American Economic Journal: Microeconomics 2022, 14(4): 122–157
https://doi.org/10.1257/mic.20200391

Games Played by Teams of Players†

By Jeongbin Kim, Thomas R. Palfrey, and Jeffrey R. Zeidel*

We develop a general framework for analyzing games where each player is a team and members of the same team all receive the same payoff. The framework combines noncooperative game theory with collective choice theory, and is developed for both strategic form and extensive form games. We introduce the concept of team equilibrium and identify conditions under which it converges to Nash equilibrium with large teams. We identify conditions on collective choice rules such that team decisions are stochastically optimal: the probability the team chooses an action is increasing in its equilibrium expected payoff. The theory is illustrated with some binary action games. (*JEL* C72, D71)

For most applications of game theory, 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

* Kim: Department of Economics, Florida State University (email: jbkimecon@gmail.com); Palfrey: Division of the Humanities and Social Sciences, California Institute of Technology (email: trp@hss.caltech.edu); Zeidel: Division of the Humanities and Social Sciences, California Institute of Technology (email: jzeidel@caltech.edu). John Asker was coeditor for this article. 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, the 2019 PECO Conference in Washington, DC, and at several departmental seminars. We thank Erkut Ozbay for his helpful discussion of the paper. We thank the reviewers for helpful comments and suggestions that improved the manuscript. Any remaining errors are our own.

† Go to https://doi.org/10.1257/mic.20200391 to visit the article page for additional materials and author disclosure statements or to comment in the online discussion forum.

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, 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, Kausel, and Kocher (2012, 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]

Given these extensive findings about group versus 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 similar team effects with respect to judgment biases and choice under uncertainty. 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: noncooperative game theory and collective choice theory.[3] Noncooperative 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, i.e., all players on the same team share the same payoff function.[4]

[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, Karni, and Levin 2010), forecasting errors (Blinder and Morgan 2005), and inefficient portfolio selection (Rockenbach, Sadrieh, and Mathauschek 2007).

[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 more distant connection with the economic theory of teams. See Marschak and Radner (1972), although the focus there is on other issues, such as communication costs, with no strategic interaction between different teams.

[4] The assumption of common values is motivated to a large extent by the many experimental studies of games played by teams of players, where all players on the same team receive exactly the same payoff. Given the extensive empirical findings in these pure common value settings, it 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. It should be clear that the focus here is *not* on applications

Collective choice theory provides an established 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]

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 *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.

where individual members of a group engage in costly private investments of different amounts for the benefit of some common outcome for the group, as for example in partnership games, voter turnout games, or more generally in public good contribution games where free riding plays a key role. Those settings already have their own extensive theoretical and empirical/experimental literatures.

[5] These errors could alternatively be interpreted as idiosyncratic additive payoff disturbances, as in quantal response equilibrium. 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, 1998), because there is no collective choice problem. That correspondence generally breaks down for teams with more than one member.

Many other collective choice rules that are included in our formulation do not have an obvious natural analog in social choice or voting theory. Various choice procedures involving direct communication between team members could be encoded in a collective choice rule. For example, an *average* rule would average the members' announced 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 ordinal-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, where—as in the theory presented here—team choices are determined by an exogenously specified collective choice rule and all members of the same team receive identical payoffs. There are two identifiable strands depending on whether the experimental task was a multiplayer game (such as the prisoner's dilemma), or a single-player decision problem (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 versus individuals by Charness and Sutter (2012) and Kugler, Kausel, and Kocher (2012), mentioned earlier.

The focus here is on the experimental studies of games rather than single-agent 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, Gomberg, and Sour (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, Kugler, and Ziegelmeyer (2004) find that teams "take" earlier than individuals in centipede games. In trust games, Kugler et al. (2007); Cox (2002); and Song (2008) 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 (2007) 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 percent lower bids by teams than individuals in Tullock contests, where individuals bid significantly above the Nash equilibrium. A similar finding is reported in Morone, Nuzzo, and

[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.

Caferra (2019) for all pay auctions. Group bidding behavior has also been investigated in auctions (Cox and Hayne 2006; Sutter, Kocher, and Strauß 2009). Most studies compare the behavior of individuals with the behavior of teams of 2 or 3 members. Variations in team size are not usually considered. An exception is Sutter (2005) in which an individual, a team with two members, and a team with four members play a beauty-contest game for four rounds. He finds that the behavior of four-member teams is closer to the Nash equilibrium action, while there is no significant difference between individuals and two-member teams.

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 fixed (i.e., nonstochastic) payoff functions. The team action is assumed to be the core of a voting rule. With this approach, existence of team equilibrium typically fails because of nonexistence of a core for many voting rules in many environments. Cason, Lau, and Mui (2019) 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, and in contrast to the results in this paper.

We first develop the formal theoretical structure of finite team games in strategic form, and provide a proof of the general existence of team equilibrium. In Section II, the effects 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 these examples illustrate how majority rule converges to Nash equilibrium with large teams, they also show that convergence is not necessarily monotone in team size; i.e., larger teams can lead to team equilibria further from Nash equilibrium. Furthermore, in mixed strategy equilibrium, individual voting probabilities within a team can be very different from the team mixed strategy equilibrium; in fact, if the Nash equilibrium is mixed, then individual voting probabilities converge to one-half, while the team equilibrium converges to the Nash equilibrium. Section III generalizes the finding in the examples in Section II that team equilibrium with large teams converges to Nash equilibrium under majority rule in 2×2 games. We prove that this *Nash convergence property* holds in *all* finite n-person games if the collective choice rules used by each team is a scoring rule. Section IV shows that under general conditions on collective choice rules, team choice behavior will satisfy two different kinds of stochastic rationality for all finite n-person games: *payoff monotonicity*, where the probability a team chooses a particular action is increasing in its equilibrium expected payoff; and *rank dependence*, where a team's choice probabilities will always be ordered by the expected payoffs of the actions. Section V generalizes the framework to extensive form games, establishes existence, and proves that the results about stochastic rationality and Nash convergence extend to arbitrary 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 VI, with a sequential prisoner's dilemma game and the four-move centipede game. Section VII discusses the results of the paper and points to some possible generalizations and extensions of the framework.

I. Team Games in Strategic Form

A team game is defined as follows. Let $\mathbf{T} = \{1, \ldots, t, \ldots, T\}$ be a collection of *teams*, where $t = \{i_1^t, \ldots, i_j^t, \ldots, i_{n^t}^t\}$, where i_j^t denotes member j of team t, and denote the *team size profile* by $n = (n^1, \ldots, n^T)$. Each team has a set of available actions, $A^t = \{a_1^t, \ldots, a_{K^t}^t\}$ and the set of action profiles is denoted $A = A^1 \times \cdots \times A^T$. The *payoff function* of the game for team t is given by $u^t : A \to \mathcal{R}$. 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 $U_k^t(\alpha) = \sum_{a^{-t} \in A^{-t}} \left[\prod_{t' \neq t} \alpha^{t'}\left(a^{t'}\right)\right] u^t\left(a_k^t, a^{-t}\right)$. 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 $\hat{U}_{ik}^t$ on α is understood.[7] We call $\hat{U}_i^t = \left(\hat{U}_{i1}^t, \ldots, \hat{U}_{iK^t}^t\right)$ i's estimated expected payoffs, and $\hat{U}^t = \left(\hat{U}_1^t, \ldots, \hat{U}_{n^t}^t\right)$ 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, $F = (F^1, \ldots, F^T)$, *admissible.* A *team collective choice rule*, C^t, is a correspondence that maps profiles of estimated expected payoffs in team t into a nonempty subset of elements of A^t. That is, $C^t : \mathcal{R}^{n^t K^t} \to \mathcal{A}^t$, where $\mathcal{A}^t$ is the set of nonempty subsets of A^t. Thus, for any α and ε^t, $C^t\left(\hat{U}^t\right) \in \mathcal{A}^t$. Different teams could be using different collective choice rules, and denote $C = (C^1, \ldots, C^T)$ the profile of collective choice rules.

For any strategic form game $G = [\mathbf{T}, A, u]$, and for any admissible F and profile of team choice rules C, call $\Gamma = [\mathbf{T}, A, n, u, F, C]$ a *team game in strategic form.* We assume that team t chooses randomly over $C^t\left(\hat{U}^t\right)$ when it is multivalued, according to the uniform distribution.[8] That is, the probability team t chooses a_k^t at $\hat{U}^t(\alpha)$ is given by the function $g_k^{C^t}$ defined as

$$g_k^{C^t}\left(\hat{U}^t\right) = \begin{cases} \dfrac{1}{\left|C^t\left(\hat{U}^t\right)\right|}, & \text{if } a_k^t \in C^t\left(\hat{U}^t\right); \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

[7] In some instances later we will write out the dependence on α explicitly to avoid ambiguity.

[8] The assumption that ties are broken fairly is made for convenience to reduce notation and to avoid artificially creating a source of bias into all collective choice rules. Ties could be broken by other means, for example by choosing the lowest index element of $C^t\left(\hat{U}^t\right)$. It is not essential to the results in the paper, except for Section IVB and Theorem 8, where neutrality of C^t is assumed.

