Narrow Framing and Risk in Games*

Fedor Sandomirskiy† Po Hyun Sung‡ Omer Tamuz§ Ben Wincelberg¶

July 15, 2024

Abstract

We study finite normal-form games under a narrow framing assumption: when players play several games simultaneously, they consider each one separately. We show that under mild additional assumptions, players must play either Nash equilibria, logit quantal response equilibria, or their generalizations, which capture players with various risk attitudes.

1 Introduction

Consider a player who plays both poker and chess at a game night. Since these are separate games, a natural modeling choice for describing her decision-making is to assume that she chooses her strategies independently in the two games. In this paper, we explore the general assumption that when players face unrelated strategic choices, these choices are made independently.

We refer to this assumption as *narrow framing* and study its implications for the choice of solution concepts in games. Despite its self-evident nature, narrow framing has far-reaching implications. Augmented with various rationality assumptions, narrow framing provides a characterization of familiar solution concepts—Nash equilibrium and logit quantal response equilibrium—and also suggests novel solution concepts, incorporating non-trivial risk attitudes in equilibrium behavior.

*The paper benefited from our colleagues' suggestions and comments. We are grateful to the audience at the Cowles Foundation conference, as well as Marina Agranov, Daniel Chen, Drew Fudenberg, Elliot Lipnowski, Fabio Maccheroni, George Mailath, Yusufcan Masatlioglu, Ellen Muir, Roger Myerson, Tom Palfrey, Luciano Pomatto, Andrew Postlewaite, Doron Ravid, Larry Samuelson, Eran Shmaya, Ran Shorrer, Charles Sprenger, and Chloe Tergiman.

†Princeton University. Email: fsandomi@princeton.edu

‡Caltech. Email: psung@caltech.edu

§Caltech. Email: tamuz@caltech.edu. Omer Tamuz was supported by a BSF award (#2018397), and a National Science Foundation CAREER award (DMS-1944153).

¶Caltech. Email: bwincelb@caltech.edu

Figure 1: Bobby Fischer playing fifty games of chess simultaneously in 1964.

As the above example hints, our narrow framing axiom is a rather mild assumption: we only suppose that people treat games separately when they are unrelated. Rather than capturing a behavioral phenomenon, narrow framing is compatible with rationality, since choosing best responses separately in two unrelated games trivially gives a best response when playing the two games simultaneously. This is in contrast with the behavioral assumption that people treat even related decisions separately, which is incompatible with rationality, and often leads to suboptimal behavior.

Narrow framing is implicitly used by experimentalists when they take models to lab-generated data without considering extraneous information about subjects' lives outside of the lab environment. If narrow framing holds, subjects' behavior inside the lab need not be affected by external considerations unknown to the experimentalist, such as their wealth or strategic interactions outside the lab. Narrow framing thus justifies experimental economics by asserting that we can meaningfully simulate isolated economic scenarios in the lab.

Assumptions about choice bracketing have been studied extensively in the context of decision problems and experimental evidence shows that individuals bracket choices across many decisions, treating them each separately. See Read, Loewenstein, Rabin, Keren, and Laibson (2000) for an extensive survey. However, we are not aware of previous studies of narrow framing in games.

Formally, we take an axiomatic approach to characterizing solution concepts of finite normal form games. A solution concept assigns to each game a set of mixed strategy profiles, which are the solutions to the game. We think of these solutions as predictions for agents' behavior in the game. Our main axiom is narrow framing. In order to impose this axiom on solution concepts, we need to clarify what it means for a game to be composed of two unrelated games. Given two games, which we call the *factor games*, involving the same players, we define the *product game* by requiring the players to choose one action in each factor game. To capture the idea that the factor games are unrelated, we think of payoffs as monetary, and define the players' payoffs in the product game to be the sums of their payoffs in the factor games.

We say that a solution concept satisfies narrow framing if, given solutions for the factor games, one of the solutions for the product game is for players to choose their actions independently in the factor games, each according to its solution. I.e., if mixed strategy profile $(p_i)_i$ is a solution for factor game 1, and $(q_i)_i$ is a solution for factor game 2, then one of the solutions for the product of these games is for each player i to choose independently from p_i in factor game 1 and from q_i in factor game 2. Note that we do not require every solution of the product game to be of this form. Rather, thinking of solutions as predictions, we require that combining predictions for the factor game provides one of the valid predictions for the product game. Accordingly, narrow framing plays the role of a consistency requirement on predictions across games.

We augment narrow framing with various rationality assumptions capturing the idea that players have rational expectations about the strategies of the other players, and that they prefer higher payoffs to lower payoffs, either in expectation or in distribution. A solution concept satisfies *expectation-monotonicity* if choice probabilities are monotone with respect to expected payoffs. More specifically, we assume that a player plays an action more frequently than another action if the expected payoff of the former action is larger, given how others are playing. This axiom includes three assumptions: first, players have rational expectations. Second, they respond to expected payoffs, and so their behavior is influenced by expected utility considerations. Third, since payoffs are monetary, they are in some sense risk neutral.

Our first main result is that the only solution concepts satisfying narrow framing and expectation-monotonicity are Nash equilibrium, logit quantal response equilibrium, and some of their refinements (Theorem 1). In this way, we demonstrate a deep connection between the two concepts. This result also provides a rationality foundation for the logit form of quantal response equilibrium, which is widely used in the experimental literature, primarily due to its tractability.

In the second part of the paper, we replace expectation-monotonicity with *distribution-monotonicity*, which only requires choice probabilities to be monotone with respect to first-order stochastic dominance of the distribution of payoffs of each action. This axiom retains the rational expectation assumption, but relaxes the expected utility and risk neutrality aspects of expectation-monotonicity.

Our first result regarding distribution-monotonicity is that it implies expectation-monotonicity when coupled with narrow framing and an additional assumption we call strategic invariance (Theorem 2). Strategic invariance means that players play identically in games that are strategically equivalent. This result uncovers a connection between strategic invariance and risk neutrality. It also implies that under narrow framing, distribution-monotonicity and strategic invariance, players play either Nash or logit quantal response equilibrium.

Dropping strategic invariance, we study solution concepts that satisfy narrow framing and distribution-monotonicity. These axiom characterize what we call *statistic response equilibria*, a novel generalization of Nash and logit quantal response equilibrium that allows for various risk attitudes (Theorem 3). Statistic response equilibria allow players to respond to not only the expectation but any monotone additive statistic (Mu et al., 2021), which can capture risk aversion, risk seeking, and mixed risk attitudes.

The class of statistic response equilibria is large, and we characterize two useful parametric subclasses. First, we show that strengthening distribution-monotonicity to an axiom that captures rational expectations and expected utilities—but not risk neutrality—yields a one parameter family in which players respond to the CARA certainty equivalents of the actions (Theorem 4). Second, using a scale-invariance axiom, we characterize a three parameter family of statistic response equilibria which we expect to be useful in estimating empirical models of games (Theorem 5). In these equilibria agents logit best respond to a convex combination of the minimum, maximum and expectation of the payoff distribution from each action.

1.1 Related literature

We contribute to a large body of literature on the axiomatic approach in economic theory. This approach has been used extensively in cooperative game theory, bargaining, and mechanism design; see surveys by Moulin (1995), Roth (2012), and Thomson (2023). Our paper concerns non-cooperative game theory, where the axiomatic approach has been applied primarily toward choosing equilibrium refinements (Harsanyi and Selten, 1988; Norde, Potters, Reijnierse, and Vermeulen, 1996; Govindan and Wilson, 2009). This approach has also been used for uniting existing solution concepts under a generalized family (Goeree and Louis, 2021), justifying maximin strategies in zero-sum games (Brandl and Brandt, 2019) and Nash equilibrium in general normal form games (Brandl and Brandt, 2024).

Quantal response equilibrium (McKelvey and Palfrey, 1995) has been empirically successful at explaining the deviations from Nash equilibrium predictions across a wide range of experiments (Goeree, Holt, and Palfrey, 2016, 2020). The axiomatic approach has been applied to define non-parametric subclasses of QRE by imposing axioms on the quantal response functions (Goeree, Holt, and Palfrey, 2005; Friedman and Mauersberger, 2022). Our approach differs from these papers, since we do not take QRE as a starting point and axiomatize solution concepts rather than response functions. Furthermore, our results differ in that they pin down the one-parameter class of logit QRE in particular, providing a novel justification for a solution concept that has been widely used to analyze empirical data (see, e.g., Goeree, Holt, and Palfrey, 2016; Wright and Leyton-Brown, 2017; Goeree, Holt, and Palfrey, 2020).

The term narrow framing is often used in a broader sense than in our paper, and includes ignoring interactions between related choices. It has been extensively studied in the context of individual decisions, where much of the literature treats narrow framing as a behavioral bias (see, e.g., Read, Loewenstein, Rabin, Keren, and Laibson, 2000; Barberis, Huang, and Thaler, 2006). Our paper belongs to a recent literature offering a rational perspective on narrow framing. Kőszegi and Matějka (2020) justify narrow framing as rational behavior in a model with costly attention and Camara (2022) offers a computational complexity justification. Sandomirskiy and Tamuz (2023) use a version of our narrow framing assumption as an axiom for single-agent decisions.

Our second main result naturally gives rise to a novel solution concept, statistic response equilibria. In these equilibria, agents respond to a *monotone additive statistic*, which were characterized by Mu, Pomatto, Strack, and Tamuz (2021). These solutions can be interpreted as reflecting various risk attitudes toward the uncertainty induced by other players' mixed strategies; we emphasize that these emerge from our axioms rather than by assuming, a priori, that players behave according to any specified risk-preference. This is in contrast with the literature on equilibrium concepts that incorporates specified risk attitudes by transforming payoffs of each game according to some utility function that reflects such an attitude (Goeree, Holt, and Palfrey, 2003; Yekkehkhany, Murray, and Nagi, 2020). In fact, the risk attitudes borne out by our characterization cannot be recovered by transforming the payoffs of games according to any utility function. In this way, we contribute to the literature on games with non-expected utility preferences (Shalev, 2000; Metzger and Rieger, 2019).

2 Solution Concepts and Narrow Framing

We consider finite normal form games played between a fixed set $N = \{1, \dots, n\}$ of agents. A game $G = (A, u)$ between these agents is defined by its finite set of action profiles $A = \prod_i A_i$ and its payoff map $u \colon A \to \mathbb{R}^N$, where the i^{th} coordinate of $u(\cdot)$, denoted $u_i(\cdot)$, is the payoff function of $i \in N$. Given a mixed strategy profile p and an action a_i for player

i, we denote by

$$\mathbb{E}[u_i(a_i, p_{-i})] = \sum_{a_{-i}} p_{-i}(a_{-i})u_i(a_i, a_{-i})$$

the expected payoff of player i for taking action a_i.

A *solution concept* S assigns to each game G a nonempty set $S(G) \subset \prod_i \Delta A_i$ of mixed strategy profiles, or *solutions*.[1] We say that a solution concept S is a *refinement* of S' if $S(G) \subseteq S'(G)$ for all games G.

Our main axiom, narrow framing, asserts that players who are engaged in multiple, unrelated games may consider each game independently. Formally, for $G = (A, u), H = (B, v)$, we define the product game $G \otimes H = (C, w)$ by

$$C_i = A_i \times B_i \quad \text{and} \quad w_i\left((a_1, b_1), \ldots, (a_n, b_n)\right) = u_i(a) + v_i(b).$$

I.e., in $G \otimes H$ players play both games and earn the sum of their payoffs from the two games. We call G and H factor games of the product game $G \otimes H$. The payoff structure of the product game captures the fact that the factor games are unrelated. To make sense of the summation of payoffs in different games, we need to think of payoffs as being quoted in the same units across all games. For simplicity, we think of payoffs as monetary.

Given mixed strategy profiles p and q of the games G and H, we define the mixed strategy profile $p \times q$ for the game $G \otimes H$ by

$$[p \times q]_i(a_i, b_i) = p_i(a_i) \cdot q_i(b_i).$$

So, if players are playing $p \times q$ in $G \otimes H$, then they are independently choosing strategies in G from p and in H from q.

Definition 1. *A solution concept S satisfies* narrow framing *if $p \in S(G)$ and $q \in S(H)$ implies $p \times q \in S(G \otimes H)$.*

When a solution concept satisfies narrow framing, then solutions of G and H can be composed into a solution of the product game $G \otimes H$ by having players choose their actions independently in the two factor games. Viewing solution concepts as predictions, our narrow framing assumption requires that independently playing the solutions to the factor games is a valid prediction for the product game. This assumption does not rule out the existence of other predictions for the product game.

