

Multi-Objective Optimization Approach to Pure Strategy Equilibria in Games With Imperfect Information

Jianan Jian^{1b}, Graduate Student Member, IEEE, Bo Zeng^{1b}, Member, IEEE,
and Zhi-Hong Mao^{1b}, Senior Member, IEEE

Abstract—Without the expected utility hypothesis, games with imperfect information can be understood as coupled multi-objective optimizations. Within this framework, equilibria of non-dominated responses are defined and their existence is studied. One class of these equilibria is formulated as Nash equilibria of normal-form games. Another class of these equilibria is formulated with payoff constraints interpreted as robustness of the games. A concrete example is presented and solution methodologies are discussed.

Index Terms—Game theory, optimization, uncertain systems.

I. INTRODUCTION

IN CONTROL theory and game theory, the information structure of a problem has a decisive effect on the properties of the solution [1], [2]. In practice, agents often receive discrete information about the state of the system, such as classification results or answers to yes-no questions. Discretization of information also facilitates computation in machine learning [3]. This kind of information is typically modeled with information sets in game theory. However, the studies of information sets are usually limited in finite or countable state spaces [4], whereas uncertain information in more general state spaces such as Euclidean spaces are typically modeled with random variables using probability theory. The probabilistic approach makes strong assumptions about how information is perceived and processed by agents, particularly the expected utility hypothesis. This letter studies information sets in general topological state spaces by proposing a solution concept called non-dominated equilibria. This novel framework also incorporates the idea of robustness in games.

Manuscript received September 14, 2020; revised November 19, 2020; accepted December 5, 2020. Date of publication December 22, 2020; date of current version April 26, 2021. Recommended by Senior Editor V. Ugrinovskii. (Corresponding author: Jianan Jian.)

Jianan Jian and Zhi-Hong Mao are with the Department of Electrical and Computer Engineering, University of Pittsburgh, Pittsburgh, PA 15260 USA (e-mail: jij52@pitt.edu; zhm4@pitt.edu).

Bo Zeng is with the Department of Industrial Engineering, University of Pittsburgh, Pittsburgh, PA 15260 USA (e-mail: bzen@pitt.edu).

Digital Object Identifier 10.1109/LCSYS.2020.3046674

In this letter, we focus on two-player static games, but the ideas and results can be easily extended to more general settings. In Section II, we propose a novel model of games with imperfect information using multi-objective optimization, where the expected utility hypothesis is replaced by non-dominance of strategies and the resulting equilibria. In Section III, we relate such equilibria with the notion of robustness by adding constraints to the problem. Section IV gives a concrete example. The general direction for solving the equilibria is given in Section V before we conclude.

II. EQUILIBRIA UNDER IMPERFECT INFORMATION

The notion of state in game theory groups all the information, other than the players' decisions, necessary to evaluate the players' payoffs [5]. In games with perfect information, the state is common knowledge among the players. The best responses of all the players in all possible states constitute a Nash equilibrium (NE) [4].

Consider a two-player static game. Denote the state space by Ω . For both players $X \in \{A, B\}$, denote their decision set by U_X , their payoff function by $J_X : \Omega \times U_A \times U_B \rightarrow \mathbb{R}$.

Definition 1: A pure strategy profile $u_A^* : \Omega \rightarrow U_A, u_B^* : \Omega \rightarrow U_B$ is an NE if $\forall u_A \in U_A, u_B \in U_B, \omega \in \Omega$, we have

$$J_A(\omega, u_A^*(\omega), u_B^*(\omega)) \geq J_A(\omega, u_A, u_B^*(\omega)) \quad (1a)$$

$$J_B(\omega, u_A^*(\omega), u_B^*(\omega)) \geq J_B(\omega, u_A^*(\omega), u_B) \quad (1b)$$

In a game with imperfect information, each agent does not have the exact value of the state, but knows that the state belongs to a certain set called the information set. This knowledge is private to each player, but all the possible sets the state may belong to, which form a partition of the state space for each agent called the information partition, are a common knowledge among all the agents.

Denote the information partition of each player by P_X . A pure strategy should be a function $u_X^* : \Omega \rightarrow U_X$ constant on each information set, or equivalently a function on the information partition $u_X^* : P_X \rightarrow U_X$ [4]. For convenience in this letter, we will use u_X^* to denote either representation depending on the context. With this notation,




Fig. 1. Consider the situation where the information partitions of the state space are illustrated above $P_A = \{E_{A1}, E_{A2}\}$ and $P_B = \{E_{B1}, E_{B2}\}$. The state of the game, indicated by the black dot, will impact the players' payoffs, but neither of them knows the exact location of the dot. Player A only knows that the dot is somewhere in the region E_{A1} , and needs to make a decision solely based on this piece of information about the state. Nevertheless, the rules of the game are clear to both players, so Player A also knows that Player B will make his decision solely based on whether the dot is in E_{B1} or E_{B2} .

we propose a direct application of NE to games without perfect information.

