

STEERING NO-REGRET LEARNERS TO OPTIMAL EQUILIBRIA

Anonymous authors

Paper under double-blind review

ABSTRACT

We consider the problem of steering no-regret-learning agents to play desirable equilibria via nonnegative payments. We first show that steering is impossible if the total budget (across all iterations) is finite, both in normal- and extensive-form games. However, we establish that *vanishing* average payments are compatible with steering. In particular, when players’ full strategies are observed at each timestep, we show that constant per-iteration payments permit steering. In the more challenging setting where only trajectories through the game tree are observable, we show that steering is impossible with constant per-iteration payments in general extensive-form games, but possible in normal-form games or if the maximum per-iteration payment may grow with time. We supplement our theoretical positive results with experiments highlighting the efficacy of steering in large games, and show how our framework relates to optimal mechanism design and information design.

1 INTRODUCTION

Any student of game theory learns that games can have multiple *equilibria* of different quality—for example, in terms of social welfare (Figure 1). How can a *mediator*—a benevolent third party—*steer* players toward an optimal one? In this paper, we consider the problem of using a mediator who can dispense nonnegative payments and offer advice to players so as to guide to a better collective outcome.

Importantly, our theory does not rest upon strong assumptions regarding agent obedience; instead, we only assume that players have *sublinear regret*, a mild assumption on the rationality of the players adopted in several prior studies (e.g., Nekipelov et al., 2015; Kolumbus & Nisan, 2022b; Camara et al., 2020). Variants of this problem have received tremendous interest in the literature (e.g., Monderer & Tennenholtz, 2004; Anshelevich et al., 2008; Schulz & Moses, 2003; Agussurja & Lau, 2009; Balcan, 2011; Balcan et al., 2013; 2014; Mguni et al., 2019; Li et al., 2020; Kempe et al., 2020; Liu et al., 2022 and references therein), but prior work either operates in more restricted classes of games or makes strong assumptions regarding player obedience. We study the steering problem in its full generality for general (imperfect-information) extensive-form games under an entire hierarchy of equilibrium concepts, and we establish a number of positive algorithmic results and complementing information-theoretic impossibilities.

Summary of Our Results Our formulation enables the mediator to 1) reward players with non-negative *payments* and 2) offer *advice*. Of course, with no constraints on the payments, the problem becomes trivial: the mediator could enforce any arbitrary outcome by paying players to play that outcome. On the other extreme, we show that if the *total* realized payments are constrained to be bounded, the decentralized steering problem is information-theoretically impossible (Proposition 3.2). Therefore, we compromise by allowing the total realized payments to be unbounded, but insist that the average payment per round is *vanishing*. Further, to justify 2) above, we show that *without advice*, steering to *mixed-Nash equilibria* is impossible already in normal-form games (Appendix D), although advice is not necessary for *pure-Nash equilibria* (Sections 4 and 5). Offering recommendations is in line with much of prior work (Appendix A), and is especially natural for correlated equilibrium concepts.

The goal of the mediator is to reach an equilibrium, either explicitly provided or provided as a principal utility function. We first assume that the mediator is provided an equilibrium. We distinguish between *realized* payments and *potential* payments. Realized payments are the payments actually dispensed




Figure 1: **Left:** An extensive-form version of a stag hunt. Chance plays uniformly at random at the root node, and the dotted line connecting the two nodes of Player 2 indicates an infoset: Player 2 cannot distinguish the two nodes. Introducing *vanishing* realized payments alters the gradient landscape, steering players to the optimal equilibrium (star) instead of the suboptimal one (opposite corner). The capital letters show the players’ initial strategies. Lighter color indicates higher welfare and the star shows the highest-welfare equilibrium. Further details are in Appendix C.

to the players. *Potential* payments are payments that players *would have* received, had they played different strategies.

We first consider the *full-feedback* (Section 5) setting where players’ payments may depend on players’ full strategies. We present steering algorithms that establish under different computational assumptions the first main result.

Theorem (Informal; precise versions in Theorem 5.2). *For both normal-form and extensive-form games, the decentralized steering problem can be solved under full feedback.*

Intuitively, the mediator sends payments in such a way as to 1) reward the player a small amount for playing the equilibrium, and 2) *compensate* the player for deviations of other players. Next, we consider the more challenging *bandit setting*, wherein only game trajectories are observed. In extensive-form games, this condition significantly restricts the structure of the payment functions, and in particular rules out the full-feedback algorithm above. We show that the decentralized steering problem under bandit feedback is information-theoretically impossible in the general case with bounded potential payments.

Theorem (Informal; precise version in Theorem 5.4). *For extensive-form games, the decentralized steering problem is impossible under bandit feedback with bounded potential payments.*

To circumvent this lower bound, we next allow the *potential* payments to depend on the time horizon, while still insisting that they vanish in the limit.

Theorem (Informal; precise version in Theorem 5.6). *For extensive-form games, if the payments may depend on the time horizon, the decentralized steering problem can be solved under bandit feedback.*

The proof of this theorem is more involved than the previous two. In particular, one might hope that the desired equilibrium can be made (strictly) dominant by adding appropriate payments as in *k*-implementation (Monderer & Tennenholtz, 2004). In extensive-form games, this is not the case: there are games where making the welfare-optimal equilibrium dominant would require payments in equilibrium, thereby inevitably leading to non-vanishing realized payments. Nevertheless, we show that steering is possible despite even without dominance. This leads to the intriguing behavior where some players may actually move *farther* from obedience before they move closer (compare Figure 1). As such, we significantly depart from the approach of Monderer & Tennenholtz (2004); we elaborate on this comparison and further related work in Appendix A.

Both previous positive results require computing an equilibrium upfront, which is both computationally expensive and not adaptive to players’ actions. We next analyze an *online* setting, where the mediator employs an online regret minimization algorithm to compute an optimal equilibrium *while* guiding the players toward it. As expected, algorithms for the online steering problem attain slightly worse rates compared to algorithms for the offline problem. The rates we obtain for the various versions of the steering problem all decay polynomially with the number of rounds, and we highlight the time dependence in Table 1. We complement our theoretical analysis by implementing and testing our steering algorithms in several benchmark games in Section 7.

