A Theory of Games Played by Teams of Players *

Jeongbin Kim† Thomas R. Palfrey‡ Jeffrey R. Zeidel§

August 7, 2019

Abstract

We propose a theory of equilibrium in games where each player in the game is team. In a team equilibrium, each member of each team has unbiased estimates of the equilibrium expected payoff for each available action for that team. A team collective choice rule aggregates these individual estimates into a team action. Existence is proven and conditions on the collective choice rule are identified that guarantee team choices to be stochastically optimal in the sense that the probability the team chooses an action is increasing in its equilibrium expected payoff. We prove that team equilibrium converges to Nash equilibrium in games played by large teams using any anonymous scoring rule, such as plurality rule or Borda count. The team game framework is generalized to games in extensive form, and team equilibrium converges to sequential equilibrium with large teams for any anonymous scoring rule. The theory is illustrated with some familiar binary action games.

JEL Classification Numbers: C72. Keywords: Team Games, Collective choice, Voting, Equilibrium

---

*The financial support of the National Science Foundation (SES-1426560) and The Ronald and Maxine Linde Institute of Economics and Management is acknowledged. We are grateful for comments from audiences at the 2018 World Meeting of the Economic Science Association in Berlin, the 2018 Southwest Experimental and Behavioral Economics Workshop at Caltech, and the 2019 PECO conference in Washington, D.C. We thank Erkut Ozbay for his helpful discussion of the paper.

†Division of the Humanities and Social Sciences, California Institute of Technology, jbkim@caltech.edu

‡Division of the Humanities and Social Sciences, California Institute of Technology, trp@hss.caltech.edu

§Division of the Humanities and Social Sciences, California Institute of Technology, jzeidel@caltech.edu

1 Introduction

For most applications of game theory in economics and political science, each ”player” of the game is actually a team of players. For reasons of analytical convenience and longstanding tradition, these teams are modeled as if they are unitary actors - i.e., single individuals. Examples abound. In spectrum auctions, the players are giant corporations such as Verizon, AT&T, and Sprint. The same is true in virtually any model used to study problems in industrial organization: oligopoly, limit pricing and entry deterrence, R&D races, and so forth. In the crisis bargaining literature aimed at understanding international conflict, the players are nation states. In the political arena key players include parties, civic organizations, campaign committees, large donor groups, commissions, panels of judges, advisory committees, etc. These “teams” range not only in size and scope but also in their organizational structure and procedures for reaching decisions.

A basic premise of the theoretical framework developed in this paper is that the unitary actor approach misses a critical component of these strategic environments, namely the collective choice problem within each competing team. This premise is not merely conjectural but is supported by a growing body of experimental work that has begun to uncover inconvenient facts pointing to important behavioral differences between games played by teams of players and games played by individual decision makers.

Many of the studies that compare group and individual behavior in games find that team play more closely resembles the standard predictions of game theory. To quote from Charness and Sutter’s JEL survey (2012, p. 158): “In a nutshell, the bottom line emerging from economic research on group decision-making is that groups are more likely to make choices that follow standard game-theoretic predictions...” Similarly, Kugler et al. (2012, p. 471) summarize the main finding of their survey in the following way: “Our review suggests that results are quite consistent in revealing that group decisions are closer to the game-theoretic assumption of rationality than individual decisions.” A similar conclusion has been reached in many individual choice experiments as well. For example a variety of judgment biases that are commonly observed in individual decision making under uncertainty are significantly reduced by group decision-making.[1]

[1]Such biases include probability matching (Schulze and Newell, 2016), hindsight bias (Stahlberg et al., 1995), overconfidence (Sniezek and Henry, 1989), the conjunction fallacy (Charness et al., 2010; Tversky and Kahneman, 1983), and forecasting errors (Blinder and Morgan, 2005, 2008), and inefficient portfolio selection (Rockenbach et al., 2007).

Given these extensive findings about group vs. individual choice in games, which cross multiple disciplines, it is perhaps surprising that these observations remain in the category of "anomalies" for which there is no existing general theoretical model that can unify these anomalies under a single umbrella. In particular, one hopes such a model might apply not only to interactive games, but also to non-interactive environments such as those experiments that have documented judgment biases and choice anomalies. This paper takes a step in that direction.[2]

Our theoretical framework of team games combines two general approaches to modeling strategic behavior and team behavior: non-cooperative game theory and collective choice theory.[3] Non-cooperative game theory provides the basic structure of a strategic form game, formalized as a set of players, action sets, and payoff functions, or more generally a game in extensive form, which includes additional features including moves by nature, order of play, and information sets. The focus here is exclusively on games played by teams in a pure common value setting.[4] The common value assumption is that all players on the team share the same payoff function.

Collective choice theory provides a theoretical structure to model the effect of different procedures or rules according to which a group of individuals produces a group decision. If all members of the team have perfectly rational expectations about the equilibrium expected payoffs in the game, then they could all agree unanimously on an optimal action choice, and the collective choice problem would be trivial. For this reason, our approach relaxes the usual assumption of perfectly rational expectations about the expected payoffs of actions. Instead, individual members' expectations about the payoff of each available action to the team are correct *on average*, but subject to unbiased errors, so that members of the same team will generally have different expectations about the payoff of each available action, which one can view as opinions, but on average these expectations are the same for all members and equal the true (equilibrium) expected payoffs of each action.[5]

[2] Charness and Sutter (2012) and others have offered some qualitative conjectures about factors that might play a role in the differences between group and individual decision making. For example, perhaps group dynamics lead to more competitive attitudes among the members, due to a sense of group membership. Or perhaps groups are better at assessing the incentives of their opponents; or groups follow the lead of the most rational member ("truth wins"). Of these conjectures, our model is closest to the last one, in the sense that the aggregation process of the diverse opinions can produce better decisions if there is a grain of truth underlying those opinions. However, this would depend on the collective decision-making procedures.

[3] There is also a connection with the economic theory of teams. See Marschak and Radner (1972).

[4] The assumption of common values seems like the natural starting point for developing a team theory of games. In principle, this could be extended to allow for heterogeneous preferences among the members of the same team, for example diverse social preferences.

[5] These errors have a similar role as the additive payoff disturbances in quantal response equilibrium.

Thus, the aggregation problem within a team arises because different members of the team have different opinions about the expected payoffs of the available actions, where these different opinions take the form of individual estimates of the expected payoff of each possible action. The collective choice rule is modeled abstractly as a function mapping a profile of team members' opinions (i.e., estimates) into a team action choice. Because the individual estimates are stochastic, this means that the action choices by a team will not be deterministic, but will be "as if" mixed strategies, with the distribution of a team's effective mixed strategy a product of both the error distribution of the individual team members' estimates and the collective choice rule that transforms these estimates into a team action choice. The equilibrium restriction is that individuals have rational expectations on average, given the mixed strategy profile of all the other teams, which results from aggregation of their members' diverse estimates via some collective choice rule for each of the other teams.

Even though the collective choice rules are modeled abstractly, many of these collective choice rules correspond to voting rules or social choice procedures that are familiar. For example, if a team has exactly two possible actions, then the *majority rule* would correspond to a collective choice rule in which the team's action choice is the one for which a majority of members estimate to have the higher expected payoff (with some tie-breaking rule in case of an even-number of members). With more than two actions, this could be extended naturally to *plurality rule*. *Weighted voting* would give certain member estimates more weight than others. At the extreme, a *dictatorial rule* would specify a particular team member, and the team action choice would be the one that is best in the opinion of only that team member. A *Borda rule* would add up the individual opinion ordinal ranks of each action and choose the one with the highest average ranking. Other collective choice rules do not have an obvious natural analog in social choice or voting theory. An *average* rule would average the member estimates of each action's expected payoff and choose the action with the highest average opinion. Thus, the notion of collective choice rules includes all familiar ordinally-based rules, but is broader in the sense that it includes rules that can depend on the cardinal values of the estimates as well.

The existing literature on games played by teams of players is extensive and growing, and essentially all focused on experimental investigations of differences between the choice behavior of teams and individuals. There are two identifiable strands depending on whether

---

In fact, if each team has only one member, the team equilibrium of the game will be a quantal response equilibrium (McKelvey and Palfrey (1995, 1996,1997), because there is no collective choice problem. That correspondence generally breaks down for teams with more than one member.

the experimental task was a multi-player game (such as the prisoners dilemma), or a single-player decision problems (such as a lottery choice task or the dictator game). There are far too many papers to describe them all here, and the interested reader should consult the surveys of experimental studies of groups vs. individuals by Charness and Sutter (2012) and Kugler et al. (2012) mentioned earlier.

The focus here is on the experimental studies of games rather than single-agent decision decision tasks, although we note that the broad finding in both classes of studies is that group decision making conforms more closely to economically rational behavior than individual decision making. The range of games studied to date is quite broad. The earliest studies were conducted by social psychologists who were interested in examining alternative hypotheses about social dynamics, based on psychological concepts such as social identity, shared self-interest, greed, and schema-based distrust (fear that the other team will defect). The consistent findings in those studies is that teams defect more frequently than individuals.[6] Bornstein and Yaniv (1998) find that teams are more rational than individuals in the ultimatum game, in the sense that proposers offer less and responders accept less. Elbittar et al. (2011) study several different voting rules in ultimatum bargaining between groups, with less clear results, but also report that proposers learn with experience to offer less. Bornstein et al. (2004) find that teams "take" earlier than individuals in centipede games. In trust games, Kugler et al. (2007), Cox (2002), and Song (2009) find that trustors give less and trustees return less. Cooper and Kagel (2005) find that teams play more strategically than individuals in a limit-pricing signaling game. Charness and Jackson (2005) compare two different voting rules for team choice in a network-formation game that is similar to the stag-hunt coordination game. They report a highly significant effect of the voting rule. Sheremeta and Zhang (2010) observe 25% lower bids by teams than individuals in Tullock contests, where individuals bid significantly above the Nash equilibrium. A similar finding is reported in Morone et al. (2019) for all pay auctions. Group bidding behavior has also been investigated in auctions (Cox and Hayne 2006, Sutter et al. 2009).

We are aware of only two other comparable theoretical models of team behavior in games. Duggan (2001) takes the opposite approach to the present paper, by assuming that members of the team share common and correct beliefs about the distribution of actions of the other teams, but have different payoff functions. The team action is assumed to be the core of a

[6] The experimental social psychology literature on the subject is extensive. See, for example, Insko et al. (1988) and several other studies by Insko and various coauthors. This literature refers to this difference between teams and individuals as a "discontinuity effect". Wildschut and Insko (2007) provide a survey of much of this literature in the context of various explanations that have been proposed.

voting rule. With this approach, existence of team equilibrium typically fails because of non-existence of a core for many voting rules in many environments. Cason et al. (2018) proposes a model specifically for the prisoner's dilemma game that incorporates homogeneous group-contingent inequity-averse preferences and common/correct beliefs. The team decision is determined by a symmetric quantal response equilibrium of the within-team majority-rule voting game, assuming all members have identical inequity averse preferences. In their model, voting behavior in the team decision process becomes more random as team size increases which can lead to behavior further from Nash equilibrium, in contrast to the experimental findings cited above.

We first develop the formal theoretical structure of finite team games in strategic form, provides a proof of the general existence of team equilibrium. The effect of changing team sizes on team equilibrium are illustrated with three examples with majority rule in 2×2 games. These effects can be rather unintuitive: while majority rule guarantees convergence to Nash equilibrium with large teams, the convergence is not necessarily monotone in team size; and individual voting probabilities within a team can be very different from the team mixed strategy equilibrium. Section 4 characterizes a class of collective choice rules for which team response functions satisfy a weak form of stochastic rationality, payoff monotonicity, whereby the probability a team chooses a particular action is increasing in its equilibrium expected payoff. This is a very broad class and includes all generalized scoring rules. We identify stronger conditions that ensure a team's choice probabilities will be ordered by the expected payoffs of the actions. Proposition 1 establishes that, under the plurality collective choice rule, every sequence of team equilibria converges to a Nash equilibrium as the size of teams grow large. Section 5 generalizes this result by showing that it is true for any anonymous scoring rule. We call this the *Nash convergence property.* Section 6 generalizes the framework to extensive form games, establishes existence, and proves that the results about stochastic rationality and Nash convergence extend to extensive form team games. In fact, every convergent sequence of team equilibrium as teams grow large necessarily converges to a *sequential equilibrium* of the game. Team equilibrium in extensive form games are illustrated in Section 7, with a simple signaling game and the 4-move centipede game. Section 8 discusses the results of the paper and points to some possible generalizations and extensions of the framework.

2 Team Games in Strategic Form

A team game is defined as follows. Let $\mathbf{T} = \{1, ..., t, ..., T\}$ be a collection of *teams*, where $t = \{i_1^t, ..., i_j^t, ..., i_{n^t}^t\}$, where i_j^t denotes member j of team t. Each team has a set of available actions, $A^t = \{a_1^t, ..., a_{K^t}^t\}$ and the set of action profiles is denoted $A = A^1 \times ... \times A^T$. The *payoff function* of the game for team t is given by $u^t : A \to \Re$. Given an action profile a, all members of team t receive the payoff $u^t(a)$. A mixed strategy for team t, α^t, is a probability distribution over A^t, and a mixed strategy profile is denoted by α. We denote the expected payoff to team t from using action a_k^t, given a mixed strategy profile of the other teams, by $U_k^t(\alpha) = \sum_{a^{-t} \in A^{-t}} \left[\prod_{t' \neq t} \alpha^{t'}(a^{t'}) \right] u^t(a_k^t, a^{-t})$. For each t, given α, each member i_j^t observes an estimate of $U_k^t(\alpha)$ equal to true expected payoff plus an estimation error term. Denote this estimate by $\hat{U}_{ik}^t = U_k^t(\alpha) + \varepsilon_{ik}^t$, where the dependence of $\widehat{U}_{ik}^t$ on α is understood. We call $\hat{U}_i^t = (\hat{U}_{i1}^t, ..., \hat{U}_{iK^t}^t)$ i's estimated expected payoffs, and $\hat{U}^t = (\hat{U}_1^t, ..., \hat{U}_{n^t}^t)$ is the profile of member estimated expected payoffs in team t. The estimation errors for members of team t, $\{\varepsilon_{ik}^t\}$, are assumed to be i.i.d. draws from a commonly known probability distribution F^t, which is assumed to have a continuous density function that is strictly positive on the real line. We also assume the distribution of estimation errors are independent across teams, and allow different teams to have different distributions. Denote any such profile of estimation error distributions, $F = (F^1, ..., F^T)$, *admissible.* A *team collective choice rule*, C^t, is a correspondence that maps profiles of estimated expected payoffs in team t into a nonempty subest of elements of A^t. That is, $C^t : \Re^{n^t K^t} \to \mathcal{A}^t$, where $\mathcal{A}^t$ is the set of nonempty subsets of A^t. Thus, for any α and ε^t, $C^t(\hat{U}^t) \in \mathcal{A}^t$. Different teams could be using different collective choice rules, and denote $C = (C^1, ..., C^T)$ the profile of collective choice rules.