Proposition 1: A pure strategy profile (u_A^*, u_B^*) is a NE if $\forall u_A \in U_A, u_B \in U_B, E_A \in P_A, E_B \in P_B, \omega \in E_A \cap E_B$, we have

$$J_A(\omega, u_A^*(E_A), u_B^*(E_B)) \geq J_A(\omega, u_A, u_B^*(E_B)) \quad (2a)$$

$$J_B(\omega, u_A^*(E_A), u_B^*(E_B)) \geq J_B(\omega, u_A^*(E_A), u_B). \quad (2b)$$

Proof: Rewrite (1) for all ω on each subset $E_A \cap E_B$. ■

In a game with perfect information, E_X are singletons, and (2) can be decoupled into a NE for each state, and the original notion of NE is recovered. In a game with imperfect information, each decision $u_X^*(E_X)$ is a solution to a multi-objective optimization. Indeed, it needs to satisfy (2- X) simultaneously for all qualified $\omega \in E_A \cap E_B$, which may be multiple states in this case. Consequently, such solution does not necessarily exist in general. This is the reason why NE are not generally defined in games with imperfect information.

The standard approach to this problem is to assume that each player is only interested in the average payoff. By replacing the payoff functions by the expected payoffs on each information set according to the beliefs each player holds about the state, Bayesian equilibria [4], [5] are defined as the resulting strategies of best response.

Definition 2: A strategy profile (u_A^*, u_B^*) is a Bayesian equilibrium if $\forall u_A \in U_A, u_B \in U_B, \omega \in \Omega$, we have

$$\mathbb{E}_A[J_A(\omega, u_A^*(\omega), u_B^*(\omega))] \geq \mathbb{E}_A[J_A(\omega, u_A, u_B^*(\omega))]$$

$$\mathbb{E}_B[J_B(\omega, u_A^*(\omega), u_B^*(\omega))] \geq \mathbb{E}_B[J_B(\omega, u_A^*(\omega), u_B)].$$

The alternative approach we propose is to consider any non-dominated solution to the multi-objective optimization (2- X) as a candidate strategy, and any pair of non-dominated strategies as a candidate equilibrium of the game. Each player can select a strategy to execute from the candidate strategies based on some criterion of their choice. With this approach, it is not required to define a probability structure on the state space, and a larger set of strategies can be justified for the players. As we demonstrate in the rest of Sections II and III, the set of non-dominated strategies includes Bayesian equilibria as well as robust equilibria with payoff constraints.

Let u_X^* be a non-dominated strategy and v_X^* a different strategy, then v_X^* cannot be both 1) no worse than u_X^* in all states and 2) strictly better than u_X^* in some state. So if v_X^* is no worse than u_X^* in all states, then it can only be equivalent to u_X^* in all states.

Definition 3: A strategy $u_A^* : P_A \rightarrow U_A$ is a non-dominated response to a strategy $u_B^* : P_B \rightarrow U_B$ if $\forall E_A \in P_A, v_A^* : P_A \rightarrow U_A$ we have that $\forall E_B \in P_B, \omega \in E_A \cap E_B$

$$J_A(\omega, u_A^*(E_A), u_B^*(E_B)) \leq J_A(\omega, v_A^*(E_A), u_B^*(E_B))$$

implies $\forall E_B \in P_B, \forall \omega \in E_A \cap E_B$

$$J_A(\omega, u_A^*(E_A), u_B^*(E_B)) = J_A(\omega, v_A^*(E_A), u_B^*(E_B)).$$

Definition 4: A pure strategy profile (u_A^*, u_B^*) is an equilibrium of non-dominated responses if u_A^* and u_B^* are non-dominated responses one to another.

It is worthwhile to compare the non-dominance defined here with other similar concepts in game theory to avoid any potential confusion. 1) The traditional notion of strategic dominance in game theory is about dominance with respect to all opponent's strategies at a fixed state. Definition 3 is about dominance with respect to all states given a fixed opponent's strategy. 2) Pareto efficiency is another solution concept in game theory based on non-dominated solutions of multi-objective optimizations. It seeks to optimize simultaneously the payoffs of all the players and does not constitute an equilibrium in general. The solution concept in Definition 4 is an equilibrium where each player seeks to optimize their payoffs for all states at the same time.

A common method to find the non-dominated solutions is the weighted sum method [6]. This method can be generalized to games with an infinite state space using the Hahn–Banach separation theorem.

Proposition 2: Suppose the sets of reachable marginal payoff functions $\{(\omega, u_B) \mapsto J_A(\omega, u_A, u_B) | u_A \in U_A\}$ and $\{(\omega, u_A) \mapsto J_B(\omega, u_A, u_B) | u_B \in U_B\}$ are convex. If (u_A^*, u_B^*) is an equilibrium of non-dominated responses, then for all $E_A \in P_A, E_B \in P_B$:

$$u_A^*(E_A) \in \arg \max_{u_A \in U_A} F_A[J_A(\cdot, u_A, u_B^*(\cdot))] \quad (3a)$$

$$u_B^*(E_B) \in \arg \max_{u_B \in U_B} F_B[J_B(\cdot, u_A^*(\cdot), u_B)] \quad (3b)$$

where F_X is a non-trivial positive continuous linear functional of functions on E_X .