Table 1: Summary of our positive algorithmic results. We hide game-dependent constants and logarithmic factors, and assume that regret minimizers incur a (typical) average regret of $T^{-1/2}$.

	Steering to Fixed Equilibrium	Online Steering
Normal Form or Full Feedback	$T^{-1/4}$ (Theorem 5.2)	$T^{-1/6}$ (Theorem 6.5)
Extensive Form and Bandit Feedback	$T^{-1/8}$ (Theorem 5.6)	<i>Open problem</i>

2 PRELIMINARIES

In this section, we introduce some basic background on extensive-form games.

Definition 2.1. An *extensive-form game* Γ with n players has the following components:

1. a set of players, identified with the set of integers $\llbracket n \rrbracket := \{1, \dots, n\}$. We will use $-i$, for $i \in \llbracket n \rrbracket$, to denote all players except i ;
2. a directed tree H of *histories* or *nodes*, whose root is denoted \emptyset . The edges of H are labeled with *actions*. The set of actions legal at h is denoted A_h . Leaf nodes of H are called *terminal*, and the set of such leaves is denoted by Z ;
3. a partition $H \setminus Z = H_{\mathbf{C}} \sqcup H_1 \sqcup \dots \sqcup H_n$, where H_i is the set of nodes at which i takes an action, and \mathbf{C} denotes the chance player;
4. for each player $i \in \llbracket n \rrbracket$, a partition \mathcal{I}_i of i 's decision nodes H_i into *information sets*. Every node in a given information set I must have the same set of legal actions, denoted by A_I ;
5. for each player i , a *utility function* $u_i : Z \rightarrow [0, 1]$ which we assume to be bounded; and
6. for each chance node $h \in H_{\mathbf{C}}$, a fixed probability distribution $c(\cdot | h)$ over A_h .

At a node $h \in H$, the *sequence* $\sigma_i(h)$ of an agent i is the set of all information sets encountered by agent i , and the actions played at such information sets, along the $\emptyset \rightarrow h$ path, excluding at h itself. An agent has *perfect recall* if $\sigma_i(h) = \sigma_i(h')$ for all h, h' in the same info set. Unless otherwise stated (Section 6), we assume that all players have perfect recall. We will use $\Sigma_i := \{\sigma_i(z) : z \in Z\}$ to denote the set of all sequences of player i that correspond to terminal nodes.

A *pure strategy* of player i is a choice of one action in A_I for each information set $I \in \mathcal{I}_i$. The *sequence form* of a pure strategy is the vector $\mathbf{x}_i \in \{0, 1\}^{\Sigma_i}$ given by $\mathbf{x}_i[\sigma] = 1$ if and only if i plays every action on the path from the root to sequence $\sigma \in \Sigma_i$. We will use the shorthand $\mathbf{x}_i[z] = \mathbf{x}_i[\sigma_i(z)]$. A *mixed strategy* is a distribution over pure strategies, and the sequence form of a mixed strategy is the corresponding convex combination $\mathbf{x}_i \in [0, 1]^{\Sigma_i}$. We will use X_i to denote the polytope of sequence-form mixed strategies of player i .

A profile of mixed strategies $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n) \in X := X_1 \times \dots \times X_n$, induces a distribution over terminal nodes. We will use $z \sim \mathbf{x}$ to denote sampling from such a distribution. The expected utility of agent i under such a distribution is given by $u_i(\mathbf{x}) := \mathbb{E}_{z \sim \mathbf{x}} u_i(z)$. Critically, the sequence form has the property that each agent's expected utility is a linear function of its own sequence-form mixed strategy. For a profile $\mathbf{x} \in X$ and set $N \subseteq \llbracket n \rrbracket$, we will use the notation $\hat{\mathbf{x}}_N \in \mathbb{R}^Z$ to denote the vector $\hat{\mathbf{x}}_N[z] = \prod_{j \in N} \mathbf{x}_j[z]$, and we will write $\hat{\mathbf{x}} := \hat{\mathbf{x}}_{\llbracket n \rrbracket}$. A Nash equilibrium is a strategy profile \mathbf{x} such that, for any $i \in \llbracket n \rrbracket$ and any $\mathbf{x}'_i \in X_i$, $u_i(\mathbf{x}) \geq u_i(\mathbf{x}'_i, \mathbf{x}_{-i})$.

3 THE STEERING PROBLEM

In this section, we introduce what we call the *steering* problem. Informally, the steering problem asks whether a mediator can always steer players to any given equilibrium of an extensive-form game.

Definition 3.1 (Steering Problem for Pure-Strategy Nash Equilibrium). Let Γ be an extensive-form game with payoffs bounded in $[0, 1]$. Let \mathbf{d} be an arbitrary pure-strategy Nash equilibrium of Γ . The mediator knows the game Γ , as well as a function $R(T) = o(T)$, which may be game-dependent, that bounds the regret of all players. At each round $t \in \llbracket T \rrbracket$, the mediator picks *payment functions* for each player, $p_i^{(t)} : X_1 \times \dots \times X_n \rightarrow [0, P]$, where $p_i^{(t)}$ is linear in \mathbf{x}_i and continuous in \mathbf{x}_{-i} , and P defines the largest allowable per-iteration payment. Then, players pick strategies $\mathbf{x}_i^{(t)} \in X_i$. Each player i then gets utility $v_i^{(t)}(\mathbf{x}_i) := u_i(\mathbf{x}_i, \mathbf{x}_{-i}^{(t)}) + p_i^{(t)}(\mathbf{x}_i, \mathbf{x}_{-i}^{(t)})$. The mediator has two desiderata.

- (S1) (Payments) The time-averaged realized payments to the players, defined as $\max_{i \in \llbracket n \rrbracket} \frac{1}{T} \sum_{t=1}^T p_i^{(t)}(\mathbf{x}^{(t)})$, converges to 0 as $T \rightarrow \infty$.
- (S2) (Equilibrium) Players' actions are indistinguishable from the Nash equilibrium \mathbf{d} . That is, the *directness gap*, defined as $\frac{1}{T} \sum_{t=1}^T \mathbb{E}_{z \sim \mathbf{x}^{(t)}} (1 - \hat{\mathbf{d}}[z])$, converges to 0 as $T \rightarrow \infty$.