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 $\bar{\hat{U}}_k^t = (1/n^t)\sum_{i \in T} \hat{U}_{ik}^t$ and define $C_{ave}^t(\hat{U}^t) = \{a_k^t \in A^t \mid \bar{\hat{U}}_k^t \geq \bar{\hat{U}}_l^t \text{ for all } l \neq k\}$ to be the set of actions that maximize $\bar{\hat{U}}^t$ at given values of α^{-t} and ε^t. This implies a team mixed strategy, g^{ave}, defined by

$$g_k^{ave}(\hat{U}^t) = \begin{cases} \dfrac{1}{|C_{ave}^t(\hat{U}^t)|}, & \text{if } a_k^t \in C_{ave}^t(\hat{U}^t); \\ 0, & \text{otherwise.} \end{cases}$$

Another example is *plurality rule*, which chooses the action for which the greatest number of team members estimate to have the highest payoff. That is, define $V_k^t(\hat{U}^t) = \#\{i \in t \mid \hat{U}_{ik}^t \geq \hat{U}_{il}^t \text{ for all } l \neq k\}$ and then $C_{pl}^t(\hat{U}^t) = \{a_k^t \in A^t \mid V_k^t(\hat{U}^t) \geq V_{k'}^t(\hat{U}^t) \text{ for all } k' \neq k\}$. Then g^{pl} is defined by

$$g_k^{pl}(\hat{U}^t) = \begin{cases} \dfrac{1}{|C_{pl}^t(\hat{U}^t)|}, & \text{if } a_k^t \in C_{pl}^t(\hat{U}^t); \\ 0, & \text{otherwise.} \end{cases}$$

A. *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 $\Gamma = [\mathbf{T}, A, n, u, F, C]$, we can define a *team response function* for team t, $P^{C^t}: \mathcal{R}^{K^t} \rightarrow \Delta A^t$, 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:

$$(2) \qquad P_k^{C^t}(U^t(\alpha)) = \int_{\epsilon^t} g_k^{C^t}(\hat{U}^t(\alpha)) dF^t(\epsilon^t),$$

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

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

THEOREM 1: *For every* Γ *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 $P_k^{C^t}$ 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^t}. Different communication mechanisms might correspond to different collective choice rules. For example, 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).[9]

More directly, the group might implement a collective choice rule such as plurality, qualified majority, or Borda count, by voting. The next section illustrates team equilibrium in 2×2 games under simple majority and qualified majority rule. The examples illustrate three different features of team equilibrium: (i) limiting properties of team equilibrium as teams become large; (ii) team equilibrium with a large team playing against a small team; and (iii) team equilibrium when teams use a supermajority collective choice rule. Most of the focus of the examples is on the first of these features, i.e., how does team equilibrium change as the team sizes increase?

The examples illustrate how an increase in team sizes can be understood conceptually in terms of two different effects. One effect, which is especially intuitive in 2×2 games where teams use majority rule, is *the consensus effect.* In this case, suppose, for some fixed α, $U_1^t(\alpha) > U_2^t(\alpha)$. Then for any admissible F, the probability that $\hat{U}_{i1}^t(\alpha) > \hat{U}_{i2}^t(\alpha)$ for an individual member of team t, which we denote by $p_1^t(\alpha)$ is greater than 1/2. Thus, *fixing* α^{-t}, under majority rule, $P_1^{Ct}\big(U^t(\alpha), n\big) = \sum_{k=((n+1)/2)}^{n} \binom{n}{k} \big[p_1^t(\alpha)\big]^k \big[1 - p_1^t(\alpha)\big]^{n-k} > 1/2$ 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.[10]

Thus, for any fixed strategy profile α, as n grows for a team, the estimated expected payoffs of different actions by individual members of a team converges on average to the "true" expected payoffs of the actions. Similarly, the individual

[9]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.

[10]If $U_1^t(\alpha) = U_2^t(\alpha)$, then $p_1^t(\alpha)$ is exactly equal to 1/2 and $P_1^{Ct}\big(U^t(\alpha), n\big) = 1/2$ for all n, so there is no consensus effect.

members' ranking of an action's estimated expected payoff converges to its true ranking, given α. Thus, for many rules, such as plurality rule, scoring rules, and the average rule, the collective choice of large teams will reflect a consensus about the relative expected payoffs of the various actions, as if all members shared common and correct beliefs about $U^t(\alpha)$. In the case of majority rule, this is simply a consensus about whether $U_1^t(\alpha)$ is greater than, less than, or equal to $U_2^t(\alpha)$.

Of course the equilibrium analysis is more complicated than this. One cannot take α as fixed when n changes, because of the second effect, *the equilibrium effect*. As n changes, the team equilibrium α^* changes, so we denote its dependence on n here by α_n^*. 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 under majority rule $p_1^{*t}(\alpha_n^*) \neq p_1^{*t}(\alpha_{n'}^*)$, which feeds back and affects each team's mixed strategy response. Thus, while the consensus effect only looks at how changes in one team's size affect the actions of that team, fixing the actions of the other teams, the equilibrium effect takes into account that as team size changes (even for a single team) the actions frequencies of *all* teams will typically change. Because it is an indirect rather than a direct effect, the equilibrium effect can produce some unintuitive consequences for some games and some voting rules, and it is possible that the equilibrium effect can dampen or even work in the opposite direction of the consensus effect.

What ultimately happens in the limit of team equilibrium with large teams depends on the relative strength of these two effects, which will typically be game-dependent and rule-dependent. In Section III, we show that these two effects interact in a way such that, for a broad class of collective choice rules, in all games every limit point of team equilibria as teams grow without bound is a Nash equilibrium of the game, which we call *Nash convergence*. Nash convergence can fail to hold if the collective choice rules are nonneutral, i.e., biased in favor or against certain actions. The next section provides an example illustrating such a failure of Nash convergence with a nonneutral rule, where a sequence of team equilibria in the prisoner's dilemma game converges to a mixed strategy in the limit.[11]

The interaction of these two effects can also lead to other surprising properties of team equilibrium. For example, because the equilibrium effect can work in the opposite direction of the consensus effect, with relatively small team sizes, the effect of increasing team size can drive the team equilibrium away from the Nash equilibrium of the underlying normal form game. This can even arise in games that are strictly dominance solvable with a unique rationalizable strategy profile. In 2×2 games where team equilibria with majority rule converge to a mixed strategy Nash equilibrium, the consensus effect essentially disappears with large teams, because the individual voting probabilities converge to $1/2$ as the expected payoffs of the

[11] Nash convergence also generally fails with large teams if a team's collective choice rule is dictatorial. This is shown in one of the examples below. A related question is whether the limit points of team equilibria are restricted by familiar refinements of Nash equilibrium, such as trembling-hand perfection and proper equilibrium. The answer is negative. In particular, there are examples of simple games with limit points of team equilibria under plurality rule that are not perfect (and hence not proper). The refinement implied by limit points of team equilibrium seems more closely related to approachability in the sense of Harsanyi (1973).

two actions approach equality. These and other phenomena that can arise in team equilibrium are illustrated with examples in the next section.

II. Team Size Effects in 2×2 Games

The model is easiest to illustrate in the simple case of 2×2 games with majority rule as the collective choice rule for each team. 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. 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 chooses 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 $\hat{U}_{ik}^t = U_k^t(\alpha) + \varepsilon_{ik}^t$ where $U_k^t(\alpha) = \alpha^{-t} u^t(a_k^t, a_1^{-t}) + (1 - \alpha^{-t}) u^t(a_k^t, a_2^{-t})$. Because there are only two actions for each team and majority rule depends only on each member's ranking of the estimated payoff of each action, 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.[12] 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 $z \in \mathcal{R}$, $H^t(z) = 1 - H^t(-z)$, implying $H(0) = 1/2$. Thus, we can write $\Delta\hat{U}_i^t(\alpha) \equiv \hat{U}_{i1}^t(\alpha) - \hat{U}_{i2}^t(\alpha) = U_1^t(\alpha) - U_2^t(\alpha) + \varepsilon_i^t$.

A team choice rule, C^t, maps each profile of individual estimated expected payoff differences, $\Delta\hat{U}^t(\alpha) = \left(\Delta\hat{U}_1^t(\alpha), \ldots, \Delta\hat{U}_{n^t}^t(\alpha)\right)$ into the nonempty subsets of A^t, and g_C^t randomly selects one of these choices with equal probability. We next illustrate team size effects under majority rule in three different kinds of 2×2 games.

A. *The Prisoner's Dilemma (PD) 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.

Suppose that both teams have n (odd) members and the same distribution of estimation error differences, H. Let α_n^* be a symmetric team equilibrium probability of either team choosing D in the game, and $p_n^* = \Pr\left(U_C^t(\alpha^*) - U_D^t(\alpha^*) < \epsilon_D^t - \epsilon_C^t\right)$ be a symmetric team equilibrium probability that any player votes for action D. Equilibrium requires the behavior of each team to solve the following equations simultaneously:

$$\alpha_n^* = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} (p_n^*)^k (1 - p_n^*)^{n-k}, \qquad p_n^* = H\left(x - (x-y)\alpha_n^*\right).$$

[12] Some collective choice rules depend on more than just the profile of estimated payoff differences. An example of such a collective choice rule is the maximin rule:

$$C_{\max\min}^t\left(\hat{U}^t\right) = \left\{a_k^t \in A^t \,|\, \min_i\left\{\hat{U}_{ik}^t\right\} \geq \min_i\left\{\hat{U}_{il}^t\right\} \forall l \neq k\right\},$$

i.e., select the action that has the highest minimum estimated expected payoff.

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)	$5+x, 3-y$	$3, 3$

If $x \neq y$, then p_n^* will generally depend on n, so there is an equilibrium effect in addition to the consensus effect. The equilibrium and consensus effects can go in opposite directions. For example, if $x > y$, the effect of increasing n is buffered by the countervailing effect that increasing α_n^* leads to a decrease in p_n^*. To see this formally, notice that the equilibrium condition for p_n^* depends on α_n^* according to $p_n^* = H\big(x - (x-y)\alpha_n^*\big)$ which is strictly decreasing in α_n^* precisely when $x > y$. Since $\big(x - (x-y)\alpha_n\big) > 0$ for all x, y, it is easy to see that as n grows large, $\lim_{n\to\infty}(\alpha_n^*) = 1$ and $\lim_{n\to\infty}(p_n^*) = H(y)$.

When $x = y$ the analysis is straightforward and intuitive, because the expected payoff difference between Defect and Cooperate for either team is $U_D^t(\alpha_n^*) - U_C^t(\alpha_n^*) = x > 0$, which does not depend on α_n^* and hence is independent of n. In any team equilibrium, the probability that any team member votes for D is $H(x) > 1/2$. Thus, in this special case, changes in the team equilibrium as n increases are entirely due to the consensus effect.

Team Equilibrium in PD Games with Different Team Sizes.—Consider the case where $x > y$, the size of the row team is fixed at 1, and the size of the column team, n, is variable.[13] In this case, for any n the equilibrium is of the form $\big(q_n^*, \alpha_n^*, p_n^*\big)$, where q_n^* denotes the single row team member's defect probability, which depends on the column team's defect probability, α_n^*, and p_n^* denotes a column team member's probability of defection, which depends on q_n^*.

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

$$q_n^* = H\big(x + (y-x)\alpha_n^*\big),$$

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

$$p_n^* = H\big(x + (y-x)q_n^*\big).$$

Since $x > y$, as noted above, the row team's defection probability, q_n^*, and the column team's defection probability, α_n^*, move in opposite directions. Similarly, q_n^* and p_n^* also move in opposite directions. Because the column team members'

[13] This example also serves as an illustration where the two teams use different collective choice rules, with the row team using a dictatorial rule and the column team using simple majority rule.

