

Existence of Equilibrium Mechanisms in Generalized Principal-Agent Problems with Interacting Teams*

Brian Roberson[†]

February 2026

Abstract

We study incentive design when multiple principals simultaneously design mechanisms for their respective teams in environments with strategic spillovers. In this environment, each principal's set of incentive-compatible mechanisms—those that satisfy their own agents' incentive compatibility constraints—depends on the mechanisms offered by the other teams. Following a classic example by Myerson (1982), such games may lack equilibrium due to discontinuities in the correspondence of incentive-compatible mechanisms. We establish general conditions for equilibrium existence by introducing a novel approach that involves tracking both the outcome distributions along the truthful-obedient path and the sets of outcome distributions achievable through unilateral deviations, thereby providing a foundation for analyzing a wide range of multi-principal mechanism design with team production and agency problems.

*I am grateful to Pierre Boyer, John Duggan, Christian Ewerhart, Navin Kartik, Wojciech Olszewski, Philip J. Reny, and Ron Siegel for helpful comments and discussions on previous drafts. I also thank participants in seminars at the Rochester Department of Political Science, the Purdue Mechanism Design Workshop, the SAET 2025 Annual Meeting, and the University of South Carolina Game Theory Conference for helpful comments.

[†]brobers@purdue.edu • Purdue University, Department of Economics

1 Introduction

In many economic environments, incentive design occurs simultaneously across multiple organizations, with spillovers from each organization’s activities affecting the others. We model such environments as interacting teams, where multiple principals simultaneously design mechanisms for their respective agents while accounting for strategic interactions across teams. Examples range from competitive settings—such as competition between firms in a market¹ or between teams in innovation contests²—to cooperative settings such as production along a supply chain. In these settings, team production is often more efficient than individual production due to complementarities among team members’ skills and efforts.³ However, the presence of adverse selection (private abilities) and moral hazard (unobservable actions) within teams complicates the assessment of individual contributions and the allocation of rewards.⁴

To illustrate the issues arising from interdependent agency concerns, consider the problem each principal faces in selecting a mechanism that specifies, among other things, how the team’s winnings are allocated. These winnings depend on the team’s own performance and may be affected by how other teams perform. As a result, the mechanism that each principal selects affects not only the choices of their own team’s agents but also the sets of incentive-compatible mechanisms available to other principals. This endogeneity—where the set of incentive-compatible mechanisms available to each principal depends on other principals’ choices—poses a fundamental challenge for modeling interacting teams with agency concerns. Indeed, Myerson (1982) provides an example⁵ with two mechanism designers in which no equilibrium exists.⁶ This problem of multi-principal interaction can be understood as a generalized game—a game in which each player’s feasible strategy set is endogenously determined by the strategies of other players. In this paper, we develop a framework that provides general conditions for equilibrium existence in generalized games of this type.

Generalized games were introduced by Debreu (1952), who establishes conditions for the existence of a (pure-strategy) Nash equilibrium, also known as a social equilibrium.⁷ In a generalized game, each player has a feasible strategy correspondence—a mapping that assigns to each profile of other players’ strategies a set of feasible strategies for that player. When this correspondence is contained within a compact, convex subset of Euclidean space, the existence results in Debreu (1952), Debreu (1982), and Banks and Duggan (2004) rely critically on its lower hemicontinuity. This property ensures that small changes in other players’ strategies do not suddenly render a previously feasible choice infeasible. As Tóbiás (2022) demonstrates, lower hemicontinuity is essential for establishing the existence of Nash equilibria in generalized games and

¹See for example D’Aspremont and Jacquemin (1988), Kamien, Muller, and Zang (1992), Raith (2003), and Antón, Ederer, Giné, and Schmalz (2023) among others.

²See for example Nitzan (1991), Esteban and Ray (2001), Nitzan and Ueda (2011), Nitzan and Ueda (2018), Balart et al. (2016), and Konishi, Sahuguet, and Crutzen (2024) among others.

³For more on this issue, see Candoğan et al. (2025), who examines innovation contests involving a finite number of individuals and/or teams with complete information and exogenous prize-sharing rules, and identifies conditions under which teams outperform individuals.

⁴In the context of a single designer designing incentives for a team facing both adverse selection and moral hazard, see McAfee and McMillan (1991). Also related are the literatures on teamwork (e.g., Admati and Perry, 1991; Yıldırım, 2006; Bonatti and Höner, 2011; Georgiadis, 2015; Bowen, Georgiadis, and Lambert, 2019; Cetemen, Hwang, and Kaya, 2020; Ozertürk and Yıldırım, 2021; and Yıldırım, 2023) and contracting with teams (e.g., Holmstrom, 1982; Itoh, 1991; Che and Yoo, 2001; Winter, 2004; and Halac, Lipnowski, and Rappoport, 2021).

⁵In particular, see Section 4 of Myerson (1982), entitled “Equilibria among several principals.”

⁶In this example, equilibrium fails to exist because the principals’ sets of feasible incentive-compatible mechanisms are not lower hemicontinuous with respect to the other principal’s choice of mechanism. For more on the role of lower hemicontinuity in generalized games, see Tóbiás (2022).

⁷For additional details, see Debreu (1982), Border (1985, Chapter 19), Facchinei and Kanzow (2010), and Dasgupta and Maskin (2015). Related work includes Tao and Xu (2024), who demonstrate the existence of pure-strategy Bayesian Nash equilibrium in generalized games in which the feasible strategy correspondence is uniformly continuous.

cannot be dispensed with.⁸ This continuity requirement plays a central role in our analysis: it is precisely the property that fails to hold in the Myerson (1982) example, leading to non-existence of equilibrium.

Our primary contribution is a novel approach to establishing the existence of a Bayesian-Nash Principals' Equilibrium (BNPE) in a generalized game with team production and agency concerns. Our methodological approach builds on two complementary strands in the literature on the existence of Bayesian-Nash equilibrium in Bayesian games.⁹ The first strand, initiated by Milgrom and Weber (1985), draws on the insight that strategies in Bayesian games can be fully characterized by the joint probability measures they induce over types and actions. Rather than working directly with strategy functions that map types to actions, this distributional approach treats each player's strategy as a probability measure on the joint type-action space. To determine when two distributional strategies are “close” to each other—and thereby characterize the continuity properties of payoff functions and best-response correspondences—Milgrom and Weber (1985) use the classical narrow topology.¹⁰ This topology captures the idea that two probability distributions are close if they assign similar probabilities to similar events, which ensures that a small change in a player's strategy induces only a small change in all players' expected payoffs. This notion of closeness is particularly well-suited for Bayesian games, where payoffs depend on the distribution of types and actions.

Building on that foundation, Kadan, Reny, and Swinkels (2017) adapt this distributional approach to mechanism design settings by associating each incentive-compatible mechanism with an “on-path” joint probability measure over rewards, outputs, actions, and types—the probability measure generated when all agents truthfully report and obediently follow the intended mechanism.¹¹ Like Milgrom and Weber (1985), Kadan, Reny, and Swinkels (2017) work with the classical narrow topology on the space of probability measures. However, focusing on the “on-path” outcome measure creates challenges when characterizing the continuity properties of the set of incentive-compatible mechanisms, since incentive compatibility fundamentally depends on the possibility of “off-path” deviations. In the single-principal case, Kadan, Reny, and Swinkels (2017) address this issue by making creative use of results on Young measure convergence, as summarized in Balder (2021) and rooted in the foundational work of Komlós (1967),¹² to demonstrate that the set of incentive-compatible mechanisms is well-behaved. Unfortunately, this elegant approach does not extend naturally to our setting with interacting teams and agency concerns.

The second strand, exemplified by Balder (1988), extends the existence results in Milgrom and Weber (1985) by shifting focus from distributional strategies to behavior strategies. In this framework, a behavior strategy specifies for each type of each player a probability measure over that player's action space—that is, the stochastically chosen action conditional on type. As in the distributional approach, Balder (1988) makes use of the classical narrow topology to examine the continuity properties of payoff functions and best-response correspondences on the space of behavior strategies. While Balder (1988) studies strategic-form Bayesian games without communication or reporting, that approach can be extended to our mechanism

⁸Tóbiás (2022) provides a detailed account of the relevant literature. See in particular the discussion in footnote 5.

⁹Our work is also related to the literature on the existence of equilibria in discontinuous games initiated by Reny (1999). The related literature is too large to comprehensively summarize here, but some notable contributions include Carmona (2009), Barelli and Meneghel (2013), McLennan, Monteiro, and Tourky (2011), He and Yannelis (2016), Bich and Laraki (2017), Carbonell-Nicolau and McLean (2018), Reny (2020), Olszewski and Siegel (2023), and Prokopovych and Yannelis (2023).

¹⁰Recall that a sequence of probability measures $\{\mu_n\}$ converges to μ in the narrow topology if and only if $\int f d\mu_n \rightarrow \int f d\mu$ for every bounded continuous function f . When the underlying space is Polish (complete separable metric), this coincides with the weak* topology on the space of probability measures.

¹¹Recent studies have made significant progress in addressing the challenges of moral hazard and adverse selection in principal-agent problems. Notably, Castro-Pires, Chade, and Swinkels (2024) introduce the concept of decoupling, a simple method for studying both moral hazard and adverse selection simultaneously, and provide tests for determining its validity. Other notable contributions include Chen, Chen, and Rietzke (2020), Ke and Xu (2023), and Gottlieb and Moreira (2022), which collectively advance our understanding of optimal contracts and mechanisms in various settings.

¹²See also Balder (1984, 1985, 1998).

design setting.

The key innovation in our approach is the combination of the behavior strategy framework of Balder (1988) with a careful extension of the Kadan, Reny, and Swinkels (2017) approach to what it means for two mechanisms to be “close” to each other. Our measure of closeness, or metric, imposes two simultaneous requirements for mechanisms to be close. First, as in Kadan, Reny, and Swinkels (2017), we require that the outcome measures along the truthful-obedient path must be similar.¹³ Second, and crucially, we require that the sets of outcome measures achievable through any unilateral (behavior strategy) deviation—whether mis-reporting types and/or disobeying recommendations—must also be similar. This ensures that as a sequence of mechanisms converges, not only do the outcome measures along the truthful-obedient path converge, but so do agents’ sets of achievable outcome measures under unilateral deviations.

For the truthful-obedient path, we use the narrow topology to measure convergence of the induced probability measures, just as in the prior literature. For the complete set of outcome measures that can arise from any unilateral (behavior strategy) deviation by an agent we use the Hausdorff metric, which measures the distance between two sets by asking how far each point in one set is from the nearest point in the other set. Two mechanisms are considered close under our robust narrow topology only if: (i) their on-path outcome measures are close in the narrow sense, and (ii) their sets of achievable outcome measures from all possible unilateral (behavior strategy) deviations are close in the Hausdorff sense. Incentive compatibility fundamentally depends on comparing what an agent gets from being truthful and obedient versus what they could achieve by deviating. By tracking both on-path outcome measures and unilateral deviation outcome measures simultaneously, our topology ensures that convergence of mechanisms implies convergence of strategic opportunities.

Our framework accommodates rich and flexible environments with multidimensional types, actions, outputs, and rewards, including various specifications of how team winnings map to feasible individual rewards. The analysis relies on four categories of assumptions. First, we require the underlying spaces of types, actions, winnings, and rewards to be compact Polish spaces—that is, complete, separable, and metrizable. This ensures that these spaces are bounded (so sequences cannot escape to infinity) and have no gaps (so every convergent sequence has a limit point within the space). Second, we require the teams’ stochastic output technologies to provide a sufficiently smooth and predictable landscape for principals to design incentives. In particular, we assume that small changes in a team’s action profile lead to small, predictable changes in the probability distribution over the team’s outputs. Third, we require the correspondence from team winnings to feasible individual rewards to be well-behaved: it must be continuous and always map to a non-empty compact set of feasible rewards. This assumption is satisfied in standard economic environments where budget constraints or technological constraints smoothly determine how team winnings can be divided among team members.¹⁴ Fourth, we assume that all players’ utility functions are bounded and continuous.

Our theory of incentive design for interacting teams contributes to two strands of the literature on contests. We generalize the literature on contests with endogenous prize sharing, which has typically focused on restricted sets of prize-sharing rules, to the case of the generalized principal-agent problem. Additionally, we extend the literature on contests featuring stochastic production as a function of effort to the case of a

¹³The standard approach in the prior literature measures closeness between mechanisms using the classical narrow topology. Under this topology, a sequence of probability measures converges if and only if expectations of all bounded continuous functions converge. Applied to mechanisms, this means that two mechanisms are considered close if they produce similar expected outcomes when all agents report their types truthfully and follow recommendations obediently.

¹⁴For example, this includes perfectly divisible monetary prizes (where the sum of individual rewards equals team winnings), pure public goods (where each team member receives the same reward equal to team winnings), and any combination of private and public rewards.

general stochastic mapping from the profiles of types and actions to the profile of team winnings. Beginning with the branch of the team contest literature with stochastic production – in the form of the Tullock contest success function – and endogenous prize sharing,¹⁵ one common setting is the case of complete information and a single prize that has both public good and private good aspects. In this case the endogenous sharing-rule allocates the private good component among the team members.¹⁶ Examples of this approach include: (i) Nitzan (1991), Nitzan and Ueda (2011), and Balart, Flamand, and Troumpounis (2016) in which the set of feasible sharing rules is given by the set of convex combinations of a relative effort component and fair division, (ii) Trevisan (2020), Simeonov (2020),¹⁷ Kobayashi and Konishi (2021) which allow for general allocations of the private good component of the prize, and (iii) Kobayashi, Konishi, and Ueda (2025) which allows for general allocations of a private good prize and the value of the prize is endogenously determined. More closely related to our focus is the extension to the case of multiple indivisible prizes, as in Crutzen, Flamand, and Sahuguet (2020) and Konishi, Sahuguet, and Crutzen (2024).

