

Hindsight and Sequential Rationality of Correlated Play

Dustin Morrill¹, Ryan D’Orazio², Reza Sarfati³, Marc Lanctot⁴,
James R. Wright¹, Amy R. Greenwald⁵, Michael Bowling^{1,4}

¹University of Alberta; Alberta Machine Intelligence Institute, Canada

²Université de Montréal; Mila, Canada

³Massachusetts Institute of Technology, United States

⁴DeepMind

⁵Brown University, United States

morrill@ualberta.ca, ryan.dorazio@mila.quebec, sarfati@mit.edu, lanctot@google.com,
james.wright@ualberta.ca, amy_greenwald@brown.edu, mbowling@ualberta.ca

Abstract

Driven by recent successes in two-player, zero-sum game solving and playing, artificial intelligence work on games has increasingly focused on algorithms that produce equilibrium-based strategies. However, this approach has been less effective at producing competent players in general-sum games or those with more than two players than in two-player, zero-sum games. An appealing alternative is to consider adaptive algorithms that ensure strong performance in hindsight relative to what could have been achieved with modified behavior. This approach also leads to a game-theoretic analysis, but in the correlated play that arises from joint learning dynamics rather than factored agent behavior at equilibrium. We develop and advocate for this *hindsight rationality* framing of learning in general sequential decision-making settings. To this end, we re-examine *mediated equilibrium* and *deviation types* in *extensive-form games*, thereby gaining a more complete understanding and resolving past misconceptions. We present a set of examples illustrating the distinct strengths and weaknesses of each type of equilibrium in the literature, and prove that no tractable concept subsumes all others. This line of inquiry culminates in the definition of the deviation and equilibrium classes that correspond to algorithms in the *counterfactual regret minimization (CFR)* family, relating them to all others in the literature. Examining CFR in greater detail further leads to a new recursive definition of rationality in correlated play that extends *sequential rationality* in a way that naturally applies to hindsight evaluation.

Introduction

Algorithms to approximate maximin or Nash equilibrium strategies have fueled major successes in the autonomous play of human-scale games like Go (Silver et al. 2016, 2017b, 2018), chess and shogi (Silver et al. 2017a, 2018), poker (Moravčík et al. 2017; Brown and Sandholm 2018, 2019) and StarCraft (Vinyals et al. 2019). With the exception of multi-player poker addressed by Brown and Sandholm (2019), all of these are two-player, zero-sum games. Less success has been achieved in games outside this space.

While both maximin and Nash equilibrium concepts apply in general-sum, multi-player games, it is not clear that

either should be fielded against unknown strategies in these situations. The maximin concept assumes all other players collude against you, which can produce overly pessimistic strategies. Nash equilibria make the less pessimistic assumption that each player is rational and independent, but strategies from such equilibria have no performance guarantees against arbitrary strategies and are hard to compute. Perhaps one reason that AI has struggled to find the same success in multi-agent, general sum games is the lack of a strong theoretical base like that provided by the maximin objective in zero-sum games.

An appealing alternative is to consider adaptive algorithms that ensure strong performance in hindsight relative to what could have been achieved with modified behavior (e.g., Hannan (1957)). We advocate for a *hindsight rationality* framing of learning in sequential decision-making settings where rational correlated play, rather than factored equilibrium behavior, arises naturally through repeated play. Hindsight rationality is the dynamic interpretation of static *mediated equilibrium* (Aumann 1974) concepts that have led to elegant decentralized learning algorithms like *regret matching* (Hart and Mas-Colell 2000) and connections between learning and statistical calibration (Foster and Vohra 1997).

Mediated equilibria in EFGs have gained increasing interest recently with the introduction of *extensive-form (coarse-)correlated equilibrium (EF(C)CE)* (Forges and von Stengel 2002; von Stengel and Forges 2008; Farina, Bianchi, and Sandholm 2020) and *agent-form (Selten 1974) (coarse-)correlated equilibrium (AF(C)CE)* (von Stengel and Forges 2008), along with related algorithms (Celli et al. 2019, 2020). However, there has yet to be a clear description of the general mediated equilibrium landscape in EFGs beyond *causal* and *action deviations* (von Stengel and Forges 2008; Dudík and Gordon 2009). There has also yet to be a clear proposal about how this line of inquiry can be used to design more competent game playing algorithms. We show how hindsight rationality operationalizes these concepts in concrete *regret minimization* objectives while preserving equilibria as descriptions of jointly rational learning.

We present a more complete picture of the relationships between the deviation types and equilibrium classes dis-

cussed in the literature and resolve the misconception that all EFCEs are AFCEs¹. A central component of this analysis is a set of examples illustrating the distinct strengths and weaknesses of each type of equilibrium. The resulting table of equilibrium relationships reveals that while some such concepts are stronger than others, no tractable concept subsumes all others.

We define *counterfactual deviations* and *counterfactual (coarse-)correlated equilibrium (CF(C)CE)* to explicitly describe CFR's underlying deviation type and equilibrium behavior. We illustrate the unique strengths and weaknesses of counterfactual deviations in examples where there is a beneficial counterfactual deviation in an EFCE and an AFCE, and where there are beneficial causal and action deviations in a CFCE.

We go beyond counterfactual equilibria to more precisely characterize CFR's behavior by building a new connection between online learning and sequential rationality (Kreps and Wilson 1982). *Observable sequential rationality* is a new extension of sequential rationality to correlated play in a way that naturally applies to hindsight evaluation. We show that CFR is observably sequentially hindsight rational with respect to counterfactual deviations, and provide an example where a *correlated equilibrium (CE)* (Aumann 1974) is not an observable sequential CFCCE. Considering that it is generally intractable to compute a CE in EFGs, it is perhaps surprising that CFR, a simple and efficient algorithm, behaves in accordance with an equilibrium class that is not subsumed by a CE.

Extensive-Form Games

As we are interested in multi-agent, sequential decision-making, we use the *extensive-form game (EFG)* framework as the basis of our technical discussions.

The set of finite-length *histories*, \mathcal{H} , defines all ways that an EFG can partially play out, *i.e.*, the game begins with the *empty history*, \emptyset , and given a history h , the successor history after taking *action* $a \in \mathcal{A}(h)$ is ha . $\mathcal{A}(h)$ is the finite set of actions available at h for the player acting at h , determined by the *player function*, $\mathcal{P} : \mathcal{H} \setminus \mathcal{Z} \rightarrow \{1, \dots, N\} \cup \{c\}$, where $N > 0$ is the number of players and c is the *chance player*. Let $\mathcal{H}_i = \{h \in \mathcal{H} \mid \mathcal{P}(h) = i\}$ be the set of histories where player i must act. We write $h \sqsubseteq h'$ to state that h is a predecessor of h' . A *terminal history*, $z \in \mathcal{Z} \subseteq \mathcal{H}$, has an empty action set and no successors, and each history must lead to at least one terminal history (since all histories have finite length), at which point rewards are given to each player according to a bounded *utility function*, $u_i : \mathcal{Z} \rightarrow [-U, U]$.