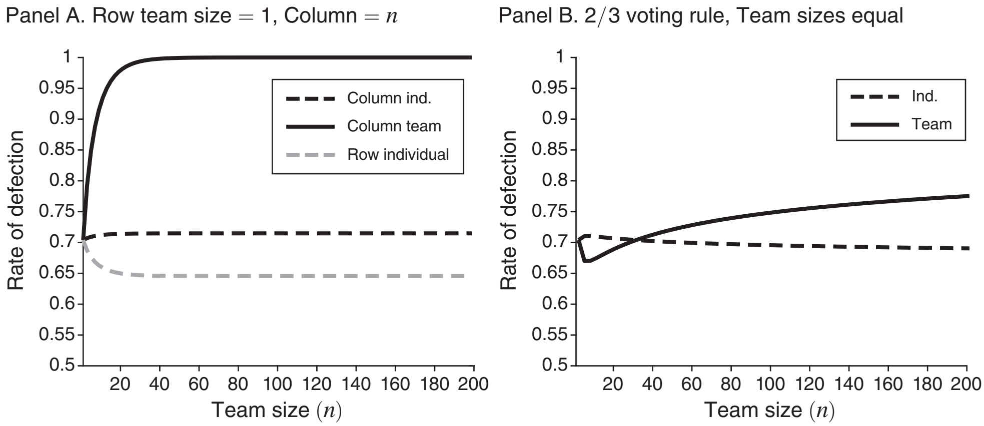

FIGURE 1. TEAM EQUILIBRIUM IN PD AS A FUNCTION OF $n(x = 5, y = 2)$

defection probability, p_n^* is strictly bounded above 0.5, α_n^* must eventually increase to 1 as n increases without bound, so the consensus effect eventually dominates the equilibrium effect. Similarly, the row team's defection probability, q_n^*, must eventually decrease to a limiting value of $H(y)$, and the column team members' defection probability, p_n^* eventually increases to a limiting value of $H\big(x + (y - x)H(y)\big)$.[14]

This is illustrated in panel A of Figure 1, which displays the team equilibrium values, $(\alpha_n^*, q_n^*, p_n^*)$, as a function of n for $x = 5$, $y = 2$ and $H(z) = 1/(1 + e^{-0.3z})$, for column team sizes up to 200 (row team size fixed at 1).

Equilibrium with PD Games with Different Voting Rules.—Even in games as simple as the PD game, the equilibrium effects in team games can be quite subtle. To illustrate this, we consider the same game as above, except where the teams are the same size and both teams use a supermajority voting rule: the team action is to defect if at least 2/3 of the team members favor defection; otherwise the team cooperates. In this case, the equilibrium conditions for a team equilibrium are

$$\alpha_n^* = \sum_{k=\lceil 2n/3 \rceil}^{n} \binom{n}{k} (p_n^*)^k (1 - p_n^*)^{n-k}, \qquad p_n^* = H\big(x + (y - x)\alpha_n^*\big),$$

where $\lceil 2n/3 \rceil$ denotes the least integer greater than or equal to $2n/3$. For the parameters used in panel A $(x = 5,\ y = 2$ and $H(z) = 1/(1 + e^{-0.3z}))$, we get the surprising result that α_n^* does not converge to 1. The logic behind this is that if α_n^* did converge to 1 then in the limit we would have $p_n^* \to H(y) = 1/(1 + e^{-0.3y}) \approx 0.65 < 2/3$, which (given the 2/3 rule) would imply that $\alpha_n^* \to 0$, a contradiction. A similar argument shows that α_n^* cannot converge to 0 either. For if α_n^* did converge to 0 then in the limit we would have $p_n^* \to H(x) = 1/(1 + e^{-0.3x}) \approx 0.82 > 2/3$, which would imply that $\alpha_n^* \to 1$, a contradiction.

[14] We conjecture that the three equilibrium values, α_n^*, q_n^*, and p_n^* each converge monotonically.

Table 2—Weak Prisoner's Dilemma (WPD) Game

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

Hence, in the limit with large n a symmetric team equilibrium must converge to a *mixed* equilibrium![15] The only way this can happen is if $p_n^* \to 2/3$. Hence, the limiting equilibrium team probability of defection is calculated from the second equilibrium condition: $p_\infty^* = 2/3 = H\big(x - (x-y)\alpha_\infty^*\big)$, which gives $\alpha_\infty^* \approx 0.89$. Panel B of Figure 1 displays the team equilibrium (p_n^*, α_n^*) as a function of n for $x = 5$, $y = 2$ and $H(z) = 1/\big(1 + e^{-0.3z}\big)$, for column team sizes up to 200. Convergence is much slower because the limiting team defect strategy is mixed.

B. *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 the prisoner dilemma game displayed in Table 1, which we call the *Weak Prisoner's 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. 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, and for this example, we assume $x = y$. The third parameter, z, is the payoff gain the column player gets from cooperating if the row player cooperates.

Let $\alpha_n^* = \big(\alpha_n^{*1}, \alpha_n^{*2}\big)$ 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, we have for the row team:

$$\alpha_n^{*1} = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} \big(p_n^{*1}\big)^k \big(1 - p_n^{*1}\big)^{n-k}, \qquad p_n^{*1} = H(x).$$

For a player on the column team, the probability of voting for D, p_n^{*2}, varies with n, as it depends directly on α_n^*. So the two equations for the column team are

$$\alpha_n^{*2} = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} \big(p_n^{*2}\big)^k \big(1 - p_n^{*2}\big)^{n-k}, \qquad p_n^{*2} = H\big(\alpha_n^{*1}(x+z) - z\big).$$

[15] There are also asymmetric team equilibria that converge to pure strategies in the limit, where one team defects and the other team cooperates.

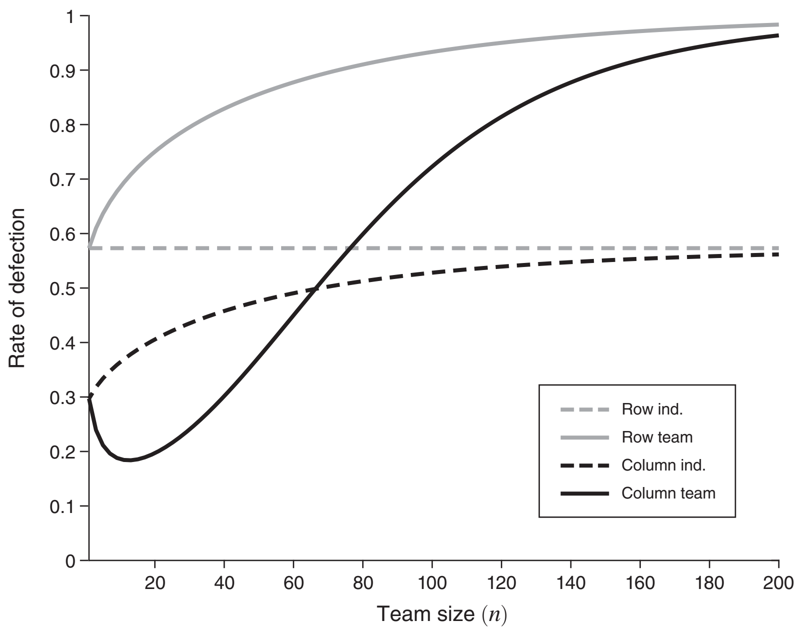

Figure 2. Team Equilibrium in Weak Prisoner's Dilemma

For any values of n, x, z, and H, a team equilibrium is given by any solution of this system of four equations.

Figure 2 illustrates how the equilibrium voting and team choice probabilities $\left(p_n^{*1}, \alpha_n^{*1}, p_n^{*2}, \alpha_n^{*2}\right)$ vary with n, for the parameters $x = 1$, $z = 8$ and $H(z) = 1/\left(1 + e^{-0.3z}\right)$ for team sizes up to 200. Notice that for relatively small teams sizes (< 20) the column team becomes *more* cooperative as it grows. This results from a combination of the consensus effect (since $p_n^{*2} < 0.5$) and the equilibrium effect (since α_n^{*1} is increasing).

C. *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 as asymmetric matching pennies (AMP) games. The payoffs are displayed in Table 3, where $a > c, b < d, w < x, y > z$:

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

$$p_n^{*1} = H\left(a\left(1 - \alpha_n^{*2}\right) + b\alpha_n^{*2} - c\left(1 - \alpha_n^{*2}\right) - d\alpha_n^{*2}\right). \tag{3}$$

TABLE 3—ASYMMETRIC MATCHING PENNIES GAME

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

The probability that a column team member estimates that R is better than L is equal to

$$p_n^{*2} = H\Big(z\big(1-\alpha_n^{*1}\big) + x\alpha_n^{*1} - y\big(1-\alpha_n^{*1}\big) - w\alpha_n^{*1}\Big). \tag{4}$$

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

$$\alpha_n^{*1} = \sum_{k=\frac{n+1}{2}}^{n} \binom{n}{k} \big(p_n^{*1}\big)^k \big(1-p_n^{*1}\big)^{n-k}, \tag{5}$$

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

The team equilibrium is obtained by solving equations (3), (4), (5), and (6) simultaneously for p_n^{*1}, p_n^{*2}, α_n^{*1}, and α_n^{*2}. The team equilibrium and the equilibrium individual voting probabilities are displayed in Figure 3 for the parameters $w = b = c = z = 0$, $d = y = 1$, $a = 5$, $x = 0.5, 1.0, 2.0$, and $H(z) = 1/\big(1+e^{-0.3z}\big)$ for team sizes up to size 200.

From these examples, we see that $\lim_{n\to\infty}\alpha_n^* = \big((y-z)/(y-z+x-w), (a-c)/(a-c+d-b)\big)$, the unique Nash equilibrium, in all three cases. As will be proved in Section III, when teams use majority rule, convergent sequences of team equilibria must converge to a Nash equilibrium, so this is a general feature of team equilibria of any 2×2 game with a unique mixed strategy equilibrium.

