A Unified Framework for Extensive-Form Game Abstraction with Bounds

Christian Kroer and Tuomas Sandholm

Carnegie Mellon University, Pittsburgh PA 15213, USA
ckroer,sandholm@cs.cmu.edu

Abstract. Abstraction has long been a key component in the practical solving of large-scale extensive-form games. Despite this, abstraction remains poorly understood. There have been some recent theoretical results but they have been confined to specific assumptions on abstraction structure and are specific to various specific disjoint types of abstraction, and specific solution concepts, for example, exact Nash equilibria or strategies with bounded immediate regret. In this paper we present a unified framework for analyzing abstractions that can express all types of abstractions and solution concepts used in prior papers with performance guarantees—while maintaining comparable bounds on abstraction quality. Moreover, our framework extends well beyond prior work. We present the first exact decomposition of abstraction error for a broad class of abstractions that encompasses abstractions used in practice. Because it is significantly more general, this decomposition has a stronger dependence on the specific strategy computed in the abstraction. We show that this dependence can be removed by making similar, though slightly weaker, assumptions than in prior work. We also show, via counterexample, that such assumptions are necessary for some games. Finally, we prove the first bounds for how ϵ-Nash equilibria computed in abstractions perform in the original game. This is important because often one cannot afford to compute an exact Nash equilibrium in the abstraction. All our results apply to general-sum n-player games.

Keywords: Extensive-form game · equilibrium finding· Nash equilibrium · abstraction · counterfactual regret minimization.

1 Introduction

Game-theoretic equilibria have played a key role in several recent advances in the ability to construct AIs with superhuman performance in games with imperfect information [5, 10, 32]. In particular these results rely on computing an approximate *Nash equilibrium* [33] for the game at hand. In typical real-world situations these games are so large that even approximate equilibria are intractable. Instead, the dominant paradigm has been to first construct some smaller *abstraction* of the game, apply an iterative algorithm for computing a Nash equilibrium in the abstraction, and map the resulting strategy back to the full game. This approach was used in the recent *Libratus* agent, which beat four top poker pros in the game of heads-ups no-limit Texas hold'em [10] (in addition to abstraction and equilibrium approximation the agent also utilized real-time subgame solving [9] and action abstraction refinement). Abstraction has also been used in trading-agent competitions [39] and security games [3, 1, 2].

In practice, abstractions are generated heuristically with no theoretical guarantees on solution quality [36, 4, 15, 16, 19, 20, 18, 21, 22, 24, 14, 6, 34]. Ideally, abstraction would be lossless, such that implementing an equilibrium from the abstract game results in an equilibrium in the full game. Gilpin and Sandholm [17] study lossless abstraction techniques for a structured class of games. Unfortunately, lossless abstraction often leads to games that are still too large to solve. Thus, one must turn to lossy abstraction. However, significant abstraction *pathologies* (*nonmonotonicities*) have been shown in games which cannot exist in single-agent settings: if an abstraction is refined, the equilibrium strategy from that new abstraction can be worse in the original game than the equilibrium strategy from a coarser abstraction [37]! Lossy abstraction remains poorly understood from a theoretical perspective. Results have been obtained only for various restricted models of abstraction. Basilico and Gatti [3] give bounds for the special game class called *patrolling security game*. Sandholm and Singh [35] provide lossy abstraction algorithms with bounds for stochastic games. Brown and Sandholm [7], Waugh et al. [38], Brown and Sandholm [8], and Čermák et al. [12] develop iterative abstraction-refinement schemes that have various forms of converge guarantees but they do not give solution-quality guarantees for the original game for strategies computed in limited-size abstractions.

Results which are for *extensive-form games (EFGs)* are most related to this work. Lanctot et al. [30] show that the *counterfactual regret minimization algorithm (CFR)* converges to an approximate NE when run on an imperfect-recall abstraction that is a *skew well-formed game* (SWF) with respect to the original game, where the error in the NE has a linear dependence on the number of information sets. Kroer and Sandholm [27] show that Nash equilibria and strategies with bounded counterfactual regret computed in *chance-relaxed SWF* (CRSWF) (a generalization of SWF that allows error in chance outcomes) are approximate NE in the original game, with a linear dependence on game-tree height. Kroer and Sandholm [25] show that NE computed in perfect-recall abstractions that satisfy conditions that are similar to those in CRSWF abstractions are approximate NE in the original game with a constant dependence on payoff error (as opposed to a linear dependence on height in Kroer and Sandholm [27] or linear dependence on information sets in Lanctot et al. [30]). Kroer and Sandholm [26] extend the results of Kroer and Sandholm [25] to continuous action spaces.