To model imperfect information, histories are partitioned into *information sets*, $\mathcal{I}_i = \{I \mid I \subseteq \mathcal{H}_i\}$, where players are unable to distinguish between histories within an information set. Thus, the set of actions at each history within an information set must, must coincide, *i.e.*, $\mathcal{A}(I) \doteq \mathcal{A}(h) = \mathcal{A}(h')$, for any $h, h' \in I$. We restrict ourselves to *perfect-*

recall information partitions, which are those that ensure players never forget the information sets they encounter during play. Under this assumption, each player's information state transition graph forms a forest (not a tree since other players may act first) and are partially ordered. We use $I \preceq I'$ to denote that $I \in \mathcal{I}_i$ is a predecessor of $I' \in \mathcal{I}_i$.

A *pure strategy* for player i maps each information set to an action. We denote the finite set of such strategies as $S_i = \{s_i : s_i(I) \in \mathcal{A}(I), \forall I \in \mathcal{I}_i\}$ and *pure strategy profiles* with a strategy for each player by $S = \times_{i \in \{1, \dots, N, c\}} S_i$. If we let $-i$ denote the set of players other than player i , we can likewise define $S_{-i} = \times_{j \neq i} S_j$ as the set of strategies for all players except i .

Randomized behavior can be represented with a *mixed strategy*, $\pi_i \in \Pi_i = \Delta^{|S_i|}$, which is a distribution over pure strategies, or a *behavioral strategy* (Kuhn 1953), which is an assignment of *immediate strategies*, $\pi_i(I) \in \Delta^{|\mathcal{A}(I)|}$, to each information set². Perfect recall ensures *realization equivalence* between the set of mixed and behavioral strategies (Kuhn 1953). That is, for every mixed strategy there is a behavioral strategy, and *vice-versa*, that applies the same weight to each terminal history. Thus, we treat mixed and behavioral strategies as interchangeable representations where $\pi_i(s_i)$ is the probability of sampling pure strategy s_i and $\pi_i(a \mid I)$ is the probability of playing action a in information set I . Noticing as well that a pure strategy is just a behavioral strategy with deterministic immediate strategies, we default to the behavioral representation unless otherwise specified.

The *reach probability function* P describes the probability of playing from one history to another, following from the familiar chain rule of probability. Let $I(h)$ denote the unique information set that contains history h . Then the probability that player i plays from h to h' , through ha when $h \neq h'$, according to strategy π_i , is

$$P(h, h'; \pi_i) = \begin{cases} 1 & \text{if } h = h' \\ 0 & \text{if } h \not\sqsubseteq h' \\ \pi_i(a \mid I(h))P(ha, h'; \pi_i) & \text{if } \mathcal{P}(h) = i \\ P(ha, h'; \pi_i) & \text{o.w.} \end{cases}$$

The probability of reaching h from the beginning of the game is $P(h; \pi_i) \doteq P(\emptyset, h; \pi_i)$, and we overload P for strategy tuples, *i.e.*, $P(h, h'; \pi_{-i}) \doteq \prod_{j \neq i} P(h, h'; \pi_j)$ is the probability that players except for i play from h to h' and $P(h, h'; \pi) \doteq P(h, h'; \pi_i)P(h, h'; \pi_{-i})$. The terminal reach probabilities of a strategy profile describes its distribution over game outcomes, so the expected utility for player i is $u_i(\pi) = \sum_{z \in \mathcal{Z}} P(z; \pi)u_i(z)$.

Factored, Correlated, and Online Play

Factored play is that where players act entirely independently from one another. A *Nash equilibrium* (Nash 1951) models jointly rational play under such factored play. It is a mixed strategy profile where no player can benefit from a unilateral deviation, *i.e.*, π where $u_i(\pi'_i, \pi_{-i}) \leq u_i(\pi)$ for all π'_i .

² Δ^d is the d -dimensional probability simplex.

¹We use the definition of EFCE as used in the majority of the literature in artificial intelligence, which differs from von Stengel's intended definition (Personal communication 2020). See Footnote 3.

In a two-player, zero-sum game, every Nash equilibrium is a pair of *maximin strategies*, which are those that maximize the player’s minimum value against any opponent. In games with more than two players or general-sum payoffs, a strategy from a Nash equilibrium does not confer the same guarantee. Additionally, a maximin strategy for player i makes the pessimistic assumption that all other players will collude against i . These deficiencies leave us wanting for a theoretical framework that is better suited to multi-player, general-sum games, upon which to design game playing algorithms.

We suggest that *hindsight rationality*, which is the idea of learning correlated play over time that approaches optimality in hindsight, can serve as such a framework. This concept arises from an old connection between *mediated equilibria* (Aumann 1974) and *online learning* (e.g., Hannan (1957)).

Mediated Equilibria

A mediated equilibrium is a generalization of Nash equilibrium to correlated play. It is a *pure strategy profile distribution–deviation set* profile pair,

$$\left(\mu \in \Delta^{|S^I|}, (\Phi_i \subseteq \{\phi \mid \phi : S_i \rightarrow S_i\})_{i=1}^N \right),$$

where there is no beneficial unilateral deviation in any player’s deviation set. The benefit to player i of a deviation is the expected utility under the deviation compared to that of following the *recommendations* (strategies) sampled for player i from the recommendation distribution μ , i.e., $\mathbb{E}_{s \sim \mu} [u_i(\phi(s_i), s_{-i}) - u_i(s)]$. A *rational* player follows the recommendations from a mediated equilibrium since their value is optimal (with respect to deviation set Φ_i) given the play of the other players. The deviation set constrains the extent to which individual strategy modifications can condition on input strategies, thereby providing a mechanism for varying the strength and character of equilibrium rationality constraints.

For example, an equilibrium recommendation distribution under the set of *external transformations*, which are those that ignore their inputs, $\Phi_{S_i}^{\text{EX}} = \{\phi \mid \exists s'_i, \phi(s_i) = s'_i, \forall s_i\}$, is a *coarse-correlated equilibrium (CCE)* (Moulin and Vial 1978). In contrast, a *correlated equilibrium (CE)* (Aumann 1974) is one where the deviation set is unconstrained, which ensures that all CEs are CCEs. The unconstrained set of deviations is called the set of *swap transformations* and is denoted $\Phi_{S_i}^{\text{SW}}$. It turns out that the smaller set of *internal transformations* (Foster and Vohra 1999), which are those that act like the identity transformation except given one particular strategy, i.e.,

$$\Phi_{S_i}^{\text{IN}} = \{\phi \mid \exists s_i, s'_i, \phi(s_i) = s'_i, \phi(s''_i) = s''_i, \forall s''_i \neq s_i\},$$

also corresponds to CE (Greenwald, Jafari, and Marks 2003). In spite of this size reduction, computing a CE in EFGs is intractable since the pure strategy space grows exponentially with the size of the game.

Online Learning

In an online learning setting, a learner repeatedly plays a game with unknown, dynamic, possibly adversarial players.