An interesting implication is that in the limit of team equilibria, *individual team members are voting randomly*. That is, $\lim_{n\to\infty}p_n^{*1} = \lim_{n\to\infty}p_n^{*2} = 1/2$, which is necessarily the case because the expected payoffs of the two strategies are equal in the Nash equilibrium limit. The argument is similar to the earlier illustration of a mixed equilibrium in the PD game with a two-thirds voting rule. Suppose to the contrary that $\lim_{n\to\infty}p_n^{*1} > 1/2$. Then $\lim_{n\to\infty}\alpha_n^{*1} = 1$, implying that the right-hand side of equation (4) converges to $H(x-w) > 1/2$ because $x - w > 0$, implying $\lim_{n\to\infty}p_n^{*2} > 1/2$, so $\lim_{n\to\infty}\alpha_n^{*2} = 1$. This in turn implies that the right-hand

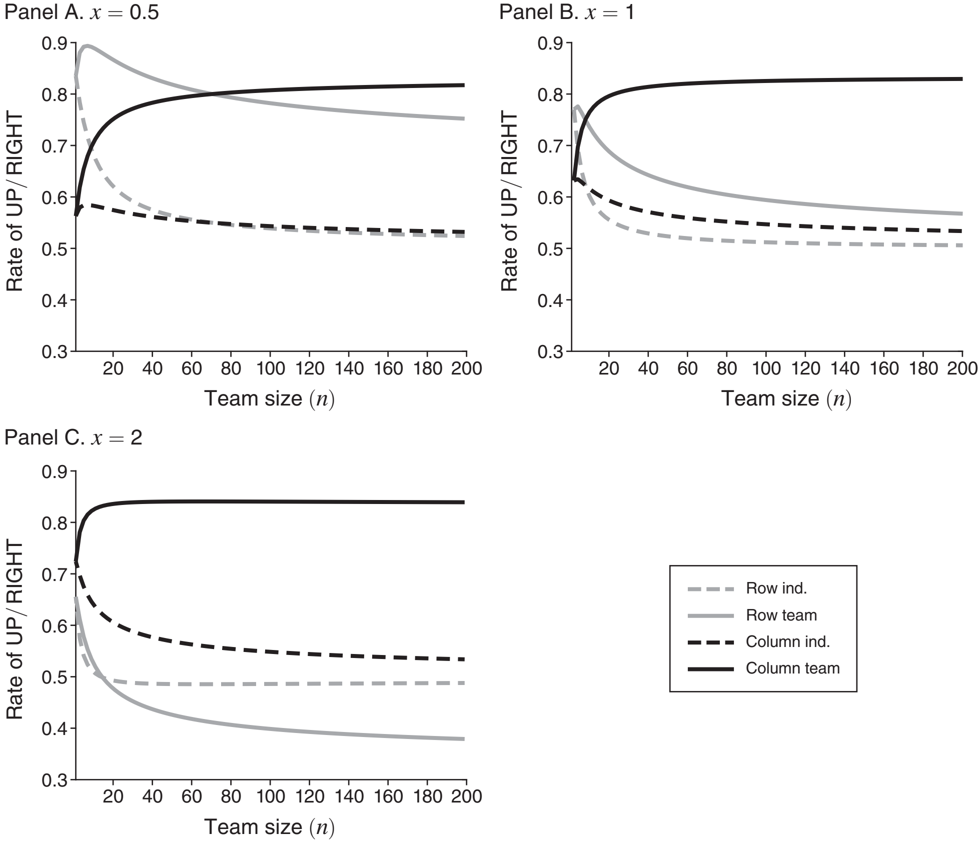

Figure 3. Team Equilibrium in Asymmetric Matching Pennies

side of equation (3) converges to $H(b-d) < 1/2$ because $b-d < 0$, and hence $\lim_{n\to\infty} p_n^{*1} < 1/2$, a contradiction. A similar contradiction arises if one were to suppose instead that $\lim_{n\to\infty} p_n^{*1} < 1/2$.

In this class of games, the equilibrium effect works in the opposite direction and thereby dampens the consensus effect, even in the limit. Without the equilibrium effect, which pushes p_n^* to exactly $1/2$ for both teams, the reinforcement effect by itself would lead to a pure strategy team actions in the limit. This is illustrated starkly in the examples shown in Figure 3, where *for all n*, $p_n^{*t} > 1/2$ for both teams, yet $\lim_{n\to\infty} \alpha_n^* \neq (1,1)$.

III. 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 equilibria converges to a Nash equilibrium. In the examples of the previous section, it was the case that the Nash convergence held if both teams used majority rule and the sizes of both teams increased without bound. This observation raises the more general question whether majority rule or 2×2 games are somehow unique in this regard, or if it is a property shared by a broader class of collective choice rules and a broader class of games. It is clearly not unique, as it is not difficult to show that the average rule also has this property. Thus, the challenge is to characterize, at least partially, the class of collective choice rules and game environments with this *Nash convergence property.*

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

DEFINITION 2: *An* Individual Scoring Function $S_i^t : A^t \times \mathcal{R}^{K^t} \to \mathcal{R}^{K^t}$, *is a function defined such that* $S_{ik}^t\left(\hat{U}_i^t\right) = s_{im}$ *whenever* $\left|\left\{a_l^t \in A^t : \hat{U}_{il}^t > \hat{U}_{ik}^t\right\}\right| = m - 1$, *for some given set of* K^t *scores* $s_{i1}^t \geq s_{i2}^t \geq \ldots \geq s_{iK^t}^t \geq 0$ *such that* $s_{i1}^t > s_{iK^t}^t$.

DEFINITION 3: *A team collective choice rule* C^t *is an* Anonymous Scoring Rule *if there exists a profile of individual scoring functions,* $\left(S_1^t, \ldots, S_{n^t}^t\right)$ *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* $\sum_{i=1}^{n^t} \sigma_k^t\left(\hat{U}_i^t\right) \geq \sum_{i=1}^{n^t} \sigma_l^t\left(\hat{U}_i^t\right)$ *for all* $l \neq k$.

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. In an anonymous scoring rule, all members of the team have the same individual scoring function. Examples of common anonymous scoring rules include: plurality rule, where $s_{i1}^t = 1$ and $s_{im}^t = 0$ for all $i \in t$ and for all $m > 1$, 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 1), so ties in an individual member's ordinal rankings are irrelevant.

THEOREM 2: *Consider an infinite sequence of team games,* $\left\{\Gamma_m\right\}_{m=1}^{\infty}$ *such that* (1) $A_m^t = A_{m'}^t = A^t$ *for all* t, m, m'; (2) $u_m^t = u_{m'}^t = u^t \; \forall t, m, m'$; (3) $n_{m+1}^t > n_m^t$ *for all* m, t; *and* (4) $C_m^t = \sigma$, *an anonymous scoring rule, for all* m, t. *Let* $\left\{\alpha_m\right\}_{m=1}^{\infty}$ *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* $\left[\mathbf{T}, A, u\right]$.

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 $U_k^t\left(\alpha^*\right) > U_l^t\left(\alpha^*\right)$, 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, for all α, and for all $a_k^t \in A^t$, $a_k^t \in C^t\left(\bar{U}^t(\alpha)\right)$ if and only if

$$\sum_{i=1}^{n^t} \sigma_k\left(\hat{U}_i^t(\alpha)\right) \geq \sum_{i=1}^{n^t} \sigma_l\left(\hat{U}_i^t(\alpha)\right) \quad \text{for all} \quad l \neq k$$

$$\Leftrightarrow$$

$$\bar{\sigma}_{kn^t}^t\left(\hat{U}^t(\alpha)\right) \geq \bar{\sigma}_{ln^t}^t\left(\hat{U}^t(\alpha)\right) \quad \text{for all} \quad l \neq k,$$

Table 4—Example of a Nash Equilibrium That Is Not a Limit of Team Equilibria with Majority Rule

	Column Team (2)	
Row Team (1)	Left (L)	Right (R)
Up (U)	1, 1	0, 0
Down (D)	0, 0	0, 0

where $\bar{\sigma}^t_{kn^t}\left(\hat{U}^t(\alpha)\right)$ denotes the average score of a^t_k among the n^t members of t at $\hat{U}^t(\alpha)$. That is, the score of a^t_k is maximal if and only if the average individual score of a^t_k is maximal. If $U^t_k(\alpha) > U^t_l(\alpha)$ then σ_k stochastically dominates σ_l, and hence their respective expected scores are strictly ordered. That is, $E\left\{\sigma_k\left(\hat{U}^t_i(\alpha)\right)\right\} > E\left\{\sigma_l\left(\hat{U}^t_i(\alpha)\right)\right\}$, where

$$E\left\{\sigma_k\left(\hat{U}^t_i(\alpha)\right)\right\} = \int \sigma_k\left(U^t(\alpha) + \epsilon^t_i\right) dF^t\left(\epsilon^t_i\right).$$

Hence, at the limiting strategy profile, α^*, we have $E\left\{\sigma_k\left(\hat{U}^t_i(\alpha^*)\right)\right\} > E\left\{\sigma_l\left(\hat{U}^t_i(\alpha^*)\right)\right\}$. Since $U^t(\alpha_m) \to U^t(\alpha^*)$ and $E\{\sigma\}$ is continuous in α, $\bar{\sigma}^t_{kn^t_m}\left(\hat{U}^t(\alpha_m)\right) \to E\left\{\sigma^t_k\left(\hat{U}^t(\alpha^*)\right)\right\}$ and $\bar{\sigma}^t_{ln^t_m}\left(\hat{U}^t(\alpha_m)\right) \to E\left\{\sigma^t_l\left(\hat{U}^t(\alpha^*)\right)\right\}$ in probability as $m \to \infty$. Therefore, $E\left\{\sigma_k\left(\hat{U}^t_i(\alpha^*)\right)\right\} > E\left\{\sigma_l\left(\hat{U}^t_i(\alpha^*)\right)\right\}$ implies that there exists $\bar{m}$ such that $\Pr\left\{\bar{\sigma}^t_{kn^t_m}\left(\hat{U}^t(\alpha_m)\right) \leq \bar{\sigma}^t_{ln^t_m}\left(\hat{U}^t(\alpha_m)\right)\right\} < (\xi/2)\,\forall m > \bar{m}$. This leads to a contradiction to the initial hypothesis that $\alpha^{t*}_l = \xi > 0$ since $\Pr\left\{\bar{\sigma}^t_{kn^t_m}\left(\hat{U}^t(\alpha_m)\right) \leq \bar{\sigma}^t_{ln^t_m}\left(\hat{U}^t(\alpha_m)\right)\right\} < (\xi/2)\,\forall m > \bar{m}$, implies that $\alpha^t_{ln^t_m} < (\xi/2)\,\forall m > \bar{m}$ and hence $\alpha^{t*}_l < \xi$. ∎

It is useful to clarify the implications and generality of the result with a few comments. First, the theorem does not imply that team equilibrium with larger teams is necessarily closer to Nash equilibrium than equilibrium with smaller teams. It is an asymptotic result for large teams. Examples in the last section show that the convergence can be nonmonotonic even in very simple games.