For any strategic form game $G = [\mathbf{T}, A, \{u^t\}_{t=1}^T]$, and for any admissible F and profile of team choice rules C, call $\Gamma = [\mathbf{T}, A, \{u^t\}_{t=1}^T, F, C]$ a *team game in strategic form.* We assume that team t chooses randomly over $C^t(\hat{U}^t)$ when it is multivalued, according to the uniform distribution. That is, the probability team t chooses a_k^t at $\hat{U}^t(\alpha)$ is given by the function g_M^t defined as:

$$\begin{aligned} g_k^{C^t}(\hat{U}^t) &= \frac{1}{|C^t(\hat{U}^t)|} \text{ if } a_k^t \in C^t(\hat{U}^t) \\ &= 0 \text{ otherwise} \end{aligned} \tag{1}$$

An example of a team choice rule is the *average rule*, which mixes uniformly over the actions with the highest average estimated payoff. That is, let $\widehat{\overline{U_k}}^t = \frac{1}{n^t}\sum_{i \in T} \hat{U}^t_{ik}$ and define $C^t_{ave}(\hat{U}^t) = \{a^t_k \in A^t | \widehat{\overline{U_k}}^t \geq \widehat{\overline{U_l}}^t$ for all $l \neq k\}$ to be the set of actions that maximize $\widehat{\overline{U}}^t$ at given values of α^t and ε^t. This implies a mixed team strategy, g^{ave}, defined by:

$$\begin{aligned} g^{ave}_k(\hat{U}^t) &= \frac{1}{|C^t_{ave}(\hat{U}^t)|} \text{ if } a^t_k \in C^t_{ave}(\hat{U}^t) \\ &= 0 \text{ otherwise} \end{aligned}$$

Another example is *plurality rule*, which chooses the action for which the greatest number of team members estimate to have to highest payoff. That is, define $V^t_k(\hat{U}^t) = \#\{i \in t | \hat{U}^t_{ik} \geq \hat{U}^t_{il}$ for all $l \neq k\}$ and then $C^t_{pl}(\hat{U}^t) = \{a^t_k \in A^t | V^t_k(\hat{U}^t) \geq V^t_{k'}(\hat{U}^t)$ for all $l \neq k\}$. Then g^{pl} is defined by:

$$\begin{aligned} g^{pl}_k(\hat{U}^t) &= \frac{1}{|C^t_{pl}(\hat{U}^t)|} \text{ if } a^t_k \in C^t_{pl}(\hat{U}^t) \\ &= 0 \text{ otherwise} \end{aligned}$$

2.1 Team response functions and team equilibrium

It is important to note that the team choices are generally stochastic (unless C is a constant function), and for any given distribution of other teams' action choices, the distribution of the mixed strategy by team t is inherited from the estimation error distribution via a team collective choice rule. It is this distribution of each team's choices under their team collective choice rule that is the object to which we ascribe equilibrium properties.

Given a team game $\Gamma = [\mathbf{T}, A, \{u^t\}^T_{t=1}, F, C]$, we can define a *team response function* for team t, $P^t : \mathbb{R}^{K^t} \to \Delta A^t$, a function that maps profiles of expected utilities for team actions to a team distribution over actions by taking an expectation of g^{C^t} over all possible realizations of ϵ^t:

$$P^t_k(U^t(\alpha)) = \int_{\epsilon^t} g^{C^t}_k(\hat{U}^t) dF^t(\epsilon^t) \tag{2}$$

where $g^C_k(\hat{U}^t)$ is defined as in equation (1). An equilibrium of a team game is a fixed point of P.

Definition 1. *A* ***team equilibrium*** *of the team game* $\Gamma = [\mathbf{T}, A, \{u^t\}^T_{t=1}, F, C]$ *is a mixed* *strategy profile* $\alpha = (\alpha^1, ..., \alpha^T)$ *such that, for every* t *and every* $k = 1, ..., K^t$, $\alpha_k^t = P_k^t(U^t(\alpha))$.

Theorem 1. *For every* Γ *and* C *a team equilibrium exists.*

Proof. This follows in a straightforward way. With the admissibility assumptions on F, the integral on the right hand side of equation (2) is well-defined for all admissible F and P_k^t is continuous in α. Brouwer's fixed point theorem then implies existence. □

It is one thing to define aggregation rules in the abstract, but quite another to model how such an aggregation rule might be implemented within a team. In a sense, team games nest a game within a game, but the definition above models the game within a game in reduced form, via the function g^C. We discuss later possible communication mechanisms that might correspond to different team choice rules. One possibility for the average rule would be a mechanism in which each player announces $\hat{U}_i^t$ to the other members of the team, and the team just takes the average, and then chooses the action that maximizes the average announced estimated expected payoff. Because it is a common value problem for the team, there is an implicit assumption of sincere reporting, and the team is choosing optimally. One might also conjecture that free form communication within a group would lead to group choice approximating the average rule, with a dynamic similar to what has been theoretically modeled as group consensus formation (McKelvey and Page 1986 and others).[7] Alternatively, the group might decide using a voting rule, such as plurality, to implement plurality rule. Before moving to the next section that illustrates the model with examples, we prove following proposition.

Proposition 1. *Consider an infinite sequence of team games,* $\{\Gamma_m\}_{m=1}^{\infty}$ *such that (1)* $A_m^t = A_{m'}^t = A^t$ *for all* t, m, m'*; (2)* $u_m^t = u_{m'}^t = u^t$ $\forall t, m, m'$*; (2)* n_m^t *is odd for all* m, t*; (3)* $n_{m+1}^t > n_m^t$ *for all* m, t*; and (4)* C *is the plurality team choice rule for all* m, t*. Let* $\{\alpha_m\}_{m=1}^{\infty}$ *be a convergent sequence of team equilibria where* $\lim_{m\to\infty} \alpha_m = \alpha$*. Then* α *is a Nash equilibrium of the strategic form game* $[\mathbf{T}, A, u]$.

[7]The formal connection between the average rule and the consensus formation literature is not direct, as that approach assumes the members of the group share a common prior. Our model of estimation errors does not specify a common prior distribution of expected payoffs from actions. Rather individual beliefs are modeled simply as unbiased point estimates of an unknown true value.

Proof. First, the statement is not vacuous because a team equilibrium of Γ_m is guaranteed to exist for all m, and there always exists at least one convergent sequence of equilibria $\{\alpha_m\}_{m=1}^{\infty}$ by the Bolzano-Weierstrass Theorem. Suppose α is not a Nash equilibrium of the strategic form game $[T, A, u]$. Then it must be the case that for some t, there exists a pair of actions $a_k^t, a_l^t \in A^t$ such that $\alpha_k^t > 0$ but $U^t(a_l^t; \alpha) - U^t(a_k^t; \alpha) > 0$.

Because $\lim_{m\to\infty} \alpha_m = \alpha$, there exists $\delta > 0$, $\overline{m} < \infty$ such that $U^t(a_l^t; \alpha_m) - U^t(a_k^t; \alpha_m) > \delta$ for all $m > \overline{m}$. For any t and any $m > \overline{m}$ denote by $q_{mkl}^t(\alpha_m)$ that probability that $\widehat{U}_{il}^t - \widehat{U}_{ik}^t > 0$. Since $U^t(a_l^t; \alpha_m) - U^t(a_k^t; \alpha_m) > \delta$ for all $m > \overline{m}$, we have $m > \overline{m}$ implies $q_{mkl}^t(\alpha_m) > H^t(\delta)$, where $H^t(\delta)$ is the cumulative distribution function of $\varepsilon_{ik}^t - \varepsilon_{il}^t$. Because ε_{ik}^t and ε_{il}^t are i.i.d draws from F^t, it follows that $H^t(0) = \frac{1}{2}$ and hence $H^t(\delta) > \frac{1}{2}$. Thus, for each $m > \overline{m}$, at α_m the probability that a majority of members of team t estimate that action α_k^t has a higher expected payoff than action α_l^t is bounded above by:

$$Q_{mkl}^t(\alpha_m) == \sum_{k=\frac{n+1}{2}}^{n_m^t} \binom{n_m^t}{k} (1 - H^t(\delta))^k (H^t(\delta))^{n_m^t - k}$$

Since $H^t(\delta) > \frac{1}{2}$, we have $\lim_{m\to\infty} Q_{mkl} = 0$. But if the probability more members rank α_k^t higher than α_l^t goes to 0, then since all the errors are i.i.d, the probability that more members rank α_k^t the highest also converges to 0, which implies $\lim_{m\to\infty} \alpha_k^t = 0$, a contradiction.

□

Accordingly, we say that plurality rule has the *Nash convergence property.* The intuition behind this result is straightforward. First, the existence of at least one convergent sequence of team equilibria follows immediately from the Bolzano-Weierstrass theorem. The fact that the limit points of such sequences must be Nash equilibria of the underlying strategic form game follows from the fact that in the limit, the expected fraction of members who rank an action first is ordered by the equilibrium expected payoff of that action. This implies that for any action that is not optimal in the limit, the probability that action wins a plurality of votes goes to zero, and hence only optimal replies can win a plurality of votes in the limit. A different proof, but one that also relies on the law of large numbers, works for the average rule as well. An interesting general problem is to characterize precisely the set of collective choice rules that satisfy the Nash convergence property.

The result for plurality rule illustrates two different effects on team equilibrium as a result of changing team sizes. One effect, which is quite intuitive, is *the reinforcement effect.* For any fixed strategy profile α, as n grows, the ranking of the estimated payoffs of different actions by individual members of a team converges on average to the "true" ranking of the actions' expected payoffs.[8] Thus, for many rules, such as plurality rule and the average rule, the choice of large teams will correspond to the choice the team would make if they all shared common and correct beliefs about a. This effect is especially easy to see in the case of majority rule when $|A^t| = 2$. In this case, suppose, for some α, $U_1^t(\alpha) > U_2^t(\alpha)$. Then for any admissible F, the probability that $\widehat{U}_{i1}^t(\alpha) > \widehat{U}_{i2}^t(\alpha)$ for an individual member of team t, which we denote by $p_1^t(\alpha)$ is greater than $\frac{1}{2}$. Thus, if n increases, *fixing* α, then under majority rule, $P_1^t(U^t(\alpha)) = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k}(p_1^t(\alpha))^k(1-p_1^t(\alpha))^{n-k} > \frac{1}{2}$ and increases in n monotonically, eventually converging to 1, because an increase in team size increases the likelihood of a majority consensus for a_1^t, and this consensus is guaranteed in the limit as n increases without bound.

Of course the analysis is more complicated than this, because, one cannot take α as fixed when n changes, because there is also a second effect, *the equilibrium effect.* As n changes, the team equilibrium α^* changes, so, even in the case where $|A^t| = 2$, it will typically be the case that for $n \neq n'$, $U_1^t(\alpha_n^*) - U_2^t(\alpha_n^*) \neq U_1^t(\alpha_{n'}^*) - U_2^t(\alpha_{n'}^*)$, so $p_1^{*t}(\alpha_n^*) \neq p_1^{*t}(\alpha_{n'}^*)$, which feeds back and affects each team's mixed strategy response. The equilibrium effect can produce some unintuitive consequences for some games and some voting rules, and it's possible that the equilibrium effect can dampen or even work in the opposite direction of the reinforcement effect. The equilibrium effect also arises when the size of one team changes while the size of all the other teams remain fixed, which can result in different equilibrium team mixing probabilities for all the teams.

One can also think of the Nash convergence result for plurality rule as being related to a sincere voting version of the Condorcet jury theorem and obviously specializes to the case where $|A^t| = 2$ for all t. The next part of the paper gives some examples to illustrate team equilibrium under majority rule in 2×2 games, and show the variety of ways that team equilibrium can be effected by sizes of the two teams. In particular these examples shed light on the different roles of the reinforcement and equilibrium effects.[9]

[8] In fact, this reinforcement, or consensus, effect is true in a stronger sense. The average estimated payoff of each action converges in probability to the true expected payoff of the action.

[9] Exploring the consequences of voting rules with more than 2 actions for a team introduce additional aggregation issues that we do not wish to consider here.

3 Team Size Effects in 2×2 Games with Majority Rule

The model is easiest to illustrate in the simple case of 2×2 games. Let $T = \{1, 2\}$, $A^t = \{a_1^t, a_2^t\}$, and the set of action profiles is $A = A^1 \times A^2$. Denote by α^t the probability that team t decides to use action a_1^t. Then member i of team t estimates the expected payoff if team t chooses action a_k^t when team $-t$ uses a mixed strategy α^{-t} by $\widehat{U}_{ik}^t = U_k^t(\alpha^{-t}) + \varepsilon_{ik}^t$ where $U_k^t(\alpha^{-t}) = \alpha^{-t} u^t(a_k^t, a_1^{-t}) + (1 - \alpha^{-t}) u^t(a_k^t, a_2^{-t})$. Because there are only two actions for each team, the notation can be simplified by letting $\varepsilon_i^t = \varepsilon_{i1}^t - \varepsilon_{i2}^t$ denote the difference in estimation errors for individual i on team t and denote by H^t the distribution of the *difference* of these estimation errors, ε_i^t. Given that each of the payoff estimation errors, ε_{i1}^t and ε_{i2}^t, are distributed according to an admissible error distribution, H^t is also admissible and symmetric around 0. That is, for all $x \in \Re$, $H^t(x) = 1 - H^t(-x)$, implying $H(0) = \frac{1}{2}$. Thus, we can write:

$$\begin{aligned} \Delta \widehat{U}_i^t &\equiv \widehat{U}_{i1}^t - \widehat{U}_{i2}^t \\ &= U_1^t(\alpha^{-t}) - U_2^t(\alpha^{-t}) + \varepsilon_i^t. \end{aligned}$$

A team choice rule, C^t, maps each profile of individual estimation error differences, $\varepsilon^t = (\varepsilon_1^t, ..., \varepsilon_{n^t}^t)$ into the nonempty subsets of A^t, and g_C^t randomly selects one of these choices with equal probability.