The assumption imposed on the payment functions in Definition 3.1 ensures the existence of Nash equilibria in the payment-augmented game (e.g., Glicksberg, 1952). Throughout this paper, we will refer to players as *direct* if they are playing actions prescribed by the target equilibrium strategy \mathbf{d} . Critically, (S2) does not require that the strategies themselves converge to the direct strategies, i.e., $\mathbf{x}_i^{(t)} \rightarrow \mathbf{d}_i$, in iterates or in averages. They may differ on nodes off the equilibrium path. Instead, the requirement defined by (S2) is equivalent to *the reach probability of every node not reached in the equilibrium \mathbf{d} converging to 0*, so that, *on path*, the players play the equilibrium. Similarly, (S1) refers to the *realized payments* $p_i^{(t)}(\mathbf{x}^{(t)})$, not the *maximum offered payment* $\max_{\mathbf{x} \in X} p_i^{(t)}(\mathbf{x})$.

For now, we will assume that a pure Nash equilibrium has been computed, and therefore our only task is to steer the agents toward it. In Section 6 we show how our steering algorithms can be directly applied to other equilibrium concepts such as *mixed* or *correlated* equilibria, and *communication* equilibria, and to the case where the equilibrium has not been precomputed.

The mediator does not know anything about how the players pick their strategies, except that they will have regret bounded by a function that vanishes in the limit and is known to the mediator. This condition is a commonly adopted behavioral assumption (Nekipelov et al., 2015; Kolumbus & Nisan, 2022b; Camara et al., 2020). The regret of Player $i \in \llbracket n \rrbracket$ in this context is defined as

$$\text{Reg}_{X_i}^T := \frac{1}{P+1} \left[\max_{\mathbf{x}_i^* \in X_i} \sum_{t=1}^T v_i^{(t)}(\mathbf{x}_i^*) - \sum_{t=1}^T v_i^{(t)}(\mathbf{x}_i^{(t)}) \right].$$

That is, regret takes into account the payment functions offered to that player.¹ The assumption of bounded regret is realistic even in extensive-form games, as various regret minimizing algorithms exist. Two notable examples are the *counterfactual regret minimization* (CFR) framework (Zinkevich et al., 2007), which yields *full-feedback* regret minimizers, and IXOMD (Kozuno et al., 2021), which yields *bandit-feedback* regret minimizers.

How large payments are needed to achieve (S1) and (S2)? If the mediator could provide totally unconstrained payments, it could enforce any arbitrary outcome. On the other hand if the total payments are restricted to be bounded, the steering problem is information-theoretically impossible:

Proposition 3.2. *There exists a game and some function $R(T) = O(\sqrt{T})$ such that, for all $B \geq 0$, the steering problem is impossible if we add the constraint $\sum_{t=1}^{\infty} \sum_{i=1}^n p_i^{(t)}(\mathbf{x}^{(t)}) \leq B$.*

(Proofs are in Appendix E unless otherwise stated.) Hence, a weaker requirement on the size of the payments is needed. Between these extremes, one may allow the *total* payment to be unbounded, but insist that the *average* payment per round must vanish in the limit.

4 STEERING IN NORMAL-FORM GAMES

We start with the example of *normal-form games*. A normal-form game, in our language, is simply an extensive-form game in which every player has one information set, and the set of histories correspond precisely to the set of pure profiles, i.e., for every pure profile \mathbf{x} , we have $\hat{\mathbf{x}}[z] = 1$ for exactly one terminal node z . This setting is, much more simple than the general extensive-form setting which we will consider in the next section. In normal-form games, the strategy sets X_i are simplices, $X_i = \Delta(A_i)$, where A_i is the action set of player i at its only decision point. In this setting, we are able to turn to a special case of a result of Monderer & Tennenholtz (2004):

Theorem 4.1 (Costless implementation of pure Nash equilibria, special case of k -implementation, Monderer & Tennenholtz, 2004). *Let \mathbf{d} be a pure Nash equilibrium in a normal-form game. Then there exist functions $p_i^* : X_1 \times \dots \times X_n \rightarrow [0, 1]$, with $p_i^*(\mathbf{d}) = 0$, such that in the game with utilities $v_i := u_i + p_i^*$, the profile \mathbf{d} is weakly dominant: $v_i(\mathbf{d}_i, \mathbf{x}_{-i}) \geq v_i(\mathbf{x}_i, \mathbf{x}_{-i})$ for every profile \mathbf{x} .*

¹The division by $1/(P+1)$ is for normalization, since $v_i^{(t)}$'s has range $[0, P+1]$.

Indeed, it is easy to check that the payment function

$$p_i^*(\mathbf{x}) := (\mathbf{d}_i^\top \mathbf{x}_i) \left(1 - \prod_{j \neq i} \mathbf{d}_j^\top \mathbf{x}_j\right),$$

which on pure profiles \mathbf{x} returns 1 if and only if $\mathbf{x}_i = \mathbf{d}_i$ and $\mathbf{x}_j \neq \mathbf{d}_j$ for some $j \neq i$, satisfies these properties. Such a payment function is *almost* enough for steering: the only problem is that \mathbf{d} is only *weakly* dominant, so no-regret players *may* play other strategies than \mathbf{d} . This is easily fixed by adding a small reward $\alpha \ll 1$ for playing \mathbf{d}_i . That is, we set

$$p_i(\mathbf{x}) := \alpha \mathbf{d}_i^\top \mathbf{x}_i + p_i^*(\mathbf{x}) = (\mathbf{d}_i^\top \mathbf{x}_i) \left(\alpha + 1 - \prod_{j \neq i} \mathbf{d}_j^\top \mathbf{x}_j\right). \quad (1)$$