A key feature of our approach is the incorporation of a general stochastic mapping from the profile of types and actions to the profile of team winnings. In the contest theory literature, it is common to allow for stochastic production as a function of effort, with the Tullock (1980) ratio-form contest success function and the Lazear and Rosen (1981) rank-order tournament being notable examples. Our approach builds upon this literature on stochastic production for individual contestants (e.g., Kirkegaard (2023); Bastani et al. (2022); Drugov and Ryvkin (2020); Ryvkin and Drugov (2020) - which extend early contributions by Fullerton and McAfee (1999)).¹⁸ Most closely related is Kirkegaard (2023), who examines the optimal contest design problem in this environment and frames the contest designer's problem as a team moral hazard problem with a finite number of agents, where the designer observes individual contestants' stochastic outputs but not their effort choices. By extending the stochastic production approach to accommodate a general form of stochastic team production that depends on the profile of types and actions, our results provide a foundation for examining a wide range of issues arising in competitive environments involving stochastic production with agency concerns.

The remainder of this paper is organized as follows. Section 2 discusses the example in Myerson (1982) where equilibrium fails to exist. Section 3 presents our general theoretical framework. Section 4 provides our assumptions, introduces the metric structure, and presents an example based on Nitzan (1991). Section 5 establishes our main results on the existence of a Bayesian-Nash Principals' Equilibrium (BNPE). Section 6 discusses implications and directions for future research.

¹⁵Also related is the strand of literature on team contests with deterministic team production, a public good prize, and private information. This includes: Barbieri and Malueg (2014), Barbieri and Malueg (2016), Eliaz and Wu (2018), Barbieri et al. (2019), and Barbieri and Topolyan (2021). See also Brookins and Ryvkin (2016), who – for the case where each team's winnings are in the form of a public good for the group – establishes the existence of equilibrium for a range of information and production configurations.

¹⁶See also Nitzan and Ueda (2014) which examines the related issue of cost-sharing rules.

¹⁷Note that the stochastic production in Simeonov (2020) is more general but allows for the Tullock CSF as a special case.

¹⁸Note that the focus in these papers differs from our focus here in that the individual contestants are assumed to have perfect information regarding the other contestants types (however, Ryvkin and Drugov (2020) features uncertainty regarding the number of contestants) and as is common in the contest theory literature, the contestants are assumed to have additively separable utility. For additional background on stochastic production in contests, see Konrad (2009) and Vojnović (2015). Also related is the innovation competition literature (e.g., Halac et al. (2017); Moscarini and Squintani (2010); Terwiesch and Xu (2008); and Taylor (1995) among others), which examines several formulations of individual stochastic production in contests.

2 Myerson (1982) Example

Consider a strategic environment with two competing teams, labeled $j = 1, 2$. Each team j consists of a principal p_j and a team member m_j . The team member has private information (their “type”), and the principal designs a mechanism to coordinate the team’s actions. In this example, we focus on adverse selection where team members only report their types.

The game unfolds across five stages, which we describe below and which are illustrated in Figure 1.

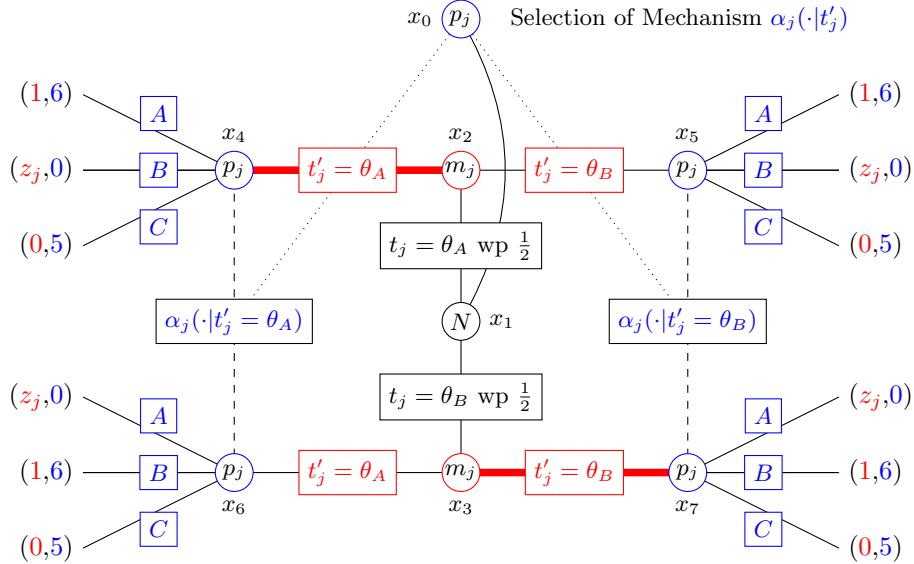
Stage 0 (Mechanism Design): Each principal p_j chooses a coordination mechanism for their team. Specifically, principal p_j selects a behavior strategy $\alpha_j(\cdot|t'_j)$ that specifies, for each possible type report $t'_j \in \{\theta_A, \theta_B\}$ from member m_j , a probability distribution over recommended actions from the set $\{A, B, C\}$.

Stage 1 (Type Realization): Nature independently draws a type $t_j \in \{\theta_A, \theta_B\}$ for each team member m_j , with $P(t_j = \theta_A) = P(t_j = \theta_B) = \frac{1}{2}$. This corresponds to node x_1 in Figure 1.

Stage 2 (Type Report): Each member m_j privately observes their type t_j (reaching node x_2 if $t_j = \theta_A$ or node x_3 if $t_j = \theta_B$) and then sends a cheap-talk report $t'_j \in \{\theta_A, \theta_B\}$ to their principal p_j . This report is private within team j : only principal p_j observes it.

Stage 3 (Action Recommendation): After observing the report t'_j , principal p_j uses the mechanism $\alpha_j(\cdot|t'_j)$ to recommend an action $a_j \in \{A, B, C\}$ (possibly stochastically). In Figure 1, this corresponds to the information sets at nodes (x_4, x_6) for report $t'_j = \theta_A$ and nodes (x_5, x_7) for report $t'_j = \theta_B$.

Stage 4 (Payoffs): The game ends with payoffs realized at the terminal nodes. Each team’s payoffs depend on three factors: their member’s true type t_j , their own recommended action a_j , and the other team’s recommended action a_{-j} . In Figure 1, principal payoffs are shown in blue and member payoffs in red.



For $j = 1$,

$$z_1 = \begin{cases} 2 & \text{if } p_2 \text{ chooses } A \text{ or } B, \\ 1 & \text{if } p_2 \text{ chooses } C. \end{cases} \quad z_2 = \begin{cases} 2 & \text{if } p_1 \text{ chooses } C, \\ 1 & \text{if } p_1 \text{ chooses } A \text{ or } B. \end{cases}$$

Figure 1: Illustration of Team j ’s Mechanism

Given a profile of mechanisms $(\alpha_1(\cdot | t'_1), \alpha_2(\cdot | t'_2))$, the team j mechanism $\alpha_j(\cdot | t'_j)$ is said to be *incentive-compatible* if it is a Bayesian equilibrium for team member m_j to truthfully report their type. That is, there exists a Bayesian equilibrium in which the reporting strategy satisfies $t'_j(t_j) = t_j$ for each $t_j \in \{\theta_A, \theta_B\}$. In Figure 1, the truthful reporting paths (x_2, x_4) and (x_3, x_7) are emphasized in red. Team j 's *incentive-compatible mechanism correspondence*, IC_j , is a set-valued map that associates each mechanism of the other team, $\alpha_{-j}(\cdot | t'_{-j})$, with the set of incentive-compatible mechanisms for team j , denoted $IC_j(\alpha_{-j}(\cdot | t'_{-j}))$.

We are now in position to define a principals' equilibrium. Consider the generalized game between principals in which each principal j 's feasible-strategy correspondence is its incentive-compatible mechanism correspondence IC_j . A profile of mechanisms $(\alpha_1(\cdot | t'_1), \alpha_2(\cdot | t'_2))$ is said to be a *principals' equilibrium* if each principal j 's mechanism $\alpha_j(\cdot | t'_j)$ is a best response to $\alpha_{-j}(\cdot | t'_{-j})$ among team j 's set of incentive-compatible mechanisms $IC_j(\alpha_{-j}(\cdot | t'_{-j}))$.

To address the issue of nonexistence of equilibrium, first note that the principals' feasible-strategy correspondences do not have closed graphs or, equivalently, are not lower hemicontinuous. We focus here on team 1's feasible-strategy correspondence, though a corresponding issue arises for team 2.

Specifically, suppose team 2's mechanism chooses C with probability 1 for both possible reports, i.e., $\alpha_2(C | t'_2) = 1$ for each $t'_2 \in \{\theta_A, \theta_B\}$. Then team 1's set of incentive-compatible mechanisms, $IC_1(\alpha_2(\cdot | t'_2))$, includes all mechanisms in which C is played with the same probability for both possible reports. However, if team 2's mechanism chooses A or B with strictly positive probability for either possible report, then team 1's set of incentive-compatible mechanisms discontinuously shrinks. Consequently, team 1's feasible-strategy correspondence does not have a closed graph, as it lacks the continuity required for feasible strategies to vary smoothly with changes in the other team's mechanism.

This discontinuity in the feasible-strategy correspondences propagates to the best-response correspondences, which also fail to have closed graphs. This failure results in the absence of a fixed point and the nonexistence of a principals' equilibrium. To see this, consider the following cycle:

- If principal 2 chooses C with probability 1 for both reports, then principal 1's best response is to choose A following report θ_A and B following report θ_B .
- But if principal 1 chooses A following θ_A and B following θ_B , then principal 2's best response is to choose A following θ_A and B following θ_B .
- If principal 2 chooses A or B with strictly positive probability for either report, then principal 1's best response is to choose C with probability 1 for both reports.
- But if principal 1 chooses C with probability 1 for both reports, then principal 2's best response is also to choose C with probability 1 for both reports.

Having illustrated how equilibrium may fail to exist, we now examine a general environment and provide conditions under which equilibrium can be shown to exist.

3 Model

We analyze a model in which teams of agents with agency concerns interact and the teams' principals (or mechanism designers) initially specify mechanisms that address their respective generalized principal-agent problems subject to feasibility constraints arising from incentive compatibility considerations. Relative to the formulation in Myerson (1982), we introduce some additional structure – which arises naturally in many

economic environments – around the mapping from the aggregate type and action profiles to the team payoff profile.

Beginning with a brief overview, consider a game that consists of N teams, each comprised of n team members,¹⁹ where an arbitrary team is denoted by $j \in \{1, 2, \dots, N\}$ and an arbitrary team member is denoted by $i \in \{1, 2, \dots, n\}$. Each team faces both adverse selection (private ability) and moral hazard (unobservable actions) within the team. The game begins with each team’s principal specifying a team mechanism. The ensuing continuation game unfolds as follows. First, team members privately learn their types. Second, the team members report their types to the team through cheap, unverifiable talk. Third, the team mechanism recommends a profile of actions to the team, and team members individually choose unobservable actions. Fourth, team winnings are stochastically determined by the actual profiles of types and actions. Finally, each team’s winnings are distributed among its team members as individual rewards according to the team’s mechanism.

Having provided a brief overview of the game, we now delve deeper into each of the stages in the game, and the following subsections provide details on the game’s progression across these four stages. We then revisit the principals’ initial problems. Finally, we conclude this section with a summary of the multi-principal extensive-form game with team production and agency concerns.

3.1 First (Private-Type) Stage

In the first stage, each team member privately observes their stochastic type, where the set of possible types is denoted by T . Let $\mathcal{P}(T)$ denote the set of all probability measures on the Borel sets of T , denoted $\mathcal{B}(T)$. Throughout the paper, we work with Borel sets and maps, omitting the terms ‘Borel’ and ‘measurable’ unless clarity requires otherwise. The joint type space T^{nN} is endowed with a probability measure $H \in \mathcal{P}(T^{nN})$. The entire joint type profile $\hat{\mathbf{t}} \in T^{nN}$ is drawn according to H , and each agent (i, j) —that is, member i of team j —is privately informed of their type $t_{i,j}$.

In the following discussion, it will also be convenient to let $\hat{\mathbf{t}}_{-i,j}$ denote the entire $(nN - 1)$ -tuple of types of all agents other than agent (i, j) —when referring to the $(n - 1)$ -tuple of types of the members of team j other than team member (i, j) we will use the notation $\mathbf{t}_{-i,j}$.

3.2 Second (Type-Reporting) Stage

Information revelation takes the form of cheap talk, and each team member makes an unverifiable report of their type. For each agent (i, j) , a second-stage type-reporting strategy is a function $t'_{i,j} : T \rightarrow T$, and the set of stage 2 type-reporting strategies is denoted by:

$$\mathcal{T} : \left\{ t'_{i,j} : T \rightarrow T \mid t'_{i,j} \text{ is Borel measurable} \right\}.$$

3.3 Third (Action) Stage

Let A denote the space of possible actions for individual team members. We now introduce the first component of a team mechanism, the recommended actions in the third (action) stage. In subsection 3.4 on the fourth (team-winnings) stage, we introduce the second component of a team mechanism, the distribution of team winnings. For any profile of reported types $\mathbf{t}'_j \in T^n$ by the n members of team j , the mechanism

¹⁹The assumption of equal team sizes is for notational convenience; our results extend directly to teams with different numbers of members.

recommends an action profile $\mathbf{a}'_j \in A^n$ drawn at the beginning of the third (action) stage from a transition probability $\alpha_j : T^n \times \mathcal{B}(A^n) \rightarrow [0, 1]$. For each reported type profile \mathbf{t}'_j , this assigns probabilities to sets of action profiles, which we write as $\alpha_j(\cdot | \mathbf{t}'_j)$.

Then, a third-stage action strategy for agent (i, j) is a function $a_{i,j} : T \times T \times A \rightarrow A$, which takes their private type $t_{i,j} \in T$, their reported type $t'_{i,j} \in T$, and their recommended action $a'_{i,j} \in A$, and maps this into a feasible action in A . The set of stage 3 action strategies is denoted by:

$$\mathcal{A} : \left\{ a_{i,j} : T \times T \times A \rightarrow A \mid a_{i,j} \text{ is Borel measurable} \right\}.$$

The actions of the individual team members are unobservable to the team, and for any agent (i, j) , a strategy $(t'^{\star}_{i,j}(\cdot), a^{\star}_{i,j}(\cdot, \cdot, \cdot))$ is said to be *honest and obedient* if in stage 2 $t'^{\star}_{i,j}(t_{i,j}) = t_{i,j}$ for almost every $t_{i,j} \in T$ and in stage 3 $a^{\star}_{i,j}(t_{i,j}, t'_{i,j}, a'_{i,j}) = a'_{i,j}$ for almost every $t_{i,j} \in T$ and $a'_{i,j} \in A$.