[1] There is a technical nuance that can be safely ignored without missing the gist of the paper: the collection of all finite set is not a set, and neither is the collection of all finite games. Hence, for a solution concept to be a well-defined correspondence, we assume that all actions available to any player in any game belong to a universal, non-empty set of actions $\mathcal{A}$. We also suppose that $\mathcal{A}$ is closed under pairing, so that $\mathcal{A} \times \mathcal{A} \subset \mathcal{A}$.

Narrow framing is satisfied by the Nash correspondence $\mathrm{Nash}(G)$ that assigns to a game G the set of all its mixed Nash equilibria. Note that not all mixed Nash equilibria in product games are products of equilibria in the factor games, but these products do appear in the solution of the product game, as required by narrow framing.

Many refinements of Nash also satisfy narrow framing. These include maximal-entropy Nash equilibria, trembling hand perfect equilibria, and welfare-maximizing equilibria. However, not all refinements are guaranteed to satisfy narrow framing, e.g., minimal-entropy Nash equilibria—which can be thought of as a natural extension of pure Nash equilibria to a non-empty correspondence—violate narrow framing as entropy can be reduced by correlating unrelated choices. Likewise, the Pareto optimal Nash equilibria do not satisfy narrow framing.[2]

Given $\lambda \geq 0$ and a game $G = (A, u)$, the logit quantal response equilibrium correspondence is given by

$$\mathrm{LQRE}_\lambda(G) = \{p \in \Delta(A) \mid p_i(a_i) \propto \exp(\lambda \mathbb{E}[u_i(a_i, p_{-i}))]), \forall (i, a_i)\}.$$ [3]

When the λ parameter is not specified, we write LQRE to refer to LQRE_λ for any $\lambda \geq 0$.

This correspondence also satisfies narrow framing. In fact, this correspondence satisfies a stronger property: every solution of a product game is a product of solutions of the factor games.

Other solution concepts that satisfy narrow framing include the rationalizable mixed strategies, welfare maximizing (or minimizing) mixed strategy profiles, level k models in which level 0 players choose uniformly, as well as cognitive hierarchy models with the same base choices. Probit QRE does not satisfy narrow framing, and more generally, neither does any QRE that is not logit.

We also consider anonymity, a simplifying assumption which does not affect the essence of our results. It captures a sense in which all the players are identical under a solution concept. Given a game $G = (A, u)$ and a permutation $\pi \colon N \to N$, define the permuted game $G_\pi = (B, v)$ by $B_i = A_{\pi(i)}$, and $v_i(a_{\pi(1)}, \ldots, a_{\pi(n)}) = u_{\pi(i)}(a_1, \ldots, a_n)$ for all $i \in N$, and $a \in A$.

Definition 2. *Say that S satisfies* anonymity *if for any permutation π and any game G with $p \in S(G)$, we have $p_\pi \in S(G_\pi)$, where $(p_\pi)_i = p_{\pi(i)}$ for all i.*

[2] A standard intuition applies: Pareto optimal allocation in sub-markets may not give rise to a Pareto optimal allocation in the market itself thanks to beneficial trades across sub-markets.

[3] We use $p_i(a_i) \propto \exp(\lambda \mathbb{E}[u_i(a_i, p_{-i}))]$ to indicate equality up to normalization of the probabilities, i.e.,

$$p_i(a_i) = \frac{\exp(\lambda \mathbb{E}[u_i(a_i, p_{-i})])}{\sum_{b \in A_i} \exp(\lambda \mathbb{E}[u_i(b, p_{-i})])}.$$

Beyond narrow framing and anonymity, we will also need mild rationality assumptions reflecting that players tend to choose actions with higher payoffs more often. These assumptions are formalized as monotonicity axioms discussed in the corresponding sections below.

3 Expectation-Monotonicity and Logit Quantal Response Equilibrium

In this section we introduce a monotonicity axiom and study it together with narrow framing.

Definition 3. *A solution concept S satisfies* expectation-monotonicity *if for every game G, $i \in N$, and $p \in S(G)$, if $\mathbb{E}[u_i(a_i, p_{-i})] > \mathbb{E}[u_i(b_i, p_{-i})]$, then $p_i(a_i) \geq p_i(b_i)$.*

This axiom includes a number of conceptual assumptions: first, it introduces a notion of rational expectations into a solution concept, in the sense that players anticipate the others' actions sufficiently well to rank expected payoffs. Second, it implies that players only care about the expectation of their payoffs, and hence is, in some sense, an expected utility axiom. And since we think of payoffs as monetary, this axiom furthermore implies that players are risk neutral. Finally, it captures a notion of rationality, because players prefer actions with higher expected payoffs.

Expectation-monotonicity is satisfied by Nash, LQRE$_\lambda$, probit QRE, and, more generally, any regular QRE (Goeree, Holt, and Palfrey, 2005), and M-equilibrium (Goeree and Louis, 2021). It is closed under refinements. The level k and cognitive hierarchy model solution concept do not satisfy it, and neither does rationalizability.

The next definition introduces a weakening of expectation-monotonicity.

Definition 4. *A solution concept S satisfies* approximate expectation-monotonicity *if there exist constants $M \geq 0$ and $\varepsilon > 0$ such that for every game G, $i \in N$, and $p \in S(G)$, if $\mathbb{E}[u_i(a_i, p_{-i})] > \mathbb{E}[u_i(b_i, p_{-i})] + M$, then $p_i(a_i) \geq \varepsilon p_i(b_i)$.*

This definition weakens expectation-monotonicity in two ways. First, choice probabilities do not have to be higher, but only at least ε as high, and for this to hold, expected payoffs do need to be higher by at least M. This property is satisfied by, for example, ε-Nash equilibrium (where no player can gain more than ε by deviating) and ε-proper equilibrium of Myerson (1978) (where any action is played at most ε times as often as a better one).

The next lemma shows that narrow framing strengthens approximate expectation-monotonicity to (exact) expectation monotonicity.

Lemma 1. *Every solution concept that satisfies narrow framing and approximate expectation-monotonicity also satisfies expectation-monotonicity.*

Proof. Let S satisfy narrow framing and approximate expectation-monotonicity. Let $G = (A, u)$, $a_i, b_i \in A_i$ and suppose, for the sake of contradiction, that there is $p \in S(G)$ with $\mathbb{E}[u_i(a_i, p_{-i})] > \mathbb{E}[u_i(b_i, p_{-i})]$ and $\frac{p_i(a_i)}{p_i(b_i)} < 1$. Then there is an $m \in \mathbb{N}$ such that

$$m\left(\mathbb{E}[u_i(a_i, p_{-i})] - \mathbb{E}[u_i(b_i, p_{-i})]\right) > M$$

and $\left(\frac{p_i(a_i)}{p_i(b_i)}\right)^m < \varepsilon$. Let p^m denote the m-fold product $p \times \cdots \times p$, and G^m denote the m-fold product $G \otimes \cdots \otimes G$. Define v as the utility function of G^m, i.e., $G^m = (A^m, v)$. Then by narrow framing, we have $p^m \in S(G^m)$ with $\mathbb{E}[v_i((a_i, \cdots, a_i), p^m_{-i})] > M + \mathbb{E}[v_i((b_i, \cdots, b_i), p^m_{-i})]$ and $\frac{p^m_i(a_i, \ldots, a_i)}{p^m_i(b_i, \ldots, b_i)} < \varepsilon$, violating approximate expectation-monotonicity. $\square$

This lemma highlights that narrow framing, far from being a behavioral assumption, is conducive to rationality.

The three properties of narrow framing, anonymity and expectation-monotonicity do not seem too restrictive, as each of them is satisfied by many well-known solution concepts. Among the examples mentioned above, Nash and LQRE_λ satisfy all three, as does trembling hand Nash, and maximum entropy Nash. The next theorem shows that any solution concept satisfying all of the three properties must return only Nash equilibria or only logit quantal response equilibria.

Theorem 1. *If S satisfies expectation-monotonicity, narrow framing, and anonymity, then S is either a refinement of* Nash *or of* LQRE_λ *for some $\lambda \geq 0$.*

Theorem 1 gives yet another piece of evidence for the importance of Nash equilibria (Brandl and Brandt, 2024). It also provides a novel justification for the particular logit form of QRE, beyond its tractability. This is a simple and important one-parameter family that has been useful for predicting outcomes of games in the lab (Goeree et al., 2016; Wright and Leyton-Brown, 2017; Goeree et al., 2020).

Furthermore, this result establishes a connection between Nash and LQRE_λ. Notice that Nash is not merely a limiting case of LQRE_λ as $\lambda \to \infty$; McKelvey and Palfrey (1995) show that limit points of LQRE are Nash equilibria, but not every Nash equilibrium can be obtained as a limit point of logit equilibria. Indeed, there is an interesting distinction between logit quantal response equilibria (and their limit points) and Nash equilibria. While all logit quantal response equilibria of a product game are products of equilibria of its factor games, there exist Nash equilibria of a product game that do not satisfy this property. In fact, Nash equilibria exhibit a very rich correlation structure: any strategy profile of a product game that marginalizes to Nash equilibria of its factor games constitutes a Nash equilibrium. Theorem 1 is proved in Appendix A.

Anonymity is not a crucial assumption for the above theorem, but it ensures that all players behave alike. Without anonymity, we get a version of LQRE with agent specific λ_i instead of common λ as well as chimeral rules where some agents use logit best response and some best-respond as in Nash equilibrium.[4] As a corollary, we do not require anonymity to characterize refinements of an anonymous solution concept: Nash equilibrium.

Corollary 1. *If S satisfies expectation-monotonicity and narrow framing, and players never play strictly dominated strategies, then S is a refinement of* Nash*.*

The assumption on strictly dominated strategies serves to distinguish Nash from LQRE. In fact, under the assumptions of expectation-monotonicity and narrow framing, *any* feature of Nash equilibrium that does not apply to LQRE will lead to a characterization of only Nash equilibrium and vice versa. We can thus weaken the assumption on strictly dominated strategies by supposing that every player plays some strategy with zero probability in some particular game.

It is worth comparing this characterization of Nash equilibrium with that of Brandl and Brandt (2024) who characterize Nash equilibrium as the unique solution concept that satisfies consequentialism—duplicates of an action are treated as the same action—, consistency—if a mixed-strategy profile is a solution to two games with the same action sets, it is a solution to any convex combination of the games—, and rationality—dominant actions are played with positive probability. Their axioms rule out all strict refinements of Nash equilibrium, whereas our axioms allow for some refinements of Nash, such as trembling hand perfect equilibrium and max welfare Nash. Recall that not all refinements of Nash satisfy narrow framing, which provides a justification for selecting some refinements of Nash over others.

Our narrow framing axiom and the consequentialism and consistency axioms of Brandl and Brandt (2024) impose coherency restrictions on solutions for related games. Interestingly, consequentialism and consistency imply a property that is a weakening of narrow framing: for any games G and H *there exist* $p \in S(G), q \in S(H)$ such that $p \times q \in S(G \otimes H)$. Narrow framing proper is not implied, but is implied if we add their rationality assumption. Our expectation-monotonicity axiom, which assumes that players' mixing probabilities are constrained by the true mixing probabilities of the other players, implies Brandl and Brandt (2024)'s rationality axiom, which makes no such "rational-expectations" assumption.[5]

Note that expectation-monotonicity only requires that actions yielding strictly higher payoffs are played weakly more often. A slight strengthening of expectation-monotonicity

[4] McKelvey et al. (2000) extended the QRE framework to allow for λ-heterogeneity.

[5] While our monotonicity assumption is constrained by rational expectations, it is only an ordinal restriction which is consistent with inaccurate beliefs, provided they do not change the ranking. See Goeree and Louis (2021).

would require that actions yielding weakly higher utility are played weakly more often. As the following corollary shows, this strengthening of expectation-monotonicity rules out Nash and its refinements.