The results in the previous paragraph are all for disparate models of abstraction, a specific solution concept, or specific algorithm. Yet they share a common structure on the assumptions needed in order to obtain theoretical results. They assume that information sets (i.e., decision points) are aggregated into larger information sets. All pairs of information sets that are aggregated together are compared by defining a mapping between subtrees under the information sets. This mapping then requires that the payoffs are similar, the distribution over chance outcomes is similar, and for pairs of leaves mapped to each other, the leaves have the same sequence of information-set-action pairs leading to them in the abstraction. Payoff and chance-outcome similarity is similar to what good practical abstraction algorithms seek to obtain. However, the requirement that information-set-action pairs are the same for leaf nodes mapped to each other is not satisfied by the best heuristic abstraction algorithms used in practice [24, 14, 6]. In this paper we develop an exact decomposition of the solution-quality error that does not require any such assumption. This is the first decomposition of solution-quality error resulting from abstraction. This decomposition depends on several quantities that prior results did not (owing to its more general and exact nature). We then show that by making a weaker variant of previous assumptions, our decomposition can recover all previous solution-quality bounds. We show via counterexample that there exist games where the assumption on information-set-action pairs is, in a sense, necessary in order to avoid large abstraction error that is not measurable by the type of technique presented here and in prior work.

Finally, we prove the first bounds for how ϵ-Nash equilibria computed in abstractions perform in the original game. This is important because often one cannot afford to compute an exact Nash equilibrium in the abstraction. All our results apply to general-sum n-player games.

2 Extensive-form games (EFGs)

An *extensive-form game (EFG)* is a game tree, where each node in the tree corresponds to some history of actions taken by the players. Each node belongs to some player, and the actions available to the player at a given node are represented by the branches. Uncertainty is modeled by having a special player, *Chance*, that moves with some predefined fixed probability distribution over actions. EFGs model imperfect information by having groups of nodes in *information sets*, where an information set is a group of nodes all belonging to the same player such that the player cannot distinguish among them. In the original game that we are trying to solve, we assume *perfect recall*, which requires that no player forgets information they knew earlier in the game. This is a natural condition since you generally cannot force players to forget information, and it would not be in their interest to do so. Formally, an *extensive-form game* Γ is a tuple $(H, Z, A, P, \pi_0, \{\mathcal{I}_i\}, \{u_i\})$. H is the set of nodes in the game tree, corresponding to sequences (or histories) of actions. H_i is the subset of histories belonging to Player i. $Z \subseteq H$ is the set of terminal histories, or *leaves*. A is the set of actions in the game. A_I denotes the set of actions available at nodes in information set I. P, the player function, maps each non-terminal history $h \in H \setminus Z$ to $\{0, \dots, n\}$, representing the player whose turn it is to move after history h. If $P(h) = 0$, the player is Chance. π_0 is a function that assigns to each $h \in H_0$ the probability of reaching h due to Chance (i.e., assuming that both players play to reach h). An information set $\mathcal{I}_i$, for $i \in \{1, \dots, n\}$, is a partition of $\{h \in H : P(h) = i\}$.

The utility function u_i maps $z \in Z$ to the utility obtained by player i when the terminal history is reached.

A *behavioral* strategy σ_i for a player i is a probability distribution over actions at each information set in $\mathcal{I}_i$. A *strategy profile* σ is a behavioral strategy for each player. The probability that σ puts on $a \in A_I$ is denoted $\sigma(I, a)$. We let $\pi^\sigma(z)$ and $\pi^\sigma(I)$ denote the probability of reaching z and I respectively, if players choose actions according to σ. We likewise let $\pi^\sigma(z|I)$ and $\pi^\sigma(\hat{I}|I)$ denote the reach probabilities conditioned on being at information set I. For a given strategy profile σ we let $\sigma_{I\to a}$ denote the same strategy except that $\sigma_{I\to a}(I, a) = 1$.

