility maximizers. More precisely, consider a sequence of purifying perturbations $\{\mu^\lambda\}_{\lambda \in \mathbb{N}}$. We say that this sequence *vanishes* when for each $i \in N$ and each neighborhood of zero, G, as $\lambda \to \infty$, $\mu_i^\lambda(G) \to 1$. If there is a convergent sequence of Bayesian Nash equilibria of the respective games $(\Gamma(u), \mu^\lambda)$ for a sequence of vanishing perturbations, the limit of the corresponding induced observable strategies is necessarily a Nash equilibrium of $\Gamma(u)$ (c.f., van Damme, 1991). Thus, any sequence of SQRFs whose corresponding perturbations are vanishing is utility maximizing in the limit (see Definition 4).

If additive randomly perturbed payoff models are unrestricted, each possible observable distribution of behavior in a normal-form game is in the empirical content of this theory. More precisely, for each $\sigma \in \Sigma(A)$ and each $u \in \mathcal{U}$, there is a purifying perturbation μ such that $\sigma \in \mathrm{BNE}(u,\mu)$ (Haile et al., 2008). If this theory is further disciplined by convergence to Nash behavior, it produces no strict refinement of the Nash equilibrium set.

Proposition 1. For each $\sigma \in \mathrm{N}(\Gamma)$, there is a sequence of full support purifying vanishing perturbations $\{\mu^\lambda\}_{\lambda\in\mathbb{N}}$ and a sequence of corresponding $\sigma^\lambda \in \mathrm{BNE}(\Gamma,\mu^\lambda)$ that converges to σ.

Thus, requiring proximity to the empirical content of unrestricted randomly perturbed payoff models does not refine the set of Nash equilibria.

McKelvey and Palfrey (1996) and Goeree et al. (2005) argue that additive randomly perturbed payoff models lack of refutability, which was most prominently pointed out by Haile et al. (2008), can be resolved by imposing consistency with payoff monotonicity, a phenomenon for which there is empirical support. Moreover, two attempts have been made with the purpose of refining the set of Nash equilibria based on monotone randomly perturbed payoff models. First, van Damme (1991) imposes permutation invariance of perturbations in Harsanyi's additive randomly perturbed models. Second, McKelvey and Palfrey (1996) propose approximation by Logistic QRE, i.e., restrict to a particular parametric family of perturbations. Both constructions are implicitly imposing that best responses are ordinally equivalent to expected utility.

The following theorem allows us to identify a sharp difference between these and other possible approaches based on monotone additive randomly perturbed payoff models and EE. It states that for each strategy space in which at least an agent has at least three actions available, one can always construct a payoff matrix so the resulting normal-form game possesses an empirical equilibrium that cannot be approximated by any additive randomly perturbed payoff model whose associated best response correspondences are weakly payoff-monotone.

Definition 7. Let $\mathcal{M} \subseteq \mathcal{B}$ be the set of perturbations μ for which for each $u \in \mathcal{U}$, each $i \in N$, and each pair $\{a_i, b_i\}$, if $B^{u_i,\mu_i}_{a_i i}(\sigma_{-i}) > B^{u_i,\mu_i}_{b_i i}(\sigma_{-i})$, then $E_{(\sigma_{-i},a_i)}u_i > E_{(\sigma_{-i},b_i)}u_i$. We refer to $\mu \in \mathcal{M}$ as a *perturbation that induces weakly payoff-monotone best responses*.

Theorem 3. Suppose that at least an agent has at least three actions available. Then, there is $u \in \mathcal{U}$ for which there is $\sigma^* \in \mathrm{EE}(u)$ and $\varepsilon > 0$ such that

$$\{\sigma : ||\sigma - \sigma^*|| < \varepsilon\} \cap \{\sigma : \exists \mu \in \mathcal{M},\ \text{s.t. } \sigma \in \mathrm{BNE}(u,\mu)\} = \emptyset.$$

Theorem 3 reveals that even though additive randomly perturbed payoff models are not refutable if unrestricted, their empirical content, if restricted by weak payoff-monotonicity of best responses as suggested by Goeree et al. (2005), is restricted beyond weak payoff-monotonicity (note that we do not even impose the necessity of perturbations to vanish). The weakly payoff-monotone behavior that is missed by the empirical content of these models can be that in the neighborhood of a Nash equilibrium. In particular, this implies that the refinements proposed by van Damme (1991) and McKelvey and Palfrey (1996) depend on their structural form of approximation and are strict subrefinements of empirical equilibrium for some games.

To gain some intuition on Theorem 3, consider three actions $A_i = \{a_1, a_2, a_3\}$ and a vector of corresponding expected utilities $U_1 < U_2 < U_3$ such that $U_3 - U_2 > U_2 - U_1$. Clearly, there are $\sigma_i \in \Delta(A_i)$ ordinally equivalent to U and in which $\sigma_i(a_2) > 1/3$ (Fig. 2). In such a distribution, the difference in frequency of play between a_2 and a_3 is smaller than the difference in frequency of play between a_1 and a_2. Even though a_2 and a_3 are farther away in utility, payoff monotonicity does not imply they have to be farther away behaviorally. These actions can be more similar in some sense not captured by the difference in utilities.

A similar point has been made by Fosgerau et al. (2020) in the related domain of random choice models. These authors observe that Luce's IIA, which implies the Logit model, is incompatible with asymmetric changes in behavior for symmetric shifts in utility. As a response, they devise a deterministically perturbed payoff model (see Sec. 4.4) that captures them. Theorem 3 goes beyond the observation that the Logit model cannot account for the asymmetries we identify in our example. It reveals that no additive randomly perturbed model can.

Since closed-form solutions are available for the Logistic QRE models, one can use calculus to prove that for none of these models $\sigma_i(a_2) > 1/3$. It is actually easier to prove this at a more general level. The Logistic QRE corresponds to an i.i.d. perturbation, which in turn is permutation invariant. So consider a permutation invariant μ_i on $\mathbb{R}^{A_i}$. Since perturbations enter linearly in payoffs, we can essentially take any $0 < x < y$ and consider the probability with which the actions are played when we draw from a distribution that obtains the permutations of perturbations $0, x, y$ with equal probability. Surprisingly, at most three of these permutations lead a_2 to be a maximizer (Fig.2). Thus, no permutation invariant perturbation, including the distribution that induces the Logit model, induces $\sigma_i(a_2) > 1/3$.

To the length of our knowledge, permutation invariance is the most general known sufficient condition that guarantees a perturbation induces payoff-monotone best responses.

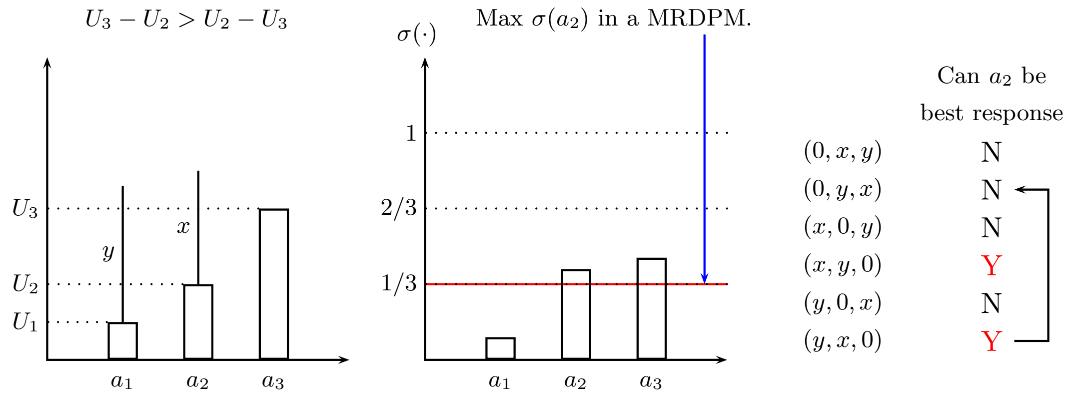

Figure 2: Permutation invariant random utility does not span all payoff-monotone behavior. The profile of (expected) utility is such that $U_1 < U_2 < U_3$ and $U_3 - U_2 > U_2 - U_1$. Let $0 < x < y$. Draws a fair dice with six permutations of $\{0, x, y\}$ as faces. Given the outcome, add the corresponding draw to the baseline utility. The probability that action a_2 is the maximizer at most $1/3$. To see this, observe (right) that a_2 can be the maximizer only when $(0, y, x)$, $(x, y, 0)$ or $(y, x, 0)$ are drawn. However, if a_2 is the maximizer when $(y, x, 0)$, it is not the maximizer when $(0, y, x)$ is drawn.