Corollary 2. *Suppose S satisfies narrow framing and anonymity. Then S is a refinement of* LQRE_λ *for some $\lambda \geq 0$ if either of the following conditions apply:*

1. *For every $G, p \in S(G)$, and $i \in N$, if $\mathbb{E}[u_i(a_i, p_{-i})] \geq \mathbb{E}[u_i(b_i, p_{-i})]$, then $p_i(a_i) \geq p_i(b_i)$.*
2. *S satisfies expectation-monotonicity and players always play totally mixed strategies.*

4 Distribution-Monotonicity and Statistic Response Equilibria

In this section we relax the assumption that S satisfies expectation-monotonicity, which supposes that players' decisions are constrained by the means of the payoff distributions induced by each action, conditional on the opponents' strategies. Instead, we assume that players' mixing probabilities are monotone with respect to first-order stochastic dominance of the conditional payoff distributions induced by each action.[6]

We write $u_i(a_i, p_{-i})$ to denote the lottery i faces when playing action a_i given that the other players play mixed-strategies according to p. Thus, $u_i(a_i, p_{-i}) >_{\mathrm{FOSD}} u_i(b_i, p_{-i})$ denotes that the distribution of $u_i(a_i, p_{-i})$ strictly first order stochastically dominates that of $u_i(b_i, p_{-i})$, which is the lottery i faces when playing b_i. Our next definition, distribution-monotonicity, is another weak concept of rationality in terms of stochastic dominance.

Definition 5. *A solution concept S satisfies* distribution-monotonicity *if for every game $G, p \in S(G)$, and $i \in N$, if $u_i(a_i, p_{-i}) >_{FOSD} u_i(b_i, p_{-i})$, then $p_i(a_i) \geq p_i(b_i)$.*

Distribution-monotonicity is implied by expectation-monotonicity, but is much weaker: it relaxes both the expected utility and the risk neutrality components of expectation-monotonicity, keeping only rational expectations and monotonicity. Unlike expectation-monotonicity, distribution-monotonicity is invariant to monotone transformations of payoffs. For example, Nash equilibrium with monotone-reparameterized payoffs (Weinstein, 2016) and risk-adjusted QRE under CRRA reparameterized payoffs (Goeree et al., 2003) satisfy distribution-monotonicity. Distribution-monotonicity is also satisfied by $S(K)$ equilibria (Osborne and Rubinstein, 1998), where players respond to independent draws from the payoff distribution induced by each action.

[6]When there is only one player, distribution-monotonicity is equivalent to expectation-monotonicity. Henceforth, we suppose there are at least two players.

Next, we explore the interaction between narrow framing, distribution-monotonicity and an additional axiom that we call *strategic invariance*. This axiom restricts behavior across *strategically equivalent* games.

Definition 6. *We say that games* (A, v) *and* (A, u) *are* strategically equivalent *if for each player* i *there exists a function* $w_i \colon A_{-i} \to \mathbb{R}$ *such that* $v_i(a) = u_i(a) + w_i(a_{-i})$.

That is, for all $a_i, b_i \in A_i$ and $a_{-i} \in A_{-i}$, $v_i(a_i, a_{-i}) - v_i(b_i, a_{-i}) = u_i(a_i, a_{-i}) - u_i(b_i, a_{-i})$, i.e., agent i's marginal payoff of switching from one action to the other is the same in the two games.

Definition 7. *A solution concept* S *satisfies* strategic invariance *if* $S(A, u) = S(A, v)$ *whenever* (A, u) *and* (A, v) *are strategically equivalent.*

Strategic equivalence is respected by Nash and LQRE, as well as many other concepts that do not have a rational expectations component, such as rationalizability and k-level reasoning.

Our next theorem shows that strategic invariance is a powerful assumption, when coupled with narrow framing.

Theorem 2. *Suppose* S *satisfies narrow framing, strategic invariance, distribution-monotonicity, and anonymity. Then it satisfies expectation-monotonicity.*

Recall that distribution-monotonicity is a pure rational expectations and monotonicity axiom, and does not have an expected utility or risk neutrality component. However, both expected utility and risk neutrality are implied by expectation-monotonicity. Theorem 2 thus shows that strategic invariance is a potent assumption that highly constrains behavior to resemble risk-neutrality.

Strategic invariance clearly implies that players display no wealth effects, since adding a constant to all payoffs does not change their behavior. To gain some intuition for why strategic invariance furthermore rules out any non-trivial risk attitudes, consider the following example of a two player game. Player 2 has two actions, a_2 and b_2, gets payoff 0 regardless of the action profile, and mixes evenly between the two actions. Player 1 has two actions, a_1 and b_1, and gets the following payoffs:

	a_2	b_2
a_1	0	2
b_1	1	1

Table 1: Player 1's utility.

In this game, both actions yield the same expected utility, but action a_1 has variance 1, whereas action b_1 has variance 0, and so would be preferred by any risk averse player. Consider now the following, strategically equivalent game:

	a_2	b_2
a_1	0	1
b_1	1	0

Table 2: Player 1's utility in a strategically equivalent game.

Here, both actions yield the same distribution of payoffs to player 1, and hence risk attitudes should not influence the choice between a_1 and b_1. Since this game is strategically equivalent to the previous, we conclude that under strategic invariance players would be indifferent between the two actions in the previous game, and so are effectively risk neutral. The proof of Theorem 2 requires an additional argument, since we cannot directly reason about players' preferences and risk attitudes, but only about their strategies in games. See Appendix B for the proof.

A consequence of Theorem 2 and Theorem 1 is that if S satisfies anonymity, narrow framing, strategic invariance and distribution-monotonicity then it is a refinement of either Nash or LQRE. This yields a motivation for these solution concepts that does not directly assume risk neutrality, highlighting the strength of the strategic invariance assumption.

We next drop the strategic invariance assumption, and explore the joint consequences of narrow framing and distribution-monotonicity, without any assumptions that lead to risk neutrality or indeed expected utility. Our next theorem shows that the only solution concepts which satisfy anonymity, narrow framing and distribution-monotonicity are *statistic response equilibria*, where players respond to a statistic of each payoff distribution. This class of equilibria will generalize Nash and LQRE, in which players evaluate actions by the expectation of the payoff distribution, to equilibria in which players evaluate distributions by statistics that may be different than the expectation and represent various risk attitudes.

We use the term *statistic* to refer to a function Φ that assigns a real number to every lottery with finitely many outcomes, and such that $\Phi(c) = c$ for the deterministic lottery yielding the constant amount c. Here, a lottery is simply a distribution over monetary payoffs. Lotteries will arise in our setting as the payoffs a player anticipates when choosing an action, given the mixed strategies of the other players. We denote lotteries by X, Y, and when they are independent denote by $X + Y$ the lottery whose distribution is the convolution of the distributions of X and Y.

Below, we define a class of statistics that are monotone with respect to first order stochastic dominance and additive for independent lotteries (Mu et al., 2021). We show that players must respond to this class of statistics.

Definition 8. *A* monotone additive statistic *is a statistic that is additive for independent lotteries and monotone with respect to first order stochastic dominance (FOSD).*

For a lottery X, let $K_a(X) = \frac{1}{a}\log \mathbb{E}\left[e^{aX}\right]$, and let $K_{-\infty}(X), K_0(X)$, and $K_\infty(X)$ denote, respectively, the minimum, expectation, and maximum of the distribution of X. Mu et al. (2021) show that all monotone additive statistics are of the form $\Phi(X) = \int_{\overline{\mathbb{R}}} K_a(X)\,\mathrm{d}\mu(a)$ for some Borel probability measure μ on $\overline{\mathbb{R}}$. Under expected utility, $K_a(X)$ is the certainty equivalent of the lottery X for an individual whose preference over monetary payoffs is represented by a utility function with a constant coefficient $(-a)$ of absolute risk aversion. Hence, $\Phi(X) = \int_{\overline{\mathbb{R}}} K_a(X)\,\mathrm{d}\mu(a)$ is a weighted average of CARA certainty equivalents across risk coefficients, which may reflect both risk-averse and risk-loving preferences as μ can place mass on both negative and positive risk coefficients.[7]

In Nash equilibrium and logit quantal response equilibrium players best respond or "better respond" to the expectation of each action's payoff distribution. In statistic response equilibrium, players respond to a monotone additive statistic of each distribution. Below we define the two classes of statistic response equilibria.

Definition 9. *Given a monotone additive statistic Φ, p is a* Nash$_\Phi$ equilibrium *of $G = (A, u)$ if for all $i \in N$, and $a_i \in A_i$*

$$\operatorname{supp} p_i \subseteq \operatorname*{argmax}_a \Phi(u_i(a, p_{-i})).$$

In a Nash$_\Phi$ equilibrium, players best respond to the other players according to Φ by randomizing over actions whose payoff distributions maximize Φ. Since Φ is a monotone additive statistic, players never play an action that is first-order stochastically dominated. Thus any solution concept that only returns Nash$_\Phi$ equilibria will satisfy distribution monotonicity.

The next definition introduces a class of statistic response equilibria in which players "better respond" to a monotone additive statistic of each distribution.

Definition 10. *Given a monotone additive statistic Φ, p is a* LQRE$_{\lambda\Phi}$ *of $G = (A, u)$ if $i \in N$, and $a_i \in A_i$*

$$p_i(a_i) \propto \exp(\lambda\Phi(u_i(a_i, p_{-i}))).$$

[7] Mu et al. (2021) provide a characterization of monotone additive statistics on the domain of all compactly supported (rather than finitely supported) lotteries. Their characterization also applies to the domain of finitely supported lotteries, since any monotone additive statistic on this restricted domain can be extended to the compactly supported lotteries. For the proof, see Appendix B.2.

Nash_Φ equilibria generalize Nash equilibria, in which players respond to the expectation, i.e., a $\mathrm{Nash}_\mathbb{E}$ equilibrium. Likewise, $\mathrm{LQRE}_{\lambda\Phi}$ equilibria generalize logit quantal response equilibria where players respond to the expectation, i.e., a $\mathrm{LQRE}_{\lambda\mathbb{E}}$.

While every game has a Nash equilibrium, the existence of a Nash_Φ equilibrium is not guaranteed for all Φ when there is more than one player.[8] For example, Nash_Φ equilibria may not exist when Φ is the minimum or maximum of a distribution. This issue does not arise for $\mathrm{LQRE}_{\lambda\Phi}$, which do exist for every game. As the next result shows, the existence of Nash_Φ equilibria is guaranteed for a large family of monotone additive statistics, namely those in which the maximum and minimum do not play a role.

Proposition 1. *There is a* Nash_Φ *equilibrium for every game if and only if* $\Phi(X) = \int_\mathbb{R} K_a(X)\,\mathrm{d}\mu(a)$. *Moreover, every game has a* $\mathrm{LQRE}_{\lambda\Phi}$ *equilibrium for every* $\lambda \geq 0$ *and monotone additive statistic* Φ.

Equivalently, for $\Phi = \int K_a\,\mathrm{d}\mu(a)$, Nash_Φ equilibria exist for every game if and only if μ places zero mass on $-\infty$ and $+\infty$, i.e., on the minimum and the maximum, while $\mathrm{LQRE}_{\lambda\Phi}$ equilibria exist for any μ on $\overline{\mathbb{R}}$. The existence of an $\mathrm{LQRE}_{\lambda\Phi}$ equilibrium is guaranteed as a fixed-point of the logit response function by Brouwer's fixed-point theorem. The proof that Nash_Φ equilibria exist when μ places no mass on the minimum or maximum does not use Kakutani's fixed point theorem, because the best-response correspondence is not convex. Rather, we apply a limiting argument taking λ to infinity to obtain a Nash_Φ equilibrium as a limit point of $\mathrm{LQRE}_{\lambda\Phi}$ equilibria. This argument does not work when μ places any positive mass on the minimum or maximum since Φ may be discontinuous at the limit point.

To prove that there may not be a Nash_Φ equilibrium when there is positive mass on the minimum or maximum, we construct an example of such games, using a variant of matching pennies. Proposition 1 is proved in Appendix C.

We call Nash_Φ and $\mathrm{LQRE}_{\lambda\Phi}$ *statistic response equilibria* as players best or better respond to the statistic Φ of distributions induced by each available action. Formally, a statistic response equilibrium (SRE) is a solution concept that returns all Nash_Φ or all $\mathrm{LQRE}_{\lambda\Phi}$ equilibria for some Φ and λ.[9]

It is easy to verify that the SRE solution concepts satisfy our axioms. Narrow framing follows from the additivity of Φ, distribution monotonicity is a consequence of the monotonicity of Φ, and anonymity holds since all players use the same Φ; as above, it can be removed, with the conclusion appropriately altered to allow different players to best

[8] When there is one player, every payoff is deterministic, so Nash_Φ equilibria coincide with Nash equilibria for all Φ.