In the examples below we consider the case where C is majority rule. Let $K^+(\varepsilon^t; \alpha^{-t}) = |\{i \in t| \ \Delta \widehat{U}_i^t > 0\}|$ and let $K^-(\varepsilon^t; \alpha^{-t}) = |\{i \in t| \Delta \widehat{U}_i^t < 0\}|$, then majority rule is defined by:λ

$$\begin{aligned} C_{mr}^t(\varepsilon^t; \alpha^{-t}) &= \{a_1^t\} \text{ if } K^+(\varepsilon^t; \alpha^{-t}) > K^-(\varepsilon^t; \alpha^{-t}) \\ &= \{a_2^t\} \text{ if } K^+(\varepsilon^t; \alpha^{-t}) < K^-(\varepsilon^t; \alpha^{-t}) \\ &= \{a_1^t, a_2^t\} \text{ if } K^+(\varepsilon^t; \alpha^{-t}) = K^-(\varepsilon^t; \alpha^{-t}) \end{aligned}$$

That is, the team choice is the action for which a majority of team members estimate to have a higher expected payoff, with ties broken randomly. We next illusstrate team size effects under majority rule in three different kinds of 2×2 games.

3.1 The Prisoner's Dilemma played by teams

Consider the family of PD games displayed in Table 1, where the two parameters $x > 0$ and $y > 0$ are, respectively the payoff gain from defecting if the other player cooperates and the payoff gain from defecting if the other player defects.

Table 1: Prisoner's dilemma game

	Column Team (2)	
Row Team (1)	Cooperate(C)	Defect(D)
Cooperate(C)	$5, 5$	$3 - y, 5 + x$
Defect(D)	$5 + x, 3 - y$	$3, 3$

Suppose that both teams have n members and the same distribution of estimation error differences, H, and let α_n^* be a symmetric team equilibrium probability of choosing D in the game.

If $x = y$ is analysis is straightforward and intuitive. The expected payoff difference between Defect and Cooperate for either team is $U_D^t(\alpha_n^*) - U_C^t(\alpha_n^*) = y > 0$.

Hence, for any player, i, on either team the probability their estimated expected payoff difference between D and C, $\Delta\widehat{U}_i^t \equiv \widehat{U}_{iD}^t - \widehat{U}_{iC}^t$, is greater than 0, is given by:

$$\Pr\{\Delta\widehat{U}_i^t > 0)\} = H(y) \tag{3}$$

Because this probability is independent of α_n^*, and is the same for players on both teams, call it, p^*, which is simply the probability that any individual on a team votes for D. Note that for any admissible H, $p^* > \frac{1}{2}$. Letting C be majority rule, given any p^* we can compute α_n^*, which is simply the probability that more than $\frac{n}{2}$ members estimate that D yields a higher expected payoff than C:

$$\alpha_n^* = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} (p^*)^k (1 - p^*)^{n-k} \tag{4}$$

Observe that if $n = 1$, this reduces to $\alpha_n^* = p^*$, which is a fixed point of the the function $p^* = H(U_D^t(p^*) - U_D^t(p^*))$. $\lim_{n\to\infty}(\alpha_n^*) = 1$. The proposition below shows that this convergence is monotone in n, i.e., $\alpha_{n+2}^* > \alpha_n^*$.

Proposition 2. $\alpha_{n+2}^* > \alpha_n^*$

Proof. Suppose there are $n = 2m+1$ players and p^* is the probability that each team member defects, and compare this to the case where there are $n+2 = 2m+3$ players. The only case where team's decision changes when two new members join the team is if the original team's decision is determined by a margin of exactly 1 vote. Otherwise, adding two members does not affect team's decision. Then, there are two cases to consider:

(1) A majority of exactly $m+1$ defects, and the two additional members cooperate.

(2) A minority of exactly m defects, and the two additional members defect.

The probability of the first event is $\binom{n}{m+1}(p^*)^{m+1}(1-p^*)^{m+2}$, while the probability for the second event is $\binom{n}{m}(p^*)^{m+2}(1-p^*)^{m+1}$. Since $p^* > \frac{1}{2}$,

$$\begin{aligned} \frac{\binom{n}{m+1}(p^*)^{m+1}(1-p^*)^{m+2}}{\binom{n}{m}(p^*)^{m+2}(1-p^*)^{m+1}} &= \frac{1-p^*}{p^*} \\ &< 1. \end{aligned}$$

Thus, $\binom{n}{m+1}(p^*)^{m+1}(1-p^*)^{m+2} < \binom{n}{m}(p^*)^{m+2}(1-p^*)^{m+1}$ and hence $\alpha^*_{n+2} > \alpha^*_n$. □

Figure 1 displays the team equilibrium as a function of n for the value of and $x = y = 2$ and $H(x) = \frac{1}{1+e^{-.5x}}$.

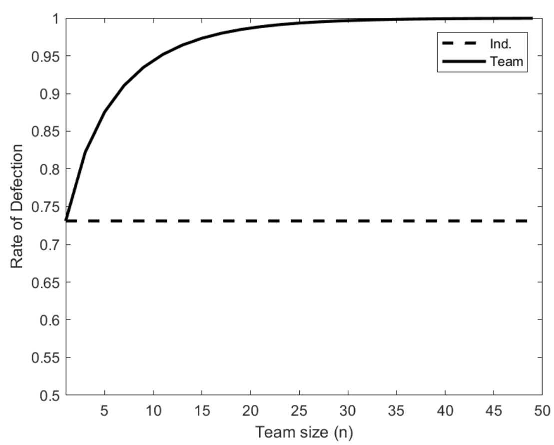

Figure 1: Team equilibrium as a function of n (x=y=2)

Thus, in PD games where $x = y$ the changes in the team equilibrium as n increases are entirely due to the reinforcement effect because p^* is independent of n. Adding new members to a team just reinforces the tendency for the majority outcome to be D. However, if $x \neq y$

this will not generally be the case, as there will be an equilibrium effect of changing n, in addition to the reinforcement effect, and these go in opposite directions. If $x > y$, the effect of increasing n is buffered by the counterveiling effect that increasing α_n^* leads to a decrease in p^*.

Consider in particular the case where the size of the row team is fixed at 1 and the size of the column team, n, is variable. In this case, for any n the equilibrium is of the form (q_n^*, α_n^*), where q_n^* denotes the single row team member's defect probability, which depends on the of the column team's defect probability, α_n^*. Note also that in general it will be the case that the defect probability of individual members of the column team, q_n^*, are no longer independent of n, since q_n^* will typically vary with n.

Thus, the equilibrium conditions are interdependent. Formally an equilibrium solves the following two equations:

$$\begin{aligned} q_n^* &= H(x + (y-x)\alpha_n^*) \\ \alpha_n^* &= \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} (p_n^*)^k (1-p_n^*)^{n-k} \end{aligned}$$

where $p_n^* = H(x + (y-x)q_n^*)$ is the probability a member of the column team defects. If $x > y$, we get the surprising result that q_n^* *decreases* in n, while, as before, α_n^* increases with n. In the limit, for large n, we get:

$$\begin{aligned} \lim_{n\to\infty} (\alpha_n^*) &= 1 \\ \lim_{n\to\infty} (q_n^*) &= H(y) \end{aligned}$$

Figure 2 displays the team equilibrium (q_n^*, α_n^*) as a function of n for the value of $x = 5$, $y = 2$ and $H(x) = \frac{1}{1+e^{-.5x}}$.

3.2 The "Weak" Prisoner's Dilemma

We next examine a class of team games where two teams, each of size n (odd), play a variation on prisoner dilemma game displayed in Table 1, which we call the *Weak Prisoners Dilemma*. While (D,D) is not a dominant strategy equilibrium, D is strictly dominant for the row player and is the unique solution to the game in two stages of iterated strict dominance.

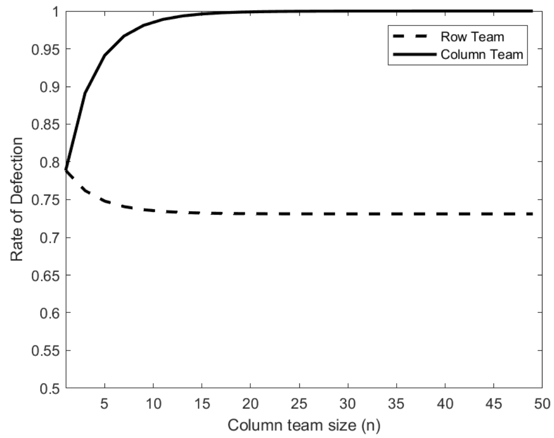

Figure 2: Team equilibrium as a function of n (x=5, y=2)

The column player is best off matching the row player's action choice. We characterize the team equilibrium for this class of games under majority rule. We have three free parameters to the game, which is displayed in Table 2. The first two, x and y, are the same as above. The third parameter, z, is the payoff gain the column player gets from cooperating if the row player cooperates.

Table 2: Weak Prisoner's Dilemma (WPD) game

	Column Team (2)	
Row Team (1)	Cooperate(C)	Defect(D)
Cooperate(C)	$5, 5$	$3 - y, 5 - z$
Defect(D)	$5 + x, 3 - y$	$3, 3$

Let $\alpha_n^* = (\alpha_n^{*1}, \alpha_n^{*2})$ be the team equilibrium probabilities of choosing D in the weak prisoner's dilemma game for row (team 1) and column (team 2). Following similar steps as in the PD example with $x = y$, we have, for the row team:

$$p_n^{*1} = H(y)$$

and

$$\alpha_n^{*1} = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} (p_n^{*1})^k (1 - p_n^{*1})^{n-k}$$

For a player on the column team, the probability of voting for D, p^{*2}, varies with n, as it depends directly on α_n^*, and is given by

$$p_n^{*2} = H(\alpha_n^{*1}(y+z) - z) \tag{5}$$

and the probability that the column team chooses D is:

$$\alpha_n^{*2} = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} (p_n^{*2})^k (1-p_n^{*2})^{n-k} \tag{6}$$

For any values of n, y, z, and λ, this system of equations can be solved by first substituting equation (3) into equation (4) to obtain α_n^{*1}, then substituting this into equation (5) to obtain p^{*2}, and finally substituting this into equation (6) to obtain α_n^{*2}.[10]

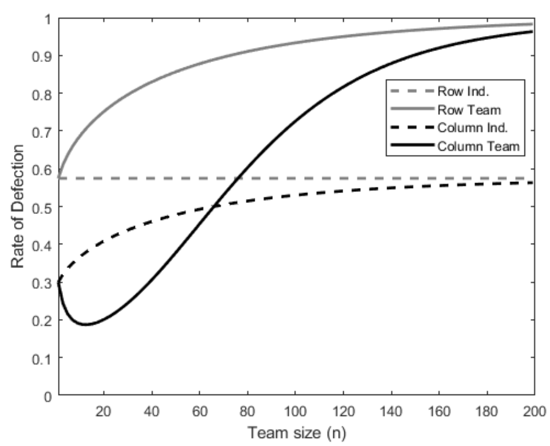

Figure 3: Team equilibrium as a function of n (x=y=1, z=8)

Figure 3 illustrates how the equilibrium voting and team choice probabilities $(p_n^{*1}, \alpha_n^{*1}, p_n^{*2}, \alpha_n^{*2})$ vary with n, for the parameters $x = y = 1, z = 8$ and $H(x) = \frac{1}{1+e^{-.3x}}$. Notice that for relatively small teams sizes (< 20) the column team becomes *more* cooperative as it grows. This results from a combination of the reinforcement effect (since $p_n^{*2} < .5$) and the equilibrium effect (since α_n^{*1} is increasing).

[10]If $x \neq y$ then a fixed point would need to be solved for, rather than using this simple iterative procedure.

3.3 Asymmetric Matching Pennies Games

We examine a class of simple team games where two teams, each of size n (odd), play a 2×2 game with a unique (mixed-strategy) Nash equilibrium, which we refer to a asymmetric matching pennies (AMP) games. The payoffs are displayed in table 3, where $a > c, b < d, w < x, y > z$:

Table 3: Asymmetric patching pennies game

	Column	
Row	L	R
U	a, w	b, x
D	c, y	d, z

In a team equilibrium α_n^*, for any distribution of estimation errors, F, the probability that an individual player on the row team estimates that U is better than D is equal to:

$$p_n^{*1} = H(a(1-\alpha_n^{*2}) + b\alpha_n^{*2} - c(1-\alpha_n^{*2}) - d\alpha_n^{*2}) \tag{7}$$

and the probability that an individual player on the column team estimates that R is better than L is equal to:

$$p_n^{*2} = H(z(1-\alpha_n^{*1}) + x\alpha_n^{*1} - y(1-\alpha_n^{*1}) - w\alpha_n^{*1}) \tag{8}$$

These are the voting probabilities. As in the previous examples, given p^{*1} and p^{*2} we can compute α_n^{*1} and α_n^{*2}, as the probability that at least $\frac{n+1}{2}$ members of the respective team estimate that U (R) yields a higher expected payoff than D (L). Hence:

$$\alpha_n^{*1} = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} (p^{*1})^k (1-p^{*1})^{n-k} \tag{9}$$

$$\alpha_n^{*2} = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} (p^{*2})^k (1-p^{*2})^{n-k} \tag{10}$$

The team equilibrium is obtained by solving equations (7), (8), (9), and (10) simultaneously for p^{*1}, p^{*2}, α_n^{*1}, and α_n^{*2}.

We know from Theorem 1 that for any admissible F, $\lim_{n\to\infty} \alpha_n^* = (\frac{y-z}{y-z+x-w}, \frac{a-c}{a-c+d-b})$.An

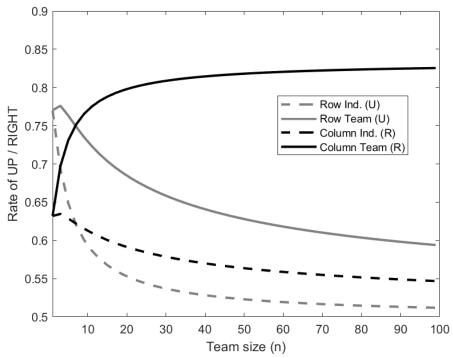

(a) Matching Pennies with a=5, x=1

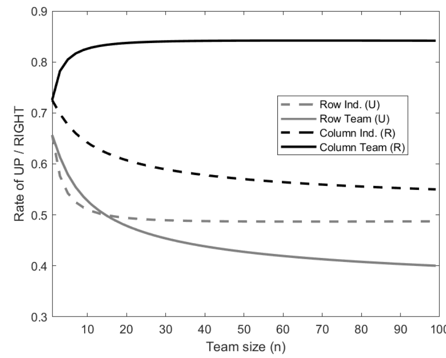

(b) Matching Pennies with a=5, x=2

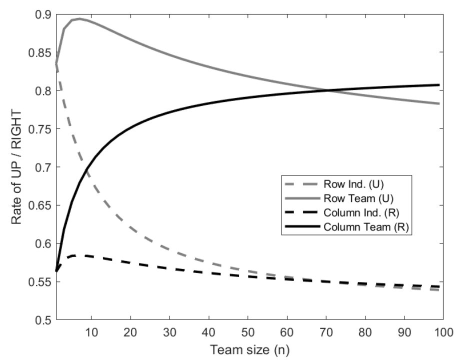

(c) Matching Pennies with a=5, x=0.5

Figure 4: Team Equilibrium in Asymmetric Matching Pennies

interesting feature of this game, or any other 2×2 game with a unique mixed strategy, is that in the limit, *individual team members are voting randomly.* That is, $\lim_{n\to\infty} p_n^{*1} = \lim_{n\to\infty} p_n^{*2} = \frac{1}{2}$, which is necessarily the case because the expected payoffs of the two strategies are equal in the Nash equilibrium limit. Thus, in this class of games, the equilibrium effect completely dominates the reinforcement effect, in the sense that, *for all* n, $U^t(a_1^t;\alpha_n^*) > U^t(a_{2'}^t;\alpha_n^*)$, and hence $p_n^{*t} > \frac{1}{2}$ for all n for both teams, yet $\lim_{n\to\infty} \alpha_n^* \neq (1,1)$.