Second, Nash convergence is an upper hemicontinuity property of the team equilibrium correspondence, but that correspondence is not generally lower hemicontinuous. Some Nash equilibria are not approachable, just as some weak Nash equilibria fail to be limit points of payoff disturbed games (Harsanyi 1973) or limit points of quantal response equilibria as the error terms vanish (McKelvey and Palfrey 1995). In the following game, (D, R) is a Nash equilibrium that cannot be approached by a sequence of large team equilibria. It is easy to see that in any team equilibrium, for any n, the probability the row team plays Up and the probability the column team plays Left is always greater than 0.5. Hence, (D, R) cannot be a limit of team equilibria.

Third, the 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. Fourth, the rate of convergence to large teams can differ across teams. Fifth, 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. Finally, we conjecture that the class is much broader than scoring rules, including many other anonymous and neutral collective choice rules. 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, such as weighted average rules.

IV. Stochastic Rationality and Team Response Functions

In our framework, individual team member estimated expected utilities obey two intuitive, normatively appealing properties of stochastic rationality. First, the probability a member of team t ranks action a_k^t as having the highest estimated expected payoff (and hence would be chosen if it were an individual choice problem) is increasing in $U_k^t(\alpha)$, the "true" equilibrium expected payoff of action k, ceteris paribus, a condition we call *payoff monotonicity*. Second, the probability a member of team t ranks action a_k^t's estimated expected payoff as highest is greater than the probability the member ranks action a_l^t's highest if and only if $U_k^t(\alpha) > U_l^t(\alpha)$, a condition we call *rank dependence*.[16] In the context of team decision-making, it is the collective choice rule in combination with the error structure that determines team choice probabilities. It is easy to see that team decision-making will not generally inherit these two properties for all collective choice rules. Given the normative appeal of these two properties, this naturally leads to the following question. Under what conditions on the collective choice rule will team response functions satisfy them? In addition to the normative appeal of payoff monotonicity and rank dependence, violations of these two properties might affect the incentives faced by team members during the team decision-making process. While our framework does not explicitly model individual choice behavior in the team decision-making process, it is plausible that stochastically irrational collective choice rules that violate payoff monotonicity or rank dependence, could hinder the ability for teams to effectively aggregate members' diverse beliefs. For example, if payoff monotonicity fails for the prescribed collective choice rule, then individual team members might profit by behaving as if their expected utility estimates for some of the actions are lower or higher than they are in truth.

This section identifies restrictions on the collective choice rule that guarantee team response functions to satisfy these two properties. We first show that payoff monotonicity of team response functions requires only two weak assumptions on the collective choice rule, unanimity and positive responsiveness. On the other hand, rank dependence holds only for a more restricted class of neutral collective choice rules. Many nonneutral collective choice rules, such as those that give a status quo

[16] This follows from the i.i.d. assumption on estimation errors. 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 individual choice probabilities in a quantal response equilibrium.

advantage to an action, will fail to satisfy rank dependence, as the last example in Section IIA demonstrates. Second, 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.

A. *Payoff Monotonicity*

For any team game, Γ, P^{C^t} depends on the strategy profile α, 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^{C^t} to be payoff monotone for all admissible F^t. The formal definition of payoff monotonicity is given below.

DEFINITION 4: *A team collective choice rule C^t satisfies* Payoff Monotonicity *if, for all a_k^t, α, α':*

$$U_k^t(\alpha) > U_k^t(\alpha') \quad \text{and} \quad U_l^t(\alpha) = U_l^t(\alpha') \ \forall l \neq k \Rightarrow$$

$$P_k^{C^t}\big(U^t(\alpha)\big) > P_k^{C^t}\big(U^t(\alpha')\big).$$

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 that a_k^t has the highest expected utility, then it is uniquely chosen by C^t.[17]

DEFINITION 5: *A team collective choice rule C^t satisfies* Unanimity *if:*

$$\hat{U}_{ik}^t > \hat{U}_{il}^t \quad \textit{for all} \quad i \in t \quad \textit{and for all} \quad l \neq k \Rightarrow C^t\big(\hat{U}^t\big) = \big\{a_k^t\big\}.$$

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 6: *A profile $\hat{U}^t$ 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 $\tilde{U}_{ik}^t \geq \hat{U}_{ik}^t$ and $\tilde{U}_{il}^t = \hat{U}_{il}^t$ for all $l \neq k$.*

[17] 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 of the strategic form game, $[T, A, u]$.

DEFINITION 7: *A team collective choice rule C^t satisfies* Positive Responsiveness *if $a_k^t \in C^t(\hat{U}^t) \Rightarrow a_k^t \in C^t(\tilde{U}^t) \subseteq C^t(\hat{U}^t)$, for all a_k^t, $\hat{U}^t$ and all monotonic transformations $\tilde{U}^t$ of $\hat{U}^t$ with respect to action a_k^t.*

This definition of Positive Responsiveness is essentially a cardinal version of the usual definition of positive responsiveness from the social choice literature. 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.

Many collective choice rules satisfy positive responsiveness. For example, 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. In this section we consider a class of collective choice rules, called *generalized* scoring rules, and show that positive responsiveness is satisfied for any such collective choice rule. A generalized scoring rule is substantially more general than the standard definition of a scoring rule in the social choice literature, which was defined in the previous section as an *anonymous* scoring rules (i.e., all the individual scoring functions are the same). Generalized scoring rules relax the anonymity requirement that all individual scoring functions are the same. It includes a wide range of non-anonymous collective choice rules, including dictatorial rules.

DEFINITION 8: *A team collective choice rule C^t is a* Generalized Scoring Rule *if there exists a profile of individual scoring functions, $(S_1^t, \ldots, S_{n^t}^t)$, such that for all $a_k^t \in A^t$ and for all $\hat{U}^t \in \mathcal{R}^{K^t n^t}$, $a_k^t \in C^t(\hat{U}^t)$ if and only if $\sum_{i=1}^{n^t} S_{ik}^t(\hat{U}_i^t) \geq \sum_{i=1}^{n^t} S_{il}^t(\hat{U}_i^t)$ for all $l \neq k$.*

PROPOSITION 1: *All generalized scoring rules satisfy positive responsiveness.*[18]

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 3: *If F^t is admissible and C^t satisfies unanimity and positive responsiveness then P^{C^t} satisfies payoff monotonicity.*

PROOF:

Let C^t satisfy positive responsiveness and unanimity and F^t admissible. Suppose that $U_k^t - U_k'^t = \delta > 0$, and $U_l^t = U_l'^t, \forall l \neq k$. Then for all realizations of

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

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_k^t. So by positive responsiveness of C^t we have that if $a_k^t \in C^t(U'^t + \epsilon^t)$ then $a_k^t \in C^t(U^t + \epsilon^t)$, and if $a_l^t \in C^t(U^t + \epsilon^t)$ then $a_l^t \in C^t(U'^t + \epsilon^t)$. So $g_k^{C^t}(U^t + \epsilon^t) \geq g_k^{C^t}(U'^t + \epsilon^t)$ for all ϵ^t, and therefore $P_k^{C^t}(U^t) \geq P_k^{C^t}(U'^t)$. To show the strict inequality, $P_k^{C^t}(U^t) > P_k^{C^t}(U'^t)$, we show that there exists a region $\beta \subset \mathcal{R}^{K^t \times N^t}$ with positive measure such that if $\epsilon^t \in \beta$, then $g_k^{C^t}(U^t + \epsilon^t) > g_k^{C^t}(U'^t + \epsilon^t)$. In particular, unanimity of C^t is used as follows to construct β such that if $\epsilon^t \in \beta$, then $g_k^{C^t}(U^t + \epsilon^t) = 1 > g_k^{C^t}(U'^t + \epsilon^t) = 0$. That is, such that a_k^t is uniquely chosen under $U^t + \epsilon^t$, and not chosen under $U'^t + \epsilon^t$. Let $\tilde{U}^t$ be an estimated expected utility profile such that all team members strictly prefer some action a_l^t to action a_k^t, all members prefer a_k^t to all other actions a_m^t (i.e., all members rank a_k^t second), and for all members we have $\tilde{U}_l^t - \tilde{U}_k^t = \delta/2$. Define:

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

Then if $\epsilon^t \in \beta$, by unanimity we have $C^t(U'^t + \epsilon^t) = \{a_l^t\}$ and $C^t(U^t + \epsilon^t) = \{a_k^t\}$, so $g_k^{C^t}(U^t + \epsilon^t) = 1 > g_k^{C^t}(U'^t + \epsilon^t) = 0$. β is an open set and hence has positive measure since the distribution of ϵ^t has full support. Therefore, $P_k^{C^t}(U^t) > P_k^{C^t}(U'^t)$, as desired. ∎

B. *Rank Dependence*

In this section we show that P^{C^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. The formal definition of rank dependence follows.

DEFINITION 9: *A team collective choice rule C^t satisfies* Rank Dependence *if, for all $a_k^t, a_l^t, \alpha, U_k^t(\alpha) > U_l^t(\alpha) \Rightarrow P_k^{C^t}(U^t(\alpha)) > P_l^{C^t}(U^t(\alpha))$.*

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.

Let $\psi : A^t \to A^t$ be any permutation of team actions. Denote by $U^{t,\psi} = (U_{\psi 1}^t, \ldots, U_{\psi(K^t)}^t)$ the permuted profile of expected utilities and by $\hat{U}^{t,\psi} \equiv (U_{\psi(1)}^t + \epsilon_{\psi(1)}^t, \ldots, U_{\psi(K^t)}^t + \epsilon_{\psi(K^t)}^t)$ the permuted profile of estimated expected utilities. We can then define neutrality formally.

DEFINITION 10: *A team collective choice rule C^t satisfies* Neutrality *if, for all a_k^t, $\hat{U}^t$, for all permutations ψ, $a_k^t \in C^t(\hat{U}^t) \Leftrightarrow a_{\psi(k)}^t \in C^t(\hat{U}^{t,\psi})$.*

Neutrality, along with admissibility of F^t, imply that when the expected payoffs of team actions are permuted, the team choice probabilities are permuted.