[9] By Proposition 1, Nash_Φ is a solution concept only for Φ that puts no mass on the maximum or minimum.

or better respond to different statistics. The next result shows that these axioms in fact characterize SRE.

Theorem 3. *Suppose S satisfies distribution-monotonicity, narrow framing, and anonymity. Then S is a refinement of some SRE.*

Theorem 3 is proven in Appendix D. As shown in the proof, Φ assigns a certainty equivalent to each payoff distribution, which players respond to. As a monotone additive statistic, Φ is a weighted average of CARA certainty equivalents across different values of coefficients a. Hence statistic response equilibria incorporate flexible risk attitudes which allow for risk-averse, risk-loving, or mixed risk attitudes.[10]

While distribution-monotonicity ensures that players' mixing probabilities are monotone with respect to first order stochastic dominance, it does not provide a way to compare any pair of distributions. Interestingly, Theorem 3 shows that with narrow framing and anonymity there is a total order, defined by the statistic Φ, that dictates how players rank *every* payoff distribution.

5 Parametric Families of Statistic Response Equilibria

In statistic response equilibria, players respond to a statistic Φ which is parameterized by a Borel probability measure μ on $\bar{\mathbb{R}}$, an infinite-dimensional parameter. In this section, we introduce additional axioms that restrict players' behavior across games and lead to simpler parametric families of SRE.

A special class of Nash_Φ equilibria are the Nash_{K_a} equilibria for $a \in \mathbb{R}$. In these equilibria, players randomize over the actions whose payoff distributions have maximal CARA certainty equivalents under risk-coefficient $-a$. Equivalently, players randomize over the actions whose payoff distributions maximize CARA expected utility. These equilibria thus coincide with the Nash equilibria of a game whose payoff function is reparameterized according to the CARA utility function

$$c_{-a}(x) = \begin{cases} \frac{e^{ax}-1}{a} & a \neq 0 \\ x & a = 0. \end{cases}$$

We characterize Nash_{K_a} and $\text{LQRE}_{\lambda K_a}$ as following from narrow framing, anonimity, and an assumption that is stronger than distribution-monotonicity but weaker than expectation monotonicity: it retains the rational expectation aspect of distribution-monotonicity

[10]The statistic $\Phi(X) = \int_{\bar{\mathbb{R}}} K_a(X)\,\mathrm{d}\mu(a)$ defines a risk attitude: if μ places mass only on negative values of a, $\Phi(X) \leq \mathbb{E}[X]$ for any lottery X, i.e., Φ reflects risk-aversion. Conversely, if μ places mass only on positive values of a, then $\Phi(X) \geq \mathbb{E}[X]$ and Φ reflects a risk-loving attitude. If μ places mass on both negative and positive values of a, then Φ reflects a mixed risk attitude, i.e. there are lotteries X and Y with $\Phi(X) < \mathbb{E}[X]$ and $\Phi(Y) > \mathbb{E}[Y]$. See Proposition 5 of Mu et al. (2021).

and adds to it an expected utility aspect, but without the risk neutrality assumption of expectation-monotonicity. Specifically, we assume that there exists a strictly increasing function $f\colon \mathbb{R} \to \mathbb{R}$ such that choice probabilities are monotone with respect to the expectation of $f \circ u_i$. This function f plays the role of a utility for monetary payoffs, and was taken to be the identity in the expectation-monotonicity axiom.

Theorem 4. *Suppose that S satisfies narrow framing, anonymity, and that there exists a strictly increasing function $f\colon \mathbb{R} \to \mathbb{R}$ such that for every game $G, p \in S(G)$, and $i \in N$, if $\mathbb{E}[(f \circ u_i)(a_i, p_{-i})] > \mathbb{E}[(f \circ u_i)(b_i, p_{-i})]$, then $p_i(a_i) \geq p_i(b_i)$. Then S is a refinement of either Nash_{K_a} or $\mathrm{LQRE}_{\lambda K_a}$ for some $a \in \mathbb{R}$ and $\lambda \geq 0$.*

It follows from this theorem that Nash_{K_a} are the only Nash_{Φ} equilibria that coincide with Nash equilibria under a reparameterization of payoffs. Note that under $\mathrm{LQRE}_{\lambda K_a}$ best responses are given by

$$p_i(a_i) \propto \mathbb{E}[\exp(a \cdot u_i(a_i, p_{-i}))]^{\lambda/a}.$$

This is different than logit responding to the transformed payoffs, as in Goeree et al. (2003). Indeed, there is no Φ other than the expectation such that $\mathrm{LQRE}_{\lambda\Phi}$ equilibria coincide with LQRE_{λ} equilibria under some reparameterization of payoffs.

The risk attitudes borne out by the characterization in Theorem 4 are pinned down by a single parameter a and correspond to coefficients of absolute risk aversion. They thus exhibit either risk-averse, risk neutral, or risk-loving attitudes. Theorem 4 is proved in Appendix E.

Our next result characterizes another simple parametric family of statistic response equilibria in which players respond to a convex combination of the minimum, maximum, and expectation of each payoff distribution, exhibiting a complex attitude toward risk. To recover this family we consider the following assumption on behavior across games.

Definition 11. *S satisfies* scale-invariance *if whenever p_i is the uniform distribution on A_i for all i, $p \in S(A, u)$ implies $p \in S(A, \alpha \cdot u)$ for all $\alpha \in (0, 1)$.*

This assumption is weak, along two dimensions. First, note that it only implies invariance when the player is completely indifferent between all the available actions.[11] Second, we only require it for scales less than unity. Intuitively, it seems plausible that if a player is indifferent between all actions, then they would still be indifferent when the stakes are made lower.

To keep the statement of the next theorem simple, we rule out Nash equilibria by assuming that players play completely mixed strategies; due to Proposition 1 removing this assumption would only add back the Nash solution concept and some of its refinements.

[11] A longer but more cumbersome name such as indifference-scale-invariance might be more appropriate.

Theorem 5. *Suppose S satisfies distribution-monotonicity, narrow framing, scale-invariance, anonymity, and players always play totally mixed strategies. Then there is $\lambda \in \mathbb{R}^3_{\geq 0}$ such that for all $G = (A, u), p \in S(G), i \in N$ and $a_i \in A_i$,*

$$p_i(a_i) \propto \exp\left(\lambda_1 \min_{a_{-i}} u_i(a_i, a_{-i}) + \lambda_2 \mathbb{E}[u_i(a_i, p_{-i})] + \lambda_3 \max_{a_{-i}} u_i(a_i, a_{-i})\right).$$

Perhaps surprisingly, scale-invariance kills all risk attitudes except for extreme risk-aversion, extreme risk-seeking, and risk neutrality (and their convex combinations). From a behavioral point of view, taking the minimum and maximum into account is consistent with these values being more salient than intermediate values. This family is simple enough to be represented by only three parameters, but is rich enough to capture mixed risk attitudes, allowing them to be risk averse in one game but risk seeking in another.

6 Conclusion

In this paper we studied solution concepts of games under narrow framing. Our solution concepts were constrained to be only mixed strategy profiles. A natural next step would be to look at joint distributions and study correlated equilibria and their generalizations.

Our results show that Nash equilibria and logit QRE (and their generalization) arise under narrow framing, as do some of their refinements. It would be interesting to understand which refinements satisfy narrow framing. In particular, it might be that logit QRE only has very few such refinements.

Along another dimension, one could restrict the set of games to interesting subclasses such as symmetric games or zero-sum games, and study narrow framing there.

A Proof of Theorem 1

We begin by introducing a number of definitions and establishing a key lemma. Say that S satisfies *expectation-neutrality* if for any $G = (A, u)$ and $p \in S(G)$, if $\mathbb{E}[u_i(a, p_{-i})] = \mathbb{E}[u_i(b, p_{-i})]$, then $p_i(a) = p_i(b)$.

The next axiom is commonly assumed in the QRE literature. It requires that players play every action with positive probability in all games.

Definition 12. *A solution concept S satisfies* interiority *if for every $G = (A, u)$, $p \in S(G)$, and $i \in N$, $p_i(a_i) > 0$ for each $a_i \in A_i$.*

The following lemma is an analog of Theorem 2 of Sandomirskiy and Tamuz (2023) for finite normal form games.

Lemma 2. *Suppose a solution concept S satisfies expectation-neutrality, narrow framing, interiority, expectation-monotonicity, and anonymity. Then for all $i, G = (A, u), p \in S(G), a_i \in A_i$,*

$$p_i(a_i) \propto \exp(\lambda \mathbb{E}[u_i(a_i, p_{-i})])$$

for some $\lambda \geq 0$.

Proof of Lemma 2. Let $i \in N$ and consider for each $x \in \mathbb{R}$ the game $G_x = (A, u)$, where $A_i = \{a_0, a_x\}$, $u_i(a_0, a_{-i}) = 0, u_i(a_x, a_{-i}) = x$ and all the other players have a single action and receive zero utility for all action profiles. Let $p_x \in S(G_x)$ and define $f(x) = \ln \frac{p_{x_i}(a_x)}{1-p_{x_i}(a_x)}$, which is well-defined by interiority. Let $x, y \in \mathbb{R}, p_x \in S(G_x), p_y \in S(G_y)$, and $p_{x+y} \in S(G_{x+y})$. By narrow framing, $p_x \times p_y \times p_{x+y} \in S(G_x \otimes G_y \otimes G_{x+y})$. By expectation-neutrality, $p_{x_i}(a_x)p_{y_i}(a_y)(1 - p_{x+y_i}(a_{x+y})) = (1 - p_{x_i}(a_x))(1 - p_{y_i}(a_y))p_{x+y_i}(a_{x+y})$. Rearranging and taking logs, we get $f(x + y) = f(x) + f(y)$, Cauchy's functional equation. Since S satisfies narrow framing and expectation-monotonicity, $p_{x_i}(a_x)$ is increasing in x, as $p_{x_i}(a_x)p_{x'_i}(a_0) \geq p_{x'_i}(a'_x)p_{x_i}(a_0)$ for $x > x'$. Thus f is nondecreasing, and $f(x) = \lambda x$, for some $\lambda \geq 0$.

Fix any $G = (B, v), q \in S(G)$, with $b, c \in B_i$. Let $x = \mathbb{E}[v_i(b, q_{-i})]$ and $y = \mathbb{E}[v_i(c, q_{-i})]$. By narrow framing, $r := q \times p_x \times p_y \in S(G \otimes G_x \otimes G_y)$. Let w denote the utility map for $G \otimes G_x \otimes G_y$. Note that $\mathbb{E}[w_i((b, a_0, a_y), r_{-i})] = \mathbb{E}[w_i((c, a_x, a_0), r_{-i})]$. By expectation-neutrality,

$$q_i(b)(1 - p_{x_i}(a_x))p_{y_i}(a_y) = q_i(c)p_{x_i}(a_x)(1 - p_{y_i}(a_y)).$$

Rearranging, we have

$$\frac{q_i(b)}{q_i(c)} = \frac{\exp f(x)}{\exp f(y)}.$$

Since c was arbitrary,

$$q_i(b) \propto \exp(f(x)) = \exp(\lambda x) = \exp(\lambda \mathbb{E}[v_i(b, q_{-i})]),$$

for some $\lambda \geq 0$. By anonymity, this holds for all $i \in N$. $\square$

Remark 6. *If the utilities in $v_i(\cdot, q_{-i})$ are deterministic, $w_i((b, a_0, a_y), r_{-i}) = w_i((c, a_x, a_0), r_{-i})$ as distributions. Thus, even under distribution-neutrality[12] we have the result $q_i(b) \propto \exp(\lambda \mathbb{E}[v_i(b, q_{-i})])$.*

We are now ready to prove Theorem 1.

[12]Distribution-neutrality is an analog of expectation-neutrality: if for any $G = (A, u)$ and $p \in S(G)$, if $u_i(a, p_{-i}) = u_i(b, p_{-i})$, then $p_i(a) = p_i(b)$.

Proof of Theorem 1. Let $j \in N$ and consider the game G where $A_j = \{0, 1\}, A_i = \{0\}$ for $i \neq j$ and each player's utility is simply their action. Let $p \in S(G)$. We first show that if $p_j(0) = 0$, then S is a refinement of Nash.