Proof: For given u_B^* , define the marginal payoff function $f_{u_A} : E_A \rightarrow \mathbb{R}$ such that $f_{u_A}(\omega) = J_A(\omega, u_A, u_B^*(\omega))$ for all $u_A \in U_A$. Let C be the closed convex cone of positive functions on E_A , then the set of marginal payoff functions dominating f_{u_A} is given by the Minkowski sum $C' = \{f_{u_A}\} + C$. The set of reachable marginal payoff functions $S = \{f_{u_A} | u_A \in U_A\}$ is convex by the assumption. If $u_A = u_A^*(E_A)$ is optimal, then $C' \cap S = \{f_{u_A}\}$. This shows that f_{u_A} is on the relative boundaries of S and C' , so there is a supporting functional F_A separating S and C' at f_{u_A} , i.e., for all $u_A \in U_A$ and for all positive function $g \in C$:

$$\begin{aligned} F_A[J_A(\cdot, u_A, u_B^*(\cdot))] &\leq F_A[J_A(\cdot, u_A^*(E_A), u_B^*(\cdot))] \\ &\leq F_A[J_A(\cdot, u_A^*(E_A), u_B^*(\cdot)) + g] \end{aligned} \quad (4)$$

i.e., $u_A^*(E_A)$ maximizes $F_A[J_A(\cdot, u_A, u_B^*(\cdot))]$ for a positive functional F_A . The same arguments apply to B . ■

In this letter, the notations defined in the proofs, such as S , C , C' , and f_{u_A} , will be used subsequently in the rest of this letter. The converse of Proposition 2 holds under the following conditions.

Proposition 3: If a pure strategy profile (u_A^*, u_B^*) satisfies (3) where F_X is a strictly positive continuous linear functional of functions on each E_X , then (u_A^*, u_B^*) is an equilibrium of non-dominated responses.

Proof: If $u_A^*(E_A)$ maximizes $F_A[J_A(\cdot, u_A, u_B^*(\cdot))]$, then F_A is a supporting functional separating S and C' . This shows that $S \cap C'$ is on the relative boundaries. When F_A is strictly positive, the second inequality in (4) is strict if $g > 0$, so $S \cap C'$ is a singleton, i.e., $u_A^*(E_A)$ is a non-dominated response. The same arguments apply to B . ■

If F_X is not strictly positive, then it cannot reflect the dominance if the margin is in its kernel. However, the existence of strictly positive functionals requires some regularity of the function space they can be defined on.

Proposition 4: Suppose Ω is a locally compact, second countable Hausdorff space. Suppose both P_X are finite partitions and both $J_X(\cdot, u_A, u_B)$ are continuous and vanish at infinity for all u_A, u_B . Then on each E_X , there is a function space H such that 1) all reachable marginal payoff functions (i.e., f_{u_A} or their counterparts for B) belong to H , and 2) strictly positive functionals F_X exist on H .

Proof: The space of continuous functions vanishing at infinity on the locally compact Hausdorff space Ω is a Banach space $C_0(\Omega)$. The Banach space $C_0(\Omega)$ is separable when Ω is second countable.

Note that all marginal payoff functions can be written as $f_{u_A} = \sum_{E_B \in P_B} f_{u_A, E_B} 1_{E_A \cap E_B}$ where $f_{u_A, E_B} = J_A(\cdot, u_A, u_B^*(E_B)) \in C_0(\Omega)$, so $f_{u_A} \in H$ where $H = \{\sum_{E_B \in P_B} h_{E_B} 1_{E_A \cap E_B} | h_{E_B} \in C_0(\Omega)\}$. Let $G = \bigoplus_{E_B \in P_B} C_0(\Omega)$. Finite direct sum of separable Banach spaces is again a separable Banach space. There is a natural surjection $\sigma : G \rightarrow H$ and we have $H \simeq G/N$ where $N = \{E_B \mapsto h_{E_B} \forall E_B : h_{E_B} = 0 \text{ on } E_A \cap E_B\}$. Clearly, N is a closed subspace of G , so H is also a separable Banach space, on which strictly positive functionals exist [7, Th. 2.6(a)]. It concludes the proof for A as H satisfies both requirements. The same arguments apply for B . ■

The finiteness of P_X is an essential assumption. It allows us to decompose the problem into finitely many subproblems and make use of nice properties of finite operations.

Solutions to (3) can be interpreted as NE among $\sum_X |P_X|$ (finitely many) players indexed by E_X . Their decision set is either U_A or U_B , and their payoff function is either $F_A[J_A(\cdot, u_A, u_B^*(\cdot))]$ or $F_B[J_A(\cdot, u_A^*(\cdot), u_B)]$. The state in the original game is integrated into the payoff, so there is no notion of state in the new game. It is simply a normal-form game. Each of the new players E_X represents the original player X in a reality where $\omega \in E_X$. Any pair (E_A, E_B) of the new players have direct interaction in their payoffs through their decisions whenever their intersection $E_A \cap E_B$ is non-empty. The existence of a NE to (3) can be given with the existing result.