On a high level, the structure of the payment function guarantees that the average strategy of any no-regret learner $i \in [n]$ should be approaching the direct strategy \mathbf{d}_i by making \mathbf{d}_i the strictly dominant strategy of player i . At the same time, it is possible to ensure that the average payment will also be vanishing by appropriately selecting parameter α . With appropriate choice of α , this is enough to solve the steering problem for normal-form games:

Theorem 4.2 (Normal-form steering). *Let $p_i(\mathbf{x})$ be defined as in (1), set $\alpha = \sqrt{\varepsilon}$, where $\varepsilon := 4nR(T)/T$, and let T be large enough that $\alpha \leq 1$. Then players will be steered toward equilibrium, with both payments and directness gap bounded by $2\sqrt{\varepsilon}$.*

5 STEERING IN EXTENSIVE-FORM GAMES

The extensive-form setting is significantly more involved than the normal-form setting, and it will be the focus for the remainder of our paper, for two reasons. First, in extensive form, the strategy spaces of the players are no longer simplices. Therefore, if we wanted to write a payment function p_i with the property that $p_i(\mathbf{x}) = \alpha \mathbb{1}\{\mathbf{x} = \mathbf{d}\} + \mathbb{1}\{\mathbf{x}_i = \mathbf{d}_i; \exists j \mathbf{x}_j \neq \mathbf{d}_j\}$ for pure \mathbf{x} (which is what was needed by Theorem 4.2), such a function would not be linear (or even convex) in player i 's strategy $\mathbf{x}_i \in X_i$ (which is a sequence-form strategy, not a distribution over pure strategies). As such, even the meaning of extensive-form regret minimization becomes suspect in this setting. Second, in extensive form, a desirable property would be that the mediator give payments conditioned only on what actually happens in gameplay, *not* on the players' full strategies—in particular, if a particular information set is not reached during play, the mediator should not know what action the player *would have* selected at that information set. We will call this the *bandit* setting, and distinguish it from the *full-feedback* setting, where the mediator observes the players' full strategies.² This distinction is meaningless in the normal-form setting: since terminal nodes in normal form correspond to (pure) profiles, observing gameplay is equivalent to observing strategies. (We will discuss this point in more detail when we introduce the bandit setting in Section 5.2.)

We now present two different algorithms for the steering problem, one in the full-feedback setting, and one in the bandit setting.

5.1 STEERING WITH FULL FEEDBACK

In this section, we introduce a steering algorithm for extensive-form games under full feedback.

Algorithm 5.1 (FULLFEEDBACKSTEER). At every round, set the payment function $p_i(\mathbf{x}_i, \mathbf{x}_{-i})$ as

$$\underbrace{\alpha \mathbf{d}_i^\top \mathbf{x}_i}_{\text{directness bonus}} + \underbrace{[u_i(\mathbf{x}_i, \mathbf{d}_{-i}) - u_i(\mathbf{x}_i, \mathbf{x}_{-i})]}_{\text{sandboxing payments}} - \underbrace{\min_{\mathbf{x}'_i \in X_i} [u_i(\mathbf{x}'_i, \mathbf{d}_{-i}) - u_i(\mathbf{x}'_i, \mathbf{x}_{-i})]}_{\text{payment to ensure nonnegativity}}, \quad (2)$$

where $\alpha \leq 1/|Z|$ is a hyperparameter that we will select appropriately.

By construction, p_i satisfies the conditions of the steering problem (Definition 3.1): it is linear in \mathbf{x}_i , continuous in \mathbf{x}_{-i} , nonnegative, and bounded by an absolute constant (namely, 3). The payment function defined above has three terms:

²To be clear, the settings are differentiated by what the *mediator* observes, not what the *players* observe. That is, it is valid to consider the full-feedback steering setting with players running bandit regret minimizers, or the bandit steering setting with players running full-feedback regret minimizing algorithms.

1. The first term is a *reward for directness*: a player gets a reward proportional to α if it plays \mathbf{d}_i .
2. The second term *compensates the player* for the indirectness of other players. That is, the second term ensures that players’ rewards are *as if* the other players had acted directly.
3. The final term simply ensures that the overall expression is nonnegative.

We claim that this protocol solves the basic version of the steering problem. Formally:

Theorem 5.2. *Set $\alpha = \sqrt{\varepsilon}$, where $\varepsilon := 4nR(T)/T$, and let T be large enough that $\alpha \leq 1/|Z|$. Then, FULLFEEDBACKSTEER results in average realized payments and directness gap at most $3|Z|\sqrt{\varepsilon}$.*

5.2 STEERING IN THE BANDIT SETTING

In FULLFEEDBACKSTEER, payments depend on full strategies \mathbf{x} , not the realized game trajectories. In particular, the mediator in Theorem 5.2 observes what the players *would have played* even at infosets that other players avoid. To allow for an algorithm that works without knowledge of full strategies, $p_i^{(t)}$ must be structured so that it could be induced by a payment function that only gives payments for terminal nodes reached during play. To this end, we now formalize *bandit steering*.

Definition 5.3 (Bandit steering problem). Let Γ be an extensive-form game in which rewards are bounded in $[0, 1]$ for all players. Let \mathbf{d} be an arbitrary pure-strategy Nash equilibrium of Γ . The mediator knows Γ and a regret bound $R(T) = o(T)$. At each $t \in \llbracket T \rrbracket$, the mediator selects a payment function $q_i^{(t)} : Z \rightarrow [0, P]$. The players select strategies $\mathbf{x}_i^{(t)}$. A terminal node $z^{(t)} \sim \mathbf{x}^{(t)}$ is sampled, and all agents observe the terminal node that was reached, $z^{(t)}$. The players get payments $q_i^{(t)}(z^{(t)})$, so that their expected payment is $p_i^{(t)}(\mathbf{x}) := \mathbb{E}_{z \sim \mathbf{x}} q_i^{(t)}(z)$. The desiderata are as in Definition 3.1.