The team equilibrium and the equilibrium voting probabilities are displayed in Figure 4 for the parameters $w = b = c = z = 0, d = y = 1, a = 5$, $x = 0.5, 1, 2$, and $H(x) = \frac{1}{1+e^{-x}}$.

4 Payoff-monotone and Rank-dependent Team Response Functions

Individual estimated expected payoffs possess two intuitive responsiveness properties that follow from the assumed structure of the estimation errors. First, the probability a member of team t ranks action a_k^t as having the highest estimated possible payoff is increasing in $U_k^t(\alpha)$, the "true" equilibrium expected payoff of action k, ceteris paribus, a condition we call *payoff monotonicity*. Related to this, and following from the i.i.d. assumption, a member of team t ranks action a_k^t's estimated expected payoff as higher than action a_l^t's estimated expected payoff if and only if $U_k^t(\alpha) > U_l^t(\alpha)$, a condition we call *rank dependence*. [11]

The more general question that we address in this section involves identifying conditions on team collective choice rules such that team response function will be stochastically responsive in a similar way that individual responses are. That is, we ask: For what class of team collective choice rules will team response functions satisfy payoff monotonicity and rank dependence? While these are not properties that are satisfied by every possible team collective choice rule, we show that payoff monotonicity requires only two weak assumptions, unanimity and positive responsivenenss, and does not require any extra conditions such as neutrality or anonymity. On the other hand, rank dependence holds only for a more restricted class of neutral collective choice rules. Many non-neutral collective choice rules, such as those that give a status quo advantage to an action, will fail to satisfy rank dependence. We show that rank dependence is satisfied for $K^t = 2$ with any collective choice rule satisfying unanimity, positive responsiveness, and neutrality, and for $K^t > 2$ with plurality rule or weighted average rules.

4.1 Payoff Monotonicity

For any team game, Γ, P^t depends on both the distribution of member's estimation errors, F^t, and the team collective choice rule, C^t. In this section we identify conditions on C^t that are sufficient for P^t to be payoff monotone for all admissible F^t. Specifically, we require team collective choice rules to satisfy two axioms: unanimity and positive responsiveness.

The first condition, unanimity, simply states that if all members of the team estimate

[11] Furthermore, the estimated expected payoffs are continuous in U^t and have full support on the real line. Thus, they have properties similar to the choice probabilities in a quantal response equilibrium.

that a_k^t has the highest expected utility, then it is uniquely chosen by C^t.[12]

Definition 2. *A team collective choice rule C^t satisfies **Unanimity** if:*

$$\begin{aligned} \widehat{U}_{ik}^t &> \widehat{U}_{il}^t \text{ for all } i \in t \text{ and for all } a_l^t \neq a_k^t \\ &\Rightarrow \\ C^t(\widehat{U}^t) &= \{a_k^t\} \end{aligned}$$

In addition to using this axiom to prove payoff monotonicity, it also guarantees that team response functions are interior, in the sense that every action is chosen with positive probability. The second axiom, positive responsiveness, requires that the team choice responds positively to all members of a team increasing their estimated expected payoff of an action, keeping all other estimated expected payoffs the same. The following definition is used in the statement of the axiom.

Definition 3. *A profile $\tilde{U}^t$ of member estimated expected utilities is a monotonic transformation of $\hat{U}^t$ with respect to action a_k^t if, for all members $i \in t$, we have*

1. $\tilde{U}_{ik}^t \geq \hat{U}_{ik}^t$

2. $\tilde{U}_{il}^t = \hat{U}_{il}^t$ for all $l \neq k$

Definition 4. *A team collective choice rule C^t satisfies **Positive Responsiveness** if for all a_k^t, $\hat{U}^t$ such that $a_k^t \in C^t(\hat{U}^t)$, then for all monotonic transformations $\tilde{U}^t$ of $\hat{U}^t$ with respect to action a_k^t, it is the case that $a_k^t \in C^t(\tilde{U}^t) \subseteq C^t(\hat{U}^t)$.*

This definition of Positive Responsiveness is essentially a cardinal version of the usual definition of positive responsiveness from the social choice literature (see Moulin 1988). It says that if an action a_k^t is chosen at some profile of estimated expected utilities, and all team members' estimates of the expected utility of that action weakly increase, ceteris paribus, then a_k^t must still be chosen, and no new actions can be added to the choice set.

[12]For the "standard" case of games played by one-person teams, unanimity implies that if $n = 1$ then every team equilibrium is equivalent to a *quantal response equilibrium (QRE)* of the strategic form game, $[T, A, u]$.

Many collective choice rules satisfy positive responsiveness. Obviously, any weighted average rule, where the team choice corresponds to the action with the highest weighted average of individual members estimates, is positively responsive. Plurality rule also clearly satisfies this condition. Plurality rule is a special case of general scoring rules, which are defined over ordinal rankings of actions. We show next that all scoring rules satisfy positive responsiveness. Formally:

Definition 5. *An* ***Individual i Scoring Function*** $S_i^t : A^t \times \mathbb{R}^{K^t} \to \mathbb{R}^{K^t}$*, is a function defined that* $S_{ik}^t(\hat{U}_i^t) = s_{im}$ *whenever* $|\{a_l^t \in A^t : \hat{U}_{il}^t > \hat{U}_{ik}^t\}| = m-1$*, for some given set of* K^t *scores* $s_{i1}^t \geq s_{i2}^t \geq ... \geq s_{iK^t}^t \geq 0$ *such that* $s_{i1}^t > s_{iK^t}^t$*. Given a profile of individual scoring functions,* $(S_1^t, ...S_{n^t}^t)$*, The* ***Team Scoring Function*** $S^t : A^t \times \mathbb{R}^{n^t K^t} \to \mathbb{R}^{K^t}$ *is defined by* $S_k^t(\hat{U}^t) = \sum_{i=1}^{n^t} S_{ik}^t(\hat{U}_i^t)$.

The individual scoring functions depend only on a team member's *ordinal* estimated expected utilities of the alternatives, and each alternative is awarded a score that is weakly increasing in its estimated expected utility rank by that team member. The individual scores for each team member are then summed to arrive at a total score for the team, and the alternatives with the highest total score are chosen. This is more general than the standard definitions of scoring rules in the social choice literature, where anonymity is typically imposed (i.e., all the individual scoring functions are the same). We do not require anonymity.

Definition 6. *A team collective choice rule* C^t *is a* ***Scoring Rule*** *if there exists a profile of individual scoring functions,* $(S_1^t, ...S_{n^t}^t)$*, such that for all* $a_k^t \in A^t$ *and for all* $\hat{U}^t \in \Re^{K^t n^t}$*,* $a_k^t \in C^t(\hat{U}^t)$ *if and only if* $S_k^t(\hat{U}^t) \geq S_l^t(\hat{U}^t)$ *for all* $l \neq k$*.*

Examples of common scoring rules include: plurality rule, where $s_{i1}^t = 1$ and $s_{im}^t = 0$ for all $i \in t$ and for all $m > 1$; dictatorship, where, for some $i^* \in t$, $s_{i^*1}^t = 1$ and $s_{i^*m}^t = 0$ for all $m > 1$ and $s_{j1}^t < \frac{1}{n+1}$ for all $j \neq i^*$; and Borda count, where $s_{im} = K^t - m$ for all $i \in t$ and for all m. We note that in our framework, every member of every team almost always has a strict order over the K^t actions (i.e., $\hat{U}_{il}^t \neq \hat{U}_{ik}^t$ for all l, k, t, i with probability one), so ties in an individual member's ordinal rankings are irrelevant.

Proposition 3. *All scoring rules satisfy positive responsiveness.*[13]

[13] If $s_{i1}^t > s_{i2}^t$ for all $i \in t$, then the scoring rule also satisfies unanimity.

Proof. Positive responsiveness follows from the fact that the value of the team score function evaluated at any alternative is weakly increasing in that alternative's estimated expected utility for each team member, and weakly decreasing in every other alternative's estimated expected utility. □

We can now state the main result of this subsection.

Theorem 2. *If F^t is admissible and C^t satisfies unanimity and positive responsiveness then P^t satisfies payoff monotonicity.*

Proof. Let C^t satisfy positive responsiveness and unanimity and F^t admissible. Suppose that $U^t_k - U'^t_k = \delta > 0$, and $U^t_l = U'^t_l$, $\forall l \neq k$. Then for all realizations of the estimation errors ϵ^t, we have that $U^t + \epsilon^t$ is a monotonic transformation of $U'^t + \epsilon^t$ with respect to a^t_k. So by positive responsiveness of C^t we have that if $a^t_k \in C^t(U'^t + \epsilon^t)$ then $a^t_k \in C^t(U^t + \epsilon^t)$, and if $a^t_l \in C^t(U^t + \epsilon^t)$ then $a^t_l \in C^t(U'^t + \epsilon^t)$. So $g^{C^t}_k(U^t + \epsilon^t) \geq g^{C^t}_k(U'^t + \epsilon^t)$ for all ϵ^t, and therefore $P^t_k(U^t) \geq P^t_k(U'^t)$.

To show the strict inequality, $P^t_k(U^t) > P^t_k(U'^t)$, we show that there exists a region $\beta \subset \mathbb{R}^{K^t \times N^t}$ with positive measure such that if $\epsilon^t \in \beta$, then $g^{C^t}_k(U^t + \epsilon^t) > g^{C^t}_k(a^t_k; U'^t + \epsilon^t)$. In particular, unanimity of C^t is used as follows to construct β such that if $\epsilon^t \in \beta$, then $g^{C^t}_k(U^t + \epsilon^t) = 1 > g^{C^t}_k(U'^t + \epsilon^t) = 0$. That is, such that a^t_k is uniquely chosen under $U^t + \epsilon^t$, and not chosen under $U'^t + \epsilon^t$.

Let $\tilde{U}^t$ be an expected utility profile such that all team members strictly prefer some action a^t_l to action a^t_k, all members prefer a^t_l to all other actions a^t_m (i.e., all members rank a^t_k second), and for all members we have $\tilde{U}^t_l - \tilde{U}^t_k = \frac{\delta}{2}$. Define

$$\beta = \{\tilde{U}^t - U'^t + \xi : \xi_k \in (0, \frac{\delta}{4}), \xi_l \in (-\frac{\delta}{4}, 0), \xi_m < 0\}$$

Then if $\epsilon^t \in \beta$, by unanimity we have $C^t(U'^t + \epsilon^t) = \{a^t_l\}$ and $C^t(U^t + \epsilon^t) = \{a^t_k\}$, so $g^{C^t}_k(U^t + \epsilon^t) = 1 > g^{C^t}_k(U'^t + \epsilon^t) = 0$. Since β has positive measure, and ϵ^t has full support, we have that $P^t_k(U^t) > P^t_k(U'^t)$, as desired.

$\square$

4.2 Rank Dependence

In this section we show that P^t satisfies rank dependence for $K^t = 2$ with any collective choice rule satisfying unanimity, positive responsiveness and neutrality, and for $K^t > 2$ with plurality rule and weighted average rules. Neutrality is an essential property for proving that team response functions satisfy rank dependence. Informally a neutral team collective choice rule is one that is not biased against or in favor of any particular action. This is analogous to the neutrality axiom from the social choice literature. Formally:

Definition 7. *A team collective choice rule C^t satisfies **Neutrality** if, for all a_k^t, $\hat{U}^t$, for all permutation functions, $\psi : A^t \to A^t$ of team actions,*

$$\begin{array}{rcl} a_k^t & \in & C^t(\hat{U}_1^t, ..., \hat{U}_{K^t}^t) \\ & \Longleftrightarrow & \\ a_{\psi(k)}^t & \in & C^t(\hat{U}_{\psi(1)}^t, ..., \hat{U}_{\psi(K^t)}^t). \end{array}$$

Violations of neutrality of C^t lead to straightforward violations of regularity of P^t. For example, collective choice rules that favor a status quo action will typically lead to such a violation. Consider $K^t = 2$ and a choice rule that selects action 1 if and only if all team members estimate its expected utility to be greater than that of action 2, and selects action 2 otherwise. For any admissible F^t, if the size of the team is large enough, a team using this choice rule will select action 2 more often than action 1 even when $U_1^t(\alpha) > U_2^t(\alpha)$.

Neutrality, along with admissibility of F^t, imply that when the expected payoffs of team actions are permuted, the team probabilities are permuted. This further implies an important property that when two actions have equal expected payoffs, the team must play these actions with equal probability.