Theorem 3 does not assume permutation invariance, it only requires that perturbations induce weakly payoff-monotone behavior. So its proof requires a much more intricate argument that exploits invariance properties induced by weak payoff-monotonicity. We present the proof in the Appendix.

4.4 Approachability by behavior in deterministically perturbed payoff models

We learn from Theorem 2 that if behavior is weakly payoff-monotone and approaches a Nash equilibrium, there is an RQRE model that fits this behavior. In this section we show that one can alternatively define empirical equilibrium by approximation by deterministically perturbed payoff models, a subfamily of RQRE models. In these models, agents are parameterized by a function that determines an agent's cost to identify optimal actions in the game. These models, also known as variational preference models, provide a deterministic interpretation of deviations from utility maximization as control costs (van Damme, 1991) and rational inattention (Fosgerau et al., 2020).

A *perturbation function* for player i is $c_i : [0, 1] \to \mathbb{R} \cup \{\infty\}$ with the following properties.

- c_i is strictly decreasing with $c_i(0) = \infty$ and $c_i(1) = 0$.
- c_i is continuously differentiable on $(0, 1]$.[12]

[12] van Damme (1991) assumes that each c_i is twice differentiable. One can easily see that Lemmas 4.2.1-4.2.5 and Theorem 4.2.6. in van Damme (1991) go through with our weaker assumption. One needs continuous differentiability in order to apply Lagrange's theorem in Lemma 4.2.3. All other results follow

- c_i is a strictly convex function.

Given a normal-form game $\Gamma(u) = (N, A, u)$ and a profile of perturbation functions $c := (c_i)_{i \in N}$ we associate a game in which players are N, agent i's action space is $\Delta(A_i)$, and payoff of action profile $\sigma \in \Sigma(A)$ is for each $i \in N$,

$$\mathrm{E}_\sigma u_i - \sum_{a_i \in A_i} c_i(\sigma_i(a_i)). \tag{2}$$

For each $\sigma_{-i} \in \Sigma(A)_{-i}$, there is a unique best response for agent i in this game, which solely depends on the profile of expected utility of the different actions in Γ (Lemma 4.2.1, van Damme, 1991). Let p^{c_i} be this function. One can easily see that this function is an RQRF (van Damme, 1991; Goeree et al., 2018a). Thus, the equilibria of the perturbed game are the RQRE with respect to $(p^{c_i})_{i \in N}$.

The deterministically perturbed payoffs model is less general than the RQRE model for it imposes some restrictions of behavior across different extended games, i.e., if one varies u.[13] For fixed u, they are behaviorally equivalent, however. Let $\sigma \in \Sigma(A)$ be interior and payoff-monotone for a payoff matrix u. It is well known that since the vector of expected utilities $(E_{\sigma_{-i}} u_i(a_i))_{a_i \in A_i}$ is ordinally equivalent to σ_i, there is a perturbation function c_i for which σ_i is the maximizer of (2) (see, for instance, Fudenberg et al., 2015). The following remark follows.

Remark 1. Let $\sigma \in \Sigma(A)$. Then, there are perturbation functions c such that σ is an equilibrium of the corresponding perturbed game for u if and only if σ is interior and payoff-monotone.

Deterministically perturbed payoff models also allow us to easily articulate the idea of convergence to utility maximization. We say that a sequence of profiles of perturbation functions, $\{c^\lambda\}_{\lambda \in \mathbb{N}}$, *vanishes*, if for each $i \in N$ and each $x \in (0, 1]$, $\lim_{\lambda \to \infty} c_i^\lambda(x) = 0$. It is well known that the behavior in games with vanishing perturbations can converge only to Nash equilibria of the underlying game (van Damme, 1991, Theorem 4.3.1).[14] Thus, for a sequence of vanishing perturbation functions, its associated sequence of RQRFs are utility maximizing in the limit.

from convexity and continuity. Our model coincides with Perturbed Utility as defined by Fudenberg et al. (2015). The greater generality of our model, compared to van Damme (1991)'s, allows us to easily construct perturbation functions hitting some specific targets of its derivative without matching the second derivative.

[13] For instance, note that adding a constant to u_i does not change the solution to the maximization of (2). Thus, perturbed payoff models are translation invariant.

[14] Technically, Theorem 4.3.1 in van Damme (1991) applies only to vanishing sequences of perturbation functions of the form εf with $\varepsilon \to 0$. One can easily see that his argument extends for a general sequence of vanishing perturbation functions as we define it because our assumption also implies that for each $x \in (0, 1]$, $(f_i^\lambda)'(x) \to 0$. We have included an explicit proof of this result in an online Appendix.

Theorem 4. Let $\sigma^* \in \mathrm{EE}(u)$. There is a sequence of perturbation functions $\{c^\lambda\}_{\lambda\in\mathbb{N}}$ that vanishes and a corresponding sequence of equilibria $\{\sigma^\lambda\}_{\lambda\in\mathbb{N}}$ that converges to σ^*.

The family of RQRFs is infinitely dimensional. Thus, Theorem 2 does not point exactly to a parametric family of models that is well specified for the analysis of experimental data that one expects to be payoff-monotone. By Theorem 3, the most obvious candidate, the Monotone SQRE model, including all its parametric incarnations, e.g., the Logistic form, is not flexible enough to account for some payoff-monotone behavior in finite games. Theorem 4 implies that the family of perturbed payoff models is enough to generate all payoff-monotone behavior. Our proof of this result is relatively simple and of independent interest, so we present it in the body of the paper. It reveals that one can generate all payoff-monotone behavior with perturbation functions c_i obtained as a logarithmic asymptote stitched to a $(|A_i| + 1)$-segment second order spline.

Proof. To prove the theorem one needs to resolve two issues. First, given an interior σ that is payoff-monotone with respect to a payoff matrix u, one should be able to construct perturbation functions c whose associated game has σ as an equilibrium. Second, one needs to guarantee that if $\{\sigma^\lambda\}_{\lambda\in\mathbb{N}}$ converges to a Nash equilibrium σ^*, the corresponding c^λs can be chosen so they vanish.

Constructing perturbation functions satisfying our admissibility requirements of monotonicity, smoothness, and convexity can be easily achieved by specifying a negative continuous strictly increasing function on $(0, 1]$ that is integrable on any subinterval $[a, 1]$ for $a > 0$ and whose integral on $(0, 1]$ is $-\infty$. Let c_i' be such a function. Then,

$$c_i(x) = -\int_x^1 c_i'(t)dt,$$

is a valid perturbation function and its derivative on $(0, 1]$ is c_i'.