The bandit steering problem is more difficult than the non-bandit steering problem in two ways. First, as discussed above, the mediator does not observe the strategies \mathbf{x} , only a terminal node $z^{(t)} \sim \mathbf{x}$. Second, the form of the payment function $q_i^{(t)} : Z \rightarrow [0, P]$ is restricted: this is already sufficient to rule out FULLFEEDBACKSTEER. Indeed, p_i as defined in (2) cannot be written in the form $\mathbb{E}_{z \sim \mathbf{x}} q_i(z) : p_i(\mathbf{x}_i, \mathbf{x}_{-i})$ is nonlinear in \mathbf{x}_{-i} due to the nonnegativity-ensuring payments, whereas every function of the form $\mathbb{E}_{z \sim \mathbf{x}} q_i(z)$ will be linear in each player’s strategy.

We remark that, despite the above algorithm containing a sampling step, the payment function is defined *deterministically*: the payment is defined as the *expected value* $p_i^{(t)}(\mathbf{x}) := \mathbb{E}_{z \sim \mathbf{x}} q_i^{(t)}(z)$. Thus, the theorem statements in this section will also be deterministic.

In the normal-form setting, the payments p_i defined by (1) already satisfy the condition of bandit steering. In particular, let z be the terminal node we have $p_i(\mathbf{x}) = \mathbb{E}_{z \sim \mathbf{x}} [\alpha \mathbb{1}\{z = z^*\} + \mathbb{1}\{\mathbf{x}_i = \mathbf{d}_i; \exists j \mathbf{x}_j \neq \mathbf{d}_j\}]$. Therefore, in the normal-form setting, Theorem 4.2 applies to both full-feedback steering and bandit steering, and we have no need to distinguish between the two. However, in extensive form, as discussed above, the two settings are quite different.

5.2.1 LOWER BOUND ON REQUIRED PAYMENTS

Unlike in the full-feedback or normal-form settings, in the bandit setting, steering is impossible in the general case in the sense that per-iteration payments bounded by any constant do not suffice.

Theorem 5.4. *For every $P > 0$, there exists an extensive-form game Γ with $O(P)$ players, $O(P^2)$ nodes, and rewards bounded in $[0, 1]$ such that, with payments $q_i^{(t)} : Z \rightarrow [0, P]$, it is impossible to steer players to the welfare-maximizing Nash equilibrium, even when $R(T) = 0$.*

For intuition, consider the extensive-form game in Figure 2, which can be seen as a three-player version of Stag Hunt. Players who play Hare (H) get a value of $1/2$ (up to constants); in addition, if all three players play Stag (S), they all get expected value 1. The welfare-maximizing equilibrium is “everyone plays Stag”, but “everyone plays Hare” is also an equilibrium. In addition, if all players are playing Hare, the only way for the mediator to convince a player to play Stag without accidentally also paying players in the Stag equilibrium is to pay players at one of the three boxed nodes. But those three nodes are only reached with probability $1/n$ as often as the three nodes on the left, so the mediator would have to give a bonus of more than $n/2$. The full proof essentially works by deriving an algorithm that the players could use to exploit this dilemma to achieve either large payments or bad convergence rate, generalizing the example to $n > 3$, and taking $n = \Theta(P)$.




Figure 2: The counterexample for Theorem 5.4, for $n = 3$. Chance always plays uniformly at random. Infosets are linked by dotted lines (all nodes belonging to the same player are in the same infoset).

5.2.2 BANDIT STEERING WITH LARGE OFF-PATH PAYMENTS

To circumvent the lower bound in Theorem 5.4, in this subsection, we allow the payment bound $P \geq 1$ to depend on both the time limit T and the game. Consider the following algorithm.

Algorithm 5.5 (BANDITSTEER). Let α, P be hyperparameters. Then, for all rounds $t = 1, \dots, T$, sample $z \sim \mathbf{x}^{(t)}$ and pay players as follows. If all players have been direct (i.e., if $\hat{\mathbf{d}}[z] = 1$), pay all players α . If at least one player has not been direct, pay P to all players who have been direct. That is, set $q_i^{(t)}(z^{(t)}) = \alpha \hat{\mathbf{d}}[z] + P d_i[z](1 - \hat{\mathbf{d}}[z])$.

Theorem 5.6. Set the hyperparameters $\alpha = 4|Z|^{1/2}\varepsilon^{1/4}$ and $P = 2|Z|^{1/2}\varepsilon^{-1/4}$, where $\varepsilon := R(T)/T$, and let T be large enough that $\alpha \leq 1$. Then, running BANDITSTEER for T rounds results in average realized payments bounded by $8|Z|^{1/2}\varepsilon^{1/4}$, and directness gap by $2\varepsilon^{1/2}$.

The proof of this result is more involved than those for previous results. One may hope that—as in FULLFEEDBACKSTEER—the desired equilibrium can be made dominant by adding payments. But this is impossible: in the smaller “stag hunt” game in Figure 1, for Player 2, Stag cannot be a weakly-dominant strategy unless a payment is given at the boxed node, which would be problematic because such payments would also appear in equilibrium, in violation of (S1). In fact, a sort of “chicken-and-egg” problem arises: (S2) requires that all players converge to equilibrium. But for this to happen, other players’ strategies must first converge to equilibrium so that i ’s incentives are as they would be in equilibrium. The main challenge in the proof of Theorem 5.6 is therefore to carefully set the hyperparameters to achieve convergence despite these apparent problems.

6 OTHER EQUILIBRIUM NOTIONS AND ONLINE STEERING

So far, Theorems 5.2 and 5.6 refer only to *pure-strategy* Nash equilibria of a game. We now show how to apply these algorithms to other equilibrium notions such as mixed-strategy or correlated equilibrium. The key insight is that many types of equilibrium can be viewed as pure-strategy equilibria in an augmented game. For example, an extensive-form correlated equilibrium of a game Γ can be viewed as a pure-strategy equilibrium of an augmented game Γ' in which the mediator samples actions (“recommendations”) and the acting player observes those recommendations. Then, in Γ' , the goal is to guide players toward the pure strategy profile of following recommendations.