On each round $1 \leq t \leq T$, the learner who acts as player i chooses a behavioral strategy, π_i^t , simultaneously with the other players who in aggregate choose π_{-i}^t . The learner is then evaluated based on the expected value they achieve under the selected strategy profile, $u_i(\pi^t)$.

In this context, we can consider deviations to the learner’s play in hindsight. A swap transformation generates a corresponding mixed strategy transformation, so we overload $\phi(\pi_i)$ as the mixed strategy that assigns probability $[\phi\pi_i](s_i) = \sum_{s'_i, \phi(s'_i)=s_i} \pi_i(s'_i)$ to each s_i (where $[\phi\pi_i] \doteq \phi(\pi_i)$). The cumulative benefit of deviating from the learner’s strategy sequence, $(\pi_i^t)_{t=1}^T$, is then

$$\rho^{1:T}(\phi) = \sum_{t=1}^T u_i(\phi(\pi_i), \pi_{-i}^t) - u_i(\pi^t). \quad (1)$$

We describe Eq. (1) as the cumulative *regret* the learner suffers for failing to modify their play according to ϕ . We denote the *maximum regret* under the set of deviations Φ as $R^T(\Phi) = \max_{\phi \in \Phi} \rho^{1:T}(\phi)$. A *no- Φ -regret* algorithm ensures its average (positive) regret for any deviation in Φ approaches zero over time.

In a two-player, zero-sum game, two competing no-regret learners generate strategy profiles that converge toward a Nash equilibrium when marginalized as $(\frac{1}{T} \sum_{t=1}^T \pi_i^t)_{i=1}^N$. However, marginalized empirical strategy profiles do not generally converge toward mediated or Nash equilibria (Shapley 1964), nor are they necessarily maximin strategies. Instead, hindsight rationality (as we discuss next) suggests an online learning approach to game playing since it makes a statement about the optimality of play that is actually observed, irrespective of how rational the other players behave.

Hindsight Rationality

The T strategy profiles deployed after T rounds of play forms an *empirical distribution of play*, $\mu^T \in \Delta^{|S^I|}$. That is, for all pure strategy profiles s , $\mu^T(s) = \frac{1}{T} \sum_{t=1}^T \pi^t(s)$, according to the probability of sampling s on each round, $\pi^t(s) = \prod_i \pi_i^t(s_i)$. If we treat the empirical distribution of play as a recommendation distribution, the benefit to player i of deviation ϕ is then exactly player i ’s average regret per round, i.e., $\mathbb{E}_{s \sim \mu^T} [u_i(\phi(s_i), s_{-i}) - u_i(s)] = \frac{1}{T} \rho^{1:T}(\phi)$. A learner participating in online play as player i is thus rational in hindsight with respect to Φ if they have no regret, where we say “in hindsight” because μ^T is constructed from past play. Furthermore, a learner with with $\varepsilon = R^T(\Phi)/T$ average regret is approximately rational in hindsight (precisely ε -rational in hindsight). Thus, a no-regret learner’s behavior approaches exact rationality in hindsight as they gain experience. We introduce the term *hindsight rational* as an equivalent alternative to “no-regret” for situations where it is useful to emphasize this shift in perspective, from learning to perform as well as strategy modifications, to learning to correlate optimally with other players.

EFG Deviation Types

Causal deviations: An *extensive-form correlated equilibrium (EFCE)* (Forges and von Stengel 2002; von Stengel and Forges 2008) is an equilibrium without beneficial *informed causal deviations* (Dudík and Gordon 2009). A causal deviation follows the recommended strategy, s_i , unless it is triggered at a particular information set, $I^!$ $\in \mathcal{I}_i$, at which point an alternative strategy, $s_i^!$, is used to play out the rest of the game.³ The immediate strategy of an informed causal deviation at a given information set is

$$[\phi s_i](I) = \begin{cases} s_i^!(I) & \text{if } s_i(I^!) = a^! \text{ and } I^! \preceq I \\ s_i(I) & \text{o.w.,} \end{cases}$$

where a trigger action, $a^!$, must be recommended in $I^!$ for the deviation to be triggered. A *blind causal deviation* is instead always triggered upon reaching $I^!$ (it does not get to “see” $s_i(I^!)$). An *extensive-form coarse-correlated equilibrium (EFCCCE)* (Farina, Bianchi, and Sandholm 2020) is an equilibrium without beneficial blind causal deviations.

Action deviations: *Action deviations* correspond to *agent-form equilibria* (von Stengel and Forges 2008). They trigger like causal deviations but they resume following the recommended strategy (and thereby re-correlate with the sampled recommendations) after a single action, as if a different agent is in control at each information set. An *informed action deviation* is defined as

$$[\phi s_i](I) = \begin{cases} s_i^!(I) & \text{if } s_i(I^!) = a^! \text{ and } I = I^! \\ s_i(I) & \text{o.w.} \end{cases}$$

and no such deviation is beneficial in an *agent-form correlated equilibrium (AFCE)*. As above, a *blind action deviation* is always triggered upon reaching the trigger information set and its equilibrium concept is *agent-form coarse-correlated equilibrium (AFCCCE)*.

Intuition: For some intuition about these deviation types, imagine a driver traveling through a city with route recommendations. A causal deviation follows the recommendations until they reach a particular part of the city and then they ignore the recommendations for the rest of the trip. An action deviation also follows the recommendations to a particular part of the city but then deviates from the recommendations at a single intersection before following them again. Later, we introduce *counterfactual deviations*, which ignore recommendations until reaching a particular part of the city and then begin following recommendations from there.

³von Stengel (Personal communication 2020) explains that it was intended for players in an EFCE to receive action recommendations at each information set regardless of earlier behavior. However, all work that we are aware of using the EFCE concept assumes the causal deviation structure codified by Dudík and Gordon (2009) where players commit to a fixed strategy after being triggered. This structure arises from von Stengel and Forges (2008)’s discussion of EFCEs in the context of recommendation distributions over reduced strategies that only specify the recommendations at information sets that can be reached by following earlier recommendations. In this case, a player receives arbitrary deterministic “uninformative” recommendations after deviating.




Figure 1: A summary of deviation types and their relationships in EFGs. The game plays out from top to bottom. Straight lines are actions, zigzags are action sequences, and triangles are decision trees. Decisions where recommendations are followed are colored black, alterations to recommendations are colored red, and trigger information is colored cyan. Arrows denote ordering from a stronger to a weaker deviation type (and therefore a subset to superset equilibrium relationship).

For a visual metaphor, see Fig. 1, where we visualize each of these deviation types alongside the previously described external and internal transformations.

Practical Implications

We now illustrate some practical differences between tractable deviation types. These differences are important because algorithms tied to more limited deviation types may not achieve as much reward as ones using stronger deviation types, and the types of mistakes an algorithm makes can depend on the structure of its associated deviation type. Moreover, even if an algorithm is not designed with a deviation type in mind, it may implicitly use one. For example, *policy gradient policy iteration (PGPI)* (Srinivasan et al. 2018a) and Monte Carlo approximations thereof (*i.e.*, standard formulations of policy gradient (Williams 1992; Sutton et al. 2000)) are implicitly tied to action deviations, as we later discuss). Thus, these results have substantial generality and widespread impact.