Proposition 5: Suppose both decision sets U_X are non-empty compact convex subsets of locally convex Hausdorff spaces, both families of functions $\{(u_A, u_B) \mapsto f_X(\omega, u_A, u_B) | \omega \in \Omega\}$ are equicontinuous, and both functions $u_A \mapsto J_A(\cdot, u_A, \cdot)$ and $u_B \mapsto J_B(\cdot, \cdot, u_B)$ are *-quasiconcave [8]. Then (3) has a solution for any positive continuous linear functional F_X of functions on each E_X .

Proof: Given the assumptions on U_X , the Debreu-Glicksberg-Fan theorem [9, Th. 1] states that a pure strategy NE exists when $(u_A, u_B^*) \mapsto F_A[J_A(\cdot, u_A, u_B^*(\cdot))]$ is a continuous function quasiconcave in u_A , together with the same condition for B .

Let $(v_A, v_B^*) \in U_A \times (U_B)^{P_B}$ and $\varepsilon > 0$. The equicontinuity condition implies that $\forall E_B \in P_B$ there is a neighborhood of the origin $\delta(E_B) \subseteq U_A \times U_B$ such that once (u_A, u_B^*) satisfies $(u_A, u_B^*(E_B)) - (v_A, v_B^*(E_B)) \in \delta(E_B)$, it also satisfies $\forall \omega \in \Omega$:

$$\|f(\omega, u_A, u_B^*(E_B)) - f(\omega, v_A, v_B^*(E_B))\| < \varepsilon/2,$$

and therefore

$$\begin{aligned} & \|f(\cdot, u_A, u_B^*(\cdot)) - f(\cdot, v_A, v_B^*(\cdot))\| \\ &= \max_{E_B \in P_B} \sup_{\omega \in E_A \cap E_B} \|f(\omega, u_A, u_B^*(E_B)) \\ &\quad - f(\omega, v_A, v_B^*(E_B))\| \leq \varepsilon/2 < \varepsilon. \end{aligned}$$

The projection $u_B^* \mapsto u_B^*(E_B)$ is continuous, so there is a neighborhood of the origin $\delta'(E_B) \subseteq U_A \times (U_B)^{P_B}$ such that whenever (u_A, u_B^*) satisfies $(u_A, u_B^*) - (v_A, v_B^*) \in \delta'(E_B)$, it also satisfies $(u_A, u_B^*(E_B)) - (v_A, v_B^*(E_B)) \in \delta(E_B)$. To conclude, as any finite intersection of neighborhoods is neighborhood, whenever

$$(u_A, u_B^*) - (v_A, v_B^*) \in \bigcap_{E_B \in P_B} \delta'(E_B).$$

We have $\|f(\cdot, u_A, u_B^*(\cdot)) - f(\cdot, v_A, v_B^*(\cdot))\| < \varepsilon$. This shows the continuity of the function $(u_A, u_B^*) \mapsto J_A(\cdot, u_A, u_B^*(\cdot))$. Composed with a continuous functional, $(u_A, u_B^*) \mapsto F_A[J_A(\cdot, u_A, u_B^*(\cdot))]$ is again continuous.

Let $\Psi : C_0(\Omega \times U_B) \rightarrow H$ be the continuous linear operator which maps the function $(\omega, u_B) \mapsto J_A(\omega, u_A, u_B)$ to the function $\omega \mapsto J_A(\omega, u_A, u_B^*(\omega))$. Then we have $F_A[J_A(\cdot, u_A, u_B^*(\cdot))] = (F_A \circ \Psi)[J_A(\cdot, u_A, \cdot)]$, which is quasiconcave in u_A by the *-quasiconcavity condition. ■

If J_X is uniformly continuous, then the equicontinuity condition is automatically satisfied. If J_X is concave in u_X , then the *-quasiconcavity condition is automatically satisfied.

Propositions 4 and 5 together give sufficient conditions on the game $(\Omega, P_X, U_X, \text{ and } J_X)$ under which equilibria of non-dominated responses exist. These conditions will be assumed for the rest of this letter.

The Riesz representation theorem states that a functional on $C_0(\Omega)$ is given by the integration on Ω with respect to a Borel measure. If this measure is σ -finite, it can be interpreted as a probability distribution of ω on each E_X . It generalizes the Bayesian equilibrium in which $P_A = P_B = \{\Omega\}$.

III. ROBUSTNESS

Proposition 2 suggests that some equilibria of non-dominated responses might be overlooked by the weighted

sum method if the convexity is not respected. A common method to overcome this limitation is the ε -constraint method [10]. In this method, player X sets a threshold ε_X of payoff to satisfy for each state, and optimizes the payoff in one particular state of priority ω_X^0 . Evaluating the payoff function at one particular state is a special case of linear functionals of the payoff function, so we can combine the ε -constraint method and the weighted sum method.