We now formalize these ideas. For this section, let Γ refer to a *mediator-augmented* game (Zhang & Sandholm, 2022), which has $n + 1$ players $i \in [n] \cup \{0\}$, where player 0 is the mediator. We will assume the *revelation principle*, which allows us to fix a target pure strategy profile \mathbf{d} that we want to make the equilibrium profile for the non-mediator players. We will write Γ^μ to refer to the n -player game in which the mediator is fixed to playing the strategy μ .

Definition 6.1. An *equilibrium in the mediator-augmented game* Γ is a strategy $\boldsymbol{\mu} \in X_0$ for the mediator such that \boldsymbol{d} is a Nash equilibrium of Γ^μ . An equilibrium $\boldsymbol{\mu}$ is *optimal* if, among all equilibria, it maximizes the mediator’s objective $u_0(\boldsymbol{\mu}, \boldsymbol{d})$.

By varying the construction of the augmented game Γ , the family of solution concepts for extensive-form games captured by this framework includes, but is not limited to, normal-form coarse correlated equilibrium (Aumann, 1974; Moulin & Vial, 1978); extensive-form correlated equilibrium (EFCE)³ (von Stengel & Forges, 2008); communication equilibrium (Forges, 1986); mechanism design; and information design/Bayesian persuasion (Kamenica & Gentzkow, 2011).

Unlike the offline setting (where the target equilibrium is given to us), in the online setting we can choose the target equilibrium. In particular, we would like to steer players toward an *optimal* equilibrium $\boldsymbol{\mu}$, without knowing that equilibrium beforehand. To that end, we add a new criterion:

- (S3) (Optimality) The mediator’s reward should converge to the reward of the optimal equilibrium. That is, the *optimality gap* $u_0^* - \frac{1}{T} \sum_{t=1}^T u_0(\boldsymbol{\mu}^{(t)}, \boldsymbol{x}^{(t)})$, where u_0^* is the mediator utility in an optimal equilibrium, converges to 0 as $T \rightarrow \infty$.

In Appendix D, we discuss why it is in some sense necessary to allow the mediator to give recommendations, not just payments, if the target equilibrium is not pure.

Since equilibria in mediator-augmented games are just strategies $\boldsymbol{\mu}$ under which \boldsymbol{d} is a Nash equilibrium, we may use the following algorithm to steer players toward an optimal equilibrium of Γ :

Algorithm 6.2 (COMPUTETHENSTEER). Compute an optimal equilibrium $\boldsymbol{\mu}$. With $\boldsymbol{\mu}$ held fixed, run any steering algorithm in Γ^μ .

As observed earlier, the main weakness of COMPUTETHENSTEER is that it must compute an equilibrium offline. To sidestep this, in this section we will introduce algorithms that compute the equilibrium in an *online* manner, while steering players toward it. Our algorithms will make use of a Lagrangian dual formulation analyzed by Zhang et al. (2023).

Proposition 6.3 (Zhang et al. (2023)). *There exists a (game-dependent) constant $\lambda^* \geq 0$ such that, for every $\lambda \geq \lambda^*$, the solutions $\boldsymbol{\mu}$ to*

$$\max_{\boldsymbol{\mu} \in X_0} \min_{\boldsymbol{x} \in X} u_0(\boldsymbol{\mu}, \boldsymbol{d}) - \lambda \sum_{i=1}^n [u_i(\boldsymbol{\mu}, \boldsymbol{x}_i, \boldsymbol{d}_{-i}) - u_i(\boldsymbol{\mu}, \boldsymbol{d}_i, \boldsymbol{d}_{-i})], \quad (3)$$

are exactly the optimal equilibria of the mediator-augmented game.

Algorithm 6.4 (ONLINESTEER). The mediator runs a regret minimization algorithm \mathcal{R}_0 over its own strategy space X_0 , which we assume has regret at most $R_0(T)$ after T rounds. On each round, the mediator does the following:

- Get a strategy $\boldsymbol{\mu}^{(t)}$ from \mathcal{R}_0 . Play $\boldsymbol{\mu}^{(t)}$, and set $p_i^{(t)}$ as defined in (2) in $\Gamma^{\boldsymbol{\mu}^{(t)}}$.
- Pass utility $\boldsymbol{\mu} \mapsto \frac{1}{\lambda} u_0(\boldsymbol{\mu}, \boldsymbol{d}) - \sum_{i=1}^n [u_i(\boldsymbol{\mu}, \boldsymbol{x}_i^{(t)}, \boldsymbol{d}_{-i}) - u_i(\boldsymbol{\mu}, \boldsymbol{d}_i, \boldsymbol{d}_{-i})]$ to \mathcal{R}_0 , where $\lambda \geq 1$ is a hyperparameter.

Theorem 6.5. *Set the hyperparameters $\alpha = \varepsilon^{2/3}|Z|^{-1/3}$ and $\lambda = |Z|^{2/3}\varepsilon^{-1/3}$, where $\varepsilon := (R_0(T) + 4nR(T))/T$ is the average regret bound summed across players, and let T be large enough that $\alpha \leq 1/|Z|$. Then running ONLINESTEER results in average realized payments, directness gap, and optimality gap all bounded by $7\lambda^*|Z|^{4/3}\varepsilon^{1/3}$.*

The argument now works with the zero-sum formulation (3), and leverages the fact that the agents’ average strategies are approaching the set of Nash equilibria since they have vanishing regrets. Thus, each player’s average strategy should be approaching the direct strategy, which in turn implies that the average utility of the mediator is converging to the optimal value, analogously to Theorem 5.2.

ONLINESTEER has a further guarantee that FULLFEEDBACKSTEER does not, owing to the fact that it learns an equilibrium online: it works even when the players’ sets of deviations, \bar{X}_i , is not known upfront. In particular, the following generalization of Theorem 6.5 follows from an identical proof.

³This requires the mediator to have imperfect recall.




Figure 3: Sample experimental results. The blue line in each figure is the social welfare (left y-axis) of the players *with* steering enabled. The green dashed line is the social welfare *without* steering. The yellow line gives the payment (right y-axis) paid to each player. The flat black line denotes the welfare of the optimal equilibrium. The panels show the game, the equilibrium concept (in this figure, always EFCE). In all cases, the first ten iterations are a “burn-in” period during which no payments are issued; steering only begins after that.