Lemma 1. *If F^t is admissible, C^t satisfies neutrality, and $U_k^t = U_l^t$, then $P_k^t(U^t) = P_l^t(U^t)$.*

Proof. For any permutation of team actions ψ, say that $U^t = (U_1^t, ..., U_{K^t}^t)$ and $U^{t,\psi} = (U_{\psi(1)}^t, ..., U_{\psi(K^t)}^t)$. By neutrality and i.i.d estimation errors across team members, we have that

$$\begin{aligned}
P_k^t(U^{t,\psi}) = \sum_{a_k^t \in B \subseteq A^t} \frac{1}{|B|}\text{Prob}[B = & C^t((U^t(a_{\psi(1)}^t) + \epsilon_{1,1}^t, ..., U^t(a_{\psi(K^t)}^t) + \epsilon_{1,K^t}^t), ... \\
& ..., (U^t(a_{\psi(1)}^t) + \epsilon_{T,1}^t, ..., U^t(a_{\psi(K^t)}^t) + \epsilon_{T,K^t}^t))]
\end{aligned} \tag{11}$$

$$\begin{aligned}
= \sum_{a_k^t \in B \subseteq A^t} \frac{1}{|B|}\text{Prob}[B = & C^t((U^t(a_{\psi(1)}^t) + \epsilon_{1,\psi(1)}^t, ..., U^t(a_{\psi(K^t)}^t) + \epsilon_{1,\psi(K^t)}^t), ... \\
& ..., (U^t(a_{\psi(1)}^t) + \epsilon_{T,\psi(1)}^t, ..., U^t(a_{\psi(K^t)}^t) + \epsilon_{T,\psi(K^t)}^t))]
\end{aligned} \tag{12}$$

$$\begin{aligned}
= \sum_{a_{\psi(k)}^t \in B \subseteq A^t} \frac{1}{|B|}\text{Prob}[B = & C^t((U^t(a_1^t) + \epsilon_{1,1}^t, ..., U^t(a_{C^t}^t) + \epsilon_{1,K^t}^t), ... \\
& ..., (U^t(a_1^t) + \epsilon_{T,1}^t, ..., U^t(a_{K^t}^t) + \epsilon_{T,K^t}^t))]
\end{aligned} \tag{13}$$

$$= P_{\psi(k)}^t(U^t)$$

The i.i.d assumption is used to go from equation (11) to equation (12), and neutrality is used to go from equation (12) to equation (13).

Suppose ψ is the pairwise permutation of actions k and l, and suppose $U^t(a_k^t) = U^t(a_l^t)$. Then $P_k^t(U^{t,\psi}) = P_{\psi(k)}^t(U^t) = P_l^t(U^t) = P_l^t(U^{t,\psi})$. So whenever $U^t(a_k^t) = U^t(a_l^t)$, it must be that $P_k^t(U^t) = P_l^t(U^t)$. $\square$

Next, considering the case of $K^t = 2$, we prove next that neutrality, together with unanimity and positive responsiveness is sufficient to guarantee that a team response function satisfies rank dependence for all admissible F^t. This is proved below.

Theorem 3. *If $K^t = 2$, F^t is admissible and C^t satisfies unanimity, positive responsiveness and neutrality, then P^t satisfies rank dependence.*

Proof. Pick any U^t such that $U_1^t > U_2^t$ and let $\delta = U_1^t - U_2^t$. Let $U'^t = (U_1^t - \delta, U_2^t)$, then by lemma 1, $P_1^t(U'^t) = P_2^t(U'^t) = \frac{1}{2}$. Since C^t satisfies positive responsiveness and unanimity, theorem 2, together with the fact that $P_1^t(U^t) + P_2^t(U^t) = 1$, implies that $P_1^t(U^t) > P_1^t(U'^t) = P_2^t(U'^t) > P_2^t(U^t)$.

□

If $K^t > 2$ the next two propositions prove rank dependence with additional restrictions on the team collective choice rule.

Theorem 4. *If F^t is admissible and C^t is plurality rule, then P^t satisfies rank dependence.*

Proof. Consider any profile of expected payoffs U^t. By neutrality and admissibility we can without loss of generality label the actions such that $U_1^t \geq U_2^t \geq ... \geq U_K^t$. By lemma 1, if $U_l^t = U_k^t$, then $P_l^t = P_k^t$.

Let $F_{i,k}$ denote the marginal distribution of ϵ_{ik} and $F_{i,-k}$ denote the joint distribution of $\epsilon_{ij}, j \neq k$. The probability that any individual team member i ranks action k highest is

$$
\begin{aligned}
p_k &= \mathbb{P}(\hat{U}_{ik}^t = U_k^t + \epsilon_{ik}^t \geq \max\nolimits_{j \neq k}\{\hat{U}_{ij}^t\}) \\
&= \int \mathbb{1}\{\hat{U}_{ik}^t \geq \max\nolimits_{j \neq k}\{\hat{U}_{ij}^t\}\} dF_i^t \\
&= \int \int \mathbb{1}\{\hat{U}_{ik}^t \geq \max\nolimits_{j \neq k}\{\hat{U}_{ij}^t\}\} dF_{i,k}^t dF_{i,-k}^t \\
&= \int (1 - F(\max\nolimits_{j \neq k}\{\hat{U}_{ij}^t\} - U_{ik}^t)) dF_{i,-k}
\end{aligned}
$$

Suppose $U_k^t > U_l^t$. Then for all $\epsilon_{i,-k}$ we have that $\max_{j \neq k}\{\hat{U}_{ij}^t\} \leq \max_{j \neq l}\{\hat{U}_{ij}^t\}_{j \neq l}$, so $p_k > p_l$.

Now, still supposing that $U_k^t > U_l^t$, and therefore that $p_k > p_l$, denote by $(n_1, ..., n_K)$ the tuple of number of team members that rank each action first for a given $\hat{U}^t$. For any choice set $B \subseteq A^t$, let $V^B = \{(n_1, ..., n_K) | \forall a_k \in B, \forall a_j, n_k \geq n_j \text{ and } \sum_{j=1}^K n_j = n\}$ be the set of feasible 'vote' totals that result in B being chosen. Then, since estimated expected utilities are independent across individuals conditional on U^t, we can write the probability of this subset being chosen as

$$\mathbb{P}(C^t(\hat{U}^t) = B) = \sum_{V^B} \frac{n!}{n_1!n_2!...n_K!} \Pi_{j=1}^K p_j^{n_j}$$

Let $\psi : \{1, ..., K\} \rightarrow \{1, ..., K\}$ be the pairwise permutation between k and l, that is, the function that maps k to l, l to k, and all else to itself. Pick any B that contains a_k and not a_l. Then $(n_1, ..., n_K) \in V^B$ if and only if $(n_{\psi(1)}, , ..., n_{\psi(K)}) \in V^{(B-\{a_k\})\cup\{a_l\}}$, the set of vote totals that results in the choice set being B, minus a_k and adding a_l. Then, since $n_k > n_l$, we have that $p_k^{n_k} p_l^{n_l} > p_k^{n_l} p_l^{n_k}$, so every term of the sum in $\mathbb{P}(C^t(\hat{U}^t) = B)$ is greater than the corresponding term in $\mathbb{P}(C^t(\hat{U}^t) = (B - \{a_k\}) \cup \{a_l\})$. So we have for all B containing a_k and not a_l, that $\mathbb{P}(C^t(\hat{U}^t) = B) > \mathbb{P}(C^t(\hat{U}^t) = (B - \{a_k\}) \cup \{a_l\})$.

Finally, define B_0 to be the subsets of A^t that contain neither a_k nor a_l, B_k the subsets containing only a_k, B_l the subsets containing only a_l and not a_k, and B_{kl} the set containing both. Then

$$\begin{aligned}
P_k^t(U^t) &= 0 \times \sum_{B \in B_0} \mathbb{P}(C^t(\hat{U}^t) = B) + \sum_{B \in B_k} \frac{1}{|B|} \mathbb{P}(C^t(\hat{U}^t) = B) \\
&\quad + 0 \times \sum_{B \in B_l} \mathbb{P}(C^t(\hat{U}^t) = B) + \sum_{B \in B_{kl}} \frac{1}{|B|} \mathbb{P}(C^t(\hat{U}^t) = B) \\
P_l^t(U^t) &= 0 \times \sum_{B \in B_0} \mathbb{P}(C^t(\hat{U}^t) = B) + 0 \times \sum_{B \in B_k} \mathbb{P}(C^t(\hat{U}^t) = B) \\
&\quad + \sum_{B \in B_l} \frac{1}{|B|} \mathbb{P}(C^t(\hat{U}^t) = B) + \sum_{B \in B_{kl}} \frac{1}{|B|} \mathbb{P}(C^t(\hat{U}^t) = B)
\end{aligned}$$

P_k^t and P_l^t share all terms of the fourth sum, so

$$P_k^t[U^t] - P_l^t[U^t] = \sum_{B \in B_k} \frac{1}{|B|} \mathbb{P}[C^t(\hat{U}^t) = B] - \sum_{B \in B_l} \frac{1}{|B|} \mathbb{P}[C^t(\hat{U}^t) = B]$$
$$P_k^t[U^t] - P_l^t[U^t] = \sum_{B \in B_k} \frac{1}{|B|} [\mathbb{P}[C^t(\hat{U}^t) = B] - \mathbb{P}[C^t(\hat{U}^t) = (B - \{a_k\}) \cup \{a_l\}]] > 0$$
$$P_k^t[U^t] > P_l^t[U^t]$$

Therefore, whenever $U_k^t > U_l^t$, we have $p_k > p_l$, which implies $P_k^t(U^t) > P_l^t(U^t)$.

□

We define a weighted average rule as follows:

Definition 8. *A team collective choice rule C^t is an* ***Weighted Average Rule*** *if there exists a profile of non-negative individual voting weights, $(w_1^t, ..., w_{n^t}^t)$ with $\sum_{i=1}^{n^t} w_i^t = 1$ such that for all $a_k^t \in A^t$ and for all $\hat{U}^t \in \Re^{K^t n^t}$, $a_k^t \in C^t(\hat{U}^t)$ if and only if $\sum_{i=1}^{n^t} w_i^t \hat{U}_{ik}^t \geq \sum_{i=1}^{n^t} w_i^t \hat{U}_{il}^t$ for all $l \neq k$.*

Theorem 5. *If F^t is admissible and C^t is a weighted average rule, then P^t satisfies rank dependence.*

Proof. Consider any profile of expected payoffs U^t, and suppose $U_k^t > U_l^t$. We have $P_k^t(U^t) = \int \mathbb{1}\{\sum_{i=1}^{n^t} w_i^t \hat{U}_{ik}^t \geq \max\{\sum_{i=1}^{n^t} w_i \hat{U}_{ij}^t\}_{j=1}^{K^t}\} dF^t$. Note that the probability that any of these weighted averages are exactly equal is 0. Now, since $F^t(y - U_k^t) < F^t(y - U_l^t)$ for all $y \in \mathbb{R}$, we have $\hat{U}_{ik}^t >_{st} \hat{U}_{il}^t$, where $>_{st}$ denotes the *strict* first stochastic order, for all members i. This order is closed under convolutions, so we have $\sum_{i=1}^{n^t} w_i^t \hat{U}_{ik}^t >_{st} \sum_{i=1}^{n^t} w_i^t \hat{U}_{il}^t$. Since $\mathbb{1}\{z > 0\}$ is increasing, non-constant and bounded, we therefore have

$$\int \mathbb{1}\{\sum_{i=1}^{n} w_i \hat{U}_{ik}^t \geq \max\{\sum_{i=1}^{n} w_i \hat{U}_{ij}^t\}_{j=1}^K\}\}dF > \int \mathbb{1}\{\sum_{i=1}^{n} w_i \hat{U}_{il}^t \geq \max\{\sum_{i=1}^{n} w_i \hat{U}_{ij}^t\}_{j=1}^K\}\}dF$$

$$P_k^t(U^t) > P_l^t(U^t)$$

□

5 Nash convergence

A collective choice rule satisfies the Nash convergence property if, as the size of all teams increases without bound, every convergent sequence of team equilibrium converges to a Nash equilibrium. The formal statement for the case of plurality rule is in Proposition 1. That earlier result, namely that plurality rule satisfies the Nash convergence property begs the more general question whether plurality is somehow unique in this regard, or if it is a property shared by a broader class of collective choice rules. It is clearly not unique, as it is easy to show that the average rule also has this property, as a straightforward application of the law of large numbers. Thus the challenge is to characterize, at least partially, the class of collective choice rules with this *Nash convergence property.*

The previous section identified conditions under which a collective choice rule will guarantee that the team response function will be well behaved, and in particular will choose better actions more often than worse actions. Intuition suggests that rank dependence is closely related to the question of Nash convergence. Is it true that every collective choice rule that generates rank dependent team responses has the Nash convergence property? Unfortunately, the answer is negative, and is easy to see with an example. Dictatorial collective choice rules generate rank dependent and payoff monotone team response functions, but obviously fail the Nash convergence property. In order to achieve Nash convergence, the collective choice rule must be responsive to the estimated expected payoffs of an ever growing number of members of the team.

We next show that all *anonymous* scoring rules satisfy the Nash convergence property.

Definition 9. *A team collective choice rule C^t is an* ***Anonymous Scoring Rule*** *if there* *exists a profile of individual scoring functions,* $(S_1^t, ...S_{n^t}^t)$ *with* $S_i^t = S_j^t = \sigma^t$ *for all* $i, j \in t$*, such that alternative* a_k^t *is chosen at* $\hat{U}^t$ *if and only if* $\sum_{i=1}^{n^t} \sigma_k^t(\hat{U}_i^t) \geq \sum_{i=1}^{n^t} \sigma_l^t(\hat{U}_i^t)$ *for all* $l \neq k$*.*

Theorem 6. *Consider an infinite sequence of team games,* $\{\Gamma_m\}_{m=1}^{\infty}$ *such that (1)* $A_m^t = A_{m'}^t = A^t$ *for all* t, m, m'*; (2)* $u_m^t = u_{m'}^t = u^t$ $\forall t, m, m'$*; (2)* $n_{m+1}^t > n_m^t$ *for all* m, t*; and (3)* $C^t = \sigma$*, an anonymous scoring rule, for all* m, t*. Let* $\{\alpha_m\}_{m=1}^{\infty}$ *be a convergent sequence of team equilibria where* $\lim_{m\to\infty} \alpha_m = \alpha^*$*. Then* α^* *is a Nash equilibrium of the strategic form game* $[\mathbf{T}, A, u]$*.*

Proof. Suppose α^* is not a Nash equilibrium. Then there is some team t and some pair of actions, a_k^t, a_l^t such that $U_k^t(\alpha^*) > U_l^t(\alpha^*)$, but $\alpha_l^{t*} = \xi > 0$.