3.4 Fourth (Team-Winnings) Stage

Let W denote the set of possible team winnings. The profile of team winnings across all N teams is determined stochastically in the fourth stage by a transition probability $\Lambda : T^{nN} \times A^{nN} \times \mathcal{B}(W^N) \rightarrow [0, 1]$. For each profile of true types $\hat{\mathbf{t}} \in T^{nN}$ and true actions $\hat{\mathbf{a}} \in A^{nN}$ across all nN agents, $\Lambda(\cdot | \hat{\mathbf{t}}, \hat{\mathbf{a}})$ assigns probabilities to sets of team winnings profiles in W^N .

To model how a team's winnings may be allocated among its members we proceed as follows. First, let I denote the space of individual rewards, with $r_{i,j} \in I$ denoting an arbitrary reward for team j member i , and $\mathbf{r}_j \in I^n$ denoting an arbitrary n -tuple of rewards for team j . Next, let $\mathcal{W} : W \rightarrow I^n$ denote the correspondence mapping team winnings to feasible profiles of individual rewards. For example, if the team j winnings w_j take the form of a (perfectly divisible) monetary prize, then feasibility requires that the sum of payments to all team members not exceed total winnings: $\sum_{i=1}^N r_{i,j} \leq w_j$ with $r_{i,j} \geq 0$ for all $i \in \{1, \dots, N\}$. Similarly, if the team j winnings w_j take the form of a public good for the team—such as a shared prize that each member values equally—then feasibility requires $r_{i,j} = w_j$ for all $i \in \{1, \dots, n\}$.

Given reported types $\mathbf{t}'_j \in T^n$, recommended actions $\mathbf{a}'_j \in A^n$, and team winnings $w_j \in W$, the team j mechanism distributes its winnings among its team members according to a transition probability

$$\kappa_j : T^n \times A^n \times W \times \mathcal{B}(I^n) \rightarrow [0, 1],$$

where I^n is the space of individual rewards. For each tuple $(\mathbf{t}'_j, \mathbf{a}'_j, w_j) \in T^n \times A^n \times W$, $\kappa_j(\cdot | \mathbf{t}'_j, \mathbf{a}'_j, w_j)$ is a probability measure that assigns a probability to each set in $\mathcal{B}(I^n)$. Furthermore, feasibility of the mechanism with respect to team winnings requires that the support of $\kappa_j(\cdot | \mathbf{t}'_j, \mathbf{a}'_j, w_j)$ is contained in $\mathcal{W}(w_j) \subseteq I^n$, where $\mathcal{W}(w_j)$ is the set of feasible individual rewards given team j winnings w_j .

3.5 Principals' Initial Problems

We now turn to the interaction between the principals in which each principal selects a feasible and incentive-compatible mechanism. To set the stage for the best-response problem faced by each principal we must first: (i) provide the definition of a team j mechanism, (ii) define the set of incentive compatible mechanisms for team j , and (iii) specify the payoff functions for the principals.

Beginning with the definition of the team j mechanism, let $\mathcal{K}(T^n, A^n)$ denote the set of transition probabilities that map the space of n -tuples of types T^n into the space of probability measures on A^n .

Similarly, let $\mathcal{K}(T^n \times A^n \times W, I^n)$ denote the set of transition probabilities mapping $T^n \times A^n \times W$ into the space of probability measures on I^n . Applying the revelation principle, a team j mechanism is described as follows.

Definition 3.1. *A team j mechanism is a pair of transition probabilities $(\alpha_j, \kappa_j) \in \mathcal{K}(T^n, A^n) \times \mathcal{K}(T^n \times A^n \times W, I^n)$ satisfying the support constraint: $\text{supp } \kappa_j(\cdot | \mathbf{t}'_j, \mathbf{a}'_j, w_j) \subseteq \mathcal{W}(w_j)$ for every $(\mathbf{t}'_j, \mathbf{a}'_j, w_j) \in T^n \times A^n \times W$.*

Let M denote the set of mechanisms that satisfy the individual-reward feasibility constraint specified in Definition 3.1.

We now turn to the incentive compatibility constraint and the set of incentive compatible mechanisms for team j . For notational simplicity, we assume all agents share a common von Neumann-Morgenstern utility function u , so that agent heterogeneity enters only through type differences.²⁰ We adopt the convention that $u_{i,j}(\hat{\mathbf{t}}, \hat{\mathbf{a}}, \mathbf{w}, \mathbf{r}_j)$ denotes the utility of agent (i, j) when evaluated at type profile $\hat{\mathbf{t}}$ and action profile $\hat{\mathbf{a}}$ —that is, the subscript (i, j) indicates which agent’s utility is being evaluated. The utility function may depend on agent (i, j) ’s own type $t_{i,j}$ and action $a_{i,j}$, their team members’ types and actions, and potentially the full profiles $\hat{\mathbf{t}}$ and $\hat{\mathbf{a}}$ to capture interdependencies across teams.

Bayesian Incentive Compatibility for team j under mechanism profile $\mathbf{m} \in M^N$ requires that each member i has no incentive to unilaterally deviate from an *honest and obedient* strategy $(t'^\star_{i,j}(\cdot), a^\star_{i,j}(\cdot, \cdot, \cdot))$.²¹ That is, for all $(t'_{i,j}(\cdot), a_{i,j}(\cdot, \cdot, \cdot)) \in \mathcal{T} \times \mathcal{A}$,

$$E_{\mathbf{m}} \left[u_{i,j}(\hat{\mathbf{t}}, \hat{\mathbf{a}'}, \mathbf{w}, \mathbf{r}_j) \right] \geq E_{\mathbf{m}} \left[u_{i,j}(\hat{\mathbf{t}}, (\hat{\mathbf{a}}'_{-i,j}, a_{i,j}(t_{i,j}, t'_{i,j}(t_{i,j}), a'_{i,j})), \mathbf{w}, \mathbf{r}_j) \right]. \quad (1)$$

The left-hand side represents agent (i, j) ’s expected utility under honest and obedient play by all agents, where $\hat{\mathbf{t}}$ denotes the true type profile and $\hat{\mathbf{a}} = \hat{\mathbf{a}'}$ denotes that the recommended action profile is obediently followed by all agents. The right-hand side captures the expected utility when agent (i, j) unilaterally deviates by misreporting their type and/or disobeying their action recommendation while all other agents remain honest and obedient. Specifically, agent (i, j) may report type $t'_{i,j}(t_{i,j})$ instead of their true type $t_{i,j}$, receive an action recommendation $a'_{i,j}$ based on this report, and then choose action $a_{i,j}(t_{i,j}, t'_{i,j}(t_{i,j}), a'_{i,j})$ which may differ from the recommendation. All other agents truthfully report their types $\hat{\mathbf{t}}_{-i,j}$ and obediently follow their recommended actions $\hat{\mathbf{a}}'_{-i,j}$, which are generated by the mechanism based on the profile of reported types (including agent (i, j) ’s report).

We now examine potential deviations in the second and third stages. We first consider the effects of a deviation in the third stage, where team member i disobeys their action recommendation. If team j member i deviates from the recommended action $(a_{i,j}(t_{i,j}, t'_{i,j}(t_{i,j}), a'_{i,j}) \neq a'_{i,j})$ with strictly positive probability, this affects the transition probability of the team winnings in the fourth stage, $\Lambda(d\mathbf{w} | \hat{\mathbf{t}}, \hat{\mathbf{a}}'_{-i,j}, a_{i,j}(t_{i,j}, t'_{i,j}(t_{i,j}), a'_{i,j}))$, with the realization of team j winnings, w_j , entering the transition probability $\kappa_j(d\mathbf{r}_j | \mathbf{t}_{-i,j}, t'_{i,j}(t_{i,j}), \mathbf{a}'_j, w_j)$ of individual rewards allocated to team j members.

Next, we consider the effects of a deviation in the second stage, where team member i misreports their type. If team member i misreports their type $(t'_{i,j}(t_{i,j}) \neq t_{i,j})$ with strictly positive probability, this affects the transition probability of the recommended action profile in the third stage, $\alpha_j(d\mathbf{a}'_j | \mathbf{t}_{-i,j}, t'_{i,j}(t_{i,j}))$. Then, the realization of the recommended action profile enters the transition probability $\Lambda(d\mathbf{w} | \hat{\mathbf{t}}, \hat{\mathbf{a}}'_{-i,j}, a_{i,j}(t_{i,j}, t'_{i,j}(t_{i,j}), a'_{i,j}))$

²⁰All results extend straightforwardly to settings with agent-specific utility functions $u_{i,j}$.

²¹Note that we abstract from individual rationality, but it is straightforward to include this additional feature.

of the team winnings in the fourth stage, and ultimately affects the team's individual rewards allocated through κ_j .

Lastly, we specify the payoff functions for the principals. Each team j principal's von Neumann-Morgenstern utility function is denoted by $\pi_j : T^{nN} \times A^{nN} \times W^N \rightarrow \mathbb{R}$ which may depend on team j 's own winnings w_j , action profile \mathbf{a}_j , and type profile \mathbf{t}_j , along with the other teams' winnings \mathbf{w}_{-j} , action profiles $\{\{\mathbf{a}_{j'}\} \mid j' \neq j\}$ and type profiles $\{\{\mathbf{t}_{j'}\} \mid j' \neq j\}$.

We denote the set of all prize-feasible and incentive-compatible mechanisms for team j in the context of the mechanism profile \mathbf{m}_{-j} by $IC_j(\mathbf{m}_{-j})$. We are now in a position to state each team j principal's problem and define our equilibrium concept, Bayesian-Nash Principals' Equilibrium (BNPE). Given team j 's set of prize feasible and incentive compatible mechanisms $IC_j(\mathbf{m}_{-j})$, the team j principal's best-response problem is:

$$\sup_{m_j \in IC_j(\mathbf{m}_{-j})} E_{\mathbf{m}} \left(\pi_j \left(\hat{\mathbf{t}}, \hat{\mathbf{a}}', \mathbf{w} \right) \right) \quad (2)$$

A Bayesian-Nash Principals' Equilibrium (BNPE) is defined as follows.

Definition 3.2 (Bayesian-Nash Principals' Equilibrium (BNPE)). *A strategy profile $\mathbf{m}^* = (m_1^*, \dots, m_N^*) \in M^N$ constitutes a **Bayesian-Nash Principals' Equilibrium** if:*

- (i) **Feasibility:** For each principal j , $m_j^* = (\alpha_j^*, \kappa_j^*) \in M$, which requires that the transition probability $\kappa_j^* \in \mathcal{K}(T^n \times A^n \times W, I^n)$ satisfies the support constraint: $\text{supp } \kappa_j(\cdot | \mathbf{t}'_j, \mathbf{a}'_j, w_j) \subseteq \mathcal{W}(w_j)$ for every $(\mathbf{t}'_j, \mathbf{a}'_j, w_j) \in T^n \times A^n \times W$.
- (ii) **Incentive Compatibility** : For each team j , $m_j^* \in IC_j(\mathbf{m}_{-j}^*)$, where:
 - Agents believe others use honest and obedient strategies, which report truthfully and follow recommendations almost surely,
 - For each team j member i , the Bayesian Incentive Compatibility condition given in equation 1 holds for all unilateral deviations $(t'_{i,j}(\cdot), a_{i,j}(\cdot, \cdot, \cdot)) \in \mathcal{T} \times \mathcal{A}$.
- (iii) **Principals Best-Respond:** Each m_j^* solves principal j 's best-response problem defined in equation 2.

3.6 Multi-Principle Interaction with Team Production and Agency Concerns

To summarize, we examine the extensive-form game involving multiple principals interacting in an environment with team production and agency concerns, denoted by $\Gamma(N, n, T, A, W, I, H, \Lambda, u, \{\pi_j\}_{j=1}^N, \mathcal{W})$, in which a set of N teams (each with n team members) compete in an environment in which the set of possible team winnings is W , and the feasible individual reward correspondence \mathcal{W} maps team winnings into a set of feasible individual rewards profiles I^n .

The game begins with each team's principal specifying a team mechanism that is feasible and incentive compatible. The four stages of the continuation game are then summarized as follows:

1. **Private-Type Stage:** All nN team members privately observe their individual types, which are jointly drawn from a probability measure H on T^{nN} .
2. **Type-Reporting Stage:** The members of each team simultaneously and privately report their types to their respective team mechanisms.

3. **Action Stage:** Each team mechanism m_j privately recommends an action to each of its members based on team j 's reported types, and then all team members simultaneously choose their actions.
4. **Team-Winnings Stage:** The profile of team winnings $\mathbf{w} \in W^N$ is drawn according to Λ , conditional on the realized type and action profiles. Within each team j , the mechanism m_j allocates the team's winnings $w_j \in W$ among its n members, where the set of feasible individual reward profiles is given by $\mathcal{W}(w_j)$.

4 Assumptions and the Metric Structure

We begin by stating our key assumptions and then develop the metric structure on the space of mechanisms.

Assumptions

Ambient Space Assumptions

Assumption 1 (Compact, Complete, Separable, Metrizable Spaces). *We assume that the sets of individual types, individual actions, team winnings, and individual rewards, denoted by T , A , W , and I respectively, are compact Polish spaces endowed with their respective Borel σ -algebras. We also assume the axiom of choice and all product sets are endowed with their product σ -algebras.*

Note that from Assumption 1, it follows that $\mathcal{P}(T)$, $\mathcal{P}(A)$, $\mathcal{P}(W)$, and $\mathcal{P}(I)$ —the spaces of probability measures on T , A , W , and I respectively—are all compact Polish spaces.