Corollary 6.6. *Suppose that each player i , unbeknownst to the mediator, is choosing from a subset $Y_i \subseteq X_i$ of strategies that includes the direct strategy d_i . Then, running Theorem 6.5 with the same hyperparameters yields the same convergence guarantees, except that the mediator’s utility converges to its optimal utility against the true deviators, that is, a solution to (3) with each X_i replaced by Y_i .*

At this point, it is very reasonable to ask whether it is possible to perform *online* steering with *bandit* feedback. In *normal-form* games, as with offline setting, there is minimal difference between the bandit and the full-feedback setting. This intuition carries over to the bandit setting: ONLINESTEER can be adapted into an online bandit steering algorithm for normal-form games, with essentially the same convergence guarantee. We defer the formal statement of the algorithm and proof to Appendix F.

The algorithm, however, fails to extend to the *extensive-form* online bandit setting, for the same reasons that the *offline* full-feedback algorithm fails to extend to the online setting.

7 EXPERIMENTAL RESULTS

We ran experiments with our BANDITSTEER algorithm (Algorithm 5.5) on various notions of equilibrium in extensive-form games, using the COMPUTETHENSTEER framework suggested by Algorithm 6.2. Since the hyperparameter settings suggested by Algorithm 5.5 are very extreme, in practice we fix a constant P and set α dynamically based on the currently-observed gap to directness. We used CFR+ (Tammelin, 2014) as the regret minimizer for each player, and precomputed a welfare-optimal equilibrium with the LP algorithm of Zhang & Sandholm (2022). In most instances tested, a small constant P (say, $P \leq 8$) is enough to steer CFR+ regret minimizers to the exact equilibrium in a finite number of iterations. Two plots exhibiting this behavior are shown in Figure 3. More experiments, as well as descriptions of the game instances tested, can be found in Appendix G.

8 CONCLUSIONS AND FUTURE RESEARCH

We established that it is possible to steer no-regret learners to optimal equilibria using vanishing rewards, even under bandit feedback. There are many interesting avenues for future research. First, is there a natural *bandit online* algorithm that combines the desirable properties of both ONLINESTEER and BANDITSTEER? Also, it is important to understand the best rates attainable for the different settings of the steering problem. Furthermore, is there a steering algorithm for which the mediator needs to know even less information about the game upfront? For example, could a mediator without knowledge of the players’ utilities still steer toward optimal equilibria? **Finally, our main behavioral assumption throughout this paper is that players incur vanishing average regret. Yet, stronger guarantees are possible when specific no-regret learning dynamics are in place; e.g., see (Vlatakis-Gkaragkounis et al., 2020; Giannou et al., 2021a;b) for recent characterizations in the presence of *strict* equilibria. Concretely, it would be interesting to understand the class of learning dynamics under which the steering problem can be solved with a finite cumulative budget.**

REFERENCES

- Lucas Agussurja and Hoong Chuin Lau. The price of stability in selfish scheduling games. *Web Intell. Agent Syst.*, 7(4):321–332, 2009.
- Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal of Computing*, 32:48–77, 2002.
- Robert Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.
- Maria-Florina Balcan. Leading dynamics to good behavior. *SIGecom Exch.*, 10(2):19–22, 2011.
- Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Improved equilibria via public service advertising. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 728–737, 2009.
- Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Circumventing the price of anarchy: Leading dynamics to good behavior. *SIAM Journal on Computing*, 42(1):230–264, 2013.
- Maria-Florina Balcan, Sara Krehbiel, Georgios Piliouras, and Jinwoo Shin. Near-optimality in covering games by exposing global information. *ACM Trans. Economics and Comput.*, 2(4):13:1–13:22, 2014.
- Mark Braverman, Jieming Mao, Jon Schneider, and Matt Weinberg. Selling to a no-regret buyer. In *ACM Conference on Economics and Computation (EC)*, pp. 523–538, 2018.
- William Brown, Jon Schneider, and Kiran Vodrahalli. Is learning in games good for the learners? *CoRR*, abs/2305.19496, 2023.
- Linda Cai, S. Matthew Weinberg, Evan Wildenhain, and Shirley Zhang. Selling to multiple no-regret buyers. *CoRR*, abs/2307.04175, 2023.
- Modibo K Camara, Jason D Hartline, and Aleck Johnsen. Mechanisms for a no-regret agent: Beyond the common prior. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 259–270. IEEE, 2020.
- In-Koo Cho and Jonathan Libgober. Machine learning for strategic inference. *arXiv preprint arXiv:2101.09613*, 2021.
- Maurizio D’Andrea. Playing against no-regret players. *Operations Research Letters*, 2023.
- Yuan Deng, Pingzhong Tang, and Shuran Zheng. Complexity and algorithms of k-implementation. In *Autonomous Agents and Multi-Agent Systems*, pp. 5–13. ACM, 2016.
- Yuan Deng, Jon Schneider, and Balasubramanian Sivan. Strategizing against no-regret learners. In *Conference on Neural Information Processing Systems (NeurIPS)*, pp. 1577–1585, 2019.
- Paul Duetting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. Multi-agent contracts. *CoRR*, abs/2211.05434, 2022.
- Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. Combinatorial contracts. In *Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 815–826. IEEE, 2021a.
- Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. The complexity of contracts. *SIAM Journal on Computing*, 50(1):211–254, 2021b.
- Gabriele Farina, Chun Kai Ling, Fei Fang, and Tuomas Sandholm. Correlation in extensive-form games: Saddle-point formulation and benchmarks. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.