LEMMA 1: *If F^t is admissible and C^t satisfies neutrality, then $P_k^{C^t}\left(U^t\right) = P_{\psi(k)}^{C^t}\left(U^{t,\psi}\right)$ for all expected utility profiles, U^t, actions, a_k^t, and permutations ψ.*

PROOF:

By neutrality of C^t, for any expected utility profile, U^t, action, a_k^t, belief error profile, ϵ^t, and permutation ψ, $g_k^{C^t}\left(\hat{U}^t\right) = g_{\psi(k)}^{C^t}\left(\hat{U}^{t,\psi}\right)$, that is the probability that a_k^t is chosen at $\hat{U}^t$ is equal to the probability that $a_{\psi(k)}^t$ is chosen at $\hat{U}^{t,\psi}$. Therefore $P_k^{C^t}\left(U^t\right) = \int_{\epsilon^t} g_k^{C^t}\left(\hat{U}^t\right)dF^t\left(\epsilon^t\right) = \int_{\epsilon^t} g_{\psi(k)}^{C^t}\left(\hat{U}^{t,\psi}\right)dF^t\left(\epsilon^t\right)$. Finally, since the estimation errors are i.i.d. $\int_{\epsilon^t} g_{\psi(k)}^{C^t}\left(\hat{U}^{t,\psi}\right)dF^t\left(\epsilon^t\right) = \int_{\epsilon^t} g_{\psi(k)}^{C^t}\left(\hat{U}^{t,\psi}\right)dF^t\left(\epsilon^{t,\psi}\right) = P_{\psi(k)}^{C^t}\left(U^{t,\psi}\right)$. ∎

A corollary to the lemma is that when two actions have equal expected payoffs, the team must play these actions with equal probability. It is easy to see that nonneutral collective choice rules can lead to violations of rank dependence. For example, collective choice rules that favor one action (e.g., a status quo action) over another will generally lead to violations, as in the last example of Section IIA with a $2/3$ voting rule. Consider $K^t = 2$ and a choice rule that selects action a_1^t if and only if all team members estimate its expected utility to be greater than that of action a_2^t, and selects action a_2^t otherwise. For any admissible F^t, if the size of the team is large enough, a team using this choice rule will select action a_2^t more often than action a_1^t even when $U_1^t(\alpha) > U_2^t(\alpha)$.

Next, for the case of $K^t = 2$, we prove 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 proven below.

THEOREM 4: *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 $U'^t = \left(U_1^t - \delta, U_2^t\right)$, then by Lemma 1, $P_1^{C^t}\left(U'^t\right) = P_2^{C^t}\left(U'^t\right) = 1/2$. Since C^t satisfies positive responsiveness and unanimity, Theorem 3, together with the fact that $P_1^{C^t}\left(U^t\right) + P_2^{C^t}\left(U^t\right) = 1$, implies that $P_1^{C^t}\left(U^t\right) > P_1^{C^t}\left(U'^t\right) = P_2^{C^t}\left(U'^t\right) > P_2^{C^t}\left(U^t\right)$. ∎

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

THEOREM 5: *If F^t is admissible and C^t is plurality rule, then P^{C^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 $U_1^t \geq U_2^t \geq \ldots \geq U_{K^t}^t$. By Lemma 1, if $U_k^t = U_l^t$, then $P_k^{C^t} = P_l^{C^t}$.

Suppose $U_k^t > U_l^t$. The probability that any team member i ranks action k highest is $p_k = \Pr\left(U_k^t + \epsilon_{ik}^t - \max_{j \neq k}\left\{U_j^t + \epsilon_{ij}^t\right\} \geq 0\right)$. Let $\psi : \left\{1, \ldots, K^t\right\} \rightarrow$

$\{1,\ldots,K^t\}$be the pairwise permutation of k and l, that is the permutation that maps k to l and l to k and all else to itself. By exchangeability of the error terms, $\left(U_1^t+\epsilon_1^t,\ldots,U_{K^t}^t+\epsilon_{K^t}^t\right)$ has the same joint distribution as $\left(U_1^t+\epsilon_{\psi(1)}^t,\ldots,U_{K^t}^t+\epsilon_{\psi(K^t)}^t\right)$, and so $U_k^t+\epsilon_{ik}^t-\max_{j\neq k}\left\{U_j^t+\epsilon_{ij}^t\right\}$ has the same distribution as $U_k^t+\epsilon_{i\psi(k)}^t-\max_{j\neq k}\left\{U_j^t+\epsilon_{i\psi(j)}^t\right\}$. Since $U_k^t > U_l^t$, we have for all ϵ_i^t, $U_{\psi(k)}^t+\epsilon_{i\psi(k)}^t < U_k^t+\epsilon_{i\psi(k)}^t$, and $\max_{j\neq k}\left\{U_{\psi(j)}^t+\epsilon_{i\psi(j)}^t\right\} \geq \max_{j\neq k}\left\{U_j^t+\epsilon_{i\psi(j)}^t\right\}$. By the full support assumption, it follows that $\Pr\left(U_k^t+\epsilon_{i\psi(k)}^t-\max_{j\neq k}\left\{U_j^t+\epsilon_{i\psi(j)}^t\right\} \geq 0\right) > \Pr\left(U_{\psi(k)}^t+\epsilon_{i\psi(k)}^t-\max_{j\neq k}\left\{U_j^t+\epsilon_{i\psi(j)}^t\right\} \geq 0\right) = p_l$, so $p_k > p_l$.

Now, still supposing that $U_k^t > U_l^t$, and therefore that $p_k > p_l$, denote by $(n_1,\ldots,n_K)$ 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 $V^B = \left\{(n_1,\ldots,n_K)\,|\,\forall a_k \in B, \forall a_j, n_k \geq n_j \text{ and } \sum_{j=1}^K n_j = n\right\}$ 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

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

Let $\psi:\{1,\ldots,K^t\} \rightarrow \{1,\ldots,K^t\}$ be the pairwise permutation between k and l as defined earlier. Pick any B that contains a_k and not a_l. Then $(n_1,\ldots,n_K) \in V^B$ if and only if $\left(n_{\psi(1)},\ldots,n_{\psi(K)}\right) \in V^{(B-\{a_k\})\cup\{a_l\}}$, 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 $\Pr\left(C^t\left(\hat{U}^t\right) = B\right)$ is greater than the corresponding term in $\Pr\left(C^t\left(\hat{U}^t\right) = \left(B-\{a_k\}\right)\cup\{a_l\}\right)$. So we have for all B containing a_k and not a_l, that $\Pr\left(C^t\left(\hat{U}^t\right) = B\right) > \Pr\left(C^t\left(\hat{U}^t\right) = \left(B-\{a_k\}\right)\cup\{a_l\}\right)$. 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,

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

$P_k^{C^t}$ and $P_l^{C^t}$ share all terms of the fourth sum, so

$$P_k^{C^t}\left[U^t\right]-P_l^{C^t}\left[U^t\right] = \sum_{B\in B_k}\frac{1}{|B|}\Pr\left[C^t\left(\hat{U}^t\right) = B\right]-\sum_{B\in B_l}\frac{1}{|B|}\Pr\left[C^t\left(\hat{U}^t\right) = B\right],$$

$$P_k^{C^t}\left[U^t\right] - P_l^{C^t}\left[U^t\right] = \sum_{B \in B_k} \frac{1}{|B|}\Big[\Pr\left[C^t\left(\hat{U}^t\right) = B\right]$$

$$- \Pr\left[C^t\left(\hat{U}^t\right) = \left(B - \{a_k\}\right) \cup \{a_l\}\right]\Big] > 0.$$

Therefore, whenever $U_k^t > U_l^t$, we have $p_k > p_l$, which implies $P_k^{C^t}\left(U^t\right) > P_l^{C^t}\left(U^t\right)$. ∎

Define a weighted average rule as follows.

DEFINITION 11: *A team collective choice rule C^t is a* Weighted Average Rule *if there exists a profile of nonnegative individual voting weights, $\left(w_1^t, \ldots, w_{n^t}^t\right)$ with $\sum_{i=1}^{n^t} w_i^t = 1$ such that for all $a_k^t \in A^t$ and for all $\hat{U}^t \in \mathcal{R}^{K^t n^t}$, $a_k^t \in C^t\left(\hat{U}^t\right)$ 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 6: *If F^t is admissible and C^t is a weighted average rule, then P^{C^t} satisfies rank dependence.*

PROOF:

Consider any profile of expected payoffs U^t, and suppose $U_k^t > U_l^t$. We have $P_k^{C^t}\left(U^t\right) = \int \mathbb{1}\left\{\sum_{i=1}^{n^t} w_i^t \hat{U}_{ik}^t \geq \max\left\{\sum_{i=1}^{n^t} w_i \hat{U}_{ij}^t\right\}_{j=1}^{K^t}\right\} dF^t$. Note that the probability that any of these weighted averages are exactly equal is 0. Now, since $F^t\left(y - U_k^t\right) < F^t\left(y - U_l^t\right)$ for all $y \in \mathcal{R}$, 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 $\sum_{i=1}^{n^t} w_i^t \hat{U}_{ik}^t >_{st} \sum_{i=1}^{n^t} w_i^t \hat{U}_{il}^t$. Since $\mathbb{1}\{z > 0\}$ is increasing, nonconstant and bounded, we therefore have

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

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

$$P_k^{C^t}\left(U^t\right) > P_l^{C^t}\left(U^t\right). \blacksquare$$

V. Team Equilibrium in Extensive Form Games

A finite extensive form game consists of a Player set $\mathbf{I} = \{1, \ldots, I\}$, an action set **A,** a set of sequences contained in A called histories, Ξ, a subset of these being terminal histories, Z, initial chance moves, b^0, a player function, ι, information sets for each player, Π^i, a feasible action function, A, that specifies the set of actions available at each information set, and payoff functions, $u = \left(u^1, \ldots, u^i, \ldots, u^I\right)$, defined on Z. Thus, an extensive form game, in shorthand, can be written as $G^{EF} = \left(\mathbf{I}, \mathbf{A}, \Xi, A, \iota, \Pi, b^0, u\right)$.