Suppose, for the sake of contradiction, that there is a player i and a game $H = (B, v)$ with $q \in S(H)$ and $a, b \in B_i$ such that $\mathbb{E}[v_i(a, q_{-i})] < \mathbb{E}[v_i(b, q_{-i})]$, while $q_i(a) > 0$. By anonymity, there is a permutation π with $\pi(i) = j$ and $p_\pi \in S(G_\pi)$, where $p_{\pi_i}(0) = p_j(0) = 0$. Let (C, w) denote $H^n \otimes G_\pi$, and let $n > \frac{1}{\mathbb{E}[v_i(b,q_{-i})] - \mathbb{E}[v_i(a,q_{-i})]}$. By narrow framing, $q^n \times p_\pi \in S(C, w)$. However, $\mathbb{E}[w_i(a^n, 1, [q^n \times p_\pi]_{-i})] < \mathbb{E}[w_i(b^n, 0, [q^n \times p_\pi]_{-i})]$, while

$$[q^n \times p_\pi]_i(a^n, 1) = (q_i(a))^n > 0 = [q^n \times p_\pi]_i(b^n, 0),$$

violating expectation-monotonicity.

For the remainder of this proof, suppose that $p_j(0) > 0$. We will show S is a refinement of LQRE_λ for some $\lambda \geq 0$. We show that S satisfies interiority. By anonymity, it is without loss of generality to suppose, toward a contradiction, that there is a game $H = (B, v)$ with $q \in S(H)$ and $a, b \in B_j$ such that $q_j(a) = 0 < q_j(b)$. Let $n > \mathbb{E}[v_j(b, q_{-j})] - \mathbb{E}[v_j(a, q_{-j})]$, and, as previously, consider that $q \times p^n \in S(H \otimes G^n)$. However,

$$\mathbb{E}[u_j(1, \dots, 1, a; [q \times p^n]_{-j})] = n + \mathbb{E}[v_j(a, q_{-j})] > \mathbb{E}[v_j(b, q_{-j})] = \mathbb{E}[u_j(0, \dots, 0, b; [q \times p^n]_{-j})],$$

while

$$[q \times p^n]_j(1, \dots, 1, a) = (p_j(1))^n q_j(a) = 0 < [q \times p^n]_j(0, \dots, 0, b) = (p_j(0))^n q_j(b),$$

violating expectation-monotonicity.

We now show that S also satisfies expectation-neutrality. By anonymity, it is without loss of generality to suppose, toward a contradiction, that there is a game $H = (B, v)$ with $q \in S(H)$ and $a, b \in B_j$ such that $\mathbb{E}[v_j(a, q_{-j})] = \mathbb{E}[v_j(b, q_{-j})]$, while $q_j(a) < q_j(b)$. By interiority we may let n such that $\left(\frac{q_j(b)}{q_j(a)}\right)^n > \frac{p_j(1)}{p_j(0)}$. By narrow framing, $q^n \times p \in S(H^n \otimes G)$. However, $\mathbb{E}[u_i(1, a, \dots, a; [q^n \times p]_{-i})] > \mathbb{E}[u_i(0, b, \dots, b; [q^n \times p]_{-i})]$, while

$$[q^n \times p]_i(1, a, \dots, a) = (q_j(a))^n p_j(1) < (q_j(b))^n p_j(0) = [q^n \times p]_i(0, b, \dots, b),$$

violating expectation-monotonicity.

Since S satisfies the hypotheses of Lemma 2, there is a $\lambda \geq 0$, such that for any $i \in N, G = (A, u)$, and $p \in S(G)$,

$$p_i(a) \propto \exp(\lambda \mathbb{E}[u_i(a, p_{-i})]).$$

$\square$

B Proof of Theorem 2

The following lemma shows that there are two possible implications of narrow framing, distribution-monotonicity, and anonymity, which can be isolated by considering whether or not players play strictly dominated strategies with positive probability.

One implication we consider in the next lemma is *distribution-neutrality*, i.e., if for any $G = (A, u)$ and $p \in S(G)$, if $u_i(a, p_{-i}) = u_i(b, p_{-i})$, then $p_i(a) = p_i(b)$.

Lemma 3. *Suppose S satisfies distribution-monotonicity, narrow framing, and anonymity. Then, either players never play FOSD-dominated actions, or S satisfies interiority and distribution-neutrality.*

Proof of Lemma 3. Let $i \in N$ and consider all pairs $(h, \ell) \in \mathbb{R}^2$ with $h > \ell$ and all games $G_{h\ell} = (A, u)$ where $A_i = \{a_h, a_\ell\}, u_i(a_h, \cdot) = h$ always, and $u_i(a_\ell, \cdot) = \ell$ always. There are two possibilities to consider:

- for all h, ℓ and any $p \in S(G_{h\ell})$, $p_i(a_\ell) = 0$, or
- there exist $h > \ell$ and $p \in S(G_{h\ell})$, with $p_i(a_\ell) > 0$.

First, we consider the former case where players never play the strictly dominated action a_ℓ. We will show that for all $i \in N, G = (A, u)$ with $p \in S(G)$, if $u_i(a_i, p_{-i}) >_{\text{FOSD}} u_i(b_i, p_{-i})$ then $p_i(b_i) = 0$, i.e., players never play an action that induces a FOSD-dominated lottery. To prove this, we introduce a definition and two lemmata.

Definition 13. *Let X and Y be compactly supported lotteries. Say X dominates Y in large numbers, denoted $X \geq_L Y$, if there exists $M \in \mathbb{N}$ such that for all $m \geq M$, $X^m >_{\text{FOSD}} Y^m$, where X^m and Y^m respectively refer to sums of m independent copies of X and Y.*

The following lemma is due to Aubrun and Nechita (2007):

Lemma 4. *Let X and Y be compactly supported lotteries with $K_a(X) > K_a(Y)$ for all $a \in \overline{\mathbb{R}}$. Then $X \geq_L Y$.*

The next lemma provides a condition under which adding independent lotteries to FOSD ranked lotteries preserves the dominance ranking.

Lemma 5. *Let X, Y, A, B be compactly supported lotteries with $X >_{\text{FOSD}} Y, \max(A) > \max(B)$ and $\min(A) > \min(B)$. Then there exist $m, n \in \mathbb{N}$ such that $X^m + A^n >_{\text{FOSD}} Y^m + B^n$.*

Proof. Since $X >_{\text{FOSD}} Y$, we have $K_a(X) > K_a(Y)$ for all $a \in \mathbb{R}$ and $K_a(X) \geq K_a(Y)$ for $a = \pm\infty$. Moreover, $K_a(A) > K_a(B)$ for $a = \pm\infty$. By continuity of K_a in a, there exists $M > 0$ such that for all $a \in \overline{\mathbb{R}} \setminus [-M, M]$, $K_a(A) > K_a(B)$. Now $t :=$

$\min_{a\in[-M,M]}(K_a(X) - K_a(Y)) > 0$, and $s := \min_{a\in[-M,M]}(K_a(A) - K_a(B))$ is finite, so there is $d \in \mathbb{N}$ with $dt + s > 0$. It thus follows from Lemma 4 that $X^d + A \geq_L Y^d + B$. The result follows from the definition of $\geq_L$. $\square$

To make use of the above lemmata, we construct a variant of matching pennies such that any solution must involve someone playing, with positive probability, an action that generates a lottery with a lower max and min than its alternative.

	a_2	b_2
a_1	$(2, 0)$	$(0, 1)$
b_1	$(-1, 1)$	$(1, 0)$

Table 3: Variant of matching pennies

Let $H = (B, v)$ denote the game in Table 3, and let $p \in S(H)$ be a solution.[13] Note that p cannot be a pure-strategy profile, since one player would be playing with probability 1 an action that generates a strictly worse outcome, violating distribution-monotonicity. If one player, i, is playing a pure strategy and their opponent, j, is mixing, then j plays an action that generates a lottery with a lower maximum and minimum with positive probability, as desired. Likewise, in a totally-mixed profile, player 1 plays, with positive probability, action b_1 which yields a lottery with a lower max and min than a_1. Without loss of generality, suppose i plays with positive probability, an action $\underline{a}_i$ that yields a lottery with a lower max and min than its alternative $\overline{a}_i$.

Let $F = (C, w)$ where $C_i = \{a_{0.5}, a_0\}$, $w_i(a_{0.5}, \cdot) = 0.5$ always, and $w_i(a_0, \cdot) = 0$ always. Note that $q_i(a_{0.5}) = 1$ for any solution $q \in S(F)$, by assumption. Since the payoffs in H are integral, a difference in the max or min of lotteries generated by a player's actions involves a difference of at least unity. By narrow framing, $p \times q \in S(H \otimes F)$, and note that $p_i(\underline{a}_i) \cdot q_i(a_{0.5}) > 0$, while $p_i(\overline{a}_i) \cdot q_i(a_0) = 0$. Note also that

$$\max(v_i(\overline{a}_i, p_j) + w_i(a_0, q_j)) > \max(v_i(\underline{a}_i, p_j) + w_i(a_{0.5}, q_j)),$$
$$\min(v_i(\overline{a}_i, p_j) + w_i(a_0, q_j)) > \min(v_i(\underline{a}_i, p_j) + w_i(a_{0.5}, q_j)).$$

Let $G = (A, u)$ be an arbitrary game with $a_Y, a_Z \in A_i$ and $r \in S(G)$ such that $u_i(a_Y, r_{-i}) >_{\text{FOSD}} u_i(a_Z, r_{-i})$. We will show that $r_i(a_Z) = 0$. Indeed, by Lemma 5 there exist $m, n \in \mathbb{N}$ such that

$$u_i(a_Y, r_{-i})^m + (v_i(\overline{a}_i, p_j) + w_i(a_0, q_j))^n >_{\text{FOSD}} u_i(a_Z, r_{-i})^m + (v_i(\underline{a}_i, p_j) + w_i(a_{0.5}, q_j))^n.$$

By narrow framing, $r^m \times p^n \times q^n \in S(G^m \otimes H^n \otimes F^n)$. Since $q_i(a_0) = 0$, and $p_i(\underline{a}_i), q_i(a_{0.5}) > 0$, by distribution-monotonicity, it must be that $r_i(a_Z) = 0$.

[13]In this game, all the other players' actions do not influence the payoffs of the two players considered and can be ignored.

Suppose instead, that there exist $h > \ell$ and $p \in S(G_{h\ell})$, with $p_i(\ell) > 0$. We show that S must then satisfy interiority. By anonymity, it is without loss of generality to suppose, toward a contradiction, that there are $i \in N, G = (A, u)$ with $q \in S(G)$ and $a, b \in A_i$ such that $q_i(a) = 0 < q_i(b)$. Let $m \in \mathbb{N}$ such that $m(h - \ell) > \max[u_i(b, q_{-i})] - \min[u_i(a, q_{-i})]$, and consider that $r := q \times p^m \in S(G \otimes G^n_{h\ell})$ by narrow framing. Let v denote the utility map of $G \otimes G^n_{h\ell}$, and note that

$$v_i(a_h, \dots, a_h, a, r_{-i}) >_{\text{FOSD}} v_i(a_\ell, \dots, a_\ell, b, r_{-i})],$$

while

$$r_i(a_h, \dots, a_h, a) = 0 < r_i(a_\ell, \dots, a_\ell, b),$$

violating distribution-monotonicity.

We next show that such an S must satisfy distribution-neutrality. By anonymity, it is without loss of generality to suppose, toward a contradiction, that there is a game $G = (A, u)$ and $i \in N$ with $q \in S(G)$ and $a, b \in A_i$, such that $u_i(a, q_{-i}) = u_i(b, q_{-i})$, while $q_i(a) < q_i(b)$. We have also assumed that there is $p \in S(G_{h\ell})$ with $p_i(a_\ell) > 0$. Let $m \in \mathbb{N}$ such that $\left(\frac{q_i(b)}{q_i(a)}\right)^m > \frac{p_i(a_h)}{p_i(a_\ell)}$. By narrow framing, $r := q^m \times p \in S(G^m \otimes G_{h\ell})$. Let v denote the utility map of $G^m \otimes G_{h\ell}$, and note that $v_i(a, \dots, a, a_h, r_{-i}) >_{\text{FOSD}} v_i(b, \dots, b, a_\ell, r_{-i})$, while

$$r_i(a, \dots, a, a_h) < r_i(b, \dots, b, a_\ell),$$

violating distribution monotonicity.