Definition 5: A pure strategy profile (u_A^*, u_B^*) is an equilibrium of ε_X -constrained F_X -optimal strategies if $\forall E_A \in P_A, \forall E_B \in P_B$:

$$u_A^*(E_A) \in \arg \max_{u_A \in U_A} F_A[J_A(\cdot, u_A, u_B^*(\cdot))] \quad (5a)$$

$$\text{s.t. } J_A(\cdot, u_A, u_B^*(\cdot)) \geq \varepsilon_A \text{ on } E_A$$

$$u_B^*(E_B) \in \arg \max_{u_B \in U_B} F_B[J_B(\cdot, u_A^*(\cdot), u_B)] \quad (5b)$$

$$\text{s.t. } J_B(\cdot, u_A^*(\cdot), u_B) \geq \varepsilon_B \text{ on } E_B$$

where ε_X is a real function on E_X and F_X is a non-trivial positive continuous linear functional of functions on E_X .

The ε_X -constraints reflect the robustness of the payoffs at the equilibrium against the uncertainty in the state. So a solution to (5) will also be called a robust equilibrium.

The ε_X -constraints restricts the set of reachable marginal functions from S to $S' = S \cap (\{\varepsilon_X\} + C)$. Even when S is not convex, which is a necessary condition for Proposition 2, set S' can be convex by carefully choosing ε_X , so each equilibrium of non-dominated responses satisfies (5) for those levels of thresholds.

Proposition 6: An equilibrium of non-dominated responses is a robust equilibrium for certain ε_X .

Proof: Given an equilibrium of non-dominated responses (u_A^*, u_B^*) , let $\varepsilon_X(\omega) = J_X(\omega, u_A^*(\omega), u_B^*(\omega))$. By non-dominance, all feasible solutions yield the same payoff, so they are also optimal. This shows u_X^* solves (5-X). ■

The converse is given as follows.

Proposition 7: A robust equilibrium for any strictly positive functionals F_X is an equilibrium of non-dominated responses.

Proof: Fix u_B^* and suppose (u_A^*, u_B^*) satisfies (5). Comparing to any strategy v_A^* that satisfies the constraint, u_A^* is a non-dominated response by the hyperplane separation theorem as in Proposition 3. Comparing to any strategy v_A^* that does not satisfy the constraint, u_A^* is strictly dominate v_A^* on the states ω on which v_A^* violates the constraint. The same arguments apply for B . ■

Define the notion of feasibility due to the ε_X -constraints.

Definition 6: A pure strategy profile (u_A^*, u_B^*) is called feasible if $J_X(\cdot, u_A^*(\cdot), u_B^*(\cdot)) \geq \varepsilon_X$ on all E_X . Denote the set of feasible strategy profiles by D . Problem (5) is called feasible if $D \neq \emptyset$.

The feasibility largely depends on the thresholds ε_X , which are the levels of robustness required by the players for their payoffs. Assume that ε_X belongs to the same class of functions as J_X , so we can assume $\varepsilon_X = 0$ by absorbing ε_X into J_X without loss of generality. In addition to the conditions assumed in Proposition 5, the following conditions guarantee the existence of a robust equilibrium.

Proposition 8: Suppose D is a non-empty polytope. Suppose the function $(u_A, u_B) \mapsto J_X(\omega, u_A, u_B)$ is sup-compact [11] and quasiconcave for all $\omega \in \Omega$. Then (5) has an equilibrium.

Proof: For E_A and u_B^* , define the sets of feasible decisions

$$U'_A(u_B^*) = \{u_A \in U_A | J_A(\cdot, u_A, u_B^*(\cdot)) \geq 0 \text{ on } E_A\}$$

and the sets of optimal decisions

$$M_A(u_B^*) = \arg \max_{u_A \in U'_A(u_B^*)} F_A[J_A(\cdot, u_A, u_B^*(\cdot))].$$

Same notations are defined for B .

The Kakutani fixed-point theorem states that an equilibrium exists if 1) D is non-empty, compact, and convex, 2) for all feasible (u_A^*, u_B^*) : the maximizers $M_A(u_B^*)$ and $M_B(u_A^*)$ are non-empty, compact, and convex, and 3) the correspondence $(u_A^*, u_B^*) \mapsto M_A(u_B^*) \times M_B(u_A^*)$ is upper hemicontinuous on D .

The feasibility of (5) implies the non-emptiness of D , $U'_A(u_B^*)$, and $U'_B(u_A^*)$. Suppose $(u_A, u_B) \mapsto J_X(\omega, u_A, u_B)$ is quasiconcave and sup-compact for all $\omega \in \Omega$, then the set D , being the intersection of their contour sets, is convex and compact. For the same reason, $U'_A(u_B^*)$ and $U'_B(u_A^*)$ are convex and compact.

Given the properties of $U'_A(u_B^*)$ as well as the continuity and the quasiconcavity of $F_A[J_A(\cdot, u_A, u_B^*(\cdot))]$ from Proposition 5, the maximum theorem states that M_A is upper hemicontinuous and $M_A(u_B^*)$ is non-empty, compact, and convex if U'_A is continuous.