Let σ be a Nash equilibrium of the game associated with $\Gamma(u)$ and c. By interiority of each p_i^c and continuous differentiability and strict convexity of each c_i, σ can be characterized by the first order conditions (Lemma 4.2.3, van Damme, 1991): for each $i \in N$ and $\{a_l, a_k\} \subseteq A_i$,

$$\mathrm{E}_{(\sigma_{-i},a_l)}u_i - \mathrm{E}_{(\sigma_{-i},a_k)}u_i = c_i'(\sigma_i(a_l)) - c_i'(\sigma_i(a_k)). \tag{3}$$

Thus, given an interior σ, we can construct c whose associated game has σ as an equilibrium by satisfying (3). If σ is interior and payoff-monotone, this is trivial. Suppose that $\sigma_i(a_1) \leq ... \leq \sigma_i(a_K)$. By payoff monotonicity, $\mathrm{E}_{(\sigma_{-i},a_1)}u_i \leq ... \leq \mathrm{E}_{(\sigma_{-i},a_K)}u_i$. Thus, for some $\varepsilon > 0$, one can define $c'(1) = 0$, $c'(\sigma_i(a_K)) = -\varepsilon$, and for each $l < K$,

$$c'(\sigma_i(a_l)) = c'(\sigma_i(a_{l+1})) + \left(\mathrm{E}_{(\sigma_{-i},a_l)}u_i - \mathrm{E}_{(\sigma_{-i},a_{l+1})}u_i\right).$$

It is easy to complete the definition of c' so it is increasing, continuous, and satisfies the integrability restrictions. For instance, let $\underline{s}$ be the minimum probability assigned by σ_i to some action. Then, one can define $c'(x) = -\underline{s}c_i'(\underline{s})/x$ for $x \in (0, \underline{s}]$ and linearly interpolate the values already defined on $[\underline{s}, 1]$. This generates c_i with a logarithmic asymptote on $(0, \underline{s}]$ and a $|A_i|$-segment second order spline on $[\underline{s}, 1]$. Guaranteeing that one can select c' small enough when σ is close to σ^* requires an additional detail.

To avoid trivialities assume that σ^* does not have full support. Let $\overline{s}$ be the maximum of the probabilities assigned by σ_i to actions not in the support of σ_i^*. Let $\overline{b}$ be an action that is assigned probability $\overline{s}$ by σ_i. Let $\underline{a}$ be an action in the support of σ_i^* that is assigned minimal probability among the actions in its support. If σ^* is a Nash equilibrium of u and σ is close to σ^*, it must be that $\underline{s} \leq \overline{s} \approx 0$, and $\sigma_i(\underline{a}) \approx \sigma_i^*(\underline{a})$. Thus, one can complete the definition of c' requiring that $c_i'(\overline{s}+\varepsilon) = c_i'(\sigma_i(\underline{a}))-\varepsilon$. By payoff monotonicity, $\mathrm{E}_{(\sigma_{-i},\overline{b})}u_i - \mathrm{E}_{(\sigma_{-i},\underline{a})} < 0$. Thus, since the expected utility of actions in the support of σ^* is the same, for arbitrary $\delta > 0$, one can choose $\varepsilon \approx 0$ so c' is monotone and for each $x \in (\overline{s} + \varepsilon, 1]$, $c(x) < \delta$. $\square$

5 Discussion and concluding remarks

We have advanced our understanding of the empirical equilibrium refinement. Empirical equilibria are defined as the Nash equilibria that are the limit of weakly payoff-monotone behavior. Alternatively, they can be characterized as those that are the limits of interior payoff-monotone behavior (Sec. 4.1). This characterization facilitates the computation of this set in applications. Empirical equilibria are also the Nash equilibria that are the limits of RQRE behavior for sequences of noisy best responses that approximate best response operators in the limit (Sec. 4.2). A particular finite dimensional family of RQRE, actually span the whole empirical content of RQRE for a fixed game (Sec. 4.4). This characterization provides a clear connection between this equilibrium refinement and one of the most popular structural theories for the analysis of data from economics experiments and for the rationalization of behavior in random choice environments.

In contrast to RQRE approximation, monotone best response additive randomly perturbed payoff models do not span the whole empirical content of RQRE for a fixed game and thus are not a basis to define EE (Sec. 4.3). The result highlights that this refinement is based only on a refutable theory, weak payoff-monotonicity, and not on implicit or less understood cardinal or structural assumptions. At the same time, we also advance our understanding of additive randomly perturbed payoff models, in particular, the Monotone SQRE model. Indeed, our results show a clear tradeoff between specification and refutability of SQRE models. If there is no reason to rule out behavior beyond payoff monotonicity, a data analysis based on a parametric version of the SQRE model that guarantees monotonicity may be misspecified. As a consequence, the interpretation of SQRE estimates requires that specification of the model must be empirically addressed.

Theorem 3 implies that for each action space in which at least one agent has at least three actions, a payoff matrix exists that admits an empirical equilibrium that is not in the closure of the empirical content of randomly perturbed payoff models whose best response correspondences are weakly payoff-monotone. This result can be generalized for the SQRE model. Trivially, the later requirement can be stated in terms of the corresponding structural QRFs requiring monotonicity directly on them. More interestingly, one can actually simply require that SQRFs admit only weakly monotone fixed points.[15]

There is a precedent for the differences we uncover in the empirical content of additive randomly perturbed payoff models and deterministically perturbed models. In the random choice domain, Fudenberg et al. (2015) show that these models are independent when there are at least *four* alternatives. These models have the same empirical content with two or three alternatives. In this domain, the researcher observes the choices of an agent among menus of alternatives. The researcher observes no payoffs, but observes frequencies of choice of the same alternative in different menus. The most popular models for the rationalization of frequencies of play in this context are the random utility model and the variational preferences model. These correspond exactly to our additive randomly perturbed payoff models and our deterministically perturbed payoff models where payoffs are replaced by utility, which is not observable and is adjusted to rationalize behavior. Thus, the source of the differences we uncover in these models is independent from those previously observed. We are noting a difference in their empirical content for a particular payoff matrix when we assume that the perturbations need to produce payoff-monotone behavior for any other payoff matrix. Note that the difference exists for a payoff matrix as long as there are at least *three* actions available to an agent.

The idea to refine Nash equilibrium with proximity to behavior rationalized by Noisy Best Response models dates back to McKelvey and Palfrey (1996) and van Damme (1991). McKelvey and Palfrey (1996) propose approachability by Logistic QRE that approach the best response operator. van Damme (1991) proposes approachability by behavior in additive

[15] More precisely, one can show, by means of a fixed point argument that if a QRF violates weak payoff monotonicity, then there is a payoff matrix for which there is a QRE that violates weak payoff monotonicity. A proof of this result is available in Velez and Brown (2019).

randomly perturbed payoff models with uniformly vanishing exchangeable permutations, and by behavior in control cost models with uniformly vanishing costs. These refinements have a precedent in Harsanyi (1973) and Rosenthal (1989). Our results reveal that EE is the refinement one obtains when one considers approachability by arbitrary RQRE models that are utility maximizers in the limit. In this sense EE articulates, in a robust way, the original spirit of the approachability refinements of McKelvey and Palfrey (1996) and van Damme (1991). Our results also reveal that this refinement can be given a solid non-parametric foundation on approachability by weakly payoff-monotone behavior (Velez and Brown (2020b)'s definition of EE), that turns out to have a simpler characterization in applications than any refinement based on additive randomly perturbed payoff models.

The idea behind Velez and Brown (2020b)'s definition of EE can be related also to the meta-theories introduced by Goeree et al. (2018a) and Goeree and Louis (2021). These theories are based on the requirement that differences in utility induce differences in behavior. This requirement generates an empirical content whose closure contains all weakly payoff-monotone behavior and can contain isolated Nash equilibrium points that are not empirical. More precisely, Proposition 3 in Goeree and Louis (2021) states a characterization of distributions that can be generated by RQRE in terms of M-equilibrium inequalities. This result is useful in applications, for M-equilibrium can be computed in finite time and its empirical content contains that of RQRE. However, this proposition only applies to generic distributions in a set of generic games and thus does not provide a basis of refinement for all finite games. The set of M-equilibrium behavior can contain isolated points that are not in the closure of the empirical content of RQRE. Importantly, these isolated points can also be Nash equilibria. This is an essential difference for the purpose of refining Nash equilibrium. For instance, the discoordination equilibrium in Example 1, which corresponds to a two-person version of the Top Trading Cycles mechanism (see Velez and Brown, 2020b), is in the empirical content of M-equilibrium.[16]

As a by product of our characterizations of EE, our results contribute to a better understanding of the empirical content of RQRE, a subject of substantial current interest.[17] Lemma 1 together with Theorem 1 provide a non-parametric characterization of the closure of this empirical content for a fixed payoff matrix. This is the first characterization of this set, in a strict sense and for all games. It is worth noting that our proof of Theorem 3

[16] Goeree and Louis (2021) define a refinement of Nash equilibrium based on their theory, which discards all weakly payoff monotone behavior. The idea behind EE is to use the proximity to MNBRE models to select the weakly dominant equilibria that need to be taken into account by a mechanism designer.