We summarize our results in a table of relationships between equilibrium concepts (Table 1). The bottom two rows and the two rightmost columns reference new equilibrium concepts that will be explained later.

Since causal deviations can trigger at the start of the game, they cannot be weaker than external ones, and von Stengel and Forges (2008) showed that they are indeed stronger. We provide another example proving this relationship in an extended matching pennies game that we later describe. It is visualized in Fig. 4.

Section 2.4 of von Stengel and Forges (2008) shows how external and causal deviations can outperform action deviations in the *In or Out* game. If the player chooses In, they choose again between In and Out. If they ever choose Out, the game ends and they receive a reward of zero. If they

\downarrow implies \rightarrow	CE	CCE	EF	AF	CF	OS-CF
CE	=	$\dot{=}$	$\dot{=}$	$\dot{=}$	$\dot{=}$	S
CCE	M	=	M	B	B	B
EF	B	$\dot{=}$	=	B	B	B
AF	I	I	I	=	I	I
CF	M	$\dot{=}$	M	M	=	S
OS-CF	M	$\dot{=}$	M	M	$\dot{=}$	=

Table 1: Equilibrium class relationships. The relationships between the coarse-correlated and correlated versions of each EFG equilibrium concept are the same, *e.g.*, the table is identical if “EF” is replaced with “EFCE” or “EFCCE”. Cyan cells highlight where the row concept implies the column concept (*e.g.*, $EF \subseteq CCE$) either by definition ($\dot{=}$) or by being the same ($=$). Otherwise, the cell is colored red and references a game example that proves the row concept does not imply the column concept (*e.g.*, $EF \not\subseteq AF$), where “I” \rightarrow In or Out, “B” \rightarrow extended battle-of-the-sexes, “M” \rightarrow extended matching pennies, and “S” \rightarrow sequential extended battle-of-the-sexes. Emphasized results were previously unclear (italicized entries, *e.g.*, $EF \not\approx AF$) or undefined (bold entries, *e.g.*, $CF \not\approx EF$).

choose In twice, they receive a reward of one. Given the recommendation (Out, Out), an action deviation can swap a single Out to an In, but not both. External and causal deviations can achieve a better value by deviating to (In, In).

However, action deviations can outperform both external and causal deviations in other scenarios, ensuring that some EFCEs and CCEs are not AFCEs. This is a different conclusion than von Stengel and Forges (2008) reach under their original EFCE definition.

Action Deviations can Outperform Causal Deviations

Consider an extended battle-of-the-sexes game where player one first privately chooses between normal or upgraded seating ($\neg U$ or U) for both players at no extra cost before choosing to attend event X or Y simultaneously with their partner. Player one prefers Y and player two prefers X, but if they attend different events, neither is happy so they both get zero. Attending the less preferred event with their partner gives that player +1 with normal seating or +2 with upgraded seating, while the other player, who is at their preferred venue, gets +2 or +3.

The joint recommendations $\{(\neg U, Y | U, Y | \neg U), (Y)\}$ and $\{(\neg U, X | U, X | \neg U), (X)\}$ gives player one +2 and +1, respectively. A distribution that assigns $1/2$ to both of these strategy profiles provides player one with an expected value of +1.5. Since both players coordinate perfectly, player one can improve their value to +2.5 if they would just choose to upgrade their seating. But there is no causal deviation that achieves more than +1.5 because causal deviations must play a fixed strategy after the deviation is triggered. An external deviation must play a fixed strategy from the start of the game, which also precludes them from taking advantage of correlation in the recom-




Figure 2: A beneficial action and counterfactual deviation for player one in an EFCE in extended battle-of-the-sexes. Black lines denote recommendations, red lines denote deviations from unobserved recommendations, and grey lines mark actions that are not recommended or used by the deviation.

mendation distribution. Figures A.2 and A.3 in the appendix enumerate all external and causal deviations in this game to prove this fully.

However, it is a valid action deviation to upgrade seating and preserve event coordination because an action deviation can re-correlate, *i.e.*, follow recommendations at information sets after an earlier deviation. The resulting deviation strategy profile distribution is uniform over $\{(U, Y | U, Y | \neg U), (Y)\}, \{(U, X | U, X | \neg U), (X)\}$, which achieves a value of +2.5. See Fig. 2 for a visualization of this deviation.

The game can be found in OpenSpiel (Lanctot et al. 2019) under the name `extended_bos.efg`. See Appendix A.1 for an illustration of this game’s extensive form, the EFCE recommendation distribution, and all of the possible deviations along with their values.

CFR and Deviation Types

The hindsight rationality framing encourages us to re-examine a prominent regret minimization algorithm for EFGs, *counterfactual regret minimization (CFR)* (Zinkevich et al. 2007b). CFR was designed to solve two-player, zero-sum games by minimizing external regret, but is it incidentally hindsight rational for causal or action deviations? Or do we need a new type of deviation to fully understand CFR?

Failure on Causal and Action Deviations

We present an extension of Shapley’s game (Shapley 1964) where CFR fails to behave according to an EFCCE or an AFCCE and therefore does not necessarily eliminate incentives for causal or action deviations.

In Shapley’s game, both players simultaneously choose between Rock, Paper, and Scissors. Rock beats Scissors, Scissors beats Paper, and Paper beats Rock, but both players lose if they choose the same item. A winning player gets +1 while losing players get -1 . Our extension is that player one privately predicts whether or not player two will choose Rock after choosing their action. If they accurately predict a Rock play, they receive a bonus, b , in addition to their usual reward, and if they accurately predict that they will not play Rock, they receive a smaller bonus, $b/3$. There is no cost for inaccurate predictions, and the second player’s decisions and payoffs are unchanged. The game can




Figure 3: The gap between CFR’s self-play empirical distribution and an extensive-form or agent-form (C)CE (E/AF(C)CE) in the extended Shapley’s game with $b = 0.003$. (Left) simultaneous-update CFR. (Right) alternating-update (Burch, Moravcik, and Schmid 2019) CFR.

be found in OpenSpiel (Lanctot et al. 2019) under the name `extended_shapleys.efg`. The game’s extensive-form is drawn in Figure A.5.

Fig. 3 shows the gap between the expected utility achieved by CFR’s self-play empirical distribution (summed across players) and an optimal causal or action deviation across iterations. In all of our experiments, optimal causal and action deviations achieved the same value. The E/AFCE lines correspond to an optimal informed deviation while the E/AFCE lines correspond to an optimal blind deviation. Notice that in all figures, the gap does not continue to decrease over time as we would expect if CFR were to minimize causal or action regret. See Figure A.6 for experiments with two other bonus values (0.3 and 30).