We will often quantify statements over the set of leaves or information sets that are reachable from some given information set I belonging to Player i, sometimes conditioned on taking a specific action $a \in A_I$. We let $\mathcal{Z}_I, \mathcal{D}_I \subset \mathcal{I}_i$ be the set of leaves and information sets reachable conditioned on being at information set I. We let Z_I and $\mathcal{C}_I \subset \mathcal{I}_i$ be the set of leaves and information sets that are reachable without Player i taking any further actions before reaching them. We let $\mathcal{Z}_I^a, \mathcal{D}_I^a, Z_I^a$ and $\mathcal{C}_I^a$ be defined analogously but conditioned on taking action $a \in A_I$.

As is usual we use the subscript $-i$ to denote exclusion of Player i, for example, σ_{-i} is the set of behavioral strategies in σ except for the strategy of Player i, and $\pi_{-i}^\sigma(z)$ is the probability of reaching leaf node z disregarding actions taken by Player i, that is, assuming that Player i plays to reach z.

3 Game abstractions

We consider abstractions that are themselves EFGs, but we do not require abstractions to have perfect recall (the leading practical abstractions are of imperfect recall [24, 14, 6]). We will use *the original game* to refer to some perfect-recall game $\Gamma = (H, Z, A, P, \pi_0, \{\mathcal{I}_i\}, \{u_i\})$ that we would like to compute a Nash equilibrium for. We use *the abstract game* to refer to some other game $\Gamma' = (H', Z', A', P', \pi'_0, \{\mathcal{I}'_i\}, \{u'_i\})$ that is an abstraction of Γ. The goal is to compute a (possibly approximate) equilibrium in the abstraction, and map the resulting strategy profile to the full game. An example is shown in Figure 1.

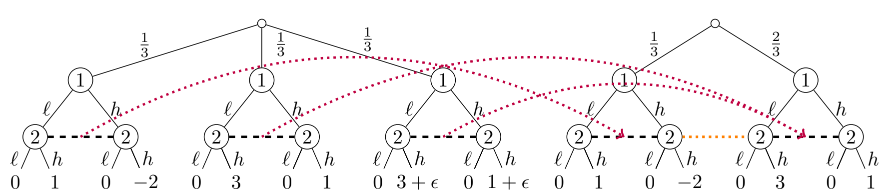

Fig. 1: Abstraction example. Left: Original EFG. Right: Abstraction (which has perfect recall in this case). Dotted red arrows denote the mapping of information sets in the original game onto information set partitions in the abstract game. The dotted orange line in the abstract game denotes an information set coarsening relative to $\mathcal{P}'$.

We model abstraction as a two-stage process. First, the full game is mapped onto the abstract game, with every original information set $I \in \mathcal{I}_i$ mapping onto some *information-set partition* I'_I in the abstraction via a function $f : \mathcal{I} \to \mathcal{P}'$ that maps $\mathcal{I}$ surjectively onto $\mathcal{P}'$. $\mathcal{P}'$ is assumed to be a partitioning of $H' \setminus Z'$ that refines $\mathcal{I}'$. Thus the information-set structure specified by $\mathcal{P}'$ can be thought of as specifying an intermediate game with (weakly) more information than Γ'; $\mathcal{P}'$ is assumed to induce a perfect-recall game[1]. In Figure 1, each of the three original information sets belonging to Player 2 map onto the same abstract information set, but the leftmost original information set maps onto the left partition, whereas the center and right information sets map onto the right partition. In the abstract game in Figure 1, Player 2 has two subsets in $\mathcal{P}'$: the left and right sides of their single information set. Actions are similarly mapped with an action mapping $g : A \to A'$ that maps each A_I surjectively onto $A_{f(I)}$. It is assumed that f respects the information-set tree structure by mapping $\mathcal{C}_I^a$ surjectively onto $\mathcal{C}_{I'_I}^{g(a)}$. The final part of the first step is a way to map leaf nodes under original information sets to leaf nodes under the corresponding abstract information set. For each information set I and action $a \in A_I$, we require a surjective leaf-node mapping ϕ_I from the set of leaf nodes reached below I, a before player i acts again, Z_I^a, onto $Z_{I'_I}^{a'}$.