Since C^t is an anonymous scoring rule, σ, we know that for all t, for all n^t, and for all $a_k^t \in A^t$, $a_k^t \in C^t(\hat{U}^t)$ if and only if

$$\begin{aligned}
\sum_{i=1}^{n^t} \sigma_k(\hat{U}_i^t) &\geq \sum_{i=1}^{n^t} \sigma_l(\hat{U}_i^t) \text{ for all } l \neq k \\
&\Leftrightarrow \\
\overline{\sigma}_{kn^t}^t(\hat{U}^t) &\geq \overline{\sigma}_{ln^t}^t(\hat{U}^t) \text{ for all } l \neq k.
\end{aligned}$$

where $\overline{\sigma}_{kn^t}^t(\hat{U}^t)$ denotes the average score of a_k^t among the n^t members of t at $\hat{U}^t$. That is, the score of a_k^t is maximal if and only if the average individual score of a_k^t is maximal. If $U_k^t > U_l^t$ then σ_k stochastically dominates σ_l, and hence $E\{\sigma_k(\hat{U}_i^t)\} > E\{\sigma_l(\hat{U}_i^t)\}$, where

$$E\{\sigma_k(\hat{U}_i^t)\} = \int \sigma_k(U^t + \epsilon_i^t) dF^t(\epsilon_i^t)$$

Hence, at the limiting strategy profile, α^*, we have $E\{\sigma_k(\hat{U}_i^t|\alpha^*)\} > E\{\sigma_l(\hat{U}_i^t|\alpha^*)\}$.

Since $U^t(\alpha_m) \to U^t(\alpha)$ and $\overline{\sigma}$ is continuous in α it follows that $\overline{\sigma}_{kn_m^t}^t(\hat{U}^t|\alpha_m) \to E\{\sigma_k^t(\hat{U}^t|\alpha^*)\}$ and $\overline{\sigma}_{ln_m^t}^t(\hat{U}^t|\alpha_m) \to E\{\sigma_l^t(\hat{U}^t|\alpha^*)\}$ in probability as $m \to \infty$. Therefore, $E\{\sigma_k(\hat{U}_i^t|\alpha^*)\} > E\{\sigma_l(\hat{U}_i^t|\alpha^*)\}$ implies $\exists\ \overline{m}$ such that $\Pr\{\overline{\sigma}_{kn_m^t}^t(\hat{U}^t|\alpha_m) \leq \overline{\sigma}_{ln_m^t}^t(\hat{U}^t|\alpha_m)\} < \frac{\xi}{2}\ \forall m > \overline{m}$. This leads to a contradiction to the initial hypothesis that $\alpha_l^{t*} = \xi > 0$ since $\Pr\{\overline{\sigma}_{kn_m^t}^t(\hat{U}^t|\alpha_m) \leq \overline{\sigma}_{ln_m^t}^t(\hat{U}^t|\alpha_m)\} < \frac{\xi}{2}\ \forall m > \overline{m}$, implies that $\alpha_{ln_m^t}^t < \frac{\xi}{2}\ \forall m > \overline{m}$ and hence $\alpha_l^{t*} < \xi$. $\square$

A few comments about the generality of the result. First, the result does not imply that equilibrium with larger teams are closer to Nash equilibrium than equilibrium with smaller teams. It is an asymptotic result for large teams. Examples in section 3 show that the convergence can be non-monotonic even in very simple games. Second, the anonymous scoring rule does not have to be the same for all teams in order to obtain Nash convergence; different teams can use different anonymous scoring rules and the result still holds. While not formally stated in the theorem, it is an obvious generalization. Second, the rate of convergence to large teams can differ across teams. Third, the result only characterizes conditions that are sufficient for the Nash convergence property; the condition is clearly not necessary as the average rule is not a scoring rule but satisfies the Nash convergence property. Fourth, we conjecture, but have not been able to prove that the class is much broader than anonymous scoring rules, including many other anonymous collective choice rules that generate payoff monotonic and rank dependent team response functions. Scoring rules operate only on the individual *ordinal* rankings of estimated expected payoffs. One imagines that there are many collective choice rules that operate on the cardinal values of the estimates and also have the Nash convergence property (in addition to the weighted average rules).

6 Team Equilibrium in Extensive Form Games

A standard definition of finite extensive form games can be found in Osborne and Rubinstein (1994), and specifies a Player set $\mathbf{I} = \{1, ..., I\}$, an action set $\mathbf{A}$, a set of sequences contained in A called histories, Ξ, a subset of these being terminal histories, Z, a player function, P, information sets for each player, Π^i, a feasible action function, A, that specifies the set of actions available at each information set, and a payoff functions, u^i, defined on Z. Thus an extensive form game, in shorthand, can be written as $\Gamma = (\mathbf{I}, \mathbf{A}, \Xi, Z, P, \Pi, A, u)$. We extend this definition to (finite) *team extensive form games*, using this notation. As in the definition of team games in strategic form, let $\mathbf{T} = \{0, 1, ..., t, ..., T\}$ be a finite collection of *teams*, $t = \{i_1^t, ..., i_j^t, ..., i_{n^t}^t\}$, where each team consists of n^t members, or individuals, and denote the *team size profile* by $n = (n^1, ..., n^T)$. Team "0" is designated *chance*. Let $\mathbf{A}$ be a finite set of *actions* and Ξ be a finite set of *histories*, that satisfy two properties: $\varnothing \in \Xi$; and $(a_1, ..., a_K) \in \Xi \Rightarrow (a_1, ..., a_L) \in \Xi$ for all $L < K$. A history $h = (a_1, ..., a_K) \in Z \subseteq \Xi$ is *terminal* if there does not exist $a \in \mathbf{A}$ such that $(h, a) \in \Xi$, and the set of terminal histories is denoted Z. The set of actions available at any non-terminal history h is determined by the function $A : \Xi \to 2^{\mathbf{A}}$, where $A(h) = \{a|(h, a) \in \Xi\}$. There is a *player function* $P : \Xi - Z \to \mathbf{T}$ that assigns each history to a unique team. Without loss of generality, assume that $P(\varnothing) = 0$ and $P(h) \neq 0$ for all $h \neq \varnothing$), and denote by b^0 the probability distribution of chance actions at $h = \varnothing$ and assume without loss of generality that $b^0(a) > 0$ for all $a \in A(\varnothing)$. For each $t \in \{0, 1, ..., t, ..., T\}$, there is an information partition Π^t of $\{h \in \Xi|P(h) = t\}$. Elements of Π^t are t's *information sets*, and are denoted H_l^t, where l indexes t's information sets. The set of available actions to t are the same in all histories that belong to the same information set. That is, if $h \in H_l^t$ and $h' \in H_l^t$ for some $H_l^t \in \Pi_t$ then $A(h) = A(h') \equiv A(H_l^t) = \{a_{l1}^t, ..., a_{lk}^t, ..., a_{lK_l^t}^t\}$ where $K_l^t = |A(H_l^t)|$.

The *payoff function* of the game for team $t \neq 0$ is given by $u^t : Z \to \Re$. Given any terminal history $z \in Z$, all members of team t receive the payoff $u^t(z)$. A behavioral strategy for team t is a function $b^t = (b_1^t, ..., b_l^t, ..., b_{L^t}^t)$, where $L^t = |\Pi^t|$ and $b_j^t : H_l^t \to \Delta A(H_l^t)$, where $\Delta A(H_l^t)$ is the set of probability distributions over $A(H_l^t)$. Denote by B the set of behavioral strategy profiles, and B^o the interior of B (i.e., the set of totally mixed behavioral strategy profiles).

Each behavioral strategy profile $b \in B^o$ determines a strictly positive *realization probability* $\rho(z|b)$ for each $z \in Z$. For any $t \neq 0$ and $b \in B^o$, define the expected payoff function for team t, $v_i : B^o \to \mathcal{R}$ by:

$$v^t(b) = \sum_{z \in Z} \rho(z|b) u_t(z).$$

Similarly, for any $t \neq 0$ and any information set $H_l^t \in \Pi^t$, and for each $a_{lk} \in A(H_l^t)$, any $b \in B^o$ determines a strictly positive *conditional realization probability* $\rho(z|H_l^t, b, a_{lk})$ for each $z \in Z$.[14] This is the probability distribution over Z, conditional on reaching H_l^t given the behavioral strategy profile b, with b_l^t replaced with the pure action a_{lk}. For each $t \neq 0$, for each $H_l^t \in \Pi^t$, and for each $a_{lk} \in A(H_l^t)$, define the *conditional payoff function* by

$$U_{lk}^t(b) = \sum_{z \in Z} \rho(z|H_l^t, b, a_{lk}) u^t(z).$$

This is the conditional payoff to t of playing the (pure) action $a_{lk} \in A(H_l^t)$ at H_l^t with probability one, and otherwise all teams (including t) playing b elsewhere.

The rest of the formal description of extensive form team games closely follows the

[14]The restriction to B^o is without loss of generality in our framework, as we will show later.

structure of team games in strategic form. At each information set H_l^t, each member of team t gets a noisy estimate of the *true* conditional payoff of each currently available action, given a behavioral strategy profile of the other teams, b^{-t}. These estimates are aggregated into a team decision via a team collective choice rule, C_l^t. We will assume C_l^t satisfies Unanimity for all t and l, and therefore expecting over all noisy estimates will lead to a totally mixed team behavioral strategy, $b_l^t \in B^o$, at that information set.

Formally, given $b \in B^o$, for every t, and each of t's information sets, $H_l^t \in \Pi^t$, at any history in H_l^t, for each $a_{lk} \in A(H_l^t)$ member $i \in t$ observes an estimate of $U_{lk}^t(b)$ equal to the true conditional expected payoff $(U_{lk}^t(b))$ plus an estimation error term. That is, $\widehat{U}_{ilk}^t(b) = U_{lk}^t(b) + \varepsilon_{ilk}^t$. We call $\widehat{U}_{il}^t(b) = (\widehat{U}_{il1}^t(b), ..., \widehat{U}_{ilK_l^t}^t(b))$ member i'sestimated conditional payoffs at H_l^t, given b. Denote by $\widehat{U}_l^t(b) = (\widehat{U}_{1l}^t(b), ..., \widehat{U}_{n^tl}^t(b))$ is the profile of member estimated conditional payoffs at $H_l^t \in \Pi^t$. The estimation errors for members of team t are i.i.d. draws from a commonly known admissible probability distribution F_l^t. We also assume the estimation errors are independent across information sets, but allow different distributions at different information sets.

A *team collective choice rule* at information set H_l^t, C_l^t, is a correspondence that maps profiles of member estimated payoffs at $H_l^t \in \Pi^t$ into a subset of elements of $A(H_l^t)$. Thus, C_l^t is a social choice correspondence. That is, $C_l^t : \Re^{n^t K_l^t} \to 2^{A(H_l^t)}$, so $C_l^t(\widehat{U}^t) \subseteq A(H_l^t)$. In principle, teams could be using different collective choice rules at different information sets, and denote $C = (C^1, ..., C^T)$, where $C^t = (C_1^t, ..., C_{L^t}^t)$. We assume that team t always mixes uniformly over $C_l^t(\widehat{U}_l^t)$. That is, the probability team t chooses a_{lk}^t at $\widehat{U}_l^t(b)$ is given by the function g^t defined as:

$$\begin{aligned} g_{lk}^{C_l^t}(\widehat{U}_l^t) &= \frac{1}{|C_l^t(\widehat{U}_l^t)|} \text{ if } a_{lk}^t \in C_l^t(\widehat{U}_l^t) \\ &= 0 \text{ otherwise} \end{aligned}$$

Given a behavioral strategy profile, b, for each realization of $(\varepsilon_{l1}^t, ..., \varepsilon_{ln^t}^t)$ at information set $H_l^t \in \Pi^t$,, team t using collective choice rule C_l^t at H_l^t is assumed to take the action $a_{lk} \in A(H_l^t)$ with probability $g_{lk}^{C_l^t}(\widehat{U}_l^t)$.

We require the profile of collective choice rules to satisfy *Unanimity*, defined similarly to Definition 3. That is, for every team $t \neq 0$ and every information set $H_l^t \in \Pi^t$, if $\widehat{U}_{ilk}^t > \widehat{U}_{ilk'}^t$ for all $a_{lk'} \in A(H_l^t) - \{a_{lk}\}$ and for all $i \in t$, then $C_l^t(\widehat{U}_l^t) = \{a_{lk}\}$. It is important to note that our assumptions about F, together with unanimity of the collective choice rules, imply that for every t, every $H_l^t \in \Pi^t$, and every $a \in A(H_l^t)$, the probability that ε_l^t is such that $\widehat{U}_{ilk}^t > \widehat{U}_{ilk'}^t$ for all $a_{lk'} \in A(H_l^t) - \{a_{lk}\}$ and for all $i \in t$ is strictly positive. Therefore the behavioral strategies implied by the team choice probabilities are always totally mixed. That is for all $b \in B^o$, for all $t \in \mathbf{T}$, for all $H_l^t \in \Pi^t$, and for all $a \in A(H_l^t)$, $P_{lk}^t(U_l^t(b)) > 0$, and hence every possible history in the game occurs with positive probability. Thus, there are no "off-path" histories, so $\rho(z|H_l^t, b, a_{lk})$ is always well-defined and computed according Bayes' rule.

Team response functions are defined in the following way. Any totally mixed behavioral strategy profile, b, implies a profile of true conditional expected continuation payoffs U_l^t for each action at each information set H_l^t . Given U_l^t, the team collective choice rule, C_l^t, and the distribution of individual estimation errors, F_l^t, together imply a behavioral strategy response to b for team t at information set H_l^t, which we denote a *team response function* for team t by P^t, where P_{lk}^t specifies the probability ε_l^t is such that a_{lk}^t is the team's action choice at H_l^t in response to b. That is:

$$P_{lk}^t(U_l^t(b)) = \int\limits_{\varepsilon_l^t} g_{lk}^{C_l^t}(\widehat{U}_{lk}^t(b))dF_l^t(\varepsilon_l^t) \tag{14}$$

where F_l^t denotes the distribution of $\varepsilon_l^t = (\varepsilon_{l1}^t, ..., \varepsilon_{ln^t}^t)$. For any extensive form game, any admissible F, and any profile of collective choice rules C call $\Gamma = [\mathbf{T}, \mathbf{A}, \Xi, A, P, \Pi, b^0, \{u^t\}_{t=1}^T, F, C]$ a *team game in extensive form*. An equilibrium of a team game in extensive form is a fixed point of P.

Definition 10. *A* ***team equilibrium*** *of the team extensive form game* $\Gamma^T = [\mathbf{T}, n, \Xi, Z, \mathbf{A}, P, \Pi, A, u, F, C]$ *is a behavioral strategy profile* b *such that, for every* $t \neq 0$, *and every information set* $H_l^t \in \Pi^t$ *and every* $a_{lk}^t \in A(H_l^t)$, $b_{lk}^t = P_{lk}^t(U^t(b))$.