Joint Winnings Transition Probability Assumptions

Assumption 2 (Narrow Continuity of Λ and Support Structure). *The transition probability $\Lambda(\cdot | \hat{\mathbf{t}}, \hat{\mathbf{a}})$ satisfies:*

1. **Pointwise Narrow Continuity in Type-Action Profiles:** The transition probability $\Lambda : T^{nN} \times A^{nN} \rightarrow \mathcal{P}(W^N)$ is narrowly (weakly) continuous in the sense that for every fixed $(\hat{\mathbf{t}}, \hat{\mathbf{a}}) \in T^{nN} \times A^{nN}$ and every sequence $(\hat{\mathbf{t}}_k, \hat{\mathbf{a}}_k) \rightarrow (\hat{\mathbf{t}}, \hat{\mathbf{a}})$, we have:

$$\Lambda(\cdot | \hat{\mathbf{t}}_k, \hat{\mathbf{a}}_k) \xrightarrow{*} \Lambda(\cdot | \hat{\mathbf{t}}, \hat{\mathbf{a}})$$

Equivalently, for every bounded continuous $f : W^N \rightarrow \mathbb{R}$,

$$\lim_{k \rightarrow \infty} \int_{W^N} f(\mathbf{w}) \Lambda(d\mathbf{w} | \hat{\mathbf{t}}_k, \hat{\mathbf{a}}_k) = \int_{W^N} f(\mathbf{w}) \Lambda(d\mathbf{w} | \hat{\mathbf{t}}, \hat{\mathbf{a}}).$$

2. **Support Structure:** There exists a family $\{\mathcal{W}_{\hat{\mathbf{t}}, \hat{\mathbf{a}}}\}_{(\hat{\mathbf{t}}, \hat{\mathbf{a}}) \in T^{nN} \times A^{nN}}$ of measurable subsets of W^N such that:

- The graph $\{(\hat{\mathbf{t}}, \hat{\mathbf{a}}, \mathbf{w}) : \mathbf{w} \in \mathcal{W}_{\hat{\mathbf{t}}, \hat{\mathbf{a}}}\}$ is measurable in $T^{nN} \times A^{nN} \times W^N$
- $\Lambda(\mathcal{W}_{\hat{\mathbf{t}}, \hat{\mathbf{a}}} | \hat{\mathbf{t}}, \hat{\mathbf{a}}) = 1$ for all $(\hat{\mathbf{t}}, \hat{\mathbf{a}})$

Feasible Rewards Correspondence

Assumption 3 (Compact-Valued, Convex-Valued, Continuous Feasible Rewards Correspondence). *The correspondence $\mathcal{W} : W \rightarrow I^n$ has nonempty, compact, convex values and is continuous (both upper and lower hemicontinuous).*

Note this assumption holds in standard cases, such as:

- *Monetary prizes:* $\mathcal{W}(w) = \{\mathbf{r} \in \mathbb{R}_+^n : \sum_{i=1}^n r_i \leq w\}$
- *Public goods:* $\mathcal{W}(w) = \{\mathbf{r} : r_i = w \ \forall i\}$

Note also that it follows from Assumption 3 that the set M (of mechanisms that satisfy the individual-reward feasibility constraint specified in Definition 3.1) is convex under pointwise convex combinations of transition probabilities. Indeed, let $(\alpha_j, \kappa_j), (\alpha'_j, \kappa'_j) \in M$ and fix $\lambda \in [0, 1]$. Define

$$\alpha_j^\lambda(\cdot | \mathbf{t}_j) := \lambda \alpha_j(\cdot | \mathbf{t}_j) + (1 - \lambda) \alpha'_j(\cdot | \mathbf{t}_j), \quad \kappa_j^\lambda(\cdot | \mathbf{t}_j, \mathbf{a}'_j, w_j) := \lambda \kappa_j(\cdot | \mathbf{t}_j, \mathbf{a}'_j, w_j) + (1 - \lambda) \kappa'_j(\cdot | \mathbf{t}_j, \mathbf{a}'_j, w_j).$$

Convex combinations of transition probabilities are transition probabilities, so α_j^λ and κ_j^λ are admissible maps. For any conditioning tuple $x = (\mathbf{t}'_j, \mathbf{a}'_j, w_j)$, feasibility of (α_j, κ_j) and (α'_j, κ'_j) implies $\text{supp } \kappa_j(\cdot | x), \text{supp } \kappa'_j(\cdot | x) \subseteq \mathcal{W}(w_j)$. Hence

$$\text{supp } \kappa_j^\lambda(\cdot | x) \subseteq \text{conv}(\text{supp } \kappa_j(\cdot | x) \cup \text{supp } \kappa'_j(\cdot | x)) \subseteq \mathcal{W}(w_j),$$

where the last inclusion uses the convexity of $\mathcal{W}(w_j)$. Thus $(\alpha_j^\lambda, \kappa_j^\lambda) \in M$, and convexity of M follows directly.

Utility Function Assumptions

Assumption 4 (von Neumann-Morgenstern Utility Functions). *The von Neumann-Morgenstern utility function u is a bounded continuous function on $T^{nN} \times A^{nN} \times W^N \times I^n$ and the von Neumann-Morgenstern utility functions $\{\pi_j\}_{j=1}^N$ are bounded continuous functions on $T^{nN} \times A^{nN} \times W^N$.*

Metric Structure

We work with two related outcome spaces. The *baseline outcome space* is

$$X := T^{nN} \times A^{nN} \times W^N \times I^{nN},$$

with generic element

$$x = (\widehat{\mathbf{t}}, \widehat{\mathbf{a}'}, \mathbf{w}, \widehat{\mathbf{r}}) \in X,$$

recording the profile of types $\widehat{\mathbf{t}} \in T^{nN}$, recommended actions $\widehat{\mathbf{a}'} \in A^{nN}$, team winnings $\mathbf{w} \in W^N$, and individual rewards $\widehat{\mathbf{r}} \in I^{nN}$.

To analyze potential deviations by a single agent, we also use an extended outcome space that records the deviating agent's type report and action choice. The *extended outcome space* for an arbitrary agent (i, j) is

$$\widetilde{X} := T^{nN+1} \times A^{nN+1} \times W^N \times I^{nN},$$

with generic element

$$\tilde{x} = (\hat{\mathbf{t}}, t'; \hat{\mathbf{a}}', a; \mathbf{w}, \hat{\mathbf{r}}) \in \tilde{X},$$

where the additional coordinates $t' \in T$ and $a \in A$ record agent (i, j) 's reported type and realized action, respectively.

For a given mechanism profile $\mathbf{m} = \{\alpha_j, \kappa_j\}_{j=1}^N \in M^N$, where $\hat{\mathbf{r}} = \{\mathbf{r}_j\}_{j=1}^N$ denotes the aggregate profile of individual rewards and $\mathbf{w} = \{w_j\}_{j=1}^N$ denotes the aggregate profile of team winnings, the joint mechanism components are constructed as:

$$\begin{aligned} \kappa^{\mathbf{m}}(d\hat{\mathbf{r}} \mid \hat{\mathbf{t}}, \hat{\mathbf{a}}', \mathbf{w}) &= \bigotimes_{j=1}^N \kappa_j(d\mathbf{r}_j \mid \mathbf{t}_j, \mathbf{a}'_j, w_j), \\ \alpha^{\mathbf{m}}(d\hat{\mathbf{a}}' \mid \hat{\mathbf{t}}) &= \bigotimes_{j=1}^N \alpha_j(da'_j \mid \mathbf{t}_j). \end{aligned}$$

Definition 4.1 (Truthful-obedient induced law). *When all agents use truthful-obedient strategies, the induced law (or probability measure) $\mu : M^N \rightarrow \mathcal{P}(X)$ on the baseline outcome space is defined by the sequential composition of transition probabilities:*

$$\mu(\mathbf{m}) := H(d\hat{\mathbf{t}}) \alpha^{\mathbf{m}}(d\hat{\mathbf{a}}' \mid \hat{\mathbf{t}}) \Lambda(d\mathbf{w} \mid \hat{\mathbf{t}}, \hat{\mathbf{a}}') \kappa^{\mathbf{m}}(d\hat{\mathbf{r}} \mid \hat{\mathbf{t}}, \hat{\mathbf{a}}', \mathbf{w}). \quad (3)$$

To build toward the robust narrow topology that our main results require, we first introduce the standard narrow topology on the space of laws $\mathcal{P}(X)$ and show how the distributional mechanism approach of Kadan, Reny, and Swinkels (2017) can be extended to our multi-team setting.

Definition 4.2 (Narrow topology on laws). *Let S be a Polish space. The narrow topology (also known as the topology of weak convergence) on $\mathcal{P}(S)$ is the coarsest topology making the maps*

$$\mathcal{P}(S) \ni \mu \longmapsto \int_S f \, d\mu \in \mathbb{R}$$

continuous for all bounded continuous functions $f \in C_b(S)$. Equivalently, a sequence $\{\mu_k\} \subset \mathcal{P}(S)$ converges to $\mu \in \mathcal{P}(S)$ in the narrow topology if and only if $\int_S f \, d\mu_k \rightarrow \int_S f \, d\mu$ for all $f \in C_b(S)$.

We now extend the distributional mechanism approach of Kadan, Reny, and Swinkels (2017) to our setting with multiple interacting teams. Each mechanism profile $\mathbf{m} = \{(\alpha_j, \kappa_j)\}_{j=1}^N \in M^N$ generates a truthful-obedient induced law $\mu(\mathbf{m}) \in \mathcal{P}(X)$. The narrow topology on $\mathcal{P}(X)$ is metrizable, with the Prokhorov metric being one natural choice.

Definition 4.3 (Prokhorov Metric). *For $\mu^1, \mu^2 \in \mathcal{P}(S)$, where S is a Polish space, define the Prokhorov metric as*

$$d_P(\mu^1, \mu^2) = \inf \left\{ \epsilon > 0 \mid \begin{array}{l} \mu^1(A) \leq \mu^2(A^\epsilon) + \epsilon \\ \text{and} \\ \mu^2(A) \leq \mu^1(A^\epsilon) + \epsilon \\ \text{for all Borel sets } A \subset S \end{array} \right\}$$

where $A^\epsilon = \{s \in S : d(s, A) < \epsilon\}$ denotes the open ϵ -neighborhood of A .

Given the Prokhorov metric, define the d_{M^N} metric between mechanism profiles $\mathbf{m}^1, \mathbf{m}^2 \in M^N$ as:

$$d_{M^N}(\mathbf{m}^1, \mathbf{m}^2) = d_P(\mu(\mathbf{m}^1), \mu(\mathbf{m}^2)).$$

The metric d_{M^N} induces the **narrow topology** on M^N , where a sequence $\{\mathbf{m}^k\}_{k \in \mathbb{N}}$ **narrowly converges** to \mathbf{m} , denoted $\mathbf{m}^k \xrightarrow{N} \mathbf{m}$, if:

$$\lim_{k \rightarrow \infty} d_{M^N}(\mathbf{m}^k, \mathbf{m}) = \lim_{k \rightarrow \infty} d_P(\mu(\mathbf{m}^k), \mu(\mathbf{m})) = 0.$$

Note that d_{M^N} treats mechanism profiles as equivalent if they generate the same truthful-obedient law, ignoring potential differences in the outcomes that arise from deviations—a limitation that becomes critical when analyzing incentive compatibility.

Given their focus on showing that the principal's objective function is lower semicontinuous in the single-team case, Kadan et al. (2017) equip M^1 with the metric d_{M^1} . A key feature of this approach is that, for any bounded continuous function $f : X \rightarrow \mathbb{R}$, the map

$$m \mapsto \int f d\mu(m)$$

is narrowly continuous with respect to $m \in M^1$. In our case, the focus will be on the correspondence of incentive-compatible mechanisms, and it will be helpful to make use of a finer topology. Towards that end, we now turn to the issue of potential deviations. We follow Balder (1988) and allow agents to use behavior strategies. A *behavior strategy* for team j member i is a pair $\sigma_{i,j} = (\sigma_{i,j}^T, \sigma_{i,j}^A)$, where:

- $\sigma_{i,j}^T \in \mathcal{K}(T, T)$ is the *type-reporting transition probability*, specifying a law over reports $t' \in T$ conditional on the agent's true type $t \in T$;
- $\sigma_{i,j}^A \in \mathcal{K}(T^2 \times A, A)$ is the *action transition probability*, specifying a law over actions $a \in A$ conditional on the agent's true type, their report, and the mechanism's action recommendation $a' \in A$.

Letting $H_{i,j} \in \mathcal{P}(T)$ denote the marginal law of agent (i, j) 's type under H , a truthful-obedient strategy may be defined in the behavior-strategy setting as follows.

Definition 4.4 (Truthful-obedient behavior strategies). *A behavior strategy $\sigma_{i,j} \in \mathcal{K}(T, T) \times \mathcal{K}(T^2 \times A, A)$ is truthful-obedient if*

$$\sigma_{i,j}^T(\cdot \mid t) = \delta_t \quad \text{for } H_{i,j}\text{-almost every } t \in T,$$

and

$$\sigma_{i,j}^A(\cdot \mid t, t, a') = \delta_{a'} \quad \text{for } H_{i,j}\text{-almost every } t \in T \text{ and all } a' \in A,$$

where δ_x denotes the Dirac probability measure concentrated at x .

Under a truthful-obedient strategy, the agent almost surely reports her type truthfully and (conditional on truthful reporting) almost surely obeys the mechanism's action recommendation. Note that the definition does not constrain off-path behavior, though such off-path choices are almost never realized under a truthful-obedient behavior strategy.