The second step in our abstraction model captures the differences between the abstract game Γ' and the game induced by using the partitioning $\mathcal{P}'$ instead. This is done by comparing the distribution over leaf nodes conditioned on being at a given $I'_I \in \mathcal{P}'$ versus the distribution conditioned on being at the corresponding abstract information set I'. In Figure 1 this would correspond to comparing the leaf nodes under e.g. the right pair of nodes in Player 2's information set in the abstraction to the leaf nodes in the overall information set for Player 2. For each partition I'_I this is done with a set-valued map $\phi_{I'_I}$ that maps the set of leaf nodes $\mathcal{Z}_{I'_I}^{a'}$ onto $\mathcal{Z}_{I'}^{a'}$ for each a' in a way such that $\{\phi_{I'_I}(z') : z' \in \mathcal{Z}_{I'_I}^{a'}\}$ specifies a partitioning of $\mathcal{Z}_{I'}^{a'}$. For a given partition I'_I, we let $\mathcal{D}_{I'_I}$ and $\mathcal{C}_{I'_I}$ be the set of descendant and child partitions, respectively, that can be reached from I'_I.

For a strategy profile σ' computed in Γ' we need a way to interpret it as strategy profiles in Γ. We present the natural extension of a *lifted strategy*, originally developed by Sandholm and Singh [35] for stochastic games, to EFGs. Intuitively, a lifted strategy $\sigma^{\uparrow\sigma'}$ is a strategy where for any abstract action a', the sum of probabilities in $\sigma^{\uparrow\sigma'}$ assigned to actions that map to a' is equal to the probability placed on a' in σ'.

Definition 1 (Strategy lifting). *Given an abstract strategy profile σ', a lifted strategy profile is any strategy profile $\sigma^{\uparrow\sigma'}$ such that for all I, all $a' \in A_{f(I)}$:* $\sum_{a \in g^{-1}(a')} \sigma^{\uparrow\sigma'}(I, a) = \sigma'(f(I), a')$.

[1] Lanctot et al. [30] and Kroer and Sandholm [27] use the notion of a *perfect-recall refinement*, which is a partitioning of each imperfect-recall information set into several perfect-recall information sets. Our definition of $\mathcal{P}'$ can be thought of as specifying a perfect-recall refinement of the abstraction.

We use the definition of counterfactual value of an information set, introduced by Zinkevich et al. [40], to reason about the value of an information set under a given strategy profile. The counterfactual value of an information set I is the expected utility of the information set, assuming that all players follow strategy profile σ, except that Player i plays to reach I. It is defined as $V_i^\sigma(I) = \sum_{z \in \mathcal{Z}_I} \pi^\sigma(z|I) u_i(z)$ when $\pi_{-i}^\sigma(I) > 0$; otherwise it is 0. Analogously, $W_i^{\sigma'} : \mathcal{I}'_i \to \mathbb{R}$ is the corresponding function for the abstract game. For the information set I_r that contains just the root node r, we have $V_i^\sigma(I_r) = V_i^\sigma(r)$, which is the value of playing the game with strategy profile σ. We assume that at the root node it is not Chance's turn to move. This is without loss of generality since we can insert dummy player nodes above a root node belonging to Chance.

We show that for an information set I, $V_i^\sigma(I)$ can be written as a sum over descendant information sets. The proof is straightforward but shown in the appendix.

Lemma 1. *For any strategy profile σ or abstract strategy profile σ', the counterfactual value of an information set I, or abstract information-set partition I'_I, can respectively be written as*

$$V_i^\sigma(I) = \sum_{a \in A_I} \sigma(I,a) \left[\sum_{\hat{I} \in \mathcal{C}_I^a} \pi_{-i}^\sigma(\hat{I}|I) V_i^\sigma(\hat{I}) + \sum_{z \in Z_I^a} \pi_{-i}^\sigma(z|I) u_i(z) \right],$$

$$W_i^{\sigma'}(I'_I) = \sum_{a' \in A_{I'}} \sigma'(I',a') \left[\sum_{\hat{I}'_{\hat{I}} \in \mathcal{C}_{I'_I}^{a'}} \pi_{-i}^{\sigma'}(\hat{I}'_{\hat{I}}|I'_I) W_i^{\sigma'}(\hat{I}'_{\hat{I}}) + \sum_{z' \in Z_{I'_I}^{a'}} \pi_{-i}^{\sigma'}(z'|I'_I) u_i(z') \right]. \tag{1}$$