Counterfactual Deviations

While CFR can fail to minimize causal or action regret, it does more than minimize external regret. We now define blind and informed *counterfactual deviations* to describe the strategy deviations that arise from the definition of counterfactual value. We show that conventional CFR with no-external-regret learners at each information set is hindsight rational for blind counterfactual deviations and that CFR with no-internal-regret learners is hindsight rational for informed counterfactual deviations. Naturally, both types of counterfactual deviations give rise to classes of counterfactual equilibria and CFR’s behavior in self-play therefore conforms to such equilibria.

Intuitively, a counterfactual deviation is one that plays from the beginning of the game to reach a particular information set and re-correlates thereafter. Counterfactual deviations re-correlate like action deviations but they are still able to deviate across sequential information sets. This ensures that the set of counterfactual deviations contains the set of external deviations. Both counterfactual deviation variants are visualized in Fig. 1.

We briefly review CFR’s definition and known theoretical properties before formally defining counterfactual deviations and equilibria and stating CFR’s relationships with these concepts.

Background: The *counterfactual value function* at I , $v_I : \mathcal{A}(I) \times \Pi \rightarrow \mathbb{R}$, is the immediate expected value for taking action a in I assuming that everyone plays according to π

except that i plays to reach I , i.e.,

$$v_I(a; \pi) = \sum_{\substack{h \in I; \\ z \in \mathcal{Z}}} \underbrace{P(h; \pi_{-i})}_{\text{Prob. of being in } h.} \underbrace{P(ha, z; \pi)u_i(z)}_{\text{Future value given } ha.}$$

We overload $v_I(\sigma; \pi) \doteq \mathbb{E}_{a \sim \sigma}[v_I(a; \pi)]$ for immediate strategies $\sigma \in \Delta^{|\mathcal{A}(I)|}$, and $v(I; \pi) \doteq v_I(\pi_i(I); \pi)$ for the counterfactual value of I .

CFR is an online algorithm so it proposes a strategy on each round t , π_i^t , is evaluated based on strategies for the other players, π_{-i}^t , and uses this feedback to propose its next strategy, π_i^{t+1} . CFR’s strategies are determined by no-external-regret learners at each I that produce immediate strategies $\pi_i^t(I) \in \Delta^{|\mathcal{A}|}$. These learners are trained on *immediate counterfactual regret* (Zinkevich et al. 2007b), i.e., $\rho_I^t(\phi) = v_I(\phi\pi_i^t(I); \pi^t) - v(I; \pi^t)$. The key result of Zinkevich et al. (2007b) is that CFR’s external regret decreases at the same rate as its immediate counterfactual regret, i.e.,

Theorem 1. (Zinkevich et al. 2007b, Theorem 3) *Let $R_I^T(\Phi_{\mathcal{A}(I)}^{\text{EX}}) = \max_{\phi \in \Phi_{\mathcal{A}(I)}^{\text{EX}}} \rho_I^{1:T}(\phi)$ be the maximum immediate counterfactual regret at each information set I . Then the incentive to deviate from a sequence of T strategies for player i , $(\pi_i^t)_{t=1}^T$, to any other pure strategy is $\frac{1}{T} R^T(\Phi_{S_i}^{\text{EX}}) \leq \frac{1}{T} \sum_{I \in \mathcal{I}_i} \left(R_I^T(\Phi_{\mathcal{A}(I)}^{\text{EX}}) \right)^+$.*

Theorem 1 follows directly from an elementary decomposition relationship that we build on (see Appendix B.1 for a simple proof):

Lemma 1. (Zinkevich et al. 2007a, Equation 13, Lemma 5) *The full regret, $R_I^T(\Phi_{S_i}^{\text{EX}})$, from information set $I \in \mathcal{I}_i$ that $(\pi_i^t)_{t=1}^T$ suffers is upper bounded by $R_I^T(\Phi_{\mathcal{A}(I)}^{\text{EX}}) + \max_{a \in \mathcal{A}(I)} \sum_{I' \in \mathcal{I}_i(I, a)} R_{I'}^T(\Phi_{S_i}^{\text{EX}})$, where $\mathcal{I}_i(I, a) \subset \mathcal{I}_i$ are the information sets that immediately follow after taking action a in I .*

Counterfactual Deviations and Intermediate Counterfactual Regret: Lemma 1 provides a one-step recursive connection between full counterfactual regret and immediate counterfactual regret. When we unroll this recursion completely from the start of the game, we arrive at Theorem 1. But what if instead we unroll this recursion an arbitrary number of steps?

Perfect recall ensures that there is a unique sequence of player i ’s information sets and actions leading to any target information set, I^\odot . Thus, if this target is n steps away from an initial information set, I_0 , then we can refer to the sequence of intermediate information sets as $(I_j)_{j=0}^{n-1}$ and the sequence of actions required to follow this sequence as $(a_j)_{j=0}^{n-1}$. Once at the target, we could consider using any action transformation ϕ^\odot . What we have just described is a strategy deviation that plays to reach a target information set and then applies an action transformation, but otherwise leaves the strategy unmodified. Denoting this strategy deviation as ϕ , we can write the counterfactual value it achieves

in terms of the usual one-step counterfactual values:

$$v(I_0; \phi(\pi_i^t), \pi_{-i}^t) = v_{I^\circ}(\phi^\circ(\pi_i^t(I^\circ)); \pi^t) + \sum_{j=0}^{n-1} \underbrace{v_{I_j}(a_j; \pi^t) - v_{I_{j+1}}(\pi^t)}_{\text{telescoping difference}}.$$

The counterfactual regret with respect to this value is intermediate between an immediate counterfactual regret where $n = 1$ and full counterfactual regret where all terminal information sets are treated as targets. With regret decomposition, we can bound this *intermediate counterfactual regret* in terms of immediate counterfactual regrets:

Theorem 2. *Let ϕ be the deviation that plays to reach I° from initial information set I_0 and uses action deviation $\phi^\circ \in \Phi^\circ \subseteq \Phi_{\mathcal{A}(I^\circ)}^{\text{sw}}$ once there, but otherwise leaves its given strategy unmodified. The instantaneous intermediate counterfactual regret with respect to ϕ is $v(I_0; \phi(\pi_i^t), \pi_{-i}^t) - v(I_0; \pi^t)$ and the cumulative intermediate counterfactual regret is bounded as*

$$\sum_{t=1}^T v(I_0; \phi(\pi_i^t), \pi_{-i}^t) - v(I_0; \pi^t) \leq R_{I^\circ}^T(\Phi^\circ) + \sum_{I^\circ \preceq I < I^\circ} R_I^T(\Phi_{\mathcal{A}(I)}^{\text{ex}}).$$

If the initial information set is one at the start of the game, we call ϕ a *counterfactual deviation*. We distinguish between two major variants—*blind counterfactual deviations* that use external deviations at the target information set and *informed counterfactual deviations* that use internal deviations there instead. We visualize both in Fig. 1 and now provide formal definitions.