□

The proof of Theorem 2 involves the set of lotteries that arise from finite normal form games. In particular, we consider the set of lotteries with finite real-valued outcomes and rational-valued CDFs. We denote this set by $\Delta_{\mathbb{Q}}$. The following remark establishes the relationship between a lottery and a random variable, which we use to construct games.

Remark 7. *Any $X \in \Delta_{\mathbb{Q}}$ can be represented as a random variable with domain $(\Omega = \{1, \dots, m\}, 2^\Omega, \mu)$, where μ is a uniform distribution, $m \in \mathbb{N}$, and $X : \Omega \to \mathbb{R}$.*

B.1 How Lotteries Arise in Games

The following lemma shows that under the assumptions of distribution-monotonicity, narrow framing, and anonymity, there are games where players must evaluate a rich set of lotteries against sure things.

Lemma 6. *Suppose S satisfies distribution-monotonicity, narrow framing, and anonymity, and let $X \in \Delta_{\mathbb{Q}}$. Then for each $i \in N$ and $r \in \mathbb{R}$, there is a game $G_{rX} = (A, u)$ with $\{a_r\} \subsetneq A_i$ and $p \in S(G_{rX})$, such that $u_i(a_r, p_{-i}) = r$ deterministically, and if $p_i(a_r) \neq 1$, then $u_i(a_i, p_{-i}) = X$ for $a_i \neq a_r$. Moreover, $u_i(a_i, p_{-i}) \geq_{FOSD} X$ for some $a_i \in A_i$.*

Proof of Lemma 6. Let $i \in N$ and $X \in \Delta_{\mathbb{Q}}$. As in Remark 7, we represent X as a random variable $X\colon \Omega \to \mathbb{R}$ where $\Omega = \{1, \ldots, m\}$ belongs to the probability space $(\Omega, 2^{\Omega}, \mu)$, and μ is the uniform distribution on Ω. For $\omega \in \Omega$, let x_ω denote $X(\omega)$. We will consider the two possibilities allowed for by Lemma 3.

First, consider the case where players never play FOSD-dominated actions. We construct $G_{rX} = (A, u)$ by $A_j = \Omega$, $A_i = \{a_r\} \cup \{f\colon A_j \to \{x_1, \ldots, x_m\} \mid f \text{ is a bijection}\}$, $u_i(a_r, \cdot) = r$ always, and $u_i(a_i, a_j) = a_i(a_j)$ for $a_j \in A_j$ and $a_i \neq a_r$. We also set $u_j = -u_i$. We think of i's action as choosing r deterministically or choosing a linear order over $\{x_1, \ldots, x_m\}$ and j choosing an index, in which case player j pays player i the value at the index j chose in i's order.

Fix $p \in S(G_{rX})$. If $p_i(a_r) = 1$, then any $f \in A_i \setminus \{a_r\}$ satisfying $f(a_j) \geq f(a'_j)$ whenever $p_j(a_j) > p_j(a'_j)$ satisfies $u_i(f, p_{-i}) \geq_{FOSD} X$, and the lemma holds. For $p_i(a_r) \neq 1$, we want to show p_j is the uniform distribution over A_j. Suppose, for the sake of contradiction, that $p_j(s) > p_j(t)$. Any $f \in A_i \setminus \{a_r\}$ $f(s) < f(t)$ is then FOSD dominated, and thus cannot be played. If, however, p_i places no weight on such f, then s is FOSD-dominated by t for player j. Thus i or j is choosing a first order stochastically dominated strategy with positive probability, which we showed cannot happen (Lemma 3). Hence, we conclude that p_j is uniform, so $u_i(a_i, p_j) = X$ for $a_i \neq a_r$.[14]

Consider then the case where S satisfies interiority and distribution-neutrality. We construct $G_{rX} = (A, u)$ by $A_i = \{a_r, a_X\}$, $A_j = \{1, \ldots, m\}$, $u_i(a_r, \cdot) = r$ always, $u_i(a_X, a_j) = x_{a_j}$, and $u_j = 0$ always. By distribution-neutrality, p_j is uniform, so $u_i(a_X, p_j) = X$. Thus, in either case, we have shown how to construct G_{rX}.

□

B.2 Representation of Monotone Additive Statistics

Mu et al. (2021) provide a characterization of monotone additive statistics on the domain of all compactly supported lotteries. The following lemma shows that we can apply their characterization to monotone additive statistics on the restricted domain of lotteries that arise from finite normal form games.

Lemma 7. *Let $\Phi\colon \Delta_{\mathbb{Q}} \to \mathbb{R}$ be a monotone additive statistic. Then*

$$\Phi(X) = \int_{\overline{\mathbb{R}}} K_a(X) \, \mathrm{d}\mu(a)$$

for some Borel probability measure μ on $\overline{\mathbb{R}}$.

[14] We have implicitly assumed that X is non-constant. For constant $X = x$, we simply let $a_x \in A_i$ with $u_i(a_x, \cdot) = x$ always.

Proof of Lemma 7. Let $\Phi\colon \Delta_{\mathbb{Q}} \to \mathbb{R}$ be a monotone additive statistic and fix a lottery X that is compactly supported.

For each $n \in \mathbb{N}$, let $\tilde{X}_n = \frac{1}{2^n}\lfloor 2^n \cdot X\rfloor$. Note that $\tilde{X}_n \leq_{\text{FOSD}} X \leq_{\text{FOSD}} \tilde{X}_n + \frac{1}{2^n}$. Define the real-valued function Ψ on the set of compactly supported lotteries by $\Psi(X) = \lim_{n\to\infty} \Phi(\tilde{X}_n)$. Since $(\tilde{X}_n)_n$ is an increasing sequence in terms of FOSD, $\Psi(X) \geq \Phi(\tilde{X}_n)$ for all n. For $X \in \Delta_{\mathbb{Q}}$, we have $\Phi(\tilde{X}_n) \leq \Phi(X) \leq \Phi(\tilde{X}_n + \frac{1}{2^n})$ by monotonicity, and $\lim_{n\to\infty} \Phi(\tilde{X}_n + \frac{1}{2^n}) = \lim_{n\to\infty} \Phi(\tilde{X}_n) + \frac{1}{2^n} = \lim_{n\to\infty} \Phi(\tilde{X}_n)$, so $\Phi(X) = \Psi(X)$, i.e., Ψ extends Φ.

Let X, Y be compactly supported lotteries. If $X \geq_{\text{FOSD}} Y$, then $\tilde{X}_n \geq_{\text{FOSD}} \tilde{Y}_n$ for all n, so $\Psi(X) \geq \Psi(Y)$, and Ψ is monotone. Moreover, if X and Y are independent, for each n, $\Phi(\tilde{X}_n + \tilde{Y}_n) \leq \Phi((\widetilde{X+Y})_n) \leq \Psi(X+Y) \leq \Phi((\widetilde{X+Y})_n) + \frac{1}{2^n} \leq \Phi(\tilde{X}_n + \tilde{Y}_n) + \frac{3}{2^n}$. Taking the limit as $n \to \infty$, we have $\Psi(X+Y) = \lim_{n\to\infty} \Phi(\tilde{X}_n + \tilde{Y}_n) = \lim_{n\to\infty} \Phi(\tilde{X}_n) + \Phi(\tilde{Y}_n) = \Psi(X) + \Psi(Y)$. Finally, any extension of Φ to a monotone additive statistic Ξ on compactly supported lotteries must satisfy $\Phi(\tilde{X}_n) \leq \Xi(X) \leq \Phi(\tilde{X}_n + \frac{1}{2^n})$ and thus the identity in the definition of Ψ, i.e., Ψ is the unique extension. Thus, by the characterization of Mu et al. (2021), $\Phi(X) = \Psi(X) = \int_{\overline{\mathbb{R}}} K_a(X)\,\mathrm{d}\mu(a)$ for some Borel probability measure μ on $\overline{\mathbb{R}}$. $\square$

B.3 Additivity for all Lotteries

Lemma 8. *Let $\Phi\colon \Delta_{\mathbb{Q}} \to \mathbb{R}$ be a monotone statistic. If Φ is additive for all lotteries, then Φ is the expectation.*

Proof of Lemma 8. We consider each finite probability space $(\Omega = \{1, \ldots, m\}, 2^{\Omega}, \mu)$, where μ is the uniform distribution on Ω. Each $X \in \Delta_{\mathbb{Q}}$ can be represented as a random variable $X\colon \Omega \to \mathbb{R}$ for a large enough m. We require that $\Phi(X+Y) = \Phi(X) + \Phi(Y)$ for any random variables $X, Y \in \mathbb{R}^{\Omega}$.

For $X\colon \Omega \to \mathbb{R}$, we may write $X = \sum_{\omega} I(\omega)X(\omega)$, where I denotes the indicator function. For each $\omega \in \mathbb{R}$, define $f_{\omega}\colon \mathbb{R} \to \mathbb{R}$ by $f_{\omega}(x) = \Phi(I(\omega)X(\omega))$, so $\Phi(X) = \sum_{\omega} f_{\omega}(X(\omega))$. It follows that each f_{ω} is a monotone additive function and is therefore linear. Thus there is $Z \in \mathbb{R}^{\Omega}$ such that $\Phi(X) = Z \cdot X$ for all X. Since Φ only depends on the distribution of X and μ is uniform, $\Phi(X) = \Phi(X \circ \pi)$ for any permutation $\pi\colon \Omega \to \Omega$, which is only possible for constant Z. Finally, since Φ is a statistic, it maps any constant random variable to its value, so $Z(\omega) = \frac{1}{m}$ for all ω. We have thus shown that for any $X \in \mathbb{R}^{\Omega}$, $\Phi(X) = \frac{1}{|\Omega|}\sum_{\omega} X(\omega) = \mathbb{E}[X]$. $\square$

The remainder of the proof of Theorem 2 relies on the existence of a monotone additive statistic Φ that players respond to (Theorem 3). Theorem 3 is proved in the subsequent section.

Proof of Theorem 2. Let $X, Y \in \Delta_{\mathbb{Q}}$. As previously, we represent X and Y as random variables on $(\Omega = \{1, \ldots, m\}, 2^{\Omega}, \mu)$, where μ is uniform. By Lemma 6, we can construct $G_{rX} = (A, u)$. We will consider the frequency that i chooses the sure thing a_r.

Since S satisfies narrow framing, distribution-monotonicity, and anonymity, by Theorem 3, S is a refinement of an SRE, T. Let $p \in S(G_{rX})$. If $T = \mathrm{LQRE}_{\lambda\Phi}$ with $\lambda = 0$, the result holds trivially. If T is any other SRE, there exists a monotone additive statistic Φ such that if $r < \Phi(X)$, then $p_i(a_r) < \frac{1}{2}$, while if $r > \Phi(X)$, then $p_i(a_r) > \frac{1}{2}$.

Consider the game (A, v), where $v(a_i, a_j) = u(a_i, a_j) + Y(a_j)$, for each $(a_i, a_j) \in A_i \times A_j$. By strategic invariance, $p \in S(A, v)$. Note also that $\Phi(v_i(a_r, p_j)) = \Phi(r + Y) = r + \Phi(Y)$, and $\Phi(v_i(a_i, p_j)) = \Phi(X + Y)$, for $a_i \neq a_r$. Since we know that $r < \Phi(X)$ implies $p_i(r) < \frac{1}{2}$, while $r > \Phi(X)$ implies $p_i(r) > \frac{1}{2}$, and probabilities are monotone with respect to Φ, it must be that $\underline{r} + \Phi(Y) \leq \Phi(X + Y) \leq \overline{r} + \Phi(Y)$, for $\underline{r} < \Phi(X) < \overline{r}$. Thus $\Phi(X + Y) = \Phi(X) + \Phi(Y)$. By Lemma 8, Φ is the expectation on $\Delta_{\mathbb{Q}}$, so by Lemma 7, Φ is the expectation.[15]

□

C Proof of Proposition 1

Proof of Proposition 1. First we show that every game has an $\mathrm{LQRE}_{\lambda\Phi}$ equilibrium for every $\lambda \geq 0$ and Φ. In order to do this, fix any $G = (A, u), i \in N, a_i \in A_i, t \in \mathbb{R} \setminus \{0\}$ and $p \in \prod_i \Delta A_i$. Let $K'_t(u_i(a_i, p_{-i})) = \frac{1}{t} \ln \mathbb{E}[\exp(t \cdot u_i(a_i, p_{-i}))]$ (i.e. $K_t(u_i(a_i, p_{-i}))$). For $t = 0$, let $K'_t(u_i(a_i, p_{-i})) = \mathbb{E}[u_i(a_i, p_{-i})]$. For $t = -\infty$, let $K'_t(u_i(a_i, p_{-i})) = \min_{a_{-i}} u_i(a_i, a_{-i})$, and let $K'_t(u_i(a_i, p_{-i})) = \max_{a_{-i}} u_i(a_i, a_{-i})$ for $t = \infty$. Define, for each monotone additive statistic $\Phi = \int_{\overline{\mathbb{R}}} K_t(u_i(a_i, p_{-i}))\, \mathrm{d}\mu(t)$, the corresponding map $\Phi'(u_i(a_i, p_{-i})) = \int_{\overline{\mathbb{R}}} K'_t(u_i(a_i, p_{-i}))\, \mathrm{d}\mu(t)$.