Definition 4.5 (Induced law under an arbitrary unilateral deviation). *When all agents except (i, j) use truthful-obedient strategies and agent (i, j) employs an arbitrary behavior strategy $\sigma_{i,j} \in \mathcal{K}(T, T) \times \mathcal{K}(T^2 \times A, A)$, the induced law $H_{i,j}^*$ on T is given by*

$A, A)$, the induced law $\tilde{\mu} : M^N \times \mathcal{K}(T, T) \times \mathcal{K}(T^2 \times A, A) \rightarrow \mathcal{P}(\tilde{X})$ on the extended outcome space is defined by:

$$\begin{aligned} \tilde{\mu}(\mathbf{m}, \sigma_{i,j}) := H(d\hat{\mathbf{t}}) \sigma_{i,j}^T(dt' | t_{i,j}) \boldsymbol{\alpha}^{\mathbf{m}}(d\hat{\mathbf{a}}' | \hat{\mathbf{t}}_{-i,j}, t') \sigma_{i,j}^A(da | t_{i,j}, t', a'_{i,j}) \\ \boldsymbol{\Lambda}\left(dw \mid \hat{\mathbf{t}}, (\hat{\mathbf{a}}'_{-i,j}, a)\right) \boldsymbol{\kappa}^{\mathbf{m}}\left(d\hat{\mathbf{r}} \mid (\hat{\mathbf{t}}_{-i,j}, t'), \hat{\mathbf{a}}', \mathbf{w}\right). \quad (4) \end{aligned}$$

Note that the sequential composition in (4) captures the following causal order:

1. Nature draws the profile of true types $\hat{\mathbf{t}}$ from prior H .
2. Agent (i, j) privately observes her true type $t_{i,j}$ and reports type $t' \in T$ according to her type-reporting transition probability $\sigma_{i,j}^T(dt' | t_{i,j})$. This is the first point where a deviation may occur.
3. The joint mechanism observes the reported type profile $(\hat{\mathbf{t}}_{-i,j}, t')$ and generates an action recommendation profile $\hat{\mathbf{a}}'$ via $\boldsymbol{\alpha}^{\mathbf{m}}(d\hat{\mathbf{a}}' | \hat{\mathbf{t}}_{-i,j}, t')$. Note that the recommendation to team j depends on agent (i, j) 's (possibly untruthful) report t' .
4. Agent (i, j) chooses action $a \in A$ according to her action transition probability $\sigma_{i,j}^A(da | t_{i,j}, t', a'_{i,j})$, where $a'_{i,j}$ is the recommended action from team j 's mechanism. This represents the second potential deviation point.
5. The winnings transition probability $\boldsymbol{\Lambda}$ realizes team winnings \mathbf{w} based on the true type profile $\hat{\mathbf{t}}$ and the action profile $(\hat{\mathbf{a}}'_{-i,j}, a)$.
6. The joint mechanism's reward transition probability $\boldsymbol{\kappa}^{\mathbf{m}}$ determines individual rewards $\hat{\mathbf{r}}$ conditional on the reported types $(\hat{\mathbf{t}}_{-i,j}, t')$, recommended actions $\hat{\mathbf{a}}'$, and realized winnings \mathbf{w} .

Note that this definition applies to *any* behavior strategy, not only to deviations. When $\sigma_{i,j}$ is truthful-obedient, the events $\{t' = t_{i,j}\}$ and $\{a = a'_{i,j}\}$ occur almost surely, so $\tilde{\mu}(\mathbf{m}, \sigma_{i,j})$ concentrates on the “on-path” subset of \tilde{X} .

The Set of Feasible Induced Laws achievable by Unilateral Deviations

For each mechanism profile $\mathbf{m} \in M^N$, define the *feasible set* of induced laws as:

$$\Phi(\mathbf{m}) := \{ \tilde{\mu}(\mathbf{m}, \sigma_{i,j}) : \sigma_{i,j} \in \mathcal{K}(T, T) \times \mathcal{K}(T^2 \times A, A) \}.$$

This set contains all laws over extended outcomes that a single agent (i, j) can induce through their choice of behavior strategy, given that all other agents play truthfully and obediently.

We now address how convergence of mechanism profiles is defined in our framework. A key element of our equilibrium existence problem is the correspondence of incentive-compatible mechanisms, which forms the endogenous feasible strategy correspondence in our generalized game. To handle this, we define a topology that distinguishes mechanism profiles based on the deviation opportunities available to agents. Our construction proceeds in two steps: we first introduce the Hausdorff metric, then use it to define the robust narrow topology on the joint mechanism space M^N , which accounts for both truthful-obedient induced laws and the feasible sets of induced laws achievable through unilateral deviations.

Definition 4.6 (Hausdorff metric on compact sets of laws). *Let S be a Polish space. Let $\mathsf{K}(\mathcal{P}(S))$ denote the family of nonempty compact subsets of $\mathcal{P}(S)$ with respect to the narrow topology, and let d_P^S denote the Prokhorov metric on $\mathcal{P}(S)$. For $C_1, C_2 \in \mathsf{K}(\mathcal{P}(S))$, the Hausdorff distance is*

$$d_H(C_1, C_2) = \max \left\{ \sup_{\mu \in C_1} \inf_{\nu \in C_2} d_P^S(\mu, \nu), \sup_{\nu \in C_2} \inf_{\mu \in C_1} d_P^S(\mu, \nu) \right\}.$$

The Hausdorff distance measures how well one compact set of laws can be approximated by another: C_1 and C_2 are close in Hausdorff distance if every law in C_1 is close (in Prokhorov distance) to some law in C_2 , and vice versa.

We are now in position to introduce the robust narrow metric, which will allow us to examine convergence of both the truthful-obedient induced laws and the feasible sets of induced laws achievable through unilateral deviations.²²

Definition 4.7 (Robust narrow metric). *Let $\mu : M^N \rightarrow \mathcal{P}(X)$ denote the truthful-obedient induced law defined in (3), let $\Phi(\mathbf{m})$ denote the set of induced laws achievable through unilateral deviations from mechanism profile \mathbf{m} , let d_P^X denote the Prokhorov metric on $\mathcal{P}(X)$, and let d_H denote the Hausdorff distance on $\mathsf{K}(\mathcal{P}(\tilde{X}))$. For mechanism profiles $\mathbf{m}^1, \mathbf{m}^2 \in M^N$, the robust narrow distance is*

$$d_{M^N}^*(\mathbf{m}^1, \mathbf{m}^2) = \max \left\{ d_H(\overline{\Phi(\mathbf{m}^1)}, \overline{\Phi(\mathbf{m}^2)}), d_P^X(\mu(\mathbf{m}^1), \mu(\mathbf{m}^2)) \right\}.$$

The robust narrow topology on M^N incorporates two distinct proximity criteria for mechanism profiles. Two mechanisms \mathbf{m}^1 and \mathbf{m}^2 are close in this topology when both of the following hold:

- *Robustness to strategic behavior.* The Hausdorff distance $d_H(\overline{\Phi(\mathbf{m}^1)}, \overline{\Phi(\mathbf{m}^2)})$ between the sets of laws attainable through unilateral deviations is small.
- *Truthful-obedient laws.* The Prokhorov distance $d_P^X(\mu(\mathbf{m}^1), \mu(\mathbf{m}^2))$ between the truthful-obedient induced laws is small.

We adopt the following notation for convergence:

- $Q_n \xrightarrow{N} Q$ denotes narrow convergence in $(\mathcal{P}(X), d_P^X)$;
- $\mathcal{A}_n \xrightarrow{H} \mathcal{A}$ denotes Hausdorff convergence in $(\mathsf{K}(\mathcal{P}(\tilde{X})), d_H)$, where d_H is induced by the Prokhorov metric $d_P^{\tilde{X}}$.

Under the robust narrow metric $d_{M^N}^*$ of Definition 4.7, a sequence $(\mathbf{m}^k)_{k \geq 1} \subset M^N$ converges to $\mathbf{m} \in M^N$ if and only if both components converge:

$$\overline{\Phi(\mathbf{m}^k)} \xrightarrow{H} \overline{\Phi(\mathbf{m})} \quad \text{and} \quad \mu(\mathbf{m}^k) \xrightarrow{N} \mu(\mathbf{m}). \quad (5)$$

Equivalently, $\mathbf{m}^k \rightarrow \mathbf{m}$ in $d_{M^N}^*$ if and only if

$$d_H(\overline{\Phi(\mathbf{m}^k)}, \overline{\Phi(\mathbf{m})}) \rightarrow 0 \quad \text{and} \quad d_P^X(\mu(\mathbf{m}^k), \mu(\mathbf{m})) \rightarrow 0.$$

²²Note that by Assumption 1, \tilde{X} is a compact metric space, so every law on \tilde{X} is automatically tight. Therefore, any subset of $\mathcal{P}(\tilde{X})$ —in particular $\Phi(\mathbf{m})$ —is uniformly tight. By Prokhorov's theorem, the closure $\overline{\Phi(\mathbf{m})}$ is compact in $(\mathcal{P}(\tilde{X}), d_P^{\tilde{X}})$, ensuring that the Hausdorff distance between such closures is well-defined.

Note that the equilibrium concept itself (Definition 3.2) does not depend on how the mechanism space M^N is topologized: a mechanism profile \mathbf{m}^* constitutes a BNPE if and only if it satisfies feasibility, incentive compatibility, and principals' best-responding. These equilibrium conditions can be verified for any given \mathbf{m}^* without reference to how mechanisms converge or how we measure distances between mechanisms.

The robust narrow topology on M^N provides a natural way to measure similarity between mechanism profiles in this environment: two mechanisms are close when they induce similar truthful-obedient outcome laws and similar sets of outcome laws achievable through unilateral deviations. This notion of closeness captures the two strategic dimensions relevant to each principal's correspondence of incentive-compatible mechanisms. As we will establish, when similarity is measured in this natural way, the best-response correspondence inherits the regularity properties needed for equilibrium existence.

Example

Before moving on to our results, we present an example – a generalized principle-agent team contest along the lines of the literature following Nitzan (1991) – that satisfies Assumptions 1-4.

Consider a contest involving N teams, each of which faces a generalized principal-agent problem while competing for a single prize that is divisible. Each team consists of n members. In the initial stage, each team's principal specifies a feasible and incentive-compatible team mechanism. We begin by describing the four stages of the continuation game, and then return to discuss the initial stage.

First (Private Type) Stage: In the first stage, the members of each team privately realize their individual types, where the *type space* is $T = [\underline{t}, \bar{t}]$, with $0 < \underline{t} < \bar{t} < \infty$, and each team member's type $t_{i,j} \in T$ is an independent draw from a common probability measure $\mu \in \mathcal{P}(T)$.

Second (Type Reporting) Stage: In the second stage, the members of each team privately report their types to the team principal, where for team j member i , $t'_{i,j} \in T$ denotes the reported type and $t_{i,j}$ denotes the true type.

Third (Action Reporting) Stage: In the third stage, the mechanism privately recommends an action from the *action space* $A = [\underline{a}, \bar{a}]$, with $0 < \underline{a} < \bar{a} < \infty$, denoted as $a'_{i,j} \in A$ for team j member i , and then the team members simultaneously choose actions, where team j member i 's actual action is denoted as $a_{i,j} \in A$.

Fourth (Team Winnings) Stage: In the fourth stage, each team's winnings are determined by competition in team outputs, where the output technology stochastically maps a team's n -tuple of types and n -tuple of actions into the team output. For a given team j profile $(\mathbf{a}_j, \mathbf{t}_j) \in A^n \times T^n$, let $P_j(\cdot | \mathbf{a}_j, \mathbf{t}_j)$ denote team j 's transition probability over the *output space* $O = [0, 1]$. The cumulative distribution function (CDF) $F_j(x | \mathbf{t}_j, \mathbf{a}_j)$ is defined as:

$$F_j(x | \mathbf{t}_j, \mathbf{a}_j) = P_j(o \leq x | \mathbf{t}_j, \mathbf{a}_j),$$

where $o \in O$ denotes the output.

Teams are ranked by their output values, where output o is valued as $V(o) = o$. In this winner-take-all contest, team winnings $W = \{0, 1\}$ are allocated according to rank, with only the highest output team receiving the prize. Team winnings take the form of a perfectly divisible monetary prize, with individual rewards in $I = [0, 1]$. Within each team, the mechanism allocates the team winnings $w \in W$ among the

n team members, where the feasible set of individual reward profiles given team winnings w is $\mathcal{W}(w) = \{\mathbf{r} \in I^n \mid \sum_{i=1}^n r_i \leq w\}$.

Given that the contest has a single prize, the probability that a team j has the highest value output and wins the prize, conditional the profile $\widehat{\mathbf{a}}, \widehat{\mathbf{t}}$, is:

$$\Lambda(w_j = 1, w_{j'} = 0 \forall j' \neq j \mid \widehat{\mathbf{a}}, \widehat{\mathbf{t}}) = \text{Prob}(j \text{ wins} \mid \widehat{\mathbf{a}}, \widehat{\mathbf{t}}) = \int_0^1 \left(\prod_{j' \neq j} F_{j'}(x \mid \mathbf{a}_{j'}, \mathbf{t}_{j'}) \right) dF_j(x \mid \mathbf{a}_j, \mathbf{t}_j). \quad (6)$$

We will assume that each distribution function $F_j(x \mid \mathbf{a}_j, \mathbf{t}_j)$ is continuous in the parameters $(\mathbf{a}_j, \mathbf{t}_j)$ for each fixed $x \in [0, 1]$.²³

Principal's Initial Problems To complete the specification of the example, for each team j member i , the von Neumann-Morgenstern utility function is given by

$$u_{i,j}(r_{i,j}, a_{i,j}, t_{i,j}) = r_{i,j} - \frac{ca_{i,j}}{t_{i,j}},$$

where $c > 0$, and for each team j principal, the von Neumann-Morgenstern utility function is given by

$$\pi_j(w_j) = w_j.$$

Satisfaction of Assumptions in the Example

With respect to the updated set of assumptions, we now verify that the example satisfies Assumptions 1–4:

Ambient Space Assumptions: The example spaces T , A , W , and I satisfy the *Ambient Space Assumptions* (Assumption 1). Specifically, these spaces are all compact Polish spaces – ensuring completeness, separability, and metrizability – and the product spaces, such as T^{nN} and A^n , inherit these properties.