Definition 1. *A blind counterfactual deviation is defined by a pair, (I°, a°) , where $I^\circ \in \mathcal{I}_i$ is a target information set that the deviation plays to reach deterministically from the start of the game and a° is the action taken at I° . The deviation follows recommendations at every other information set. Formally,*

$$[\phi s_i](I(h)) = \begin{cases} a & \text{if } \exists h' \in I^\circ, a, ha \sqsubseteq h'a^\circ \\ s_i(I(h)) & \text{o.w.} \end{cases}$$

Definition 2. *An informed counterfactual deviation is defined by a triple, (I°, a', a°) , where $I^\circ \in \mathcal{I}_i$ is a target information set that the deviation plays to reach deterministically from the start of the game. If a' is the recommendation at I° , then a° is played, otherwise the recommendation is followed. The deviation follows recommendations at every other information set. Formally,*

$$[\phi s_i](I(h)) = \begin{cases} a & \text{if } \exists h' \in I^\circ, a, ha \sqsubseteq h' \\ a^\circ & \text{if } h \in I^\circ, s_i(I^\circ) = a' \\ s_i(I(h)) & \text{o.w.} \end{cases}$$

The following conclusions immediately follow from Theorem 2, Definition 1, and the definition of CFR:

Corollary 1. *CFR is no-regret/hindsight rational with respect to blind counterfactual deviations. Its blind counterfactual deviation regret bound is the same as that for external deviations (see, e.g., Theorem 1).*

Corollary 2. *CFR with no-internal-regret learners is no-regret/hindsight rational with respect to informed counterfactual deviations. Its cumulative informed deviation regret is upper bounded by*

$$\sum_{I \in \mathcal{I}_i} \left(R_I^T \left(\Phi_{\mathcal{A}(I)}^{\text{in}} \cup \Phi_{\mathcal{A}(I)}^{\text{ex}} \right) \right)^+.$$

The counterfactual deviations we have just described correspond to two new equilibrium concepts:

Definition 3. *A recommendation distribution is a counterfactual coarse-correlated equilibria (CFCCE) if there are no beneficial blind counterfactual deviations.*

Definition 4. *A recommendation distribution is a counterfactual correlated equilibria (CFCE) if there are no beneficial informed counterfactual deviations.*

See the CF row and column in Table 1 for a summary of how counterfactual equilibria relate to other equilibrium concepts.

Practical properties: Since counterfactual deviations can re-correlate, they achieve the same improved value in the extended battle-of-the-sexes game as action deviations. Counterfactual deviations also subsume external deviations, which is not true of action deviations, so the In or Out game does not pose a problem either. However, counterfactual deviations are limited in that they cannot correlate before deviating.

Consider an extended version of matching pennies where player one privately decides whether their goal is to Match (M) with other player or to Not Match (\neg M). If player one achieves their goal, they get +1 while the other player gets -1 and *vice-versa* otherwise.

The joint recommendations $\{(\neg\text{M}, \text{H} | \text{M}, \text{T} | \neg\text{M}), (\text{H})\}$ give player one +1 while $\{(\text{M}, \text{H} | \text{M}, \text{T} | \neg\text{M}), (\text{T})\}$ yields -1 , for an average of zero. Since this time the decision holding player one back is the second action in the sequence at the Match information set where they should always play T, there is no better counterfactual deviation. But deviating to T in the Match information set, while following the Match or Not Match recommendation is a valid causal or action deviation. The resulting deviation strategy profile distribution is uniform over $\{(\neg\text{M}, \text{T} | \text{M}, \text{T} | \neg\text{M}), (\text{H})\}, \{(\text{M}, \text{T} | \text{M}, \text{T} | \neg\text{M}), (\text{T})\}$, which achieves a value of +1. See Fig. 4 for a visualization of this deviation.

The game can be found in OpenSpiel (Lanctot et al. 2019) under the name `extended.mp.efg`. See Appendix A.3 for an illustration of this game's extensive form, the CFCE recommendation distribution, and all of the possible deviations along with their values.

From this example, we can see why a causal or action deviation could outperform a counterfactual deviation in the extended Shapley's game. Under the usual Shapley dynamics where players settle into a cycle that




Figure 4: A beneficial causal and action deviation for player one in a CFCE in extended matching pennies. Black lines denote recommendations, red lines denote deviations from unobserved recommendations, and grey lines mark actions that are not recommended or utilized by the deviation.

puts 1/6 probability on Rock–Paper, Rock–Scissors, Paper–Rock, Paper–Scissors, Scissors–Rock, Scissors–Paper while avoiding Rock–Rock, Paper–Paper, and Scissors–Scissors, it is beneficial to predict Rock whenever you do not play Rock yourself. However, counterfactual deviations cannot condition behavior at one information set according to play at a previous information set by definition, thereby preventing any counterfactual deviation from taking advantage of this fact. This also explains why the performance of optimal causal and action deviations is the same in this game: the only causal deviations that are not counterfactual deviations in a two action game are in fact action deviations.

To be clear, it is not that a player considering a counterfactual deviation forgets that they are in an information set where they have chosen to play Rock. It is instead that their appraisal of whether to predict Rock or not Rock is based on the average utility they receive for that prediction given that the other player chooses Rock, Paper, and Scissors equally often, which is what the player would observe if they were to always choose Rock. Counterfactual deviations do not take into account the correlation information that may be present in previous decisions. CFR fails to minimize regret for action or causal deviations in this case because it evaluates its actions in exactly the same way through its use of counterfactual values and regrets.

Observable Sequential Rationality

It turns out that the equilibrium class associated with CFR is smaller still than CFCCE. So far, we have examined properties of CFR that emerge from regret decomposition, but we have only used bounds on regrets that begin at the start of the game. Regret decomposition, however, applies at every subtree within an information set forest in a way that resembles sequential rationality (Kreps and Wilson 1982). We introduce *observable sequential rationality* and *observable sequential equilibrium* to define CFR’s refinement of CFCCE.

Sequential rationality is based on examining the reward that follows from a particular information set I ,

$$u_i(\pi | \xi_I) \doteq \mathbb{E}_{h \sim \xi_I} \left[\sum_{\substack{z \supseteq h \\ z \in \mathcal{Z}}} P(h, z; \pi) u_i(z) \right],$$

where histories are distributed according to a belief distribution, $\xi_I : I \rightarrow [0, 1]$. An assessment, that is, a system of beliefs and strategy pair, (ξ, π_i) , is sequentially rational if $u_i(\phi(\pi_i, \pi_{-i}) | \xi_I) - u_i(\pi | \xi_I) \leq 0$ for each information set $I \in \mathcal{I}_i$ and external deviation ϕ .

Different types of equilibria have been proposed that reuse the optimality constraint described by sequential rationality but construct systems of beliefs in different ways. Typically, beliefs must be consistent, which means that each belief distribution is defined as

$$\xi_I : h \rightarrow \frac{P(h; \pi_{-i})}{\sum_{h' \in I(h)} P(h'; \pi_{-i})},$$

as derived by applying Bayes’ rule according to π_{-i} .⁴ However, if I is unreachable, *i.e.*, the probability of reaching each history in I is zero, then ξ_I is undefined and it is unclear what beliefs should be held at I . The method of constructing ξ_I for unreachable positions differentiates various types of sequentially-refined equilibria. Instead, we reuse the notion of sequential rationality under consistent beliefs without constructing beliefs at unreachable information sets at all.