[17] Melo et al. (2019) develop nonparametric tests of the SQRE hypothesis for research environments with variable payoff matrices. Friedman and Mauersberger (2022) characterize the empirical content of RQRE in 2×2 games with the additional requirement of symmetry. Friedman and Goncalves (2023) also characterize symmetric RQRE in binary-action games with a continuum of types.

can be modified to produce a fixed payoffs non-parametric characterization of the empirical content of MSQRE for games in which agents have at most three actions. It is interesting to pursue the development of specification tests based on these characterizations.

Finally, it is also interesting but beyond the scope of this paper to advance applications of EE and to evaluate its support in data. Velez and Brown (2020b) study full robust implementation in EE and contrast these results with the accumulated experimental evidence on dominant strategy mechanisms; Brown and Velez (2020) and Velez and Brown (2020a) theoretically and empirically study full implementation with complete information and obtain characterizations of EE for popular partnership dissolution auctions.

6 Appendix

Proof of Theorem 1. Payoff monotone distributions are weakly payoff-monotone. Thus we only need to prove that an empirical equilibrium is always the limit of interior payoff-monotone distributions. Let μ be weakly payoff-monotone for $u \in \mathcal{U}$. Let $\varepsilon > 0$. We prove that there is an interior γ that is payoff-monotone for u such that $||\mu - \gamma|| < \varepsilon$. This implies that $\sigma \in \mathrm{N}(u)$ is the limit of a sequence of weakly payoff-monotone distributions for u if and only if it is the limit of a sequence of interior payoff-monotone distributions for u.

For each $i \in N$, each $\zeta \in (0,1)$, and each profile of distributions $\beta \in \Sigma(A)$, let

$$f_i^\zeta(\beta) := (1-\zeta)\mu_i + \zeta l^\lambda((\mathrm{E}_{(\beta_{-i},a_i)}u_i)_{a_i \in A_i}),$$

where l^λ is the Logistic QRF defined in (1).

Let γ^ζ be a fixed point of f^ζ, that exists because f^ζ is continuous. Let $\{a_i, \hat{a}_i\} \subseteq A_i$. Suppose first that $\mu_i(a_i) = \mu_i(\hat{a}_i)$. We know that

$$l^\lambda_{a_i}((\mathrm{E}_{(\gamma^\zeta_{-i},b_i)}u_i)_{b_i \in A_i}) \geq l^\lambda_{\hat{a}_i}((\mathrm{E}_{(\gamma^\zeta_{-i},b_i)}u_i)_{b_i \in A_i}),$$

if and only if $\mathrm{E}_{(\gamma^\zeta_{-i},a_i)}u_i \geq \mathrm{E}_{(\gamma^\zeta_{-i},\hat{a}_i)}u_i$. Thus, $\mathrm{E}_{(\gamma^\zeta_{-i},a_i)}u_i \geq \mathrm{E}_{(\gamma^\zeta_{-i},\hat{a}_i)}u_i$ if and only if $\gamma^\zeta(a_i) \geq \gamma^\zeta(\hat{a}_i)$. Suppose then that $\mu_i(a_i) > \mu_i(\hat{a}_i)$. Since μ is weakly payoff-monotone for u, $\mathrm{E}_{(\mu_{-i},a_i)}u_i > \mathrm{E}_{(\mu_{-i},\hat{a}_i)}u_i$. Since as $\zeta \to 0$, $\gamma^\zeta \to \mu$, there is $c > 0$ such that for each $\zeta < c$, $\gamma_i^\zeta(a_i) > \gamma_i^\zeta(\hat{a}_i)$ and $\mathrm{E}_{(\gamma^\zeta_{-i},a_i)}u_i > \mathrm{E}_{(\gamma^\zeta_{-i},\hat{a}_i)}u_i$. Thus, for each pair $\{a_i, \hat{a}_i\} \subseteq A_i$, there is $c > 0$ such that for each $\zeta < c$, $\mathrm{E}_{(\gamma^\zeta_{-i},a_i)}u_i \geq \mathrm{E}_{(\gamma^\zeta_{-i},\hat{a}_i)}u_i$ if and only if $\gamma^\zeta(a_i) \geq \gamma^\zeta(\hat{a}_i)$. Since $\Gamma(u)$ has finite action spaces, there is $c > 0$ such that for each $\zeta < c$, each $i \in N$, and each pair $\{a_i, \hat{a}_i\} \subseteq A_i$, $\mathrm{E}_{(\gamma^\zeta_{-i},a_i)}u_i \geq \mathrm{E}_{(\gamma^\zeta_{-i},\hat{a}_i)}u_i$ if and only if $\gamma^\zeta(a_i) \geq \gamma^\zeta(\hat{a}_i)$. $\square$

Proof of Proposition 1. Let $u \in \mathcal{U}$ and $\sigma \in \mathrm{N}(u)$. We prove that there is a sequence of vanishing purifying perturbations $\{\gamma^\lambda\}_{\lambda\in\mathbb{N}}$ and a corresponding sequence of Bayesian Nash equilibria in the perturbed games whose observable strategy distributions converge to σ. We construct perturbations with satisfying Harsanyi (1973)'s requirements that can be easily modified to induce structural QRE models.

We first construct perturbations with compact support. For each Lebesgue measurable and bounded set with non-empty interior $E \subseteq \mathbb{R}^A$, let $U(E)$ be the uniform distribution on E, i.e., the normalized Lebesgue measure on it. Let $V_\varepsilon \subseteq \mathbb{R}^A$ be the open ball centered at zero with radius $\varepsilon > 0$ and for each $\lambda \in \mathbb{N}$ let

$$T_\lambda^{a_i} := \{\eta_i \in \mathbb{R}^A : \forall a_i' \in A_i \setminus \{a_i\}, \forall a_{-i} \in A_{-i}, \eta_i(a_{-i}, a_i) > \eta_i(a_{-i}, a_i') + 1/(2|A_{-i}|\lambda)\},$$

and

$$\mu_i^\lambda := \sum_{a_i \in A_i} \sigma_i(a_i) U(V_{1/\lambda} \cap T_\lambda^{a_i}).$$

To understand the structure of μ_i^λ, let a_i be in the support of $\sigma \in \mathrm{N}(u)$. Observe that for each $a_i' \in A_i$ and each realization $\eta_i \in T_\lambda^{a_i}$,

$$\textstyle\sum_{a_{-i}\in A_{-i}} (u_i(a_{-i}, a_i) + \eta_i(a_{-i}, a_i))\sigma_{-i}(a_{-i}) - \sum_{a_{-i}\in A_{-i}} (u_i(a_{-i}, a_i') + \eta_i(a_{-i}, a_i'))\sigma_{-i}(a_{-i}) \geq \min_{i\in N} 1/(2|A_{-i}|\lambda).$$

This implies that $\sigma_i = B_i^{u_i,\mu^\lambda}(\sigma)$. Thus, $\sigma \in \mathrm{BNE}(u, \mu^\lambda)$. Clearly, the sequence $\{\gamma^\lambda\}_{\lambda\in\mathbb{N}}$ is vanishing. Thus, there is a sequence of vanishing perturbations for which σ itself is an element of $\mathrm{BNE}(u, \mu^\lambda)$.

We now construct a sequence of full support perturbations with associated equilibria converging to σ. The construction is again based on μ^λ.