Joint Winnings Transition Probability Assumptions: Given the explicit formula for Λ , each outcome \mathbf{w} corresponds to exactly one team j winning (that is, $w_j = 1$ and $w_{j'} = 0$ for all $j' \neq j$), satisfying the support structure in Assumption 2. The probability assigned to team j winning is a continuous function of the type-action profile $(\widehat{\mathbf{t}}, \widehat{\mathbf{a}})$. Since the measure $\Lambda(\cdot \mid \widehat{\mathbf{t}}, \widehat{\mathbf{a}})$ is supported on finitely many points (the N possible winners) and each point's probability varies continuously with $(\widehat{\mathbf{t}}, \widehat{\mathbf{a}})$, the transition probability satisfies pointwise narrow continuity as required in Assumption 2.

²³Note that for the stochastic output model in which for a team j 's profile $(\mathbf{a}_j, \mathbf{t}_j) \in A^n \times T^n$:

$$F_j(x \mid \mathbf{a}_j, \mathbf{t}_j) = P_j(o \leq x \mid \mathbf{a}_j, \mathbf{t}_j) = x^{\left(\sum_{i=1}^n t_{i,j} a_{i,j}\right)},$$

we can substitute this into the general winning probability expression in equation (6), and it follows that the prize is awarded via a ratio-form CSF,

$$\Lambda(w_j = 1, w_{j'} = 0 \forall j' \neq j \mid \widehat{\mathbf{a}}, \widehat{\mathbf{t}}) = \text{Prob}(j \text{ wins} \mid \widehat{\mathbf{a}}, \widehat{\mathbf{t}}) = \frac{\sum_{i=1}^n t_{i,j} a_{i,j}}{\sum_{j'=1}^N \sum_{i=1}^n t_{i,j'} a_{i,j'}}.$$

Feasible Rewards Correspondence Assumptions: The feasible individual rewards correspondence $\mathcal{W} : W \rightarrow I^n$, defined by

$$\mathcal{W}(w) = \left\{ \mathbf{r} \in I^n : \sum_{i=1}^n r_i \leq w \right\}$$

for $w \in \{0, 1\}$, satisfies the *Feasible Rewards Correspondence Assumptions* (Assumption 3). Indeed, $\mathcal{W}(w)$ is nonempty (always contains $\mathbf{0} = (0, \dots, 0) \in I^n$), convex-valued, and compact-valued for each w . Because w takes only finitely many values, both upper and lower hemicontinuity hold trivially.

Utility Function Assumptions: Lastly, the example clearly satisfies the *Utility Function Assumptions* (Assumption 4). Specifically, the von Neumann-Morgenstern utility function

$$u_{i,j}(r_{i,j}, a_{i,j}, t_{i,j}) = r_{i,j} - \frac{ca_{i,j}}{t_{i,j}},$$

is a continuous function that is bounded on the compact domain $I \times A \times T$. Similarly, the von Neumann-Morgenstern utility function

$$\pi_j(w) = w$$

is a bounded continuous function on W .

5 Results

Our main result is to show that there exists a Bayesian-Nash Principals' equilibrium (BNPE) of the multi-principle interaction with team production and agency concerns team game $\Gamma(N, n, T, A, W, I, H, \Lambda, u, \{\pi_j\}_{j=1}^N, \mathcal{W})$.

In equilibrium, each team j makes use of a mechanism (α_j, κ_j) that maximizes the expected payoff of the team j principal subject to: (i) prize feasibility of the mechanism (α_j, κ_j) , and (ii) Bayesian incentive compatibility of the mechanism (α_j, κ_j) .

5.1 Existence of Equilibrium

Our main result is stated as follows.

Theorem 1. *Endow M^N with the $d_{M^N}^*$ metric. If $IC(\mathbf{m}) := \prod_{j=1}^N IC_j(\mathbf{m}_{-j})$ admits a selection, then under Assumptions 1–4, for each team j :*

- IC_j is nonempty, continuous, compact-valued, and convex-valued as a correspondence on the space of the other teams' mechanism profiles, and
- each team j principal's expected payoff is continuous on M^N and quasi-concave in m_j .

Therefore, there exists a BNPE.

Before sketching the key arguments behind Theorem 1, we note that a sufficient condition for IC to admit a selection is the existence of at least one mechanism that is always incentive compatible, as in the example in Myerson (1982). This ensures non-emptiness of the feasible set at every mechanism profile.

The proof of Theorem 1 establishes several properties of the incentive-compatible mechanism correspondence IC_j under the robust narrow topology. We focus our discussion here on the continuity of IC_j as a

correspondence, which is the most technically demanding component. The robust narrow topology plays a central role in establishing continuity by simultaneously tracking convergence of on-path induced laws and agents' deviation opportunities.

For completeness, we briefly outline how the remaining properties are established, with full details provided in the Appendix: IC_j is compact-valued by the compactness of the underlying strategy spaces together with the closed graph property established below; IC_j is convex-valued by the linearity of agents' expected payoffs; continuity of each principal's expected payoff follows along similar lines to the continuity of IC_j , since payoffs are defined via integration against the truthful-obedient induced law; and quasi-concavity of the principals' expected payoffs follows from Balder (1988, Theorem 3.1).

Now we turn to sketching the proof of the continuity of IC_j . To verify incentive compatibility, we must ensure that truthful reporting and obedient action-taking weakly dominate all possible unilateral combinations of misreporting and disobedience. Our approach makes use of our two outcome spaces: the baseline space X containing all payoff-relevant variables, and the extended space \tilde{X} that additionally tracks one agent's strategic choices. The truthful-obedient induced law $\mu(\mathbf{m})$ lives in $\mathcal{P}(X)$, while the set of induced laws under all possible unilateral deviations by a given agent, $\Phi(\mathbf{m})$, lives in $\mathcal{P}(\tilde{X})$.

This leads to a measurement problem: agent (i, j) 's utility depends only on payoff-relevant variables in X , yet her deviation possibilities generate distributions over the extended outcome space \tilde{X} . To compare an agent's truthful-obedient payoff with her payoffs under unilateral deviations, we employ a projection $pr_{i,j} : \tilde{X} \rightarrow X$ that maps each extended outcome to its payoff-relevant components. This projection discards agent (i, j) 's type report and her recommended action, retaining only the variables that affect her utility.

Since $pr_{i,j}$ is continuous, it induces a pushforward operation on probability measures: any law $\tilde{\mu} \in \mathcal{P}(\tilde{X})$ over strategic choices maps to a law $(pr_{i,j})_* \tilde{\mu} \in \mathcal{P}(X)$ over payoff-relevant outcomes, defined by²⁴

$$(pr_{i,j})_* \tilde{\mu}(A) = \tilde{\mu}(pr_{i,j}^{-1}(A))$$

for measurable sets $A \subset X$. This operation computes the marginal law on X by integrating out agent (i, j) 's type report and recommended action. The pushforward preserves the probability structure while discarding strategically-chosen but payoff-irrelevant variables. This operation is key to our approach—it makes the deviation set $\Phi(\mathbf{m}) \subseteq \mathcal{P}(\tilde{X})$ comparable with the truthful-obedient law $\mu(\mathbf{m}) \in \mathcal{P}(X)$.

Using this pushforward, we can express agent (i, j) 's *incentive compatibility slack* entirely in terms of laws on X :²⁵

$$f_{i,j}(\mathbf{m}) = \int_X u_{i,j} d\mu(\mathbf{m}) - \sup_{\nu \in (pr_{i,j})_* \overline{\Phi(\mathbf{m})}} \int_X u_{i,j} d\nu. \quad (7)$$

This formulation depends on two objects: the truthful-obedient law $\mu(\mathbf{m}) \in \mathcal{P}(X)$ and the projected feasible set $(pr_{i,j})_* \overline{\Phi(\mathbf{m})} \subseteq \mathcal{P}(X)$ containing all outcome laws achievable through strategic deviations. The structure of equation (7) reveals why the robust narrow metric is precisely tailored to ensure continuity of $f_{i,j}$.

²⁴Note that $pr_{i,j}^{-1}$ denotes the preimage operation. It captures all the different configurations in \tilde{X} (with varying reports $t'_{i,j}$ and recommendations $a'_{i,j}$) that lead to the same payoff-relevant outcome in X .

²⁵Regarding the closure $\overline{\Phi(\mathbf{m})}$ in the IC slack condition: Since $u_{i,j}$ is continuous and bounded on the compact space X , the payoff functional $\mathcal{U}_{i,j}(\mu) = \int u_{i,j} d\mu$ is continuous on $\mathcal{P}(X)$. Moreover, the pushforward map $(pr_{i,j})_*$ is continuous. The composition $\tilde{\mu} \mapsto \mathcal{U}_{i,j}((pr_{i,j})_* \tilde{\mu})$ is therefore continuous on $\mathcal{P}(\tilde{X})$. By continuity,

$$\sup_{\tilde{\mu} \in \Phi(\mathbf{m})} \mathcal{U}_{i,j}((pr_{i,j})_* \tilde{\mu}) = \sup_{\tilde{\mu} \in \overline{\Phi(\mathbf{m})}} \mathcal{U}_{i,j}((pr_{i,j})_* \tilde{\mu}),$$

so the IC condition is equivalent whether we use $\Phi(\mathbf{m})$ or its closure. The closure ensures a maximizer exists (by compactness) without affecting the constraint.

The key insight is that the robust narrow topology accounts for both terms in (7). Its narrow topology component ensures that expectations of bounded continuous functions vary continuously in \mathbf{m} when evaluated at individual measures like $\mu(\mathbf{m})$. Its Hausdorff component guarantees that the correspondence $\mathbf{m} \mapsto \overline{\Phi(\mathbf{m})}$ varies continuously—a property that follows from the equivalence between Hausdorff convergence and Kuratowski convergence for compact subsets in Polish spaces. Since \widetilde{X} is compact Polish, the Hausdorff metric topology coincides with the Fell topology (Beer, 1993, Corollary 5.1.11), which in turn implies both upper and lower hemicontinuity of $\mathbf{m} \mapsto \overline{\Phi(\mathbf{m})}$ (Beer and Rodríguez-López, 2010).

More specifically, the two components of the robust narrow topology correspond directly to the two terms in (7):

- **Prokhorov distance** between truthful-obedient laws $\mu(\mathbf{m}^1)$ and $\mu(\mathbf{m}^2)$ ensures continuity of the first term, $\int_X u_{i,j} d\mu(\mathbf{m})$.
- **Hausdorff distance** between deviation sets $\overline{\Phi(\mathbf{m}^1)}$ and $\overline{\Phi(\mathbf{m}^2)}$ ensures continuity of the second term, the supremum over all outcome laws achievable through unilateral deviations.

Together, these components guarantee that $f_{i,j}(\mathbf{m})$ is continuous in \mathbf{m} under the robust narrow metric via a Maximum Theorem argument: when the constraint correspondence is continuous and the objective is bounded and continuous, the supremum function is continuous. We formally establish this as Lemma 2 in the Appendix.

6 Conclusion

When multiple principals simultaneously design mechanisms for their respective teams, the interdependence of their choices creates a generalized game where each principal’s feasible set of incentive-compatible mechanisms depends on others’ choices. As Myerson (1982) demonstrates, such games may lack equilibrium due to discontinuities in feasible strategy correspondences. In this paper, we establish equilibrium existence for multi-principal mechanism design in environments with strategic spillovers. Our framework—which accommodates multidimensional types, actions, outputs, and rewards in environments with both adverse selection and moral hazard—provides a robust and versatile foundation for analyzing strategic environments involving multiple interacting principals. This includes settings with competing platforms, interfirm contracting networks, and hierarchical organizations where multiple decision-makers design complementary incentive schemes. By establishing when equilibria exist in these economically important environments, our results open new avenues for analyzing how strategic interaction shapes the design of incentive systems.

7 Appendix: Proof of Theorem 1

We begin with an overview of the proof of Theorem 1, and then in the following subsections address each of the steps of the proof in detail.

Overview

The proof demonstrating the existence of a Bayesian-Nash Principal’s equilibrium mechanism profile \mathbf{m} is summarized as follows. First, we show that the principals’ expected payoff functions $\{E_{\pi_j}(\mathbf{m})\}_{j=1}^N$ are continuous in \mathbf{m} , with respect to the d_M^* metric, and quasi-concave in m_j . Second, we show that the

correspondence $IC_j : M^{N-1} \rightarrow M$ is continuous with respect to the metric on M^{N-1} inherited from d_{MN}^* , with compact and convex values. Third, applying Berge's Maximum Theorem, it follows that the best-response correspondence $BR_j : M^{N-1} \rightarrow M$ is upper hemicontinuous with respect to the same metric, with nonempty, compact, and convex values. Finally, we apply the Kakutani-Fan-Glicksberg fixed-point Theorem to the best-response correspondence to demonstrate the existence of a Bayesian-Nash Principals' equilibrium profile of incentive-compatible mechanisms \mathbf{m} .

We begin with the proof that the principals' expected utility functionals are continuous and quasi-concave

Continuity and Quasi-concavity of Principal's Expected Utility

For a given mechanism profile $\mathbf{m} = \{\alpha_j, \kappa_j\}_{j=1}^N \in M^N$, note that the expected utility functional for the team j principal may be written as

$$V_{\pi_j}(\mathbf{m}) = \int \pi_j d\mu(\mathbf{m}). \quad (8)$$

Lemma 1. *The functional $V_{\pi_j}(\mathbf{m})$, defined in equation (8) is continuous in $\mathbf{m} = (m_j, \mathbf{m}_{-j})$ and concave in $m_j = \{\alpha_j, \kappa_j\}$.*

The proof of the continuity and quasiconcavity of $V_{\pi_j}(\mathbf{m})$ follows along similar lines as Balder (1988, Theorem 3.1), which makes use of his Theorem 2.2 and Lemma 3.2.

Beginning with the continuity of $V_{\pi_j}(\mathbf{m})$, by Definition 4.7, the metric d_{MN}^* metrizes the robust narrow topology so that $\mathbf{m}^k \rightarrow \mathbf{m}$ implies $\mu(\mathbf{m}^k) \xrightarrow{N} \mu(\mathbf{m})$ (narrow convergence). Since π_j is a bounded continuous function, it follows that $\int \pi_j d\mu(\mathbf{m}^k) \rightarrow \int \pi_j d\mu(\mathbf{m})$, i.e., $V_{\pi_j}(\mathbf{m}^k) \rightarrow V_{\pi_j}(\mathbf{m})$ and therefore $V_{\pi_j}(\mathbf{m})$ is continuous in \mathbf{m} .