A natural extension to the mediated execution setting is to define a sequentially rational recommendation distribution, $\mu \in \Delta^{|\mathcal{S}|}$, as one where the benefit of a deviation according to a system of beliefs ξ^{s-i} consistent with mediator-sampled strategies $s_{-i} \in \mathcal{S}_{-i}$ is non-positive,

$$\mathbb{E}_{s \sim \mu} [u_i(\phi(s_i), s_{-i}) | \xi_I] - u_i(s | \xi_I) \leq 0, \quad (2)$$

at every information set. Notice that at every reachable information set, Eq. (2) differs from full counterfactual regret only by the Bayes normalization constant (Srinivasan et al. 2018a), since counterfactual values are weighted by the probability that the other players play to each history. At unreachable information sets, the counterfactual value function is always zero so a sequential rationality-like constraint based on counterfactual values is weaker only in that all potential behavior is acceptable at unreachable positions. However, since an unreachable position corresponds to a situation that was never experienced, the difference to a hindsight analysis is entirely inconsequential. We call this condition *observable sequential rationality* because it is indistinguishable from sequential rationality under Bayesian beliefs according to all outcomes that could be observed while playing with players who always follow their recommendations.

Definition 5. A recommendation distribution, $\mu \in \Delta^{|\mathcal{S}|}$, is *observably sequentially rational with respect to a set of deviations*, $\Phi \subseteq \Phi_{\mathcal{S}_i}^{\text{sw}}$, if and only if the maximum benefit for every deviation, $\phi \in \Phi$, according to the counterfactual value function at every information set, $I \in \mathcal{I}_i$, is non-positive,

$$\max_{\phi \in \Phi} \mathbb{E}_{s \sim \mu} [v(I; \phi(s_i), s_{-i}) - v(I; s)] \leq 0.$$

Two corollaries of Theorem 2 and Definition 5 are:

Corollary 3. CFR is *observably sequentially hindsight rational with respect to blind counterfactual deviations*. Its

⁴Note that π_i cannot skew the history distribution at player i ’s information set because of perfect recall.

empirical play approaches exact rationality at the same rate as its average blind counterfactual regret vanishes, which by Corollary 1 is the same as its average external regret bound (see, e.g., Theorem 1).

Corollary 4. *CFR in self-play converges to an observable sequential CFCCE at the same rate as its average blind counterfactual regret vanishes, which by Corollary 1 is the same as its average external regret bound (see, e.g., Theorem 1).*

An analogous result applies to CFR with no-internal-regret learners, informed counterfactual deviations, and observable sequential CFCE. Corollaries 3 and 4 provide the most thorough characterization of CFR’s behavior (both for a single player and in self-play) to date and both apply to multi-player, general sum games.

To show how observable sequential rationality differs from non-sequential rationality in practice, we provide an example in Appendix A.4 where a CE is not an observable sequential CFCCE. See the OS-CF row and column in Table 1 for a summary of how observable sequential counterfactual equilibria relate to the other equilibrium concepts.

CFR for Action Deviations

We showed that CFR does not behave according to an EFCCE or AFCCE, but could we modify CFR to do so? Modifying CFR to address causal deviations requires recently developed non-trivial modifications (Celli et al. 2020), but we can easily modify CFR for action deviations by weighting counterfactual regrets by reach probabilities.

Since by definition an action deviation applies an action transformation at a single trigger information set, the benefit of such a deviation is simply the average reach-probability-weighted immediate regret, no decomposition required. Therefore, employing no-regret learners on the reach-probability-weighted counterfactual value functions in each information set ensures that the average reach-probability-weighted immediate regret becomes non-positive.

Theorem 3. *The CFR-like algorithm that trains a no-regret learner at each information set on the reach-probability-weighted counterfactual value functions, $\hat{v}_I(\cdot; \pi^t) : a \rightarrow P(I; \pi_i^t)v_I(a; \pi^t)$, is no-regret/hindsight rational with respect to action deviations. If the learners minimize external regret, then the algorithm minimizes blind action deviation regret, and if the learners minimize internal regret, then the algorithm minimizes informed action deviation regret. At all times, the algorithm’s action deviation regret cannot be more than the maximum regret suffered by any single learner.*

Celli et al. (2020)’s algorithm is also no-regret with respect to informed action deviations because it uses no-internal-regret learners trained on reach-probability-weighted immediate counterfactual values. Theorem 3 gives an algorithm that is weaker in that it does not minimize causal, counterfactual, or external regret, but it does not require action sampling, nor does it require multiple learners at every information set. It remains to investigate this trade-off experimentally.

PGPI and *actor-critic policy iteration (ACPI)* describe the optimization problems underlying standard policy gradient and actor-critic methods (Srinivasan et al. 2018a). (Euclidean) projected PGPI and ACPI using tabular representations are equivalent to *state-local generalized infinitesimal gradient ascent (GIGA(s))* (Srinivasan et al. 2018b, Definition 2) and are no-external-regret with respect to reach-probability-weighted counterfactual value functions at each information set (Srinivasan et al. 2018b, Lemma 4). These algorithms are also therefore hindsight rational with respect to blind action deviations. Thus, Theorem 3 establishes a new connection between CFR and reinforcement learning algorithms via action deviations, PGPI, and ACPI.

Conclusions

Complicated general-sum, multi-agent games, perhaps unsurprisingly, resist static solutions. Learning online and analyzing agent behavior in hindsight is a dynamic alternative. This hindsight rationality view returns equilibria to a descriptive role, as originally introduced within the field of game theory, rather than the prescriptive role it often holds in the field of artificial intelligence. Agents prescriptively aim to reduce regret for deviations in hindsight and thereby learn to behave rationally over time, while equilibria describe the consequences of interactions between such agents.

The mediated equilibrium and deviation landscapes of EFGs are particularly rich because there is substantial space for deviation types that have intermediate power and computational requirements between external and internal deviations. To develop these EFG deviation and equilibrium types, we re-examined causal and action deviations along with their corresponding equilibrium concepts, extensive-form and agent-form equilibria. We showed how action deviations can outperform causal deviations, which dispels a common misunderstanding that an EFCE is always an AFCE.

We showed that CFR’s empirical play does not converge toward an EFCCE or AFCCE in self-play and thus CFR is not hindsight rational for causal or action deviations. Instead, we defined blind and informed counterfactual deviations, and observable sequential rationality to more precisely characterize CFR’s behavior. CFR is observably sequentially hindsight rational with respect to blind counterfactual deviations and CFR with internal learners has the same property with respect to informed counterfactual deviations. In self-play, these algorithms converge toward observable sequential CFCCEs or CFCEs, respectively.