Fix $\lambda \in \mathbb{N}$, and let $\delta > 0$, $K_\delta := \{\sigma' \in \Sigma(A) : ||\sigma - \sigma'|| \leq \delta\}$, and $\sigma' \in K_\delta$. Since $\sigma \in \mathrm{N}(u)$, for each $\eta_i \in T_\lambda^{a_i}$ and each $a_i' \in A_i$,

$$\begin{aligned}
&\textstyle\sum_{a_{-i}\in A_{-i}} (u_i(a_{-i}, a_i) + \eta_i(a_{-i}, a_i))\sigma'_{-i}(a_{-i}) - \sum_{a_{-i}\in A_{-i}} (u_i(a_{-i}, a_i') + \eta_i(a_{-i}, a_i'))\sigma'_{-i}(a_{-i}) = \\
&\textstyle\sum_{a_{-i}\in A_{-i}} (u_i(a_{-i}, a_i) - u_i(a_{-i}, a_i'))\sigma_{-i})(a_{-i}) + \sum_{a_{-i}\in A_{-i}} (u_i(a_{-i}, a_i) - u_i(a_{-i}, a_i'))(\sigma'_{-i} - \sigma_{-i})(a_{-i}) + \\
&(\eta_i(a_{-i}, a_i) - \eta_i(a_{-i}, a_i'))\sigma'(a_{-i}) \geq \\
&\textstyle\{\sum_{a_{-i}\in A_{-i}} (u_i(a_{-i}, a_i) - u_i(a_{-i}, a_i'))(\sigma'_{-i}(a_{-i}) - \sigma_{-i}(a_{-i}))\} + 1/(2|A_{-i}|\lambda) \geq \\
&M \max_{a_l\in A_i} (\sigma_i'(a_l) - \sigma_i(a_l)) + \min_{i\in N} 1/(2|A_{-i}|\lambda),
\end{aligned}$$

for some $M < 0$ that depends only on u. Thus, by choosing δ close to zero we can guarantee that when perturbations follow μ^λ, each σ_i is the observable best response to a distribution that is δ close to σ. Fix this value of $\delta < 1/\lambda$ and denote it by $\delta(\lambda)$. Fix, γ, an arbitrary

	Player 1				
	a_1	a_2	$\dots$	a_{K-1}	a_K
a_{-1}^*	5	5	$\dots$	5	5
$A_{-1} \setminus \{a_{-1}^*\}$	1	2	$\dots$	2	4

Table 3: Game $\Gamma(u) := (N, A, u)$, $N := \{1, ..., n\}$, $A_1 := \{a_1, ..., a_K\}$ with $K \geq 3$, and $|A_{-1}| \geq 2$. The table shows the payoff of agent 1. Each agent $j > 1$ gets a payoff of 1 if she plays a_j^* and zero otherwise.

full support Borel probability measure on $\mathbb{R}^A$. Let $\gamma^\lambda := (1 - \delta(\lambda))\mu^\lambda + \delta(\lambda)\gamma$. Clearly, perturbations $\{\gamma^\lambda\}_{\lambda \in \mathbb{N}}$ vanish as $\lambda \to \infty$. Finally, observe that for each $\sigma' \in \Sigma(A)$,

$$||\sigma - (B_i^{u_i,\gamma^\lambda}(\sigma'_{-i}))_{i \in N}|| = ||\sigma - ((1 - \delta(\lambda))\sigma + \delta(\lambda)\psi(\sigma'))|| = \delta(\lambda)||\sigma - \psi(\sigma')|| \leq \delta(\lambda),$$

where $\psi(\sigma') = (B_i^{u_i,\gamma}(\sigma'_{-i}))_{i \in N}$. Thus, $(B_i^{u_i,\gamma^\lambda}(\cdot))_{i \in N} : B_\delta \to B_\delta$. Since B_δ is compact and convex, by Brouwer's fixed point theorem, it has a fixed point in B_δ. Thus, for each $\lambda \in \mathbb{N}$, there is $\sigma^\lambda \in \mathrm{BNE}(u, \gamma^\lambda)$ such that $||\sigma - \sigma^\lambda|| \leq \delta(\lambda)$. □

We now prove Theorem 3. The game that allows us to prove this result is defined in Table 3. In this game each agent $i \neq 1$ has a strictly dominant action. The profile of these strictly dominant actions for these agents is a_{-1}^*. Agent 1 is indifferent among all actions if all other agents play their strictly dominant action. Agent 1 has three different types of actions. Action a_1, which is weakly dominated by actions $a_2, ..., a_{K-1}$, which are all payoff equivalent. All actions $\{a_1, ..., a_{K-1}\}$ are weakly dominated by a_K for this agent. The essential feature of this game is that given any σ_{-1} for which $\sigma_{-i}(a_{-i}^*) < 1$, the difference in agent 1's expected payoff between actions a_K and a_{K-1} is greater than the difference in expected payoff between actions a_2 and a_1.

Lemma 2. Let $\Gamma(u)$ be the game in Table 3. There is $\sigma \in \mathrm{N}(u)$ that belongs to the closure of $\{\gamma : \gamma$ is weakly payoff-monotone for $u\}$ and in which each agent $j \neq 1$ plays the strictly dominant action and $\sigma_1(a_1) < 1/K < \sigma_1(a_2) = \cdots = \sigma_1(a_{K-1}) < \sigma_1(a_K)$.

Proof of Lemma 2. Clearly $\mathrm{N}(u)$ is the set of distributions in which each $j \neq 1$ plays the dominant action and agent 1 arbitrarily randomizes. Let σ be such that each agent $j \neq 1$ plays the strictly dominant action with certainty, and $\sigma_1(a_1) < \sigma_1(a_2) = \cdots = \sigma_1(a_{K-1}) < \sigma_1(a_K)$. Then, $\sigma \in \mathrm{N}(u)$. Let $\lambda \in \mathbb{N}$ and σ^λ be the convex combination that places $(1 - 1/\lambda)$ weight on σ and $1/\lambda$ on a uniform distribution. Clearly as $\lambda \to \infty$, $\sigma^\lambda \to \sigma$. Thus, there is $\Lambda \in \mathbb{N}$ such that for each $\lambda \geq \Lambda$, σ^λ is ordinally equivalent to σ. Since for each $\lambda \in \mathbb{N}$, σ^λ is interior, $E_{(\sigma_{-i}^\lambda, a_1)} u_i < E_{(\sigma_{-i}^\lambda, a_2)} u_i = \cdots = E_{(\sigma_{-i}^\lambda, a_{K-1})} u_i < E_{(\sigma_{-i}^\lambda, a_K)} u_i$, and for each $j \neq i$, if $a_j^* \in A_j$ is this agent's dominant action, $E_{(\sigma_{-j}^\lambda, a_j^*)} u_j > E_{(\sigma_{-j}^\lambda, a_j)} u_j$, and for each pair of actions $\{a_j, a_j'\} \subseteq A_j$ that are not dominant, $E_{(\sigma_{-j}^\lambda, a_j)} u_j =$

$E_{(\sigma^{\lambda}_{-j},a'_j)}u_j$. Thus, σ^{λ} is weakly payoff-monotone for u. Thus, σ belongs to the closure of $\{\gamma : \gamma$ is weakly payoff-monotone for $u\}$. Thus, for each $1/K < \alpha < 1/(K-1)$, there is $\sigma \in \mathrm{N}(u)$ that belongs to the closure of $\{\gamma : \gamma$ is weakly payoff-monotone for $u\}$ and such that $0 = \sigma_1(a_1) < \alpha = \sigma_1(a_2) = \cdots = \sigma_1(a_{K-1}) < (1-(K-2)\alpha) = \sigma_1(a_K)$. □

Lemma 2 states that there is weakly payoff-monotone behavior for u that is arbitrarily close to an empirical equilibrium of $\Gamma(u)$ in which agent 1 plays actions $\{a_2, ..., a_{K-1}\}$ with probability greater than $1/K$. The following proposition identifies restrictions on distributions generated by additive randomly perturbed payoff models whose best response operators are weakly payoff-monotone. The proof of Theorem 3 is completed by showing, based on these restrictions, that it is impossible for these models to generate behavior close to the equilibrium identified in Lemma 2.

Proposition 2. Let $\Gamma(u)$ be the game in Table 3 and $\mu \in \mathcal{M}$. Let $\sigma_{-i} \in \Delta(A_{-i})$ be such that $E_{(\sigma_{-i},a_1)}u_i < E_{(\sigma_{-i},a_2)}u_i = \cdots = E_{(\sigma_{-i},a_{K-1})}u_i < E_{(\sigma_{-i},a_K)}u_i$, and $E_{(\sigma_{-i},a_K)}u_i - E_{(\sigma_{-i},a_{K-1})}u_i > E_{(\sigma_{-i},a_2)}u_i - E_{(\sigma_{-i},a_1)}u_i$. Then, $B^{u_i,\mu_i}_{ia_{K-1}}(\sigma_{-i}) \leq 1/K$.