Next, the quasi-concavity of V_{π_j} in m_j can be demonstrated as follows. For a given mechanism profile $\mathbf{m} \in M^N$, recall the joint action recommendation mechanism component is constructed as:

$$\boldsymbol{\alpha}^{\mathbf{m}}\left(d\widehat{\mathbf{a}}' \mid \widehat{\mathbf{t}}\right) = \bigotimes_{j=1}^N \alpha_j\left(d\mathbf{a}'_j \mid \mathbf{t}_j\right),$$

where $\mathbf{m} = \{\alpha_j, \kappa_j\}_{j=1}^N$.

Note that the equation (8) expected utility functional for the team j principal may be rewritten as

$$V_{\pi_j}(\mathbf{m}) = \int_{W^N \times A^{nN} \times T^{nN}} \pi_j(\widehat{\mathbf{t}}, \widehat{\mathbf{a}}', \mathbf{w}) \boldsymbol{\Lambda}(d\mathbf{w} \mid \widehat{\mathbf{t}}, \widehat{\mathbf{a}}') \boldsymbol{\alpha}^{\mathbf{m}}(d\widehat{\mathbf{a}}' \mid \widehat{\mathbf{t}}) H(d\widehat{\mathbf{t}}). \quad (9)$$

For any $m_j = (\alpha_j, \kappa_j)$, $m'_j = (\alpha'_j, \kappa'_j) \in M$, and $\lambda \in [0, 1]$, define the joint convex combination

$$m_j^\lambda = (\alpha_j^\lambda, \kappa_j^\lambda) := (\lambda \alpha_j + (1 - \lambda) \alpha'_j, \lambda \kappa_j + (1 - \lambda) \kappa'_j).$$

For a given mechanism profile $\mathbf{m}_{-j} \in M^{N-1}$, we construct $\boldsymbol{\alpha}^{\mathbf{m}_{-j}}$ as:

$$\boldsymbol{\alpha}^{\mathbf{m}_{-j}}\left(d\widehat{\mathbf{a}}'_{-j} \mid \widehat{\mathbf{t}}_{-j}\right) = \bigotimes_{j' \neq j} \alpha_{j'}\left(d\mathbf{a}'_{j'} \mid \mathbf{t}_{j'}\right).$$

Then:

$$V_{\pi_j}(m_j^\lambda, \mathbf{m}_{-j}) = \int \pi_j(\widehat{\mathbf{t}}, \widehat{\mathbf{a}}', \mathbf{w}) \boldsymbol{\Lambda}(d\mathbf{w} \mid \widehat{\mathbf{t}}, \widehat{\mathbf{a}}')[\lambda \alpha_j + (1 - \lambda) \alpha'_j](d\mathbf{a}'_j \mid \mathbf{t}_j) \boldsymbol{\alpha}^{\mathbf{m}_{-j}}(d\widehat{\mathbf{a}}'_{-j} \mid \widehat{\mathbf{t}}_{-j}) H(d\widehat{\mathbf{t}}). \quad (10)$$

Since the integral in equation (10) is linear in $[\lambda\alpha_j + (1 - \lambda)\alpha'_j]$ and π_j is bounded (Assumption 4), it follows that:

$$V_{\pi_j}(m_j^\lambda, \mathbf{m}_{-j}) = \lambda V_{\pi_j}(m_j, \mathbf{m}_{-j}) + (1 - \lambda) V_{\pi_j}(m'_j, \mathbf{m}_{-j}), \quad (11)$$

establishing that $m_j \mapsto V_{\pi_j}(m_j, \mathbf{m}_{-j})$ is affine, and hence concave and quasiconcave, on the convex set M .

IC Correspondence: Continuity, Compactness, and Convex-Valuedness

Notation Summary

For the reader's convenience, we review the key notation used in this subsection of the appendix. Complete definitions and all assumptions appear in Sections 3–4 of the main text.

Consider an agent (i, j) facing a mechanism profile \mathbf{m} . The agent observes her true type and chooses a type report, then observes her recommended action and chooses an action. Recall that:

- X denotes the *baseline outcome space*, consisting of all agents' types, implemented actions, team winnings, and individual rewards. This is the payoff-relevant state space.
- Each agent (i, j) 's utility $u_{i,j} : X \rightarrow \mathbb{R}$ depends only on outcomes in the baseline space X .
- $\mu(\mathbf{m}) \in \mathcal{P}(X)$ denotes the truthful-obedient induced law on the baseline outcome space.
- \tilde{X} denotes the *extended outcome space*, which augments X by additionally recording one agent's type report and the mechanism's action recommendation.
- $\Phi_{i,j}(\mathbf{m}) \subseteq \mathcal{P}(\tilde{X})$ denotes the set of induced laws on the extended outcome space achievable by an agent (i, j) behavior strategy, while all other agents remain truthful and obedient.
- $\text{pr}_{i,j} : \tilde{X} \rightarrow X$ denotes the projection that maps the extended outcome space to the baseline outcome space by replacing agent (i, j) 's recommended action with her actual chosen action and discarding all type reports and action recommendations.
- The pushforward measure is defined via the preimage of sets: for any measurable set $B \subseteq X$,

$$[(\text{pr}_{i,j})_* \tilde{\mu}](B) = \tilde{\mu}(\text{pr}_{i,j}^{-1}(B)),$$

where $\text{pr}_{i,j}^{-1}(B) = \{\tilde{x} \in \tilde{X} : \text{pr}_{i,j}(\tilde{x}) \in B\}$. By marginalizing over components (reports and recommendations) that are not payoff-relevant, the pushforward induces the law on the (payoff-relevant) baseline outcome space implied by the law $\tilde{\mu}$ on the extended outcome space.

IC Slack

Fix a mechanism profile $\mathbf{m} \in M^N$. Recall from equation (7) that agent (i, j) 's *incentive compatibility slack* under \mathbf{m} is:

$$f_{i,j}(\mathbf{m}) = \int_X u_{i,j} d\mu(\mathbf{m}) - \sup_{\nu \in (\text{pr}_{i,j})_* \overline{\Phi(\mathbf{m})}} \int_X u_{i,j} d\nu. \quad (7)$$

The first term on the right-hand side of equation (7) is agent (i, j) 's expected utility from truth-telling and obedience. The second term on the right-hand side of equation (7) is her maximal achievable utility

over all possible deviations. The mechanism profile \mathbf{m} is incentive compatible for agent (i, j) if and only if $f_{i,j}(\mathbf{m}) \geq 0$.

Continuity of Team j 's Correspondence of IC Mechanisms

We begin by establishing that each team j 's IC correspondence is continuous with respect to the robust narrow metric d_M^* on mechanism profiles. We then turn to establishing that it is compact-valued and convex-valued.

Lemma 2. *The incentive compatibility slack $f_{i,j} : M^N \rightarrow \mathbb{R}$ is continuous with respect to the robust narrow metric $d_{M^N}^*$.*

Proof. The proof of Lemma 2 consists of two steps.

Step 1 (truthful-obedient utility): The first term in (7), $\int_X u_{i,j} d\mu(\mathbf{m})$, is continuous in \mathbf{m} because by Definition 4.7, the metric $d_{M^N}^*$ metrizes the robust narrow topology so that $\mathbf{m}^k \rightarrow \mathbf{m}$ implies $\mu(\mathbf{m}^k) \xrightarrow{N} \mu(\mathbf{m})$ (narrow convergence). Since $u_{i,j}$ is a bounded continuous function, it follows that $\int_X u_{i,j} d\mu(\mathbf{m}^k) \rightarrow \int_X u_{i,j} d\mu(\mathbf{m})$.

Step 2 (unilateral-deviation utility): For the second term in (7), by Definition 4.7, $\mathbf{m}^k \rightarrow \mathbf{m}$ implies $d_H(\overline{\Phi(\mathbf{m}^k)}, \overline{\Phi(\mathbf{m})}) \rightarrow 0$ (Hausdorff convergence). Since \tilde{X} is compact Polish, the Hausdorff metric topology coincides with the Fell topology (Beer, 1993, Corollary 5.1.11), and for compact subsets, convergence in the Fell topology is equivalent to Kuratowski convergence (Beer and Rodríguez-López, 2010). This implies that the correspondence $\mathbf{m} \mapsto \overline{\Phi(\mathbf{m})}$ is both upper and lower hemicontinuous.

Since $\text{pr}_{i,j} : \tilde{X} \rightarrow X$ is continuous, the pushforward map $(\text{pr}_{i,j})_* : \mathcal{P}(\tilde{X}) \rightarrow \mathcal{P}(X)$ is continuous and preserves this hemicontinuity, so that

$$d_H\left((\text{pr}_{i,j})_* \overline{\Phi(\mathbf{m}^k)}, (\text{pr}_{i,j})_* \overline{\Phi(\mathbf{m})}\right) \rightarrow 0.$$

Since $u_{i,j}$ is bounded and continuous, the supremum functional $\sup_{\nu \in (\text{pr}_{i,j})_* \overline{\Phi(\mathbf{m})}} \int_X u_{i,j} d\nu$ is continuous with respect to Hausdorff convergence. Therefore,

$$\sup_{\nu \in (\text{pr}_{i,j})_* \overline{\Phi(\mathbf{m}^k)}} \int_X u_{i,j} d\nu \rightarrow \sup_{\nu \in (\text{pr}_{i,j})_* \overline{\Phi(\mathbf{m})}} \int_X u_{i,j} d\nu.$$

Note that this argument is essentially an application of the Maximum Theorem (Ok, 2007, Chapter E, Section 3), which guarantees continuity of the maximum when the constraint correspondence is continuous and the objective function is bounded and continuous. As established in footnote 24, the maximum exists and equals the supremum due to compactness of $\overline{\Phi(\mathbf{m})}$ and continuity of the objective.

Since both terms in (7) converge, it follows that $f_{i,j}(\mathbf{m}^k) \rightarrow f_{i,j}(\mathbf{m})$. Therefore $f_{i,j}(\mathbf{m})$ is continuous in \mathbf{m} , and that completes the proof. \square

Next, note that team j 's correspondence of IC mechanisms is given by:

$$\text{IC}_j(\mathbf{m}_{-j}) = \{m_j \in M : f_{i,j}(m_j, \mathbf{m}_{-j}) \geq 0 \text{ for all } i \text{ in team } j\}.$$

Since IC_j is defined by continuous functions (by Lemma 2) and weak inequalities, we have the following result.

Lemma 3. *The correspondence $\text{IC}_j : M^{N-1} \twoheadrightarrow M$ is continuous with respect to the metric on M^{N-1} inherited from $d_{M^N}^*$.*

Compact-valuedness of the IC correspondence.

We now turn to the compact-valuedness of the IC correspondence.

Lemma 4. *For each $\mathbf{m}_{-j} \in M^{N-1}$, the set $\text{IC}_j(\mathbf{m}_{-j}) \subseteq M$ is compact with respect to d_M^* .*

Lemma 4 follows from the fact that for each \mathbf{m}_{-j} , the sublevel set

$$S_i(\mathbf{m}_{-j}) := \{m_j \in M : f_{i,j}(m_j, \mathbf{m}_{-j}) \geq 0\}$$

is closed in M by continuity of $f_{i,j}(\cdot, \mathbf{m}_{-j})$ (Lemma 2). Hence

$$\text{IC}_j(\mathbf{m}_{-j}) = \bigcap_{i \in \text{team } j} S_i(\mathbf{m}_{-j})$$

is an intersection of finitely many closed sets, thus closed in M . Since M is compact, every closed subset is compact; therefore $\text{IC}_j(\mathbf{m}_{-j})$ is compact.

Convex-valuedness of the IC correspondence.

Next, we turn to the convex-valuedness of the IC correspondence.

Lemma 5. *For each $\mathbf{m}_{-j} \in M^{N-1}$, the set $\text{IC}_j(\mathbf{m}_{-j}) \subseteq M$ is convex.*

To see this, fix \mathbf{m}_{-j} and suppose $m_j, m'_j \in \text{IC}_j(\mathbf{m}_{-j})$. For any $\lambda \in [0, 1]$, interpret the convex combination $m''_j = \lambda m_j + (1 - \lambda)m'_j$ as a mechanism that implements m_j with probability λ and m'_j with probability $1 - \lambda$. Then

$$f_{i,j}(\mathbf{m}'') = \lambda f_{i,j}(\mathbf{m}) + (1 - \lambda)f_{i,j}(\mathbf{m}') \geq 0$$

whenever both $f_{i,j}(\mathbf{m}) \geq 0$ and $f_{i,j}(\mathbf{m}') \geq 0$. This holds for all agents (i, j) in team j , confirming that $m''_j \in \text{IC}_j(\mathbf{m}_{-j})$.

Best-Response Correspondence

Given that each principal j 's objective function is continuous and quasi-concave and each IC_j is continuous and nonempty compact-valued, it follows from Berge's Maximum Theorem (Aliprantis and Border, 2006, Theorem 17.31) that each team j principal's best-response correspondence is upper hemicontinuous with compact convex values.

Theorem 2 (Berge's Maximum Theorem). *Let $\phi : X \twoheadrightarrow Y$ be a continuous correspondence between topological spaces X and Y with nonempty compact values. Let $V : \text{Gr}(\phi) \rightarrow \mathbb{R}$ be a continuous function, where $\text{Gr}(\phi) = \{(x, y) \in X \times Y \mid y \in \phi(x)\}$ is the graph of the correspondence.*

Define the value function $m : X \rightarrow \mathbb{R}$ by:

$$m(x) = \sup_{y \in \phi(x)} V(x, y),$$

and the correspondence of maximizers $\beta : X \rightarrow Y$ by:

$$\beta(x) = \{y \in \phi(x) \mid V(x, y) = m(x)\}.$$

Then:

1. The value function $m(x)$ is continuous on X .
2. The argmax correspondence $\beta(x)$ has nonempty compact values.
3. $\beta(x)$ is convex-valued for all $x \in X$ if $V(x, y)$ is quasi-concave in y and $\phi(x)$ has convex values.
4. If Y is Hausdorff, then the argmax correspondence $\beta(x)$ is upper hemicontinuous.