Table 1 summarizes all of the equilibrium relationships investigated in this work.

One avenue for future work is to modify CFR to simultaneously handle each type of deviation and show the practical benefits of such an algorithm. Additionally, investigating observable sequential rationality with non-counterfactual deviations is another avenue for interesting future study opened up by our work.

Acknowledgments

Dustin Morrill and Michael Bowling are supported by the Alberta Machine Intelligence Institute (Amii) and NSERC. James Wright is supported by a Canada CIFAR AI Chair at Amii. Amy Greenwald is supported in part by NSF Award CMMI-1761546. Thanks to Bernhard von Stengel for answering our questions about the EFCE concept. Thanks to anonymous reviewers and Julien Perolat for constructive comments and suggestions.

References

- Aumann, R. J. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1(1): 67–96.
- Brown, N.; and Sandholm, T. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science* 359(6374): 418–424.
- Brown, N.; and Sandholm, T. 2019. Superhuman AI for multiplayer poker. *Science* 365(6456): 885–890.
- Burch, N.; Moravčík, M.; and Schmid, M. 2019. Revisiting CFR+ and alternating updates. *Journal of Artificial Intelligence Research* 64: 429–443.
- Celli, A.; Marchesi, A.; Bianchi, T.; and Gatti, N. 2019. Learning to correlate in multi-player general-sum sequential games. In *Advances in Neural Information Processing Systems*, 13076–13086.
- Celli, A.; Marchesi, A.; Farina, G.; and Gatti, N. 2020. No-regret learning dynamics for extensive-form correlated equilibrium. *Advances in Neural Information Processing Systems* 33.
- Dudík, M.; and Gordon, G. J. 2009. A Sampling-Based Approach to Computing Equilibria in Succinct Extensive-Form Games. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI-2009)*.
- Farina, G.; Bianchi, T.; and Sandholm, T. 2020. Coarse Correlation in Extensive-Form Games. In *Thirty-Fourth AAAI Conference on Artificial Intelligence, February 7-12, 2020, New York, New York, USA*.
- Forges, F.; and von Stengel, B. 2002. Computationally Efficient Coordination in Games Trees. THEMA Working Papers 2002-05, THEMA (THéorie Economique, Modélisation et Applications), Université de Cergy-Pontoise.
- Foster, D. P.; and Vohra, R. 1999. Regret in the on-line decision problem. *Games and Economic Behavior* 29(1-2): 7–35.
- Foster, D. P.; and Vohra, R. V. 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior* 21(1-2): 40.
- Greenwald, A.; Jafari, A.; and Marks, C. 2003. A general class of no-regret learning algorithms and game-theoretic equilibria. In *Proceedings of the 2003 Computational Learning Theory Conference*, 1–11.
- Hannan, J. 1957. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games* 3: 97–139.
- Hart, S.; and Mas-Colell, A. 2000. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica* 68(5): 1127–1150.
- Kreps, D. M.; and Wilson, R. 1982. Sequential equilibria. *Econometrica* 50(4): 863–894.
- Kuhn, H. W. 1953. Extensive Games and the Problem of Information. *Contributions to the Theory of Games* 2: 193–216.
- Lanctot, M.; Lockhart, E.; Lespiau, J.-B.; Zambaldi, V.; Upadhyay, S.; Pérolat, J.; Srinivasan, S.; Timbers, F.; Tuyls, K.; Omidshafiei, S.; Hennes, D.; Morrill, D.; Muller, P.; Ewalds, T.; Faulkner, R.; Kramár, J.; Vylder, B. D.; Saeta, B.; Bradbury, J.; Ding, D.; Borgeaud, S.; Lai, M.; Schrittwieser, J.; Anthony, T.; Hughes, E.; Danihelka, I.; and Ryan-Davis, J. 2019. OpenSpiel: A Framework for Reinforcement Learning in Games. *CoRR* abs/1908.09453. URL <http://arxiv.org/abs/1908.09453>.
- Moravčík, M.; Schmid, M.; Burch, N.; Lisý, V.; Morrill, D.; Bard, N.; Davis, T.; Waugh, K.; Johanson, M.; and Bowling, M. 2017. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science* 356(6337): 508–513.
- Moulin, H.; and Vial, J.-P. 1978. 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.
- Nash, J. 1951. Non-Cooperative Games. *The Annals of Mathematics* 54(2): 286–295.
- Selten, R. 1974. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *Economics* .
- Shapley, L. 1964. Some Topics in Two-Person Games. In *Advances in Game Theory*. Princeton University Press.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* 529: 484–489.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; Lillicrap, T.; Simonyan, K.; and Hassabis, D. 2018. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* 362(6419): 1140–1144.
- Silver, D.; Hubert, T.; Schrittwieser, J.; Antonoglou, I.; Lai, M.; Guez, A.; Lanctot, M.; Sifre, L.; Kumaran, D.; Graepel, T.; Lillicrap, T. P.; Simonyan, K.; and Hassabis, D. 2017a. Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. *CoRR* abs/1712.01815.
- Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; Chen, Y.; Lillicrap, T.; Hui, F.; Sifre, L.; van den Driessche, G.; Graepel, T.; and Hassabis, D. 2017b. Mastering the game of Go without human knowledge. *Nature* 550: 354–359.

Srinivasan, S.; Lanctot, M.; Zambaldi, V.; Pérolat, J.; Tuyls, K.; Munos, R.; and Bowling, M. 2018a. Actor-Critic Policy Optimization in Partially Observable Multiagent Environments. In *Advances in Neural Information Processing Systems*.

Srinivasan, S.; Lanctot, M.; Zambaldi, V.; Pérolat, J.; Tuyls, K.; Munos, R.; and Bowling, M. 2018b. Actor-Critic Policy Optimization in Partially Observable Multiagent Environments. *CoRR* abs/1810.09026.

Sutton, R. S.; McAllester, D.; Singh, S.; and Mansour, Y. 2000. Policy Gradient Methods for Reinforcement Learning with Function Approximation. In *Advances in Neural Information Processing Systems 12*, 1057–1063. MIT Press.

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; Mathieu, M.; Dudzik, A.; Chung, J.; Choi, D. H.; Powell, R.; Ewalds, T.; Georgiev, P.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* 575(7782): 350–354.

von Stengel, B.; and Forges, F. 2008. Extensive-form correlated equilibrium: Definition and computational complexity. *Mathematics of Operations Research* 33(4): 1002–1022.

Williams, R. J. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8(3): 229–256. ISSN 1573-0565. doi:10.1007/BF00992696. URL <https://doi.org/10.1007/BF00992696>.

Zinkevich, M.; Johanson, M.; Bowling, M.; and Piccione, C. 2007a. Regret Minimization in Games with Incomplete Information. Technical Report TR07-14, University of Alberta.

Zinkevich, M.; Johanson, M.; Bowling, M. H.; and Piccione, C. 2007b. Regret Minimization in Games with Incomplete Information. In *Advances in Neural Information Processing Systems*, 1729–1736.