Theorem 7. *For every team game in extensive form a team equilibrium exists and is in totally mixed behavioral strategies.*

Proof. For each t, and each $H_l^t \in \Pi^t$ and each $a \in A(H_l^t)$, define:

$$\underline{b}_{lk}^t = \inf_{b \in B^o} P_{lk}^t(U_l^t(b))$$

By Unanimity of C_l^t, $P_{lk}^t(U_l^t(b)) > 0$ so $\underline{b}_{lk}^t \geq 0$. Furthermore, we have $\underline{b}_{lk}^t > 0$ since $U_l^t(b)$ is uniformly bounded in B^o for all t and l, and hence $P^t_{lk}(U^t_l(b))$ is bounded strictly away from 0 for all t, l and k and for all $b \in B^o$. Define:

$$\overline{B}^o = \{b \in B | b^t_{lk} \geq \underline{b}^t_{lk} \text{ for all } t, l \text{ and } k\} \subset B^o.$$

Since $\overline{B}^o$ is compact and convex and P^t_{lk} is a continuous function for all t, l and k, by Brouwer's fixed point theorem there exists a $b \in \overline{B}^o$ such that $b^t_{lk} = P^t_{lk}(U^t(b))$ for all t, l and k. □

A few observations about team equilibria in extensive form games are worth noting. First, the results in Section 4 about conditions for team response functions to satisfy payoff monotonicity and rank dependence carry through to extensive form team games, as applied to the behavioral strategies of each team. This is formally stated as follows.

Theorem 8. *For each t, Π^t, and $H^t_l \in \Pi^t$, if F^t_l is admissible and C^t_l satisfies unanimity and positive responsiveness then P^t_l satisfies payoff monotonicity. Furthermore, if C^t_l also satisfies neutrality and $|A(H^t_l)| = 2$, or if $|A(H^t_l)| > 2$ and C^t_l is plurality rule or a weighted average rule, then P^t_l satisfies rank dependence.*

Proof. The proof is essentially the same as the proof of Theorems 2, 3, and 4. □

Second, a stronger version of the Nash convergence property holds for team extensive form games. It is stronger because any limit point of a sequence of team equilibria in an extensive form game when teams become large is not just a Nash equilibrium, but must also be sequentially rational. This follows because equilibria in team games are always in B^o, so all information sets are on the equilibrium path and continuation payoffs are always computed simply using Bayes' rule. This is formally stated as follows.

Theorem 9. *Consider an infinite sequence of team extensive form games, $\{\Gamma^T_m\}^\infty_{m=1}$, where $\Gamma^T_m = [\mathbf{T}, n, \Xi, Z, \mathbf{A}, P, \Pi, A, u, F, C]$,where m indexes an increasing sequence of team sizes, with all the other characteristics of the game being the same. That is, $n^t_{m+1} > n^t_m$ for all m, t. Suppose C^t_{lm} is an anonymous scoring rule for all m, t, l and let $\{b^*_m\}^\infty_{m=1}$ be a convergent sequence of team equilibria where $\lim_{m\to\infty} b^*_m = b^*$. Then b^* is a sequential equilibrium strategy of the corresponding extensive form game $\Gamma = (\mathbf{I}, \Xi, Z, \mathbf{A}, P, \Pi, A, u)$.*

Proof. The proof is essentially the same as the proof of Theorem 6 except to additionally show that limit points are sequentially rational. First, the statement is not vacuous because for each m there exists at least one team equilibrium, b_m^*, and hence exists at least one convergent sequence of equilibria $\{b_m^*\}_{m=1}^{\infty}$ by the Bolzano-Weierstrass Theorem, with $\lim_{m\to\infty} b_m^* = b^*$. What we need to show is that there exist consistent beliefs, μ^* (i.e., assignments of a probability distribution over the histories at each information set that satisfy the Kreps-Wilson (1982) consistency criterion) such that, under those beliefs b^* specifies optimal behavior at each information set. That is, b^* is sequentially rational given μ^* and μ^* is consistent with b^*.

We first show that μ^* is consistent with b^*. Because F^t has full support for all t and plurality rule satisfies unanimity, it follows that $b_m^* >> 0$ for all m. That is, for every m, t, H_l^t, and $a_{lk}^t \in A(H_l^t)$, $b_{mlk}^{t*} > 0$. Consequently every history occurs with positive probability, so, by Bayes' rule, for each information set H_l^t, and for all m, μ_{lm}^* is uniquely defined, where μ_{lm}^* denotes the equilibrium beliefs over the histories in the information set H_l^t. Since μ varies continuously with b there is a unique limit, $\mu^* = \lim_{m\to\infty} \mu_m^*$. Since $\mu_m^* >> 0$, and $\lim_{m\to\infty} b_m^* = b^*$ it follows that μ^* and consistent beliefs under b^*.

What remains to be shown is that b_l^{t*} is optimal for all t and for all H_l^t. The proof is virtually the same as the proof of Theorem 6, so we omit it. □

Third, extensive form team games include games of incomplete information where teams have private information. For example, each team may have private information about u^t. We work through an example of such a game in the next section (A Simplified Poker Game). In principle, the approach could be extended to allow individual team members to have private information in the form of payoff disturbances, which would introduce preference heterogeneity within a team. For example, in the centipede game, different players on the same team may have different degrees of altruism and hence disagree on which is the best action, even without any error in the estimated expected payoffs. This leads naturally to incentives for strategic behavior the the collective choice procedure (for example, strategic voting), and would complicate matters considerably.

Fourth, there is a rough connection between team equilibria and the extensive form version of quantal resopnse equilibrium (Mckelvey and Palfrey 1998), with the difference being that the disturbances are not private value payoffs, but estimation errors, so all members of a team have common values. As in the agent model of quantal response equilibrium, in team games it is assumed that estimated expected continuation payoffs of each member of t at H_l^t are not observed until H_l^t is reached. If instead, each member of t observed all it's estimates at the beginning of the game, this would complicate the model somewhat, and would require specifying additional consistency conditions on the entire profile of a member's estimates.

A last observation is that an alternative way to model team equilibrium in extensive form games would be to apply the theory developed in Section 3 of this paper to the normal form or reduced normal form of the game. However, the team equilibria are *not* invariant to inessential transformations of the extensive form, so equilibria of the normal form or the reduced normal form will in general not be observationally equivalent to team equilibria in behavioral strategies derived from the extensive form. Such theoretical differences could suggest interesting testable implications of the team equilibrium framework.

7 Examples of Team Size Effects in Extensive Form Games with Majority Rule.

7.1 The Simplified Poker Game

There are two teams, 1 and 2. Team 1 begins by drawing a card, which can be either a high card or low card, with a common knowledge prior π probability of drawing the high card. After seeing the card, 1 chooses whether to bet (B) or fold (F). If 1 folds, 1 loses the 'ante', a and 2 earns a, that is the payoff is $(-a, a)$. If 1 bets, 2 can either call or fold. If 2 folds, the payoff is $(a, -a)$. If 2 calls, then if 1 has the high card 1 wins the bet + ante, $(b + a)$ and B loses the same amount, that is the payoff is $(a + b, -(a + b))$. If 1 has the low card then the payoffs are $(-(a + b), a + b)$.

There is a unique Nash equilibrium to this game, in which 1 bets whenever 1 has the high card and mixes between betting and folding with the low card. If 1 bets, then 2 mixes between calling and folding. The equilibrium probability of betting with the low card p^* and calling q^* are given by

$$p^* = \frac{\pi}{1-\pi} \frac{b}{2a+b}$$
$$q^* = \frac{2a}{2a+b}$$

for any set of parameters such that $a > 0, b > 0, \pi \leq \frac{2a+b}{2a+2b}$.[15] For any such set of parameters, team equilibria converge to this equilibrium at any value of λ. However, the comparative statics of how the team equilibrium mixing probabilities change with group size is strongly affected by the game parameters, π, a, b.

In fact, there is a significant bias for small teams, and the direction of bias depends on the underlying parameters of the game. For example, when $\pi = 1/2$, $a = 1$, $b = 10$, and $\lambda = 1/4$, then Team 1 under-bets with the low card, and team 2 over-calls. Thus convergence to p^* is from below, while convergence to q^* is from above, with less under-betting and over-calling as team size increases.

However, if we reverse the magnitudes of the ante and bet amounts to $a = 5$ and $b = 1$, then we get the opposite direction of convergence: teams over-bet and under-call at low team sizes. Thus convergence to p^* is from *above*, while convergence to q^* is from *below*, with less over-betting and under-calling as team size increases. This is illustrated in Figure 5 for (odd) team sizes ranging from 1 to 59.

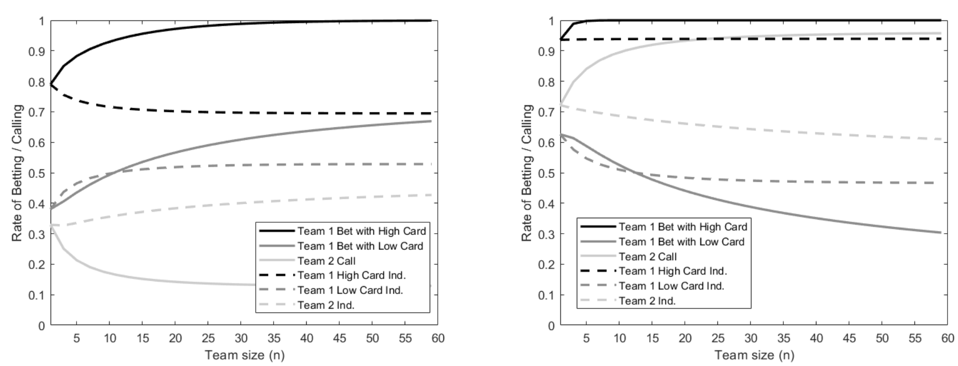

(a) Poker Game with $a = 1$, $b = 10$, $\lambda = 1/4$

(b) Poker Game with $a = 5$, $b = 1$, $\lambda = 1/4$

Figure 5: Team Equilibrium in the Poker Game

[15] If $\pi > \frac{2a+b}{2a+2b}$ the Nash equilibrium is for 1 to always bet and 2 to always fold.

7.2 Sequential Weak Prisoner's Dilemma

Next, we examine a sequential version of the Weak Prisoner's Dilemma, the simultaneous version of which we analyzed in section 3 (Table 2). In the sequential game, the equilibrium outcome depends on the order of moves. If column moves first, the unique subgame perfect Nash equilibrium is for both teams to defect. This is because after either cooperate or defect is chosen by column, row is better off choosing to defect. Since column optimizes by matching row's action, column should choose defect.

However, when row moves first, the unique subgame perfect equilibrium is for both teams to cooperate. This is because in equilibrium team 2 will choose whichever action team 1 chooses, defect after defect and cooperate after cooperate, and so team 1 will choose cooperate since $C > D$.

Figure 6 displays the team equilibrium for the two versions of this game with $x = y = 1$, and $z = 8$, the same parameter values used for the analysis of the simultaneous version in Section 3. The left panel shows the team equilibria if team 2 (column) moves first, and the right panel for the case where team 1 (row) moves first, for (odd) n ranging from 1 to 199, and $H(x) = \frac{1}{1+e^{-x/8}}$. The solid light gray line is the team 2's (first mover) defect probability, and the dashed light gray line is team 2's individual member's defect probability. The solid black line the defect probability of team 1 after either cooperation or defection by team 2, and the dashed black line is team 1's individual member's defect probability after either cooperation or defection by team 2.[16]

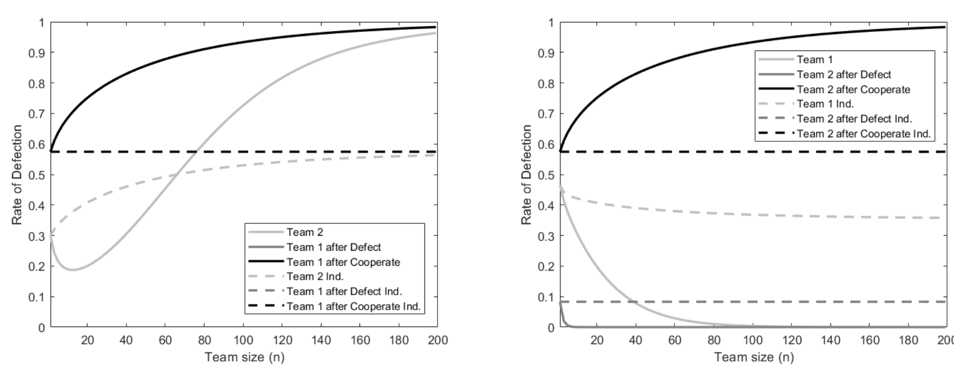

(a) Sequential WPD, team 2 moves first (b) Sequential WPD, team 1 moves first

Figure 6: Team Equilibrium in the Sequential WPD

When team 2 moves first, team 1's individual voting probabilities for defect are inde-

[16]The black and dark gray curves coincide in the left panel (a) of Figure 6, so both appear as black.

pendent of n, because $x = y = 1$, so the payoff difference between defect and cooperate for team 1 is the same. In fact, they are the same as in the simultaneous version studied in Section 3, namely $H(y)$. So in this case, for any $H(\cdot)$ and for any n, the team equilibria of the simultaneous move version and team 2 first mover sequential version of the game are identical.

However, when team 1 moves first the team equilibria converge rapidly to the action profile (C, C). When team 1 chooses cooperate, the difference in expected utility to team 2 between choosing defect and cooperate is $-z < 0$, so the individual voting probabilities are $H(-z) < \frac{1}{2}$, and when team 1 chooses defect, the difference is y, so the individual voting probabilities in this case are $H(y) > \frac{1}{2}$.

These voting probabilities are constant for all n, so team 2's individual voting probability of defection after defection converges to 1 and defection after cooperation converges to 0. Therefore, as n increases, the expected utility difference between defection and cooperation for team 1 decreases, and the team 1 individual voting probabilities and team probabilities decrease and converge to 0 probability of defection.

7.3 Centipede Game

Finally, we analyze team equilibria in a 4-move centipede game with exponentially increasing payoffs. At every outcome of this game, there is a high payoff and a low payoff, which initially equal 4 and 1, respectively. Two teams (1 and 2) take turns choosing to take or pass in sequence, starting with team 1. If team 1 chooses take the game ends, and team 1 receives the higher payoff while team 2 receives the lower payoff. If team 1 chooses pass, the two payoffs are doubled and team 2 gets to choose take or pass. This continues for up to 4 moves (fewer if one of the teams takes before the 4 node of the game), with the payoffs doubling after each pass. If pass is chosen at the last node, the game ends and team 1 receives 64 while team 2 receives 16. The two teams alternately play at most two nodes each in this game.

Since this is a finite game, the unique subgame perfect Nash equilibrium can be solved for by backward induction: choose take at every node.