Let $\lambda \geq 0$ and define $T \colon \prod_i \Delta A_i \to \prod_i \Delta A_i$ by $T_i(q)(a_i) \propto \exp(\lambda \Phi'(u_i(a_i, q_{-i})))$. Since K'_t for all $t \in \overline{\mathbb{R}}$ are continuous in q. Since $\prod_i \Delta A_i$ is convex and compact, T has a fixed-point q^* by Brouwer's fixed-point theorem. Since T maps every mixed strategy profile to a totally mixed strategy profile, q^* must be totally mixed. We thus have

$$q^*_i(a_i) \propto \exp(\lambda \Phi'(u_i(a_i, q^*_{-i}))) = \exp(\lambda \Phi(u_i(a_i, q^*_{-i}))),$$

as Φ and Φ' agree when i's opponents play totally mixed strategy profiles.

For each $\lambda \in \mathbb{N}$, let q^λ denote an $\mathrm{LQRE}_{\lambda\Phi}$ equilibrium. Then if $\Phi(u_i(a_i, p_{-i})) = \int_{\mathbb{R}} K_t(u_i(a_i, p_{-i}))\, \mathrm{d}\mu(t)$, the existence of a Nash_Φ equilibrium is guaranteed as a limit point of $\{q^\lambda\}_{\lambda \in \mathbb{N}}$, which exists by the Bolzano-Weierstrass theorem. Any such limit point q is Nash_Φ equilibrium. Indeed, let $i \in N$ and $a_i, b_i \in A_i$ such that $\Phi(u_i(a_i, q_{-i})) >$

[15] The uniqueness of the probability measure μ in the representation of Φ was shown by Mu et al. (2021).

$\Phi(u_i(b_i, q_{-i}))$. Since $\Phi(u_i(q))$ is continuous in q, there are $\varepsilon > 0$ and $\Lambda \in \mathbb{N}$ such that for all $\lambda \geq \Lambda$, $\Phi(u_i(a_i, q^\lambda_{-i})) > \Phi(u_i(b_i, q^\lambda_{-i})) + \varepsilon$. Since $q^\lambda_i(b_i) < \frac{\exp(\lambda\Phi(u_i(b_i,q^\lambda_{-i})))}{\exp(\lambda\Phi(u_i(a_i,q^\lambda_{-i})))} < \exp(-\lambda\varepsilon)$, as $\lambda \to \infty$, $q^\lambda_i(b_i) \to 0$, hence $q_i(b_i) = 0$.

Finally, we show how to construct a game, when there are at least two players, that has no Nash$_\Phi$ equilibrium when Φ places a positive weight on the minimum or maximum. For such a Φ, let $\varepsilon = \mu(-\infty) + \mu(+\infty)$ and consider the game in table 4. Since pure Nash$_\Phi$ equilibria coincide with pure Nash equilibria for all Φ, it is easy to see that the game has no pure equilibria. Likewise, there are no equilibria where either player plays a pure strategy, since the best responses to pure strategies in this game are pure for all Φ. In particular any supposed Nash$_\Phi$ equilibrium q would have player 2 playing a totally mixed strategy. We thus have

$$\begin{aligned}\Phi(u_1(a_1, q_2)) - \Phi(u_1(b_1, q_2)) &= (\mu(-\infty) + \mu(+\infty)) \cdot \frac{1}{\varepsilon} \\ &+ \int_{\mathbb{R}} K_t(u_1(a_1, q_2))\,\mathrm{d}\mu(t) - \int_{\mathbb{R}} K_t(u_1(b_1, q_2))\,\mathrm{d}\mu(t) \geq 1 - \mu(\mathbb{R}) \cdot 1 = \varepsilon > 0.\end{aligned}$$

This contradicts the assumption that q is a totally mixed Nash$_\Phi$ equilibrium, which would require that $\Phi(u_1(a_1, q_2)) = \Phi(u_1(b_1, q_2))$.

	a_2	b_2
a_1	$(1 + \frac{1}{\varepsilon}, 0)$	$(0, 1)$
b_1	$(-\frac{1}{\varepsilon}, 1)$	$(1, 0)$

Table 4: Variant of matching pennies for which extremal Nash$_\Phi$ equilibria do not exist.

$\square$

D Proof of Theorem 3

We use Δ to denote the set of all lotteries with finite real-valued outcomes.

Lemma 9. *Suppose a solution concept S satisfies distribution-neutrality, interiority, narrow framing, distribution-monotonicity, and anonymity. Then there is a monotone additive statistic Φ such that for all $i, G = (A, u), p \in S(G), a_i \in A_i$,*

$$p_i(a_i) \propto \exp(\lambda\Phi(u_i(a_i, p_{-i})))$$

for some $\lambda \geq 0$.

Proof of Lemma 9. Fix $i \in N$. We show that for any $X \in \Delta$ there is a game $G_X = (A, u)$, with $A_i = \{a_0, a_X\}$ and a $p_X \in S(G_X)$ such that $u_i(a_0, p_{X-i}) = 0$ always and $u_i(a_X, p_{X-i}) = X$.

Indeed, let $j \neq i$ and $X \in \Delta$ with outcomes $x_1, \ldots, x_m$, and let $\lambda \geq 0$ be identified by Remark 6. For $\lambda = 0, G = (A, u), p \in S(G)$, we show that p_i is the uniform distribution, so that the lemma holds. Indeed, if $p_i(a_i) > p_i(b_i)$, then consider $q \in S(G_{h\ell})$, so that by narrow framing, $p \times q \in S(G \otimes G_{h\ell})$. Note that by Remark 6, q_i is uniform. We will thus have $p_i(a_i) \cdot q_i(\ell) > p_i(b_i) \cdot q_i(h)$, which violates distribution-monotonicity for $h - \ell$ large enough.

Consider then the case where $\lambda > 0$. Define $G_X = (A, u)$ by $A_j = \{b_1, \ldots, b_m\}$ and $u_j(b_\ell, \cdot) = \frac{1}{\lambda} \ln(\mathbb{P}(X = x_\ell))$ always. Remark 6 applies as $u_j(\cdot, p_{X-j})$ is a deterministic function, so $p_{Xj}(b_\ell) = \mathbb{P}(X = x_\ell)$ for all $\ell = 1, \ldots, m$. We thus define $u_i(a_X, b_\ell) = x_\ell$ for $\ell = 1, \ldots, m$, and $u_i(a_0, \cdot) = 0$ always. Thus $u_i(a_X, p_{Xj}) = X$, while $u_i(a_0, p_{Xj}) = 0$ always.

The rest of this proof matches that of Lemma 2:

For each $X \in \Delta$ and $p_X \in S(G_X)$ as defined above, let $f(X) = \ln \frac{p_{Xi}(a_X)}{1 - p_{Xi}(a_X)}$, which is well-defined by interiority. Let $X, Y \in \Delta$ be independent lotteries and let $p_X \in S(G_X), p_Y \in S(G_Y)$, and $p_{X+Y} \in S(G_{X+Y})$. By narrow framing, $p_X \times p_Y \times p_{X+Y} \in S(G_X \otimes G_Y \otimes G_{X+Y})$. By distribution-neutrality, $p_{Xi}(a_X) p_{Yi}(a_Y)(1 - p_{X+Yi}(a_{X+Y})) = (1 - p_{Xi}(a_X))(1 - p_{Yi}(a_Y)) p_{X+Yi}(a_{X+Y})$. Rearranging and taking logs, we get $f(X+Y) = f(X) + f(Y)$. Since S satisfies narrow framing and distribution-monotonicity, $p_{Xi}(a_X)$ is increasing in X, as $p_{Xi}(a_X) p_{X'i}(a_0) \geq p_{X'i}(a_{X'}) p_{Xi}(a_0)$ for $X >_{\text{FOSD}} X'$. Hence, f is non-decreasing and additive for lotteries. Finally, define $g \colon \mathbb{R} \to \mathbb{R}$ by $g(x) = f(x)$ for deterministic lotteries yielding x for sure. Then $g(x+y) = f(x+y) = f(x) + f(y) = g(x) + g(y)$. Since f is monotone, g must be monotone, so there is a $\lambda \in [0, \infty)$ such that $f(x) = g(x) = \lambda x$ for all $x \in \mathbb{R}$. Hence, f is a scaled monotone additive statistic, i.e., $f(X) = \lambda \Phi(X)$, for some monotone additive statistic Φ.

Fix any $G = (B, v), q \in S(G)$ with $b, c \in B_i$. Let $X = v_i(b, q_{-i})$ and $Y = v_i(c, q_{-i})$. By narrow framing, $r := q \times p_X \times p_Y \in S(G \otimes G_X \otimes G_Y)$. Let w denote the utility map for $G \otimes G_X \otimes G_Y$. Note that $w_i((b, a_0, a_Y), r_{-i}) = w_i((c, a_X, a_0), r_{-i})$. By distribution-neutrality,

$$q_i(b)(1 - p_{Xi}(a_X)) p_{Yi}(a_Y) = q_i(c) p_{Xi}(a_X)(1 - p_{Yi}(a_Y)).$$

Rearranging, we have

$$\frac{q_i(b)}{q_i(c)} = \frac{\exp f(X)}{\exp f(Y)}.$$

Since c was arbitrary,

$$q_i(b) \propto \exp f(X) = \exp(\lambda \Phi[v_i(b, q_{-i})]).$$

By anonymity, this holds for all $i \in N$.

$\square$

We now show that either the first case of Theorem 3 holds or the above lemma applies and the second case of Theorem 3 holds.

Proof of Theorem 3. Let $i \in N$ and consider the two possibilities allowed for Lemma 3.

Suppose we are in the former case where agents never play FOSD-dominated actions. Define $\Phi \colon \Delta \to \mathbb{R}$ by $\Phi(X) = \sup\{r \in \mathbb{R} \colon \exists p \in S(G_{rX}) \text{ with } p_i(a_r) < 1\}$, which is finite since players never play FOSD-dominated strategies. Note that $\Phi(c) = c$ for any $c \in \mathbb{R}$, i.e., Φ is a statistic. We now show that whenever $r < \Phi(X)$, and $p \in S(G_{rX})$ we will have $p_i(a_r) < 1$. Indeed, let s with $r < s \leq \Phi(X)$ such that there is $q \in S(G_{sX})$ with $q_i(a_s) < 1$. By narrow framing, $p \times q \in S(G_{rX} \otimes G_{sX})$. By distribution-monotonicity, $p_i(a_r) \cdot (1 - q_i(a_s)) \leq (1 - p_i(a_r)) \cdot q_i(a_s)$. Since we assumed $q_i(a_s) < 1$, we cannot have $1 - p_i(a_r) = 0$.

We next show that only maximizers of Φ can be played with positive probability. Let $G = (A, u)$, with $a, b \in A_i, o \in S(G)$, and $\Phi(u_i(a, o_{-i})) > \Phi(u_i(b, o_{-i}))$. Let $X = u_i(a, o_{-i}), Y = u_i(b, o_{-i})$, and fix r, s satisfying $\Phi(Y) < r < s < \Phi(X)$. Let $p \in S(G_{sX})$ and $q \in S(G_{rY})$. By narrow framing, $o \times p \times q \in S(G \otimes G_{sX} \otimes G_{rY})$, and by distribution-monotonicity,

$$o_i(b) \cdot (1 - p_i(a_s)) \cdot q_i(a_r) \leq o_i(a) \cdot p_i(a_s) \cdot (1 - q_i(a_r)).$$

By definition of Φ, we have $q_i(a_r) = 1$. Since $1 - p_i(a_s) > 0$ and $q_i(a_r) = 1$, it must be that $o_i(b) = 0$, as desired.