- Francoise Forges. An approach to communication equilibria. *Econometrica: Journal of the Econometric Society*, pp. 1375–1385, 1986.
- Rupert Freeman, David M. Pennock, Chara Podimata, and Jennifer Wortman Vaughan. No-regret and incentive-compatible online learning. In *International Conference on Machine Learning (ICML)*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3270–3279. PMLR, 2020.
- Angeliki Giannou, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Panayotis Mertikopoulos. On the rate of convergence of regularized learning in games: From bandits and uncertainty to optimism and beyond. In *Conference on Neural Information Processing Systems (NeurIPS)*, pp. 22655–22666, 2021a.
- Angeliki Giannou, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Panayotis Mertikopoulos. Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information. In *Conference on Learning Theory (COLT)*, volume 134 of *Proceedings of Machine Learning Research*, pp. 2147–2148. PMLR, 2021b.
- I. L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *American Mathematical Society*, 3(1):170–174, 1952.
- Shengyuan Hu, Dung Daniel Ngo, Shuran Zheng, Virginia Smith, and Zhiwei Steven Wu. Federated learning as a network effects game, 2023.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6): 2590–2615, 2011.
- David Kempe, Sixie Yu, and Yevgeniy Vorobeychik. Inducing equilibria in networked public goods games through network structure modification. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20*, pp. 611–619. International Foundation for Autonomous Agents and Multiagent Systems, 2020.
- Robert D. Kleinberg, Katrina Ligett, Georgios Piliouras, and Éva Tardos. Beyond the nash equilibrium barrier. In *Innovations in Computer Science - ICS 2011*, pp. 125–140. Tsinghua University Press, 2011.
- Yoav Kolumbus and Noam Nisan. How and why to manipulate your own agent: On the incentives of users of learning agents. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2022a.
- Yoav Kolumbus and Noam Nisan. Auctions between regret-minimizing agents. In Frédérique Laforest, Raphaël Troncy, Elena Simperl, Deepak Agarwal, Aristides Gionis, Ivan Herman, and Lionel Médini (eds.), *WWW '22: The ACM Web Conference 2022*, pp. 100–111. ACM, 2022b.
- Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Symposium on Theoretical Aspects in Computer Science*, 1999.
- Tadashi Kozuno, Pierre Ménard, Remi Munos, and Michal Valko. Learning in two-player zero-sum partially observable markov games with perfect recall. *Advances in Neural Information Processing Systems*, 34:11987–11998, 2021.
- H. W. Kuhn. A simplified two-person poker. In H. W. Kuhn and A. W. Tucker (eds.), *Contributions to the Theory of Games*, volume 1 of *Annals of Mathematics Studies*, 24, pp. 97–103. Princeton University Press, Princeton, New Jersey, 1950.
- Jiayang Li, Jing Yu, Yu Marco Nie, and Zhaoran Wang. End-to-end learning and intervention in games. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Kai Li, Wenhan Huang, Chenchen Li, and Xiaotie Deng. Exploiting a no-regret opponent in repeated zero-sum games. *Journal of Shanghai Jiaotong University (Science)*, 2023.
- Boyi Liu, Jiayang Li, Zhuoran Yang, Hoi-To Wai, Mingyi Hong, Yu Nie, and Zhaoran Wang. Inducing equilibria via incentives: Simultaneous design-and-play ensures global convergence. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2022.

- Yishay Mansour, Mehryar Mohri, Jon Schneider, and Balasubramanian Sivan. Strategizing against learners in bayesian games. In *Conference on Learning Theory (COLT)*, volume 178 of *Proceedings of Machine Learning Research*, pp. 5221–5252. PMLR, 2022.
- David Mguni, Joel Jennings, Emilio Sison, Sergio Valcarcel Macua, Sofia Ceppi, and Enrique Munoz de Cote. Coordinating the crowd: Inducing desirable equilibria in non-cooperative systems. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '19, pp. 386–394. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- Dov Monderer and Moshe Tennenholtz. K-implementation. *J. Artif. Intell. Res.*, 21:37–62, 2004.
- H. Moulin and J.-P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3-4):201–221, 1978.
- Roger B Myerson. Multistage games with communication. *Econometrica: Journal of the Econometric Society*, pp. 323–358, 1986.
- Denis Nekipelov, Vasilis Syrgkanis, and Éva Tardos. Econometrics for learning agents. In *ACM Conference on Economics and Computation (EC)*, pp. 1–18, 2015.
- Ioannis Panageas and Georgios Piliouras. Average case performance of replicator dynamics in potential games via computing regions of attraction. In *ACM Conference on Economics and Computation (EC)*, pp. 703–720. ACM, 2016.
- Tim Roughgarden. *Selfish routing and the price of anarchy*. MIT Press, 2005.
- Tim Roughgarden. Intrinsic robustness of the price of anarchy. *Journal of the ACM*, 62(5):32:1–32:42, 2015.
- Tim Roughgarden and Okke Schrijvers. Online prediction with selfish experts. In *Conference on Neural Information Processing Systems (NeurIPS)*, pp. 1300–1310, 2017.
- Andreas S. Schulz and Nicolás E. Stier Moses. On the performance of user equilibria in traffic networks. In *Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 86–87, 2003.
- Oskari Tammelin. Solving large imperfect information games using CFR+. *arXiv preprint arXiv:1407.5042*, 2014.
- Emmanouil-Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, Thanasis Lianeas, Panayotis Mertikopoulos, and Georgios Piliouras. No-regret learning and mixed nash equilibria: They do not mix. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- Bernhard von Stengel and Françoise Forges. Extensive-form correlated equilibrium: Definition and computational complexity. *Mathematics of Operations Research*, 33(4):1002–1022, 2008.
- Brian Hu Zhang and Tuomas Sandholm. Polynomial-time optimal equilibria with a mediator in extensive-form games. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
- Brian Hu Zhang, Gabriele Farina, Andrea Celli, and Tuomas Sandholm. Optimal correlated equilibria in general-sum extensive-form games: Fixed-parameter algorithms, hardness, and two-sided column-generation. In *ACM Conference on Economics and Computation (EC)*, pp. 1119–1120. ACM, 2022.
- Brian Hu Zhang, Gabriele Farina, Ioannis Anagnostides, Federico Cacciamani, Stephen McAleer, Andreas Haupt, Andrea Celli, Nicola Gatti, Vincent Conitzer, and Tuomas Sandholm. Computing optimal equilibria and mechanisms via learning in zero-sum extensive-form games. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- Martin Zinkevich, Michael Bowling, Michael Johanson, and Carmelo Piccione. Regret minimization in games with incomplete information. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2007.