We extend this definition to (finite) *team extensive form games*, modifying the notation in the following ways. As in the definition of team games in strategic form, let $\mathbf{T} = \{0, 1, \ldots, t, \ldots, T\}$ be a finite collection of *teams*, $t = \{i_1^t, \ldots, i_j^t, \ldots, i_{n^t}^t\}$, where each team consists of n^t members, or individuals, and denote the *team size profile* by $n = (n^1, \ldots, n^T)$. 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 $(a_1, \ldots, a_K) \in \Xi \Rightarrow (a_1, \ldots, a_L) \in \Xi$ for all $L < K$. A history $h = (a_1, \ldots, a_K) \in Z \subseteq \Xi$ 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 $A(h) = \{a \mid (h, a) \in \Xi\}$. There is a *team function* $\iota : \Xi - Z \to \mathbf{T}$ that assigns each history to a unique team. Without loss of generality, assume that $\iota(\varnothing) = 0$ and $\iota(h) \neq 0$ 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 $t \in \{0, 1, \ldots, t, \ldots, T\}$, there is an information partition Π^t of $\{h \in \Xi \mid \iota(h) = t\}$. 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 $A(h) = A(h') \equiv A(H_l^t) = \{a_{l1}^t, \ldots, a_{lk}^t, \ldots, a_{lK_l^t}^t\}$ where $K_l^t = |A(H_l^t)|$.

The *payoff function* of the game for team $t \neq 0$ is given by $u^t : Z \to \mathcal{R}$. 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 $b^t = (b_1^t, \ldots, b_l^t, \ldots, b_{L^t}^t)$, 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* $\rho(z \mid b)$ 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

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

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* $\rho(z \mid H_l^t, b, a_{lk})$ for each $z \in Z$.[19] 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

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

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 1, and otherwise all teams (including t) playing b elsewhere.

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

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

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, $\hat{U}_{ilk}^t = U_{lk}^t(b) + \varepsilon_{ilk}^t$, where the dependence of $\hat{U}_{ilk}^t$ on b is understood. We call $\hat{U}_{il}^t = (\hat{U}_{il1}^t, \ldots, \hat{U}_{ilK_l^t}^t)$ member i's estimated conditional payoffs at H_l^t, given b. Denote by $\hat{U}_l^t = (\hat{U}_{1l}^t, \ldots, \hat{U}_{n^tl}^t)$ 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.[20]

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, $C_l^t : \mathcal{R}^{n^t K_l^t} \to 2^{A(H_l^t)}$, so $C_l^t(\hat{U}^t) \subseteq A(H_l^t)$. In principle, teams could be using different collective choice rules at different information sets, and denote $C = (C^1, \ldots, C^T)$, where $C^t = (C_1^t, \ldots, C_{L^t}^t)$. We assume that team t always mixes uniformly over $C_l^t(\hat{U}_l{}^t)$. That is, the probability team t chooses a_{lk}^t at $\hat{U}_l{}^t(b)$ is given by the function g^t defined as

$$g_{lk}^{C_l^t}(\hat{U}_l^t) = \begin{cases} \dfrac{1}{|C_l^t(\hat{U}_l^t)|}, & \text{if } a_{lk}^t \in C_l^t(\hat{U}_l{}^t); \\ 0, & \text{otherwise.} \end{cases}$$

Given a behavioral strategy profile, b, for each realization of $(\varepsilon_{l1}^t, \ldots, \varepsilon_{ln^t}^t)$ 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 $g_{lk}^{C_l^t}(\hat{U}_l^t)$.

We require C_l^t to satisfy *Unanimity* for all t and l, defined analogously to Definition 3. That is, for every team $t \neq 0$ and every information set $H_l^t \in \Pi^t$, if $\hat{U}_{ilk}^t > \hat{U}_{ilk'}^t$ for all $a_{lk'} \in A(H_l^t) - \{a_{lk}\}$ and for all $i \in t$, then $C_l^t(\hat{U}_l{}^t) = \{a_{lk}\}$. 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 $\hat{U}_{ilk}^t > \hat{U}_{ilk'}^t$ 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 $a \in A(H_l^t)$, $P_{lk}^{C_l^t}(U_l^t(b)) > 0$, and hence every possible history in the game occurs with positive probability. Thus, there are no "off-path" histories, so $\rho(z \mid H_l^t, b, a_{lk})$ is always well defined and computed according to Bayes' rule.

[20] This specification of errors can be interpreted in terms of an agent model the extensive form game.

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 $P^{C_l^t}$, where $P_{lk}^{C_l^t}$ 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,

$$P_{lk}^{C_l^t}\big(U_l^t(b)\big) = \int_{\varepsilon_l^t} g_{lk}^{C_l^t}\big(\hat{U}_{lk}^t(b)\big) dF_l^t\big(\varepsilon_l^t\big), \tag{7}$$

where F_l^t denotes the distribution of $\varepsilon_l^t = \big(\varepsilon_{l1}^t, \ldots, \varepsilon_{ln^t}^t\big)$.

For any extensive form game, any admissible F, and any profile of collective choice rules C call $\Gamma^{EF} = \big[\mathbf{T}, \mathbf{A}, n, \Xi, A, \iota, \Pi, b^0, u, F, C\big]$ a *team game in extensive form*. An equilibrium of a team game in extensive form is a fixed point of P.

DEFINITION 12: *A* team equilibrium *of the team extensive form game* Γ^{EF} *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* $a_{lk}^t \in A\big(H_l^t\big)$, $b_{lk}^t = P_{lk}^t\big(U^t(b)\big)$.

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 $a \in A\big(H_l^t\big)$, define $\underline{b}_{lk}^t = \inf_{b \in B^o} P_{lk}^t\big(U_l^t(b)\big)$. By Unanimity of C_l^t, $P_{lk}^t\big(U_l^t(b)\big) > 0$ 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, from Unanimity, $P_{lk}^t\big(U_l^t(b)\big)$ is bounded strictly away from 0 for all t, l and k and for all $b \in B^o$. Define $\bar{B}^o = \big\{b \in B \,|\, b_{lk}^t \geq \underline{b}_{lk}^t \text{ for all } t, l \text{ and } k\big\} \subset B^o$. Since $\bar{B}^o$ is compact and convex and P_{lk}^t is a continuous function for all t, l and k, by Brouwer's fixed point theorem there exists a $b \in \bar{B}^o$ such that $b_{lk}^t = P_{lk}^t\big(U^t(b)\big)$ for all t, l and k. ∎

A few observations about team equilibria in extensive form games are worth noting. First, the results in Section III 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* $H_l^t \in \Pi^t$, *if* F_l^t *is admissible and* C_l^t *satisfies unanimity and positive responsiveness then* $P_l^{C_l^t}$ *satisfies payoff monotonicity. Furthermore, if* C_l^t *also satisfies neutrality and* $\big|A\big(H_l^t\big)\big| = 2$, *or if* $\big|A\big(H_l^t\big)\big| > 2$ *and* C_l^t *is plurality rule or a weighted average rule, then* $P_l^{C_l^t}$ *satisfies rank dependence.*

PROOF:

The proof is essentially the same as the proof of Theorems 3, 4, and 5. ∎

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.[21] This is formally stated as follows.

THEOREM 9: *Consider an infinite sequence of team extensive form games,* $\left\{\Gamma_m^{EF}\right\}_{m=1}^{\infty}$, *where* $\Gamma_m^{EF} = \left[\mathbf{T}, \mathbf{A}, n_m, \Xi, A, \iota, \Pi, b^0, u, F, C\right]$, *where* m *indexes an increasing sequence of team sizes, with all the other characteristics of the game being the same. That is,* $n_{m+1}^t > n_m^t$ *for all* m, t. *Suppose* C_l^t *is an anonymous scoring rule for all* t, l *and let* $\left\{b_m^*\right\}_{m=1}^{\infty}$ *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* $\left(\mathbf{I}, \mathbf{A}, \Xi, A, \iota, \Pi, b^0, u\right)$.

PROOF:

The proof is essentially the same as the proof of Theorem 2 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 $\left\{b_m^*\right\}_{m=1}^{\infty}$ 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 $b_m^* \gg 0$ for all m. That is, for every m, t, H_l^t, and $a_{lk}^t \in A\left(H_l^t\right)$, $b_{mlk}^{t*} > 0$. 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 $\mu_m^* \gg 0$, 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 2, 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, or may have imperfect information about the path of play. Classic applications include signaling games.

Fourth, there is a rough connection between team equilibria and the extensive form version of quantal response equilibrium (McKelvey and Palfrey 1998), with the main differences being that team size is fixed at $n = 1$ in quantal response equilibrium and

[21] Not all sequential equilibria are approachable as limit points of team equilibria in extensive form. The example of non-approachability provided earlier applies here as well.

that the disturbances in team equilibrium are 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 its estimates at the beginning of the game, this would lead to a different formulation of the model.

A last observation is that an alternative way to model team equilibrium in extensive form games would be to represent an extensive form game by its normal form or reduced normal form, and then apply the theory developed in Section I of this paper for games in strategic form. However, team equilibrium 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 are suggestive of possible testable implications of the team equilibrium framework.

VI. Examples of Extensive Form Games

A. *Sequential Weak Prisoner's Dilemma Games*

Next, we examine a sequential version of the Weak Prisoner's Dilemma, the simultaneous version of which we analyzed in Section II (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 4 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 II. Panel A shows the team equilibria if team 2 (column) moves first, and panel B for the case where team 1 (row) moves first, for (odd) n ranging from 1 to 199, and $H(x) = 1/(1 + e^{-0.3x})$. The solid light gray line is 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 is 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 defect probability after either cooperation or defection by team 2.[22]

When team 2 moves first, team 1's individual voting probabilities for defect are independent of n, because $x = y = 1$, so the payoff difference between defect and

[22] The black and dark gray curves coincide in panel A of Figure 4, so both appear as black.