We next show that Φ is monotone with respect to first-order stochastic dominance. Let $r \in \mathbb{R}$ and $X, Y \in \Delta$ with $X \geq_{\text{FOSD}} Y$. Let $p \in S(G_{rX}), q \in S(G_{rY})$. By narrow framing, $p \times q \in S(G_{rX} \otimes G_{rY})$, and by distribution-monotonicity, $p_i(a_r) \cdot (1 - q_i(a_r)) \leq (1 - p_i(a_r)) \cdot q_i(a_r)$, or $p_i(a_r) \leq q_i(a_r)$. It then follows from the definition of Φ that $\Phi(X) \geq \Phi(Y)$.

It remains to be shown that Φ is additive for independent variables. Indeed, let $X, Y \in \Delta$ be independent, and let $r > \Phi(X), s > \Phi(Y)$. Fix $t > r + s$ and let $o \in S(G_{rX}), p \in S(G_{sY}), q \in S(G_{t\,X+Y})$, and note that $o \times p \times q \in S(G_{rX} \otimes G_{sY} \otimes G_{t\,X+Y})$, by narrow framing. By distribution-monotonicity,

$$o_i(a_r) \cdot p_i(a_s) \cdot (1 - q_i(a_t)) \leq (1 - o_i(a_r)) \cdot (1 - p_i(a_s)) \cdot q_i(a_t).$$

Now $(1 - o_i(a_r)) = (1 - p_i(a_s)) = 0$, while $o_i(a_r) = p_i(a_s) = 1$, so $q_i(a_t) = 1$, meaning $\Phi(X + Y) \leq t$. Since r, s and t can be chosen so that t is arbitrarily close to $\Phi(X) + \Phi(Y)$, it follows that $\Phi(X + Y) \leq \Phi(X) + \Phi(Y)$.

Likewise, we can choose $r < \Phi(X), s < \Phi(Y)$ and $t < r+s$. By distribution-monotonicity we will have

$$(1 - o_i(a_r)) \cdot (1 - p_i(a_s) \cdot q_i(a_t) \leq o_i(a_r) \cdot p_i(a_s) \cdot (1 - q_i(a_t)).$$

Now $(1 - o_i(a_r)) = (1 - p_i(a_s)) = 1$, while $o_i(a_r) = p_i(a_s) = 0$,so $q_i(a_t) = 0$. This means $\Phi(X + Y) \geq t$. Since r, s and t can be chosen so that t is arbitrarily close to $\Phi(X) + \Phi(Y)$, it follows that $\Phi(X + Y) \geq \Phi(X) + \Phi(Y)$. We conclude that $\Phi(X + Y) = \Phi(X) + \Phi(Y)$.

Since S satisfies anonymity, we have now shown that if this case holds, there is a monotone additive statistic Φ such that for all games G, $p \in S(G), i \in N$, and $a_i \in A_i$,

$$\operatorname{supp} p_i \subseteq \operatorname*{argmax}_a \Phi(u_i(a, p_{-i})).$$

Suppose then, S satisfies interiority and distribution-neutrality. By Lemma 9, there exists a monotone additive statistic Φ and $\lambda \geq 0$ such that for all $i, G = (A, u), p \in S(G), a_i \in A_i$,

$$p_i(a_i) \propto \exp(\lambda\Phi(u_i(a_i, p_{-i}))).$$

□

E Proof of Theorem 4

Proof. By Theorem 3, S is a refinement of either Nash_Φ or $\text{LQRE}_{\lambda\Phi}$ for some monotone additive statistic Φ and $\lambda \geq 0$. For $\text{LQRE}_{\lambda\Phi}$ where $\lambda = 0$, the result is trivial. For any other SRE with Φ, we define the statistic Ψ by $\Psi(X) = f^{-1}(\mathbb{E}[f(X)])$. Since f is strictly increasing, choice probabilities are maximized for lotteries with maximal Φ and Ψ in any solution to any game. It follows from Lemma 6 that Ψ and Φ agree on $\Delta_{\mathbb{Q}}$, so by Lemma 7 $\Psi = \Phi$. Since $f \circ \Psi$ is an expected utility over lotteries, Ψ satisfies independence, i.e., for all compactly supported lotteries X, Y, Z and all $\beta \in (0, 1)$, $X \succsim Y$ implies $\beta X + (1 - \beta)Z \succsim \beta Y + (1 - \beta)Z$. Hence, the result follows from Proposition 8 of Mu et al. (2021). □

F Proof of Theorem 5

Since S satisfies distribution-monotonicity, narrow framing, and anonymity, and players always play totally mixed strategies, by Theorem 3, S is a refinement of some $\text{LQRE}_{\lambda\Phi}$. Scale invariance ensures that Φ belongs to the class of positively homogenous monotone additive statistics, which we characterize in the following lemma.

Lemma 10. *Suppose that $\Phi \colon \Delta \to \mathbb{R}$ is a monotone additive statistic such that $\Phi(\alpha X) = \alpha\Phi(X)$ for all $X \in \Delta$ and some $\alpha \geq 0$. Then Φ is a convex combination of the minimum, the maximum and the expectation.*

Proof of Lemma 10. Let $\beta > 0$ and $\Phi(X)$ be a monotone additive statistic. By Lemma 7, $\Phi(X) = \int K_a(X)\,\mathrm{d}\mu(a)$. Then

$$\begin{aligned}\Phi(\beta X) &= \int \frac{1}{a} \log \mathbb{E}\mathrm{e}^{a\beta X}\,\mathrm{d}\mu(a)\\ &= \int \frac{\beta}{a\beta} \log \mathbb{E}\mathrm{e}^{a\beta X}\,\mathrm{d}\mu(a)\\ &= \beta \int K_{a\beta}(X)\,\mathrm{d}\mu(a)\\ &= \beta \int K_a(X)\,\mathrm{d}(\beta_*\mu)(a).\end{aligned}$$

Denote $\Psi(X) = \int K_a(X)\,\mathrm{d}(\beta_*\mu)(a)$, and note that this is also a monotone additive statistic. Then $\Phi(\beta X) = \beta\Psi(X)$.

Suppose $\Phi(\beta X) = \beta\Phi(X)$ for all X and some $\beta > 0$. Hence $\beta\Phi(X) = \beta\Psi(X)$ for all X, and so $\Phi = \Psi$. By Lemma 5 of Mu et al. (2021) it follows that $\mu = \beta_*\mu$. Since a finite measure on $\mathbb{R}$ can only be invariant to rescaling if it is the point mass at 0, it follows that $\mu(\{-\infty, +\infty, 0\}) = 1$. □

It is straightforward to see that if μ is supported on $\{-\infty, +\infty, 0\}$, then Φ satisfies $\Phi(\alpha X) = \alpha\Phi(X)$ for *all* $\alpha \geq 0$. We proceed with the proof of Theorem 5.

Proof of Theorem 5. Since S satisfies distribution-monotonicity, narrow framing, and anonymity, and players always play totally mixed strategies, by Theorem 3, S is a refinement of some $\mathrm{LQRE}_{\lambda\Phi}$. Let $X \in \Delta$ and consider the game $G_{rX} = (A, u)$ where $r = \Phi(X)$. Any $p \in S(G_{rX})$ must satisfy $p_i(a_r) = p_i(a_X)$, and by scale invariance this is also the case for each $p \in S(A, \alpha \cdot u)$ for any $\alpha \geq 0$. It follows that $\Phi(\alpha X) = \Phi(\alpha \cdot r) = \alpha \cdot r = \alpha\Phi(X)$. By Lemma 10, Φ is a convex combination of the minimum, the maximum and the expectation.

For any scaling factor $\alpha \geq 0$, we may write the scaled convex combination of the minimum, maximum and expectation as $\lambda_1 \min + \lambda_2\mathbb{E} + \lambda_3 \max$ for some $\lambda \in \mathbb{R}^3_{>0}$, giving us the desired representation. □

References

G. Aubrun and I. Nechita. Catalytic majorization and ℓ_p norms. *arXiv preprint arXiv:quant-ph/0702153v2*, Jun 2007. URL https://arxiv.org/abs/quant-ph/0702153v2.

N. Barberis, M. Huang, and R. H. Thaler. Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American economic review*, 96 (4):1069–1090, 2006.

F. Brandl and F. Brandt. Justifying optimal play via consistency. *Theoretical Economics*, 14(4):1185–1201, 2019.

F. Brandl and F. Brandt. An axiomatic characterization of nash equilibrium. *Theoretical Economics*, 2024.

M. K. Camara. Computationally tractable choice. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 28–28, 2022.

E. Friedman and F. Mauersberger. Quantal response equilibrium with symmetry: Representation and applications. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 240–241, 2022.

J. K. Goeree and P. Louis. M equilibrium: A theory of beliefs and choices in games. *American Economic Review*, 111(12):4002–4045, 2021. doi: 10.1257/aer.20201683.

J. K. Goeree, C. A. Holt, and T. R. Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1):97–113, 2003.

J. K. Goeree, C. A. Holt, and T. R. Palfrey. Regular quantal response equilibrium. *Experimental economics*, 8:347–367, 2005.

J. K. Goeree, C. A. Holt, and T. R. Palfrey. *Quantal response equilibrium: A stochastic theory of games*. Princeton University Press, 2016.

J. K. Goeree, C. A. Holt, and T. R. Palfrey. Stochastic game theory for social science: A primer on quantal response equilibrium. In *Handbook of Experimental Game Theory*, pages 8–47. Edward Elgar Publishing, 2020.

S. Govindan and R. Wilson. Axiomatic theory of equilibrium selection for games with two players, perfect information, and generic payoffs. Technical Report 2008, Working Paper, 2009.

J. C. Harsanyi and R. Selten. A general theory of equilibrium selection in games. *MIT Press Books*, 1, 1988.

B. Kőszegi and F. Matějka. Choice simplification: A theory of mental budgeting and naive diversification. *The Quarterly Journal of Economics*, 135(2):1153–1207, 2020.

R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and economic behavior*, 10(1):6–38, 1995.

R. D. McKelvey, T. R. Palfrey, and R. A. Weber. The effects of payoff magnitude and heterogeneity on behavior in $2\times$ 2 games with unique mixed strategy equilibria. *Journal of Economic Behavior & Organization*, 42(4):523–548, 2000.

L. P. Metzger and M. O. Rieger. Non-cooperative games with prospect theory players and dominated strategies. *Games and Economic Behavior*, 115:396–409, 2019.

H. Moulin. *Cooperative microeconomics: a game-theoretic introduction*. Princeton University Press, 1995.

X. Mu, L. Pomatto, P. Strack, and O. Tamuz. Monotone additive statistics. *arXiv preprint arXiv:2102.00618*, 2021.

R. B. Myerson. Refinements of the nash equilibrium concept. *International journal of game theory*, 7:73–80, 1978.

H. Norde, J. Potters, H. Reijnierse, and D. Vermeulen. Equilibrium selection and consistency. *Games and Economic Behavior*, 12(2):219–225, 1996.

M. J. Osborne and A. Rubinstein. Games with procedurally rational players. *American Economic Review*, pages 834–847, 1998.

D. Read, G. Loewenstein, M. Rabin, G. Keren, and D. Laibson. Choice bracketing. *Elicitation of preferences*, pages 171–202, 2000.

A. E. Roth. *Axiomatic models of bargaining*, volume 170. Springer Science & Business Media, 2012.

F. Sandomirskiy and O. Tamuz. Decomposable stochastic choice. *arXiv preprint arXiv:2312.04827*, 2023.

J. Shalev. Loss aversion equilibrium. *International Journal of Game Theory*, 29:269–287, 2000.

W. Thomson. *The Axiomatics of Economic Design: An introduction to theory and methods. Volume 1*, volume 1. Springer Nature, 2023.

J. Weinstein. The effect of changes in risk attitude on strategic behavior. *Econometrica*, 84(5):1881–1902, 2016.

J. R. Wright and K. Leyton-Brown. Predicting human behavior in unrepeated, simultaneous-move games. *Games and Economic Behavior*, 106:16–37, 2017.

A. Yekkehkhany, T. Murray, and R. Nagi. Risk-averse equilibrium for games. *arXiv preprint arXiv:2002.08414*, 2020.