The upper hemicontinuity of U'_A results from the closed graph theorem [12, Th. 17.11] given the continuity and the boundedness of the function $(u_A, u_B^*) \mapsto J_A(\omega, u_A, u_B^*(\omega))$ for all ω . The lower hemicontinuity of U'_A results from [13, Th. 3] given that D is a polytope and $\{(u_A, u_B^*) | J_A(\cdot, u_A, u_B^*(\cdot)) \geq 0 \text{ on } E_A\}$ is convex. The same arguments apply for B . Finally, the product of compact-valued hemicontinuous correspondences $(u_A^*, u_B^*) \mapsto M_A(u_B^*) \times M_B(u_A^*)$ is again compact-valued upper hemicontinuous [12, Th. 17.28]. ■

Being a polytope is a strong requirement for the feasible set D . This means that in general J_X needs to be piecewise linear on a polytope mesh of $\Omega \times U_A \times U_B$.

Proposition 9: Suppose J_X satisfies the properties in Proposition 8. Suppose there are finite partitions \mathcal{P} of $E_A \cap E_B$ and a finite partition \mathcal{Q} of $U_A \times U_B$ such that the closures of $\mathcal{E} \in \mathcal{P}$ and $\mathcal{F} \in \mathcal{Q}$ are all convex polytopes and J_X is affine on each $\mathcal{E} \times \mathcal{F}$, then D is a polytope.

Proof: The closure of $\mathcal{C}_{\mathcal{E}, \mathcal{F}} := \{(u_A, u_B) \in \mathcal{F} | \forall \omega \in \mathcal{E}: J_A(\omega, u_A, u_B) \geq 0\}$ is a convex polytope as it is a contour set of an affine function restricted on a convex polytope. The set $\mathcal{C} := \{(u_A, u_B) \in U_A \times U_B | \forall \omega \in E_A \cap E_B: J_A(\omega, u_A, u_B) \geq 0\} = \bigcup_{\mathcal{F}} \bigcap_{\mathcal{E}} \mathcal{C}_{\mathcal{E}, \mathcal{F}}$ is again a convex polytope as it is convex and compact from Proposition 8. Finally, the set $D = \bigcap_{E_A, E_B} \{(u_A^*, u_B^*) \in \times U_A^{|P_A|} \times U_B^{|P_B|} | (u_A^*(E_A), u_B^*(E_B)) \in \mathcal{C}\}$ is a convex polytope as the existence of a polytope partition \mathcal{Q} requires that U_A, U_B are polytopes. ■

If J_X is affine, then the graph of J_X is a polyhedron. If U_X and Ω are bounded then J_X is a polytope. The graph Intersection the half space is again a polytope. This polytope

projecting onto U_X is again a polytope. This shows that piecewise linear approximation is not only powerful for the purpose of digital computation, but also crucial to guaranteeing the theoretical existence of a solution.

A related notion is the robust games introduced in [14]. We adapt this concept to our context.

Proposition 10: Suppose both decision sets U_X are metric spaces and both payoff functions J_X are quasiconcave in u_X . Then there is a solution to

$$u_A^*(E_A) \in \arg \max_{u_A \in U_A} \inf_{\omega \in E_A} J_A(\omega, u_A, u_B^*(\omega)) \quad (6a)$$

$$u_B^*(E_B) \in \arg \max_{u_B \in U_B} \inf_{\omega \in E_B} J_B(\omega, u_A^*(\omega), u_B). \quad (6b)$$

Proof: For each $E_A \in P_A$, define the following functions $U_A \times (U_B)^{P_B} \rightarrow \mathbb{R}$.

$$\psi_\omega : (u_A, u_B^*) \mapsto J_A(\omega, u_A, u_B^*(\omega))$$

$$\psi_{E_B} = \inf\{\psi_\omega | \omega \in E_A \cap E_B\}$$

$$\psi = \inf\{\psi_\omega | \omega \in E_A\} = \min\{\psi_{E_B} | E_B \in P_B\}$$

Given the assumptions on U_X , the Debreu-Glicksberg-Fan theorem states that a pure strategy NE exists if ψ is a continuous function quasiconcave in u_A , together with the same condition for B .

As the payoff function $J_A \in C_0(\Omega \times U_A \times U_B)$ is bounded, the family of functions $\{\psi_\omega | \omega \in E_A \cap E_B\}$ is uniformly bounded. It is also equicontinuous by the conditions in Proposition 5. As a result, the infimum ψ_{E_B} is a bounded and continuous function [15, Problem 4.2.10]. The finite family of bounded continuous functions $\{\psi_{E_B} | E_B \in P_B\}$ is also equicontinuous and uniformly bounded, so the minimum ψ is again continuous.

Similarly, the function ψ_{E_B} , being an infimum of a family of quasiconcave functions in u_A , is quasiconcave [16]. Then ψ , the minimum of ψ_{E_B} , is again quasiconcave. ■

Definition 7: Problem (6) is called a robust game.

In Definition 5, each player requires the positivity of the worst-case payoff. In Definition 7, each player requires the optimality of the worst-case payoff. However these two exigencies have different significance. Positivity constraints have the purpose of benefiting both players. A feasible strategy profile leaves both players satisfied at all states. This can be considered as a weaker version of the Pareto efficiency. While Pareto efficiency is generally not obtained at a NE without forming coalitions, the positivity constraints are requisites to the equilibrium (5), so that the players both understand that a deal can be reached only if they keep each other happy. This illustrates that non-dominated equilibria can achieve Pareto efficiency to some extent through payoff constraints. The formulation (5) makes less sense given perfect information. Because without uncertainty, maximizing payoff and seeking a minimum payoff guarantee are somehow overlapping. Nevertheless, the constraints can serve to eliminate inefficient equilibria when multiple equilibria exist.