In Figure 7 we display the team equilibrium majority-rule choice probabilities and individual voting probabilities for this game for (odd) n ranging from 1 to 119, and $H(x) = \frac{1}{1+e^{-x/8}}$.

In the left panel are the team probabilities and in the right panel are the individual voting probabilities, with probability of taking and voting for take on the y axis and team size on the x axis.

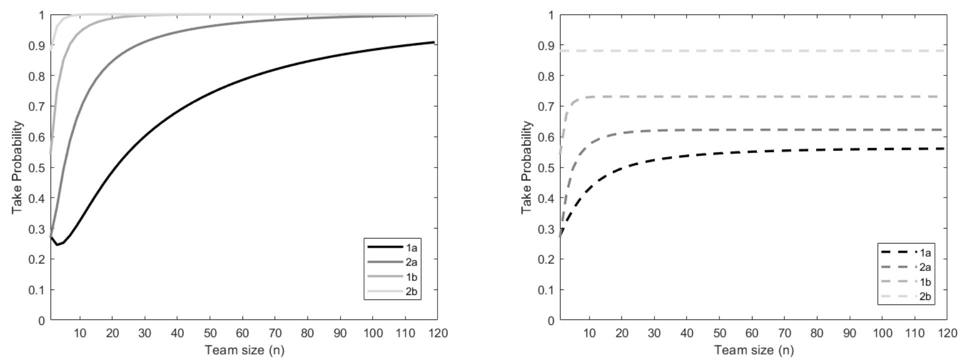

(a) Centipede Game, Team Take Probabilities

(b) Centipede Game, Individual Voting Probabilities

Figure 7: Team Equilibrium in the Centipede Game

At the final node of the game, observe that the individual voting probabilities are fixed at $H(16) \approx 0.88$. As n increases, the team probability of taking at this node converges rapidly and monotonically to 1. Voting probabilities at the early nodes are influenced by the team choice probabilities at future nodes. As n increases, the probability that the opposing team will take at future nodes increases, decreasing the continuation value of passing. This causes the voting probabilities for take to increase. Since, if the opposing team takes at the next node with probability 1, it is better to take at the current node than to pass, the individual voting probabilities converge to values strictly above 1/2, and so majority rule ensures that all team take probabilities eventually converge to 1.

8 Discussion and Conclusions

This paper proposes and develops a theory of games played by teams of players. In many, perhaps even most, applications of non-cooperative game theory in the social sciences, the players in the game are not actually individuals, but collections of individuals - teams - who collectively decide and implement a strategy in the game. Experimental economists and psychologists have observed significantly different outcomes of laboratory games in which strategy choices are made by teams compared with games in which strategy choices are made by individuals. Our goal of developing a rigorous theoretical framework to understand how games are played by teams is motivated by the combination of the observation that the actual players are teams in most applications of interest, combined with laboratory findings that consistently demonstrate that teams do not make strategic choices the same as individuals. As is the case for the laboratory implementations of games played by teams of players, we assume all members of the same team face the same payoff function.

The general framework we develop here combines the non-cooperative approach to model the strategic interaction between teams with a collective choice approach to the decision making process within teams. The individual members of each team have correct beliefs on average about the expected payoffs to each available team strategy, given the distribution of strategy profiles being used by the other teams in the game. These beliefs are correct on average, but they are subject to error. Because of this, these beliefs will generally be different for the different members of a teams, and a team collective choice rule maps the profile of error-prone beliefs of all members of the team into a strategy decision. Given an error structure and a collective choice rule, this induces a probability distribution over strategy choices for each team. A team equilibrium is a profile of mixed strategies, one for each team, with the property that, the collective choice rule of each team will generate its equilibrium mixed strategy, given the distribution of beliefs of the individual members of the team, each of which are correct on average.

The approach is initially developed for finite games in strategic form. Four main results are proved for strategic form team games. First, team equilibria generally exist. Second, we identify a broad class of collective choice rules such that team response functions are payoff monotone, in the sense that the probability a team chooses a particular action is increasing in the true expected payoff of that action. We show that unanimity and positive responsiveness are sufficient conditions on the team collective choice rule for the team response function to be payoff monotone for all admissible error distributions. Third, we identify stronger conditions on the collective choice rule that guarantee rank dependence, i.e., the property that team choice probabilities are ordered by the actions' true expected payoffs for all admissible distributions of estimation errors. In the case of $K^t = 2$ the only additional condition that is needed is neutrality, which implies that all scoring rules generate rank dependent team response functions. For $K^t > 2$ we show that plurality rule and weighted average rules generate rank dependence. Fourth, we show that all anonymous scoring rules satisfy the Nash convergence property. That is, as team sizes become large, all convergent sequences of team equilibrium with anonymous scoring rules converge to a Nash equilibrium.

Team equilibria for games in strategic form are illustrated for several 2×2 games, where the collective choice rule is majority rule and the number of team members is odd. These examples illustrate two distinct effects of changing team size on outcomes. The first is the reinforcement effect. If the probability any individual on a team choose one of the strategies is $p > 1/2$, then the probability a majority of the individuals on the team choose that strategy is greater than p, and is strictly increasing in the size of the team. The second effect is the equilibrium effect, which arises because in equilibrium, p will generally vary with n, and this equilibrium effect can go in the opposite direction from the reinforcement effect. This equilibrium effect is especially important in games with a mixed strategy equilibrium, where it comes to dominate the reinforcement effect in games with large teams.

The second half of the paper extends the framework to finite games in extensive form. Individuals are assumed to have correct beliefs on average at every information set about the expected continuation value of each available action at that information set. The results for strategic form team games about payoff monotonicity and rank dependence of team response functions and Nash convergence also apply to extensive form team games, with the latter result strengthened to show that limit points of team equilibrium in extensive form games are sequential equilibria. Team equilibrium in extensive form games is illustrated using examples of games of perfect information and games of imperfect information where teams have private types.

We hope this framework is a useful starting point for the further exploration and deeper understanding of how teams of players play games. There are many open questions that deserve further study and we mention a few. One is to generalize the class of collective choice rules that have the Nash convergence property. We identified one broad class of such collective choice rules (anonymous scoring rules), but we are well aware that there are other rules with this property. An interesting related question is how to extend the results about rank dependence with more than two actions to more general collective choice rules.

The framework can be expanded in several interesting directions. Many games of significant interest have infinite strategy spaces, including oligopoly, auction, and bargaining models in economics and spatial competition models in political science. In principle, the framework might be able to accommodate such an extension, for example by using finite approximations to the strategy space, but specific applications might face computational challenges. Similarly, some Bayesian games of interest have a continuum of types; allowing for a continuum of types would seem to be a feasible extension if the action spaces are finite.

The incorporation of behavioral biases and preferences (loss aversion, judgement biases, social preferences, etc.) should be straightforward, provided the effects are homogeneous across members of the group. A more challenging direction would be to extend the common values approach of this paper to allow for heterogeneity of team members' preferences in addition to the heterogeneous beliefs assumed here.

Finally, it should be noted that there are alternative approaches to modeling team games that are outside the scope of the framework presented here. For example, one might try to formalize the notion of "truth wins" - i.e., the idea that the team will adopt the choice favored by the most rational member of the group. This would require some formal notion of how to rank the rationality of the group members (such as level-k), coupled with a theory of persuasion, whereby the more rational members are able to change the beliefs of less rational members. Another alternative approach, which is more in the spirit of implementation theory and mechanism design, is to model the internal team decision making process as non-cooperative game rather than an abstract collective choice rule. This would create a nested game-within-a-game structure. In such an approach, it might be possible to model the deliberation process by introducing an explicit communication structure, which would add another layer of complexity.

References

[1] Blinder, Alan S., and John Morgan (2005). Are two heads better than one? Monetary policy by committee. *Journal of Money, Credit and Banking*, 789-811.

[2] Blinder, Alan S., and John Morgan (2008). Do monetary policy committees need leaders? A report on an experiment. *American Economic Review*, 98(2): 224-29.

[3] Bornstein, Gary, Tamar Kugler, and Anthony Ziegelmeyer (2004). Individual and Group Decisions in the Centipede Game: Are Groups More Rational Players? *Journal of Experimental Social Psychology* 40(5): 599-605.

[4] Bornstein, Gary, Ilan Yaniv (1998). Individual and Group Behavior in the Ultimatum Game: Are Groups More "Rational" Players? *Experimental Economics* 1:101-108.

[5] Casari, Marco, Jingjing Zhang, and Christine Jackson (2016). "Same process, different outcomes: group performance in an acquiring a company experiment." *Experimental Economics*, 19(4):764-91.

[6] Cason, Timothy N., Sau-Him Paul Lau, and Vai-Lam Mui (2018). Prior Interaction, Identity, and Cooperation in the Inter-Group Prisoner's Dilemma. *Working paper*.

[7] Charness, Gary and Matthew Jackson (2007). Group play in games and the role of consent in network formation. *Journal of Economic Theory*, 136:417-445.

[8] Charness, Gary, Edi Karni, and Dan Levin (2007). "Individual and Group Decision Making under Risk: An Experimental Study of Bayesian Updating and Violations of First-Order Stochastic Dominance." *Journal of Risk and Uncertainty*, 35(2):129-48.

[9] Charness, Gary, Edi Karni, and Dan Levin (2010). On the conjunction fallacy in probability judgment: New experimental evidence regarding Linda. *Games and Economic Behavior*, 68(2):551-556.

[10] Charness, Gary and Matthias Sutter (2012). "Groups Make Better Self-interested Decisions." *Journal of Economic Perspectives*, 26(3):157-176.

[11] Cooper, David J., and John H. Kagel (2005). Are Two Heads Better Than One? Team versus Individual Play in Signaling Games. *American Economic Review*, 95(3):477-509.

[12] Cox, James C (2002). Trust, Reciprocity, and Other-Regarding Preferences: Groups vs. Individuals and Males vs. Females. In *Advances in Experimental Business Research*, edited by Rami Zwick and Amnon Rapoport, 331-50. Dordrecht: Kluwer Academic Publishers.

[13] Cox, James C. and Stephen C. Hayne (2006). Barking up the right tree: Are small groups rational agents? *Experimental Economics*, 9:209-222.

[14] Duggan, John (2001). Non-cooperative games among groups. *Typescript*.

[15] Elbittar, Alexander, Andrei Gomberg, and Laura Sour (2011). Group Decision-Making and Voting in Ultimatum Bargaining: An experimental Study. *B.E. Journal of Economic Analysis and Policy (Contributions)*, 11(1): Article 53.

[16] Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2005), "Regular Quantal Response Equilibrium."*Experimental Economics*, 8(4):347-67.

[17] Insko C.A., Hoyle R.H., Pinkley R.L., Hong G.Y., Slim R.M., Dalton B, Lin Y-H.W., Ruffin P.P., Dardis G.J., Bernthal P.R. (1988). Individual-group discontinuity: the role of a consensus rule. *Journal of Experimental Social Psychology*, 24:505-519.

[18] Kocher, Martin G. and Matthias Sutter (2005). "The Decision Maker Matters: Individual versus Group Behavior in Experimental Beauty-Contest Games." *Economic Journal* 115(500): 200-223.

[19] Kreps, David and Robert Wilson (1982). "Sequential Equilibria." *Econometrica*, 50:863-894.

[20] Kugler, Tamar, Gary Bornstein, Martin G. Kocher, and Matthias Sutter (2007). Trust between Individuals and Groups: Groups are Less Trusting Than Individuals But Just as Trustworthy. *Journal of Economic Psychology*, 28(6): 646-57.

[21] Kugler, Tamar, Edgar E. Kausel, and Martin G. Kocher (2012). Are Groups More Rational Than Individuals? A Review of Interactive Decision Making in Groups. *Wiley Interdisciplinary Reviews - Congnitive Science*, 3(4):471-82.

[22] Marschak, Jacob and Roy Radner (1972). Economic Theory of Teams, *Cowles Foundation Monograph 22,* Yale University Press: New Haven.

[23] McKelvey, Richard D. and Talbot Page (1986). "Common Knowledge, Consensus, and Aggregate Information." *Econometrica*, 54(1): 109-127.

[24] McKelvey, Richard D. and Thomas R. Palfrey (1995). "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior*, 10:6-38.

[25] McKelvey, Richard D. and Thomas R. Palfrey (1996). "A Statistical Theory of Equilibrium in Games." *Japanese Economic Review*, 47:186-209.

[26] McKelvey, Richard D. and Thomas R. Palfrey (1998). "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics*, 1:9-41.

[27] Morone, Andrea, Simone Nuzzo, and Rocco Caferra (2019). The Dollar Auction Game: A Laboratory Comparison Between Individuals and Groups. *Group Decision and Negotiation*, 28:79-98

[28] Moulin, Hervé (1988). *Axioms of Cooperative Decisions Making*, Cambridge University Press: Cambridge.

[29] Osborne, Martin and Ariel Rubinstein (1994). *A Course in Game Theory.* MIT Press: Cambridge.

[30] Rockenbach, Bettina, Abdolkarim Sadrieh, and Barbara Mathauschek (2007). Teams Take Better Risks. *Journal of Economic Behavior and Organization*, 63:412-422.

[31] Schulze, Christin, and Ben R. Newell (2016). More heads choose better than one: Group decision making can eliminate probability matching. *Psychonomic Bulletin and Review*, 23(3): 907-914.

[32] Sheremeta, Roman M., and Jingjing Zhang (2010). Can Groups Solve the Problem of Over- bidding in Contests? *Social Choice and Welfare*, 35(2):175-97.

[33] Sniezek, Janet A., and Rebecca A. Henry (1989). Accuracy and confidence in group judgment. *Organizational Behavior and Human Decision Processes*, 43(1):1-28.

[34] Song, Fei (2008). Trust and Reciprocity Behavior and Behavioral Forecasts: Individuals versus Group-Representatives. *Games and Economic Behavior*, 62(2):675-96.

[35] Stahlberg, Dagmar, Frank Eller, Anne Maass, and Dieter Frey (1995). We knew it all along: Hindsight bias in groups. *Organizational Behavior and Human Decision Processes*, 63(1):46-58.

[36] Sutter, Matthias, Simon Czermak, and Francesco Feri (2013). Strategic Sophistication of Individuals and Teams: Experimental Evidence. *European Economic Review*, 64:395-410.

[37] Sutter, Matthias, Martin Kocher, and S. Strauss (2009). Individuals and teams in auctions. *Oxford Economic Papers*, 61(2):380-394

[38] Tversky, Amos, and Daniel Kahneman (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4):293.

[39] Wildschut, Tim and Chester A. Insko (2007). Explanations of interindividual-intergroup discontinuity: A review of the evidence. *European Review of Social Psychology*, 18:175-211.