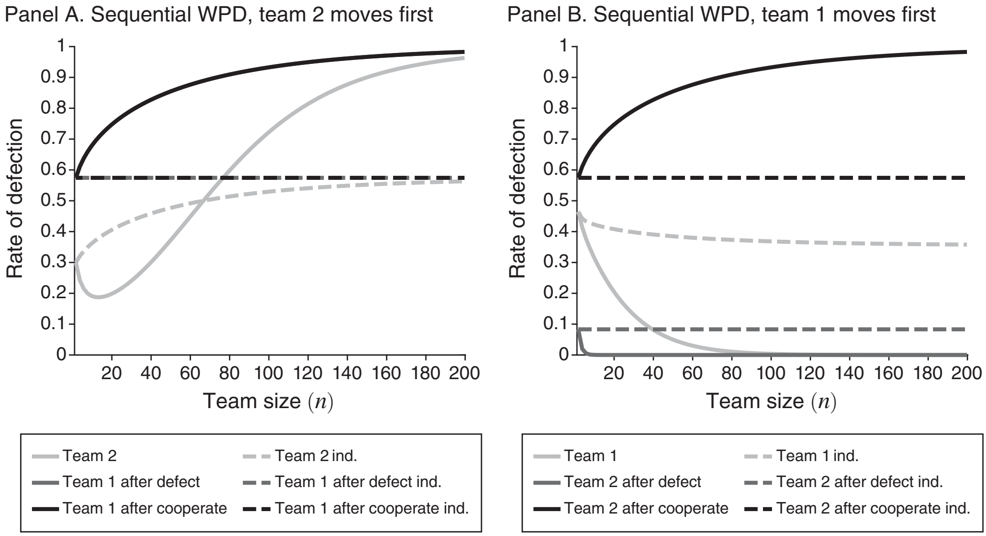

Figure 4. Team Equilibrium in the Sequential WPD

cooperate for team 1 is the same. In fact, they are the same as in the simultaneous version studied in Section II, 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 $H(-z) < 1/2$, and when team 1 chooses defect, the difference is y, so the individual voting probabilities in this case are $H(y) > 1/2$.

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.

B. *Centipede Games*

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. The game ends if team 1 chooses take, 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.

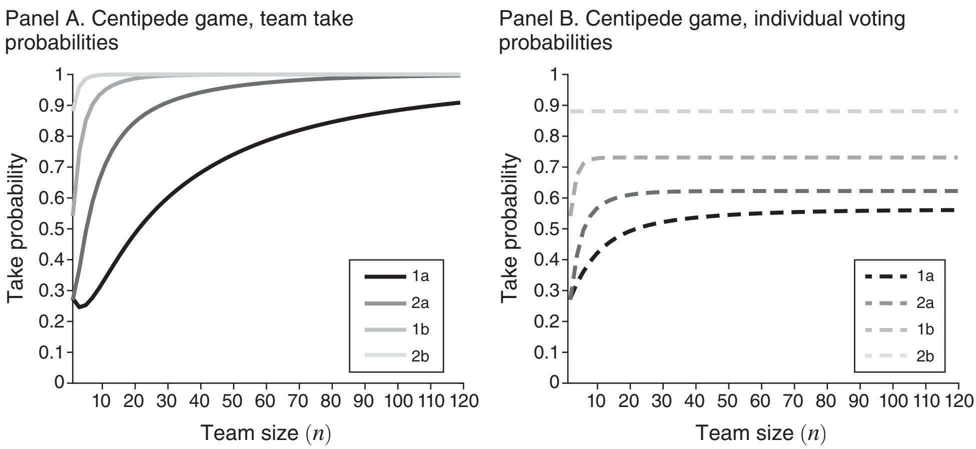

FIGURE 5. TEAM EQUILIBRIUM IN THE CENTIPEDE GAME

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 5, 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 $H(x) = 1/(1 + e^{-x/8})$. In panel A are the team probabilities and in panel B are the individual voting probabilities, with probability of taking and voting for take on the y-axis and team size on the x-axis.

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 and hence the team equilibrium probability of taking increases at every node as team size grows, which is consistent with experimental findings (Bornstein, Kugler, and Ziegelmeyer 2004). 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.

VII. Discussion and Conclusions

This paper proposed and developed a theory of games played by teams of players. The framework combines the noncooperative 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. A team collective choice rule maps the profile of members' beliefs into a team 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.

The approach is initially formulated for finite games in strategic form. Four main results are proved for strategic form team games. First, team equilibria generally exist. Second, we show that all anonymous scoring rules satisfy the Nash convergence property: as team sizes become large, all limit points of team equilibria are Nash equilibria. Counterexamples are constructed to illustrate non-Nash limit points if teams do not use anonymous and/or neutral scoring rules. Third, we identify two weak conditions, unanimity and positive responsiveness, that are sufficient for team response functions to satisfy payoff monotonicity, in the sense that the probability a team chooses a particular action is increasing in the true expected payoff of that action. Fourth, 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.

Team equilibria for games in strategic form are illustrated for several 2×2 games, where the collective choice rule is majority rule. These examples illustrate two distinct effects of changing team size on outcomes. The first is the consensus effect. If the probability any individual on a team chooses 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 consensus effect. Some of the examples suggest possible team games that might be interesting to study in the laboratory, where increasing team size can push equilibrium outcomes further away from Nash equilibrium.

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.

We are hopeful that this framework is a useful starting point for the further exploration and 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 many other rules with this property. One tempting conjecture is that Nash convergence obtains if all teams use a collective choice rule that satisfies unanimity, payoff monotonicity, rank dependence, neutrality, and anonymity. However it turns out that this is not sufficient. Consider the collective choice rule that uniquely selects action a_k^t if and only if a_k^t is unanimously ranked highest by all team members, and otherwise the collective choice rule selects the entire set,

A^t. This satisfies all of the properties above, but in the limit the team's strategy will always converge to uniform mixing over A^t, regardless of the mixed strategies of the other teams.

An interesting related question is how to extend the results about rank dependence with more than two actions to more general collective choice rules. We conjecture that rank dependence holds generally for collective choice rules that satisfy unanimity, positive responsiveness, and neutrality.

The framework can be expanded in several interesting directions. This paper assumed the collective choice rule of each team was exogenous, but as several of the examples suggest, teams may have preferences over collective choice rules in particular games. That is, one could pose the question, given the collective choice rule of the other teams, what would be an optimal collective choice rule for my team in the sense that the resulting team equilibrium with this profile of collective choice rules gives members the highest expected utility? Would optimal collective choice rules satisfy payoff monotonicity and rank dependence? Taking this a step further, one could define an equilibrium in collective choice rules and study its properties in different games.

Another direction to extend the framework would be to allow for a broader class of games than the finite games studied here. 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, such as auctions, 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, judgment biases, social preferences, etc.) would be straightforward, provided the effects are homogeneous across members of the group.

There are alternative approaches to modeling team games and extensions of the present approach that are beyond 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 a noncooperative game rather than an abstract collective choice rule. This would create a nested game-within-a-game structure, which would add another layer of complexity.

REFERENCES

Blinder, Alan S., and John Morgan. 2005. "Are Two Heads Better than One? Monetary Policy by Committee." *Journal of Money, Credit, and Banking* 37 (5): 789–811.

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.

Bornstein, Gary, and Ilan Yaniv. 1998. "Individual and Group Behavior in the Ultimatum Game: Are Groups More 'Rational' Players?" *Experimental Economics* 1: 101–08.

Cason, Timothy N., Sau-Him Paul Lau, and Vai-Lam Mui. 2019. "Prior Interaction, Identity, and Cooperation in the Inter-Group Prisoner's Dilemma." *Journal of Economic Behavior and Organization* 166: 613–29.

Charness, Gary, and Matthew O. Jackson. 2007. "Group Play in Games and the Role of Consent in Network Formation." *Journal of Economic Theory* 136 (1): 417–45.

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): 551–56.

Charness, Gary, and Matthias Sutter. 2012. "Groups Make Better Self-Interested Decisions." *Journal of Economic Perspectives* 26 (3): 157–76.

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.

Cox, James C. 2002. "Trust, Reciprocity, and Other-Regarding Preferences: Groups versus Individuals and Males versus Females." In *Experimental Business Research*, edited by Rami Zwick and Amnon Rapoport, 331–50. Boston: Kluwer Academic Publishers.

Cox, James C., and Stephen C. Hayne. 2006. "Barking up the Right Tree: Are Small Groups Rational Agents?" *Experimental Economics* 9: 209–22.

Duggan, John. 2001. "Non-cooperative Games among Groups." Unpublished.

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* 11 (1).

Harsanyi, John C. 1973. "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed Strategy Equilibrium." *International Journal of Game Theory* 2: 1–23.

Insko, Chester A., Rick H. Hoyle, Robin L. Pinkley, Gui-Young Hong, Randa M. Slim, Bret Dalton, Yuan-Huei Lin et al. 1988. "Individual-Group Discontinuity: The Role of a Consensus Rule." *Journal of Experimental Social Psychology* 24 (6): 505–19.

Kreps, David M., and Robert Wilson. 1982. "Sequential Equilibria." *Econometrica* 50 (4): 863–94.

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.

Kugler, Tamar, Edgar E. Kausel, and Martin G. Kocher. 2012. "Are Groups More Rational than Individuals? A Review of Interactive Decision Making in Groups." *WIREs Cognitive Science* 3 (4): 471–82.

Marschak, Jacob, and Roy Radner. 1972. *Economic Theory of Teams*. New Haven, CT: Yale University Press.

McKelvey, Richard D., and Talbot Page. 1986. "Common Knowledge, Consensus, and Aggregate Information." *Econometrica* 54 (1): 109–27.

McKelvey, Richard D., and Thomas R. Palfrey. 1995. "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10 (1): 6–38.

McKelvey, Richard D., and Thomas R. Palfrey. 1996. "A Statistical Theory of Equilibrium in Games." *Japanese Economic Review* 47 (2): 186–209.

McKelvey, Richard D., and Thomas R. Palfrey. 1998. "Quantal Response Equilibria for Extensive Form Games." *Experimental Economics* 1: 9–41.

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.

Rockenbach, Bettina, Abdolkarim Sadrieh, and Barbara Mathauschek. 2007. "Teams Take Better Risks." *Journal of Economic Behavior and Organization* 63 (3): 412–22.

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): 907–14.

Sheremeta, Roman M., and Jingjing Zhang. 2010. "Can Groups Solve the Problem of Overbidding in Contests?" *Social Choice and Welfare* 35 (2): 175–97.

Sniezek, Janet A., and Rebecca A. Henry. 1989. "Accuracy and Confidence in Group Judgment." *Organizational Behavior and Human Decision Processes* 43 (1): 1–28.

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

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.

Sutter, Matthias. 2005. "Are Four Heads Better Than Two? An Experimental Beauty-Contest Game with Teams of Different Size." *Economics Letters* 88 (1): 41–46.

Sutter, Matthias, Martin G. Kocher, and Sabine Strauß. 2009. "Individuals and Teams in Auctions." *Oxford Economic Papers* 61 (2): 380–94.

Wildschut, Tim, and Chester A. Insko. 2007. "Explanations of Interindividual-Intergroup Discontinuity: A Review of the Evidence." *European Review of Social Psychology* 18 (1): 175–211.