Proof of Proposition 2. Recall that in our notation $A_i = \{1, ..., K\}$. For each $k \in \{1, ..., K\}$ let $v_k := E_{(\sigma_{-i},a_k)}u_i$, fore each $x \in \mathbb{R}^A$, $x_k := \sum_{a_{-i}\in A_{-i}} x_{(a_{-i},a_k)}\sigma_{-i}(a_{-i})$, and

$$X_k := \{x \in \mathbb{R}^A : \{k\} = \arg\max_{l=1,...,K} x_l\}.$$

Let $\mu \in \mathcal{M}$. Consider $\bar{u}_i \in \mathbb{R}^A$ for which agent i has equal payoff from each action profile. Since $\mu \in \mathcal{M}$, $B^{\bar{u}_i,\mu_i}_i(\sigma_{-i}) = (1/K, ..., 1/K)$. Thus, for each $k = 1, ..., K$, $\mu_i(X_k) = 1/K$.

Since $\mu \in \mathcal{B}$, for each measurable set $G \subseteq \mathbb{R}^A$,

$$\mu_i(G) = \sum_{l=1}^{K} \mu_i(G \cap X_k). \tag{4}$$

Let $G := \{x \in \mathbb{R}^A : \{K-1\} = \arg\max_{l=1,...,K} v_l + x_l\}$. Then, $B^{u_i,\mu_i}_{ia_{K-1}}(\sigma_{-i}) = \mu_i(G)$. Since $v_1 < v_2 = \cdots = v_{K-1} < v_K$,

$$\mu_i(G) = \mu_i(G \cap X_1) + \mu_i(G \cap X_{K-1}). \tag{5}$$

Let $D_1 := v_2 - v_1$ and $D_2 := v_K - v_{K-1}$. Consider $u'_i \in \mathbb{R}^A$ for which $y' := E_{(\sigma_{-i},\cdot)}u'_i$ is such that $y'_1 = \cdots = y'_{K-1} < y'_K := y'_{K-1} + D_2$ (one can simply make payoffs be independent of the action of the other agents). Since $\mu_i \in \mathcal{M}$, $B^{u'_i,\mu_i}_{ia_1}(\sigma_{-i}) = B^{u'_i,\mu_i}_{ia_{K-1}}(\sigma_{-i})$. Moreover,

$$\begin{aligned}
&B^{u'_i,\mu_i}_{ia_{K-1}}(\sigma_{-i}) = \mu_i\left(\{x \in X_{K-1} : x_{K-1} > x_K + D_2\}\right), \\
&B^{u'_i,\mu_i}_{ia_1}(\sigma_{-i}) = \mu_i\left(\{x \in X_1 : x_1 > x_K + D_2\}\right).
\end{aligned}$$

Thus,

$$\mu_i\left(\{x \in X_{K-1} : x_{K-1} > x_K + D_2\}\right) = \mu_i\left(\{x \in X_1 : x_1 > x_K + D_2\}\right).$$

Since $\mu_i \in \mathcal{B}_i$ and $\mu_i(X_1) = \mu_i(X_{K-1})$,

$$\mu_i\left(\{x \in X_{K-1} : x_{K-1} < x_K + D_2\}\right) = \mu_i\left(\{x \in X_1 : x_1 < x_K + D_2\}\right). \tag{6}$$

We claim that

$$\begin{aligned}
&\mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-2}, x_K\} < x_{K-1}\}\right) \\
&= \mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-1}\} < x_K\}\right).
\end{aligned} \tag{7}$$

Consider $u''_i \in \mathbb{R}^A$ for which $y'' := E_{(\sigma_{-i},\cdot)}u''_i$ is such that $y''_1 < y''_2 = \cdots = y''_K := y''_1 + D_1$. Observe that

$$B^{u''_i,\mu_i}_{ia_{K-1}}(\sigma_{-i}) = \mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-2}, x_K\} < x_{K-1}\}\right) + \mu_i(X_{K-1}),$$

and

$$B^{u''_i,\mu_i}_{ia_K}(\sigma_{-i}) = \mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-1}\} < x_K\}\right) + \mu_i(X_K).$$

Since $\mu \in \mathcal{M}$, $B^{u''_i,\mu_i}_{ia_K}(\sigma_{-i}) = B^{u''_i,\mu_i}_{ia_{K-1}}(\sigma_{-i})$. Thus, since $\mu_i(X_{K-1}) = \mu_i(X_K)$, (7) follows.

Since $D_2 \geq D_1$, by monotonicity of measures with respect to set inclusion

$$\begin{aligned}
&\mu_i\left(\{x \in X_1 : x_1 < x_K + D_2\}\right) \\
&\geq \mu_i\left(\{x \in X_1 : x_1 < x_K + D_1\}\right) \\
&\geq \mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-1}\} < x_K\}\right).
\end{aligned}$$

Replacing (6) and (7) in the first and last expressions of the inequality above yields,

$$\begin{aligned}
&\mu_i\left(\{x \in X_{K-1} : x_{K-1} < x_K + D_2\}\right) \\
&\geq \mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-2}, x_K\} < x_{K-1}\}\right).
\end{aligned}$$

Now,

$$\mu_i(G \cap X_{K-1}) = \mu_i(X_{K-1}) - \mu_i\left(\{x \in X_{K-1} : x_{K-1} < x_K + D_2\}\right),$$

and by monotonicity of measures with respect to set inclusion,

$$\begin{aligned}\mu_i(G \cap X_1) &= \mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-2}, x_K + D_2\} < x_{K-1}\}\right)\\ &\leq \mu_i\left(\{x \in X_1 : \max\{x_1 - D_1, x_2, ..., x_{K-2}, x_K\} < x_{K-1}\}\right).\end{aligned}$$

Thus,

$$\mu_i(G \cap X_1) + \mu_i(G \cap X_{K-1}) \leq \mu_i(X_{K-1}) = 1/K.$$

Thus, by (5),

$$B^{u_i,\mu_i}_{ia_{K-1}}(\sigma_{-i}) = \mu_i(G) = \mu_i(G \cap X_1) + \mu_i(G \cap X_{K-1}) \leq 1/K.$$

□

Proof of Theorem 3. Let $\Gamma(u)$ be the game in Table 3. By Lemma 2, there is $\sigma^* \in \mathrm{N}(u)$ that belongs to the closure of $\{\gamma : \gamma$ is weakly payoff-monotone for $u\}$ in which each agent $j \neq 1$ plays the strictly dominant action and $\sigma^*_1(a_1) < 1/K < \sigma^*_1(a_2) = \cdots = \sigma^*_1(a_{K-1}) < \sigma^*_1(a_K)$. Thus, there is $\varepsilon > 0$ for which for each $\sigma \in \Delta$ such that $||\sigma - \sigma^*|| < \varepsilon$, $\sigma_1(a_1) < 1/K < \sigma_1(a_2)$. Let σ that is weakly payoff-monotone for u and $||\sigma - \sigma^*|| < \varepsilon$. Since $\sigma_1(a_1) < 1/K < \sigma_1(a_2)$ and σ is weakly payoff-monotone for u, we have that $E_{(\sigma_{-i},a_1)}u_i < E_{(\sigma_{-i},a_2)}u_i$. Thus, $\sigma_{-i}(a^*_{-i}, a_1) < 1$, for otherwise $E_{(\sigma_{-i},a_1)}u_i = E_{(\sigma_{-i},a_2)}u_i$. Thus, $E_{(\sigma_{-i},a_2)}u_i - E_{(\sigma_{-i},a_1)}u_i = (1 - \sigma_{-i}(a^*_{-i}, a_1)) > 0$ and $E_{(\sigma_{-i},a_K)}u_i - E_{(\sigma_{-i},a_{K-1})}u_i = 2(1 - \sigma_{-i}(a^*_{-i}, a_1))$. Thus, $E_{(\sigma_{-i},a_K)}u_i - E_{(\sigma_{-i},a_{K-1})}u_i > E_{(\sigma_{-i},a_2)}u_i - E_{(\sigma_{-i},a_1)}u_i$. Thus, there is no $\mu \in \mathcal{M}$ such that $\sigma \in \mathrm{BNE}(u, \mu)$, for otherwise by Proposition 2, $\sigma_1(a_2) \leq 1/K$. Thus, $\{\sigma : ||\sigma - \sigma^*|| < \varepsilon\} \cap \{\sigma : \exists \mu \in \mathcal{M},\ \text{s.t. } \sigma \in \mathrm{BNE}(u, \mu)\} = \emptyset$. □