We will show results for three different solution concepts that come up in practice. An ϵ-*Nash equilibrium* is a strategy profile σ such that $V_i^\sigma(r) \geq V_i^{\sigma'}(r) - \epsilon$ for all players i and $\sigma' = (\sigma_{-i}, \sigma'_i)$. In other words, each player can gain at most ϵ by deviating to any other strategy σ'_i. This is what is computed by approaches based on first-order methods [23, 29, 28]. A *Nash equilibrium* is an ϵ-Nash equilibrium where $\epsilon = 0$. Finally, a strategy profile σ has bounded counterfactual regret if for all $i, I \in \mathcal{I}$, and $a \in A_I$, $V_i^{\sigma_{I \to a}}(I) \leq V_i^\sigma(I) + r(I)$. Strategy profiles with bounded counterfactual regret are important because regret minimization algorithms for EFGs converge by producing strategies with low $r(I)$ [40, 31, 11, 13, 10].

4 Measuring differences between the original game and the abstract game

Our goal is to show a decomposition of the utility difference between the original game and the abstract game when using a lifted strategy. In order to do this, we need a way to measure differences between the original and abstract game. We measure payoff differences between nodes as

$$\Delta_i^R(z, z') = u_i(z) - u_i(z').$$

We measure leaf-node reach-probability differences conditioned on reaching a given information set I versus its corresponding abstract information set-partition I'_I as follows

$$\Delta^P_{-i}(z'|I,a,\sigma,\sigma') = \sum_{z \in \phi_I^{-1}(z') : z \in Z_I^a} \pi^\sigma_{-i}(z|I) - \pi^{\sigma'}_{-i}(z'|I'_I), \quad \text{for } z' \in Z_{I'}^{a',\Rrightarrow}.$$

We will also need to measure the difference in probability of reaching information set partitions, conditioned on being at the preceding information set partition belonging to the same player,

$$\Delta^P_{-i}(\hat{I}'_{\hat{I}}|I,a,\sigma,\sigma') = \sum_{\tilde{I} \in f^{-1}(\hat{I}'_{\hat{I}})} \pi^\sigma_{-i}(\tilde{I}|I,a) - \pi^{\sigma'}_{-i}(\hat{I}'_{\hat{I}}|I'_I).$$

Note that while the set $f^{-1}(\hat{I}'_{\hat{I}})$ can include information sets $\tilde{I}$ that do not come after I, a, such information sets are irrelevant since $\pi^\sigma_{-i}(\tilde{I}|I,a) = 0$.

We now prove a technical lemma that will be used as the primary tool for inductively proving that strategies from abstractions have bounded regret.

Lemma 2. *For any information set $I, I' = f(I)$ and pair of strategy profiles σ, σ', assume there is a bound $\operatorname{diff}(\hat{I}, f(\hat{I}))$ such that $V_i^\sigma(\hat{I}) - W_i^{\sigma'}(\hat{I}'_{\hat{I}}) \leq \operatorname{diff}(\hat{I}, f(\hat{I}))$ for all $\hat{I} \in \mathcal{C}_I^a, a \in A_I$, and $\sigma'(I', a') = \sum_{a \in g^{-1}(a')} \sigma(I, a)$. Then*

$$\begin{aligned} V_i^\sigma(I) - W_i^{\sigma'}(I'_I) \leq \sum_{a \in A_I} \sigma(I,a) \Bigg[& \sum_{z \in Z_I^{a,\Rrightarrow}} \pi^\sigma_{-i}(z|I) \Delta_i^R(z, \phi_I(z)) \\ + \sum_{z' \in Z_{I'_I}^{g(a),\Rrightarrow}} & \Delta^P_{-i}(z'|I,a,\sigma,\sigma') u_i(z') + \sum_{\hat{I} \in \mathcal{C}_I^a} \pi^\sigma_{-i}(\hat{I}|I) \operatorname{diff}(\hat{I}, f(\hat{I})) \\ + \sum_{\hat{I}'_{\hat{I}} \in \mathcal{C}_{I'_I}^{g(a)}} & \Delta^P_{-i}(\hat{I}'_{\hat{I}}|I,a,\sigma,\sigma') W_i^{\sigma'}(\hat{I}'_{\hat{I}}) \Bigg] \end{aligned}$$

The above holds with equality if $V_i^\sigma(\hat{I}) - W_i^{\sigma'}(\hat{I}'_{\hat{I}}) = \operatorname{diff}(\hat{I}, f(\hat{I}'))$ for all $\hat{I} \in \mathcal{C}_I^a$ and $a \in A_I$.