On the other hand, an equilibrium to the robust game (6) is simply a NE. It only ensures that the players are satisfied at the worst state. The equilibrium may suffer from Prisoner's dilemma and be inefficient in general. Suppose (6) admits an

equilibrium who has non-negative values on all E_X . Then then this strategy profile is feasible to (5) and it is also an equilibrium to (5) with F_X being the evaluation map at the worst state on E_X . However, this functional is not strictly positive, so this strategy profile may not be an equilibrium of non-dominated responses.

IV. EXAMPLE

Alice and Bob run for mayor in a municipal with a population of n . The state of the race can be characterized by a vector $\omega \in [-1, 1]^n$ whose components represent the political leaning of each voter ranging from entirely in favor of Alice ($\omega_i = 1$) to entirely in favor of Bob ($\omega_i = -1$). Their campaigns seek to influence voters' opinions by allocating their limited resources to each voter. This constitutes the following zero-sum game

$$\text{Alice: } \max_{u_A \geq 0} \left\{ \sum_{i=1}^n \sigma(\omega_i + u_{Ai} - u_{Bi}) : \sum_{i=1}^n u_{Ai} \leq R_A \right\}$$

$$\text{Bob: } \min_{u_B \geq 0} \left\{ \sum_{i=1}^n \sigma(\omega_i + u_{Ai} - u_{Bi}) : \sum_{i=1}^n u_{Bi} \leq R_B \right\}$$

where σ is a sigmoid-shaped function. This generalizes the Colonel Blotto game [17], [18], a classical model of electoral resource allocation which does not include such state.

Polls are conducted by the campaigns to assess ω . Suppose each poll aims at answering a binary question about some statistic of the favorability rating, the results of which may be private to the candidate but the polling efforts are commonly known. Then all the polls together form an information partition for each candidate. On each information set, the functional F_X represents each candidate's own strategy or preference regarding their constituents. This offers a more practical interpretation of the payoff than the expectation with respect to a Bayesian belief about ω does. Furthermore, it may be possible to translate a candidate's preference from one election to another but their belief about the states in different elections may be unrelated.

For the purpose of illustration, let $n = 2$. Alice conducted two polls: whether voter 1 favors her greater than -25% , and whether voter 2 favors her twice as much as voter 1 does. Bob conducted one poll: whether the voters favor Alice more than 25% on average. To decide among non-dominated strategies, Alice would like to conserve her lead in voter 2 or flip voter 1 when possible, and Bob would like to flip either voter when possible. This case is illustrated by Fig. 2. Note that the game is no longer zero-summed because F_A and F_B are different.

Non-dominated equilibria with payoff constraints may provide more appealing scenarios in this example. These equilibria are achieved given the premise that both candidates can reach some level satisfaction in their preferred goals through their individual strategies which are nonetheless non-dominated. This illustrates a potential path to understand and develop cooperation in competitive games using our proposed framework.

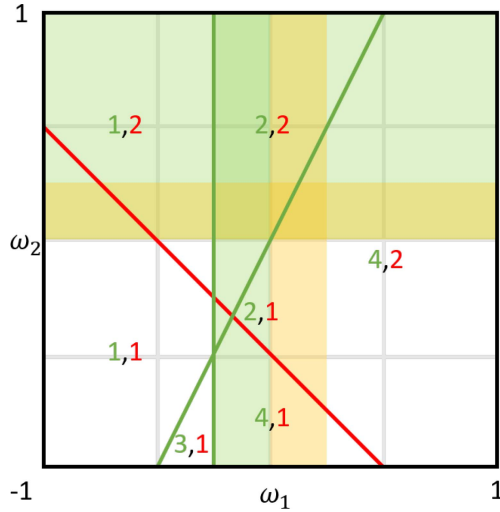


Fig. 2. State space of the game. The solid lines (green for Alice, red for Bob) represent the polls which partition the state space into information sets labeled 1-4 by Alice or 1-2 by Bob. The color blocks represent where F_A or F_B takes greater values. For instance, to conserve her lead in voter 2, Alice puts greater weight where $\omega_2 \geq 0$; to flip voter 1, she puts greater weight where ω_1 is negative by a small margin. These are the states Alice wants to have a dominance at.

V. SOLUTIONS

Solving games is a hard problem in general. In this section, we provide some general directions to finding an equilibrium.

The equilibria (3) and (5) can be found by solving the KKT conditions. Suppose Ω is a Euclidean space and the measures associated with F_X are absolutely continuous with respect to the Lebesgue measure. Then $F_A[J_A(\cdot, u_A, u_B^*(\cdot))]$ can be written as

$$\sum_{E_B \in P_B} \int_{E_A \cap E_B} w_A(\omega) J_A(\omega, u_A, u_B^*(E_B)) d\omega$$

where w_A is a positive function on Ω . To evaluate the integral, we discretize the state space, yielding a sum

$$\sum_{E_B \in P_B} \sum_{\omega_k \in L(E_A, E_B)} w_A(\omega_k) J_A(\omega_k, u_A, u_B^*(E_B))$$

where $L(E_A, E_B)$ is a lattice on $E_A \cap E_B$.