Proof of Theorem 2. Since best response operators in perturbed games are RQRFs, Theorem 2 is a corollary of Theorem 4. □

References

Brown, A. L., Velez, R. A., 2020. Empirical bias and efficiency of alpha-auctions: experimental evidence, mimeo Texas A&M University.
URL https://arxiv.org/abs/1905.03876

Fosgerau, M., Melo, E., de Palma, A., Shum, M., 2020. Discrete choice and rational inattention: A general equivalence result. International Economic Review 61 (4), 1569–1589.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/iere.12469

Friedman, E., Goncalves, D., 2023. Quantal response equilibrium with a continuum of types: Characterization and nonparametric identification.
URL https://www.dropbox.com/s/jvz4vah3fq6qkr3/ContinuousQRE.pdf?dl=0

Friedman, E., Mauersberger, F., 2022. Quantal response equilibrium with symmetry: Representation and applications.
URL https://www.dropbox.com/s/tqj2vpbeoowitz5/SQRE.pdf?dl=0

Fudenberg, D., Iijima, R., Strzalecki, T., 2015. Stochastic choice and revealed perturbed utility. Econometrica 83 (6), 2371–2409.
URL http://www.jstor.org/stable/43866415

Goeree, J., Holt, C. A., Louis, P., Palfrey, T. R., Rogers, B., 2018a. Rank-dependent choice equilibrium: A non-parametric generalization of qre. In: The Handbook of Research Methods and Applications in Experimental Economics. Forthcoming.

Goeree, J. K., Holt, C. A., Palfrey, T. R., 2005. Regular quantal response equilibrium. Experimental Economics 8 (4), 347–367.
URL http://dx.doi.org/10.1007/s10683-005-5374-7

Goeree, J. K., Holt, C. A., Palfrey, T. R., 2016. Quantal Response Equilibrium: A Stochastic Theroy of Games. Princeton Univ. Press, Princeton, NJ.

Goeree, J. K., Louis, P., December 2021. M equilibrium: A theory of beliefs and choices in games. American Economic Review 111 (12), 4002–45.
URL https://www.aeaweb.org/articles?id=10.1257/aer.20201683

Goeree, J. K., Louis, P., Zhang, J., 2018b. Noisy introspection in the 11–20 game. The Economic Journal 128 (611), 1509–1530.
URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecoj.12479

Golman, R., 2011. Quantal response equilibria with heterogeneous agents. J Econ Theory 146 (5), 2013 – 2028.
URL https://doi.org/10.1016/j.jet.2011.06.007

Govindan, S., Reny, P. J., Robson, A. J., 2003. A short proof of Harsanyi's purification theorem. Games Econ Behavior 45 (2), 369 – 374, special Issue in Honor of Robert W. Rosenthal.
URL http://www.sciencedirect.com/science/article/pii/S0899825603001490

Haile, P. A., Hortaçsu, A., Kosenok, G., 2008. On the empirical content of quantal response equilibrium. Amer Econ Review 98 (1), 180–200.
URL http://www.jstor.org/stable/29729968

Harsanyi, J. C., Dec 1973. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. Int J Game Theory 2 (1), 1–23.
URL https://doi.org/10.1007/BF01737554

Jackson, M. O., 1991. Bayesian implementation. Econometrica 59 (2), 461–477.
URL http://www.jstor.org/stable/2938265

Jackson, M. O., 1992. Implementation in undominated strategies: A look at bounded mechanisms. The Review of Economic Studies 59 (4), 757–775.
URL http://www.jstor.org/stable/2297996

Maskin, E., 1999. Nash equilibrium and welfare optimality. Review Econ Studies 66, 83–114, first circulated in 1977.
URL http://www.jstor.org/stable/2566947

McKelvey, R. D., Palfrey, T. R., 1995. Quantal response equilibria for normal form games. Games and Economic Behavior 10 (1), 6–38.
URL http://dx.doi.org/10.1006/game.1995.1023

McKelvey, R. D., Palfrey, T. R., 1996. A statistcial theory of equilibrium in games. Japanese Econ Review 47 (2), 186–209.

Melo, E., Pogorelskiy, K., Shum, M., 2019. Testing the quantal response hypothesis. International Economic Review 60 (1), 53–74.

Myerson, R. B., Jun 1978. Refinements of the nash equilibrium concept. International Journal of Game Theory 7 (2), 73–80.
URL https://doi.org/10.1007/BF01753236

Nash, J., 1951. Non-cooperative games. Annals of Mathematics 54 (2), 286–295.
URL http://www.jstor.org/stable/1969529

Palfrey, T. R., Srivastava, S., 1989. Implementation with incomplete information in exchange economies. Econometrica 57 (1), 115–134.
URL http://www.jstor.org/stable/1912575

Palfrey, T. R., Srivastava, S., 1991. Nash implementation using undominated strategies. Econometrica 59 (2), 479–501.
URL http://www.jstor.org/stable/2938266

Rosenthal, R. W., Sep 1989. A bounded-rationality approach to the study of noncooperative games. Int J Game Theory 18 (3), 273–292.
URL https://doi.org/10.1007/BF01254292

Selten, R., Mar 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4 (1), 25–55.
URL https://doi.org/10.1007/BF01766400

Tumennasan, N., 2013. To err is human: Implementation in quantal response equilibria. Games and Economic Behavior 77 (1), 138 – 152.
URL http://www.sciencedirect.com/science/article/pii/S0899825612001522

van Damme, E., 1991. Stability and Perfection of Nash Equilibria. Springer Berlin Heidelberg, Berlin, Heidelberg.
URL https://link.springer.com/book/10.1007/978-3-642-58242-4

Velez, R. A., Brown, A. L., 2019. The paradox of monotone structural QRE, mimeo, Texas A&M University.
URL https://arxiv.org/abs/1905.05814

Velez, R. A., Brown, A. L., 2020a. Empirical bias of extreme-price auctions: analysis, mimeo, Texas A&M University.
URL https://arxiv.org/abs/1905.08234

Velez, R. A., Brown, A. L., 2020b. Empirical strategy-proofness, mimeo, Texas A&M University.
URL https://arxiv.org/abs/1907.12408

Appendix not for publication

Lemma 3 (van Damme, 1991)**.** Let $\{f^\lambda\}_{\lambda\in\mathbb{N}}$ be a sequence of perturbation functions that vanishes and $\{\sigma^\lambda\}_{\lambda\in\mathbb{N}}$ a corresponding convergent sequence of equilibria of the games associated with $\Gamma(u)$ and f^λ. Then, $\{\sigma^\lambda\}_{\lambda\in\mathbb{N}}$ converges to a Nash equilibrium of $\Gamma(u)$.