We can now apply Berge's Maximum Theorem to establish properties of each team j principal's best-response correspondence. From Lemmas 3 and 4, the incentive-compatible correspondence $IC_j : M^{N-1} \rightarrow M$ is continuous with nonempty compact values. From Lemma 1, the principal's payoff function V_{π_j} is continuous. Since M is Hausdorff, parts (1), (2), and (4) of Berge's theorem imply that each team j principal's best-response correspondence is upper hemicontinuous with nonempty compact values. Moreover, since $IC_j(\mathbf{m}_{-j})$ is convex-valued (Lemma 5) and V_{π_j} is quasi-concave in m_j (Lemma 1), part (3) implies that the best-response correspondence also has convex values. Thus, the (aggregate) principals' best-response correspondence is upper hemicontinuous with nonempty compact convex values.

Existence of Bayes-Nash Principal's Equilibrium

Lastly, we can apply the Kakutani-Fan-Glicksberg fixed-point theorem (Aliprantis and Border, 2006, Theorem 17.55) to the best-response correspondence to demonstrate the existence of a Bayesian-Nash Principals' equilibrium.

Corollary 1. *Let K be a nonempty compact convex subset of a locally convex Hausdorff space, and let the correspondence $\Psi : K \rightarrow K$ have a closed graph and nonempty convex values. Then, the set of fixed points of Ψ is compact and nonempty.*

Given that each principal's best-response correspondence is upper hemicontinuous with nonempty compact convex values, the aggregate best-response correspondence satisfies the conditions of the Kakutani-Fan-Glicksberg fixed-point theorem. Therefore, there exists a Bayesian-Nash Principals' equilibrium.

References

Anat R. Admati and Motty Perry. Joint projects without commitment. *The Review of Economic Studies*, 58(2):259–276, 1991.

Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer, Berlin, Heidelberg, 3rd edition, 2006.

Miguel Antón, Florian Ederer, Mireia Giné, and Martin Schmalz. Common ownership, competition, and top management incentives. *Journal of Political Economy*, 131(5):1294–1355, 2023.

Pau Balart, Sabine Flamand, and Orestis Troumpounis. Strategic choice of sharing rules in collective contests. *Social Choice and Welfare*, 46:239–262, 2016.

Erik J. Balder. A general approach to lower semicontinuity and lower closure in optimal control theory. *SIAM Journal on Control and Optimization*, 22(4):570–598, 1984.

Erik J. Balder. An extension of Prohorov’s theorem for transition probabilities with applications to infinite-dimensional lower closure problems. *Rendiconti del Circolo Matematico di Palermo*, 34:427–447, 1985.

Erik J. Balder. Generalized equilibrium results for games with incomplete information. *Mathematics of Operations Research*, 13(2):265–276, 1988.

Erik J. Balder. Lectures on Young measure theory and its applications in economics. Technical Report 1052, Rijksuniversiteit Utrecht. Mathematisch Instituut, 1998.

Erik J. Balder. New fundamentals of Young measure convergence. In *Calculus of Variations and Optimal Control*, pages 24–48. Chapman and Hall/CRC, 2021.

J. Banks and J. Duggan. Existence of Nash equilibria on convex sets, 2004. Manuscript, California Institute of Technology and University of Rochester.

Stefano Barbieri and David A. Malueg. Group efforts when performance is determined by the “best shot”. *Economic Theory*, 56:333–373, 2014.

Stefano Barbieri and David A. Malueg. Private-information group contests: Best-shot competition. *Games and Economic Behavior*, 98:219–234, 2016.

Stefano Barbieri and Iryna Topolyan. Private-information group contests with complementarities. *Journal of Public Economic Theory*, 23(5):772–800, 2021.

Stefano Barbieri, Dan Kovenock, David A. Malueg, and Iryna Topolyan. Group contests with private information and the “weakest link”. *Games and Economic Behavior*, 118:382–411, 2019.

Paulo Barelli and Idione Meneghel. A note on the equilibrium existence problem in discontinuous games. *Econometrica*, 81(2):813–824, 2013.

Spencer Bastani, Thomas Giebe, and Oliver Gürler. Simple equilibria in general contests. *Games and Economic Behavior*, 134:264–280, 2022.

Gerald Beer. *Topologies on Closed and Closed Convex Sets*. Number 268 in Mathematics and its Applications. Springer, Dordrecht, 1993.

Gerald Beer and Jesús Rodríguez-López. Topologies associated with kuratowski-painlevé convergence of closed sets. *Journal of Convex Analysis*, 17:805–826, 2010.

Philippe Bich and Rida Laraki. On the existence of approximate equilibria and sharing rule solutions in discontinuous games. *Theoretical Economics*, 12(1):79–108, 2017.

Alessandro Bonatti and Johannes Hörner. Collaborating. *American Economic Review*, 101(2):632–663, 2011.

Kim C. Border. *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge university press, 1985.

T. Renee Bowen, George Georgiadis, and Nicolas S. Lambert. Collective choice in dynamic public good provision. *American Economic Journal: Microeconomics*, 11(1):243–298, 2019.

Philip Brookins and Dmitry Ryvkin. Equilibrium existence in group contests. *Economic Theory Bulletin*, 4(2):265–276, 2016.

Sıdika Tunç Candoğan, C. Gizem Korpeoglu, and Christopher S. Tang. Implications of team submissions in open innovation contests. *IMA Journal of Management Mathematics*, 36(1):1–19, 2025.

Oriol Carbonell-Nicolau and Richard P. McLean. On the existence of Nash equilibrium in Bayesian games. *Mathematics of Operations Research*, 43(1):100–129, 2018.

Guilherme Carmona. An existence result for discontinuous games. *Journal of Economic Theory*, 144(3):1333–1340, 2009.

Henrique Castro-Pires, Hector Chade, and Jeroen Swinkels. Disentangling moral hazard and adverse selection. *American Economic Review*, 114(1):1–37, 2024.

Doruk Cetemen, Ilwoo Hwang, and Ayça Kaya. Uncertainty-driven cooperation. *Theoretical Economics*, 15(3):1023–1058, 2020.

Yeon-Koo Che and Seung-Weon Yoo. Optimal incentives for teams. *American Economic Review*, 91(3):525–541, 2001.

Bo Chen, Yu Chen, and David Rietzke. Simple contracts under observable and hidden actions. *Economic Theory*, 69:1023–1047, 2020.

Benoît S.Y. Crutzen, Sabine Flamand, and Nicolas Sahuguet. A model of a team contest, with an application to incentives under list proportional representation. *Journal of Public Economics*, 182:104109, 2020.

Partha Dasgupta and Eric Maskin. Debreu’s social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 112(52):15769–15770, 2015.

Claude D’Aspremont and Alexis Jacquemin. Cooperative and noncooperative r & d in duopoly with spillovers. *The American Economic Review*, 78:1133–1137, 1988.

Gerard Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.

Gerard Debreu. Existence of competitive equilibrium. In K. Arrow and M. Intriligator, editors, *Handbook of Mathematical Economics*, volume 2. North-Holland, New York, 1982.

Mikhail Drugov and Dmitry Ryvkin. How noise affects effort in tournaments. *Journal of Economic Theory*, 188:105065, 2020.

Kfir Eliaz and Qinggong Wu. A simple model of competition between teams. *Journal of Economic Theory*, 176:372–392, 2018.

Joan Esteban and Debraj Ray. Collective action and the group size paradox. *American Political Science Review*, 95(3), 2001.

Francisco Facchini and Christian Kanzow. Generalized nash equilibrium problems. *Annals of Operations Research*, 175(1):177–211, 2010.

Richard L Fullerton and R. Preston McAfee. Auctioning entry into tournaments. *Journal of Political Economy*, 107(3):573–605, 1999.

George Georgiadis. Projects and team dynamics. *The Review of Economic Studies*, 82(1):187–218, 2015.

Daniel Gottlieb and Humberto Moreira. Simple contracts with adverse selection and moral hazard. *Theoretical Economics*, 17(3):1357–1401, 2022.

Marina Halac, Navin Kartik, and Qingmin Liu. Contests for experimentation. *Journal of Political Economy*, 125(5):1523–1569, 2017.

Marina Halac, Elliot Lipnowski, and Daniel Rappoport. Rank uncertainty in organizations. *American Economic Review*, 111(3):757–786, 2021.

Wei He and Nicholas C. Yannelis. Existence of equilibria in discontinuous Bayesian games. *Journal of Economic Theory*, 162:181–194, 2016.

Bengt Holmstrom. Moral hazard in teams. *The Bell Journal of Economics*, pages 324–340, 1982.

Hideshi Itoh. Incentives to help in multi-agent situations. *Econometrica*, pages 611–636, 1991.

Ohad Kadan, Philip J. Reny, and Jeroen M. Swinkels. Existence of optimal mechanisms in principal-agent problems. *Econometrica*, 85(3):769–823, 2017.

Morton I. Kamien, Eitan Muller, and Israel Zang. Research joint ventures and r&d cartels. *The American Economic Review*, 82:1293–1306, 1992.

Rongzhu Ke and Xinyi Xu. The existence of an optimal deterministic contract in moral hazard problems. *Economic Theory*, 76(2):375–416, 2023.

René Kirkegaard. Contest design with stochastic performance. *American Economic Journal: Microeconomics*, 15(1):201–238, 2023.

Katsuya Kobayashi and Hideo Konishi. Effort complementarity and sharing rules in group contests. *Social Choice and Welfare*, 56(2):205–221, 2021.

Katsuya Kobayashi, Hideo Konishi, and Kaoru Ueda. Prize-allocation rules in generalized team contests. *Economic Theory*, 79:151–179, 2025.

Janos Komlós. A generalization of a problem of steinhaus. *Acta Mathematica Academiae Scientiarum Hungaricae*, 18(1-2):217–229, 1967.

Hideo Konishi, Nicolas Sahuguet, and Benoît S.Y. Crutzen. Allocation rules of indivisible prizes in team contests. *Economic Theory*, 78(1):69–100, 2024.

Kai A. Konrad. *Strategy and Dynamics in Contests*. Oxford University Press, 2009.

Edward P. Lazear and Sherwin Rosen. Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, 89:841–864, 1981.

R. Preston McAfee and John McMillan. Optimal contracts for teams. *International Economic Review*, 32(3):561–577, 1991.

Andrew McLennan, Paulo K. Monteiro, and Rabee Tourky. Games with discontinuous payoffs: a strengthening of Reny’s existence theorem. *Econometrica*, 79(5):1643–1664, 2011.

Paul R. Milgrom and Robert J. Weber. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10(4):619–632, 1985.

Giuseppe Moscarini and Francesco Squintani. Competitive experimentation with private information: The survivor’s curse. *Journal of Economic Theory*, 145(2):639–660, 2010.

Roger B. Myerson. Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics*, 10(1):67–81, 1982.

Shmuel Nitzan. Collective rent dissipation. *The Economic Journal*, 101(409):1522–1534, 1991.

Shmuel Nitzan and Kaoru Ueda. Prize sharing in collective contests. *European Economic Review*, 55(5):678–687, 2011.

Shmuel Nitzan and Kaoru Ueda. Intra-group heterogeneity in collective contests. *Social Choice and Welfare*, 43:219–238, 2014.

Shmuel Nitzan and Kaoru Ueda. Selective incentives and intragroup heterogeneity in collective contests. *Journal of Public Economic Theory*, 20:477–498, 2018.

Efe A. Ok. *Real Analysis with Economic Applications*. Princeton University Press, 2007.

Wojciech Olszewski and Ron Siegel. Equilibrium existence in games with ties. *Theoretical Economics*, 18(2):481–502, 2023.

Saltuk Ozerturk and Huseyin Yildirim. Credit attribution and collaborative work. *Journal of Economic Theory*, 195:105264, 2021.

Pavlo Prokopovych and Nicholas C. Yannelis. On monotone pure-strategy Bayesian-Nash equilibria of a generalized contest. *Games and Economic Behavior*, 140:348–362, 2023.

Michael Raith. Competition, risk, and managerial incentives. *The American Economic Review*, 93:1425–1436, 2003.

Philip J. Reny. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica*, 67(5):1029–1056, 1999.

Philip J. Reny. Nash equilibrium in discontinuous games. *Annual Review of Economics*, 12(1):439–470, 2020.

Dmitry Ryvkin and Mikhail Drugov. The shape of luck and competition in winner-take-all tournaments. *Theoretical Economics*, 15(4):1587–1626, 2020.

Dimitar Simeonov. *Individual and team incentives in contests with heterogeneous agents*. PhD thesis, Boston College, 2020.

Yuan Tao and Huifu Xu. Generalized Bayesian Nash equilibrium with continuous type and action spaces. *arXiv preprint arXiv:2405.19721*, 2024.

Curtis R. Taylor. Digging for golden carrots: An analysis of research tournaments. *The American Economic Review*, pages 872–890, 1995.

Christian Terwiesch and Yi Xu. Innovation contests, open innovation, and multiagent problem solving. *Management Science*, 54(9):1529–1543, 2008.

Áron Tóbiás. Equilibrium non-existence in generalized games. *Games and Economic Behavior*, 135:327–337, 2022.

Francesco Trevisan. Optimal prize allocations in group contests. *Social Choice and Welfare*, 55:431–451, 2020.

Gordon Tullock. Efficient rent seeking. In James M. Buchanan, Robert D. Tollison, and Gordon Tullock, editors, *Towards a Theory of the Rent-Seeking Society*, pages 97–112. Texas A&M University Press, College Station, TX, 1980.

Milan Vojnović. *Contest Theory: Incentive Mechanisms and Ranking Methods*. Cambridge University Press, 2015.

Eyal Winter. Incentives and discrimination. *American Economic Review*, 94(3):764–773, 2004.

Huseyin Yildirim. Getting the ball rolling: Voluntary contributions to a large-scale public project. *Journal of Public Economic Theory*, 8(4):503–528, 2006.

Huseyin Yildirim. Who fares better in teamwork? *The RAND Journal of Economics*, 54(2):299–324, 2023.