Suppose J_X are continuously differentiable in u_X , and the convex compact decision sets can be specified as $U_X = \{u_X \in \mathbb{R}^{m_X} | g_X(u_X) \geq 0\}$ where g_X are continuously differentiable functions. Then we can solve the KKT conditions and verify the second-order conditions. If J_X and g_X are non-differentiable, we can use subgradients instead in the KKT conditions. If J_X satisfies Proposition 9, the piecewise linear nature of the objective function results in a mixed-integer linear program for each player.

VI. CONCLUSION

This letter applies and generalizes multi-objective games [19], [20], [21] to provide an alternative solution concept for games with imperfect information in general topological state spaces. Two class of equilibria are studied using the weighted sum method and the ϵ -constraint method, given the parameters F_X and ϵ_X , intuitively interpreted as preferences and robustness. The theoretical results of this letter will guide future work on modeling and understanding human behaviors beyond the expected utility theory.

REFERENCES

- [1] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM J. Control*, vol. 6, no. 1, pp. 131–147, 1968.
- [2] G. A. Akerlof, "The market for 'lemons': Quality uncertainty and the market mechanism," *Quart. J. Econ.*, vol. 84, no. 3, pp. 488–500, 1970.
- [3] I. S. Lee and H. Y. Lau, "Adaptive state space partitioning for reinforcement learning," *Eng. Appl. Artif. Intell.*, vol. 17, no. 6, pp. 577–588, 2004.
- [4] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA, USA: MIT Press, 1994.
- [5] J. C. Harsanyi, "Games with incomplete information played by 'Bayesian' players, I–III: Part I. The basic model," *Manag. Sci.*, vol. 14, no. 3, pp. 159–182, 1967.
- [6] S. Wang, "An existence theorem of a Pareto equilibrium," *Appl. Math. Lett.*, vol. 4, no. 3, pp. 61–63, 1991.
- [7] J. Glück and M. R. Weber, "Almost interior points in ordered Banach spaces and the long-term behaviour of strongly positive operator semigroups," 2019. [Online]. Available: arXiv:1901.03306.
- [8] V. Jeyakumar, W. Oettli, and M. Natividad, "A solvability theorem for a class of quasiconvex mappings with applications to optimization," *J. Math. Anal. Appl.*, vol. 179, no. 2, pp. 537–546, 1993.
- [9] P. Dasgupta and E. Maskin, "The existence of equilibrium in discontinuous economic games, I: Theory," *Rev. Econ. Stud.*, vol. 53, no. 1, pp. 1–26, 1986.
- [10] K. Miettinen, *Nonlinear Multiobjective Optimization*, vol. 12. New York, NY, USA: Springer, 2012.
- [11] F. Luque-Vásquez and O. Hernández-Lerma, "A counterexample on the semicontinuity of minima," *Proc. Amer. Math. Soc.*, vol. 123, no. 10, pp. 3175–3176, 1995.
- [12] C. Aliprantis, K. Border, and K. Border, *Infinite Dimensional Analysis: A Hückhiker's Guide*, (Studies in Economic Theory). Heidelberg, Germany: Springer, 1999.
- [13] P. Maćkowiak, "Some remarks on lower hemicontinuity of convex multivalued mappings," *Econ. Theory*, vol. 28, no. 1, pp. 227–233, 2006.
- [14] M. Aghassi and D. Bertsimas, "Robust game theory," *Math. Program.*, vol. 107, nos. 1–2, pp. 231–273, 2006.
- [15] P. N. De Souza and J.-N. Silva, *Berkeley Problems in Mathematics*. New York, NY, USA: Springer, 2004.
- [16] J. P. Crouzeix, "Quasi-concavity," in *The New Palgrave Dictionary of Economics*, M. M. J. Eatwell and P. Newman, Eds. London, U.K.: Palgrave Macmillan, 1987.
- [17] T. Adamo and A. Matros, "A blotto game with incomplete information," *Econ. Lett.*, vol. 105, no. 1, pp. 100–102, 2009.
- [18] R. Golman and S. E. Page, "General blotto: Games of allocative strategic mismatch," *Public Choice*, vol. 138, nos. 3–4, pp. 279–299, 2009.
- [19] S. Wang, "Existence of a Pareto equilibrium," *J. Optim. Theory Appl.*, vol. 79, no. 2, pp. 373–384, 1993.
- [20] X. P. Ding, "Constrained multiobjective games in general topological space," *Comput. Math. Appl.*, vol. 39, nos. 3–4, pp. 23–30, 2000.
- [21] Q. H. Ansari and Z. Khan, "On existence of Pareto equilibria for constrained multiobjective games," *Southeast Asian Bull. Math.*, vol. 27, no. 6, pp. 1–10, 2003.