Proof of Lemma 3. Let $\{\sigma^\lambda\}_{\lambda\in\mathbb{N}}$ be a convergent sequence such that for each $\lambda\in\mathbb{N}$, σ^λ is an equilibrium of the game associated with $\Gamma(u)$ and f^λ. Let $\sigma := \lim_{\lambda\to\infty}\sigma^\lambda$. Let $i\in N$ and $a_i\in A_i$ be a best response to σ_{-i} for agent i in $\Gamma(u)$. Suppose that $a_k\in A_i$ is not a best response to σ_{-i} for agent i. We prove that $\sigma_i(a_k)=0$. Since as $\lambda\to\infty$, $\sigma^\lambda\to\sigma$, we also have that $\sigma_i^\lambda(a_i)\to\sigma_i(a_i)$, $\sigma_i^\lambda(a_k)\to\sigma_i(a_k)$, $\mathrm{E}_{(\sigma_{-i}^\lambda,a_i)}u_i\to\mathrm{E}_{(\sigma_{-i},a_i)}u_i$, and $\mathrm{E}_{(\sigma_{-i}^\lambda,a_k)}u_i\to\mathrm{E}_{(\sigma_{-i},a_k)}u_i$. Thus, there is $\Lambda\in\mathbb{N}$ such that for each $\lambda\geq\Lambda$, $\sigma_i^\lambda(a_k)\leq\sigma_i^\lambda(a_i)$. Suppose first that $\sigma_i(a_i)=0$. Since $\sigma_i^\lambda(a_i)\to 0$, $\sigma_i^\lambda(a_k)\to 0$. Suppose then that $\sigma_i(a_i)>0$. By (van Damme, 1991, Theorem 4.2.6), for each $\lambda\in\mathbb{N}$,

$$\mathrm{E}_{(\sigma_{-i}^\lambda,a_l)}u_i-\mathrm{E}_{(\sigma_{-i}^\lambda,a_k)}u_i=(f^\lambda)_i'(\sigma_i^\lambda(a_l))-(f^\lambda)_i'(\sigma_i^\lambda(a_k)).$$

The left side of the expression above converges to a positive number. Since $\sigma_i^\lambda(a_l)\to\sigma_i(a_l)>0$ and $\{f^\lambda\}_{\lambda\in\mathbb{N}}$ vanishes, $(f^\lambda)_i'(\sigma_i^\lambda(a_l))\to 0$. Thus, $\sigma_i^\lambda(a_k))\to 0$, for otherwise there is a subsequence of $\{\sigma_i^\lambda(a_k))\}_{\lambda\in\mathbb{N}}$ that converges in the interior of $(0,1]$. If this is so the right side of the equation above converges to zero. This is a contradiction. □

The following proposition formally states our claims in Example 3

Proposition 3. Consider the game Γ_c in Table 2. Then, for each $c>0$,

$$\begin{aligned}
&\mathrm{N}(\Gamma_c)=\{(a_1,b_1),(a_2,b_2),(a_3,b_3)\},\\
&\mathrm{T}(\Gamma_c)=\mathrm{U}(\Gamma_c)=\{(a_1,b_1),(a_2,b_2)\},\\
&\mathrm{P}(\Gamma_c)=\{(a_1,b_1)\}.
\end{aligned}$$

Moreover,

$$\mathrm{EE}(\Gamma_c)=\begin{cases}\{(a_1,b_1)\} & \text{if } \min\{c_1,c_2\}\leq 1,\\ \{(a_1,b_1),(a_2,b_2)\} & \text{Otherwise.}\end{cases}$$

Proof. We first prove that $\mathrm{N}(\Gamma_c)=\{(a_1,b_1),(a_2,b_2),(a_3,b_3)\}$. Let $c:=(c_1,c_2)$ such that $c_1>0$ and $c_2>0$. One can easily see that the action profiles (a_1,b_1), (a_2,b_2), and (a_3,b_3) are the only pure strategy Nash equilibria of Γ_c. Now, let $\sigma\in\mathrm{N}(\Gamma_c)$. If $\sigma_1(a_2)>0$, then $\sigma_2(b_3)=0$. Then, $\sigma_1(a_3)=0$. It follows that either σ is equal to (a_1,b_1) or (a_2,b_2). Symmetry implies the same is true when $\sigma_2(b_2)>0$. Thus, suppose that σ_1 is not a pure strategy. Suppose that

$\sigma_1(a_1) > 0$, $\sigma_1(a_2) = 0$, and $\sigma_2(b_2) = 0$. Then $\sigma_2(b_3) = 0$. Thus, $\sigma = (a_1, b_1)$. A symmetric argument shows that if $\sigma_2(b_1) > 0$, $\sigma_1(a_2) = 0$, and $\sigma_2(b_2) = 0$, then $\sigma = (a_1, b_1)$.

It is well known that at each perfect equilibrium no agent plays a weakly dominated strategy. Clearly, a_3 and b_3 are weakly dominated for players 1 and 2, respectively. Thus, $\mathrm{T}(\Gamma_c) \subseteq \{(a_1, b_1), (a_2, b_2)\}$. Now, let $t := \min\{c_1, c_2\}$, $\varepsilon := \min\{t, t/(3c_1), t/(3c_2), 1/3\}$, and for each $\lambda \in \mathbb{N}$, σ^λ be the strategy profile for which $\sigma_1^\lambda(a_1) := \varepsilon c_2/(2\lambda t)$ and $\sigma_1^\lambda(a_3) := \varepsilon/(\lambda t)$; and $\sigma_2^\lambda(b_1) := \varepsilon c_1/(2\lambda t)$ and $\sigma_2^\lambda(b_3) := \varepsilon/(\lambda t)$. Then,

$$
\begin{aligned}
&E_{(\sigma_2^\lambda, a_1)} u_1^c = \varepsilon c_1/(2\lambda t) - (7 + c_1)\varepsilon/(\lambda t) = -7\varepsilon/(\lambda c_1) - \varepsilon c_1/(2\lambda t),\\
&E_{(\sigma_2^\lambda, a_2)} u_1^c = -7\varepsilon/(\lambda c_1),\\
&E_{(\sigma_2^\lambda, a_2)} u_1^c = -(7 + c_1)\varepsilon/(2\lambda) - 7(1 - \varepsilon/\lambda - \varepsilon/(\lambda c_1)) - 7\varepsilon/(\lambda c_1).
\end{aligned}
$$

Thus, a_2 is the unique best response to σ_2^λ for agent 1. Symmetry implies that b_2 is the unique best response to σ_1^λ for agent 2. Since σ_1^λ places probability at most $1/\lambda$ in both a_1 and a_3; σ_2^λ places probability at most $1/\lambda$ in both b_1 and b_3; and as $\lambda \to \infty$, $\sigma^\lambda \to (a_2, b_2)$, we have that $(a_2, b_2) \in \mathrm{T}(\Gamma_c)$.

Let $\Lambda > 2$ be such that for each $\lambda \geq \Lambda$, $1 - 1/(2\lambda) - 1/(3\lambda^2) > \max\{c_1/(3\lambda^2), c_2/(3\lambda^2), 1/\lambda\}$. Let $\lambda \geq \Lambda$ and σ^λ be the symmetric profile of strategies such that $\sigma_1^\lambda(a_2) := 1/(2\lambda)$ and $\sigma_1^\lambda(a_3) := 1/(3\lambda^2)$. Thus, $E_{(\sigma_2^\lambda, a_1)} u_1^c - E_{(\sigma_2^\lambda, a_2)} u_1^c = 1 - 1/(2\lambda) - 1/(3\lambda^2) - c_1/(3\lambda^2) > 0$. Clearly, $E_{(\sigma_2^\lambda, a_2)} u_1^c > E_{(\sigma_2^\lambda, a_3)} u_1^c$. Similarly, $E_{(\sigma_1^\lambda, b_1)} u_2^c - E_{(\sigma_1^\lambda, b_2)} u_2^c > 0$ and $E_{(\sigma_1^\lambda, b_2)} u_2^c > E_{(\sigma_1^\lambda, b_3)} u_2^c$. Since $\sigma_1^\lambda(a_1) > \sigma_1^\lambda(a_2)/\lambda$, $\sigma_1^\lambda(a_2) > \sigma_1^\lambda(a_3)/\lambda$, $\sigma_2^\lambda(b_1) > \sigma_2^\lambda(b_2)/\lambda$, and $\sigma_2^\lambda(b_2) > \sigma_2^\lambda(b_3)/\lambda$; and as $\lambda \to \infty$, $\sigma^\lambda \to (a_1, b_1)$, we have that $(a_1, b_1) \in \mathrm{P}(\Gamma_c) \subseteq \mathrm{T}(\Gamma_c)$. $\square$