We now introduce a shorthand for denoting the utility difference attributable to differences between a given information set I and its abstract counterpart $f(I)$. This is the utility difference that would arise from recursively applying Lemma 2 to information sets.

$$\begin{aligned} \operatorname{Mdiff}(I, \sigma, \sigma'_{-i}) \stackrel{\text{def}}{=} \sum_{a \in A_I} \sigma(I,a) \Bigg[& \sum_{z \in Z_I^{a,\Rrightarrow}} \pi^\sigma_{-i}(z|I) \Delta_i^R(z, \phi_I(z)) \\ & + \sum_{z' \in Z_{I'_I}^{g(a),\Rrightarrow}} \Delta^P_{-i}(z'|I,a,\sigma,\sigma') u_i(z') \\ & + \sum_{\hat{I} \in \mathcal{C}_I^a} \pi^\sigma_{-i}(\hat{I}|I) \operatorname{Mdiff}(\hat{I}, \sigma, \sigma'_{-i}) + \sum_{\hat{I}'_{\hat{I}} \in \mathcal{C}_{I'_I}^{g(a)}} \Delta^P_{-i}(\hat{I}'_{\hat{I}}|I,a,\sigma,\sigma') W_i^{\sigma'}(\hat{I}'_{\hat{I}}) \Bigg] \end{aligned}$$

It follows from Lemma 2 that the players' values in any lifted strategy profile in the original game are close to the players' values of the corresponding abstract strategy profile:

Lemma 3. *Given any abstract strategy profile* σ'*, any lifted strategy profile* $\sigma^{\uparrow\sigma'}$ *achieves utility*

$$W_i^{\sigma'}(r') = V_i^{\sigma^{\uparrow\sigma'}}(r) - \mathrm{Mdiff}(r, \sigma^{\uparrow\sigma'}, \sigma'_{-i})$$

Next we derive an expression for the difference between an abstract information set and any subset in its partitioning. We will need a way to measure the difference between an information set I' and any partition I'_I. For reach probability, we let

$$\Delta^P(z'|I'_I, \sigma') = \pi^{\sigma'}(z'|I'_I) - \sum_{\hat{z}' \in \phi_{I'_I}(z')} \pi^{\sigma'}(\hat{z}'|I') \tag{2}$$

be the difference between the probability of arriving at z' conditioned on a strategy σ' and being in partition I'_I of I' and the probability of arriving at any leaf node $z' \in \phi_{I'_I}^{-1}(z')$ conditioned on the same strategy σ' and being in I'. For reward differences we let the utility difference between a leaf node $\hat{z}' \in Z_{I'}$ and its corresponding leaf node $z' = \phi_{I'_I}^{-1}(z')$ in $Z_{I'_I}$ be

$$\Delta_i^R(\hat{z}'|I'_I) = u_i(z') - \delta_{I'_I} u_i(\hat{z}') \tag{3}$$

These terms allow us to measure the difference between the value $W_i^{\sigma'}(I')$ and $W_i^{\sigma'}(I'_I)$ for any information set I' and any I'_I in its partition. We let $\mathrm{Pdiff}(I'_I, \sigma')$ denote this difference.

Lemma 4. *For any player* i*, abstract strategy profile* σ'*, information set* I' *and any* I'_I *in its partition,*

$$W_i^{\sigma'}(I'_I) - \delta_{I'_I} W_i^{\sigma'}(I') = \sum_{\hat{z}' \in \mathcal{Z}_{I'}} \pi^{\sigma'}(\hat{z}'|I') \Delta_i^R(\hat{z}'|I'_I) + \sum_{z' \in \mathcal{Z}_{I'_I}} \Delta^P(z'|I'_I, \sigma') u_i(z') \stackrel{\mathrm{def}}{=} \mathrm{Pdiff}(I'_I, \sigma')$$

5 An exact decomposition of abstraction error

Our first theorem shows that an ϵ-Nash equilibrium in the abstract game maps to an ϵ'-Nash equilibrium in the original game, where ϵ' depends on the difference terms introduced in the previous section. We say that the abstract game has a *cycle* if there exists a sequence of information sets $I'_1, \dots, I'_k$ such that for all $j \neq k$ there exist nodes $h'_j \in I'_j, h'_{j+1} \in I'_{j+1}$ such that h'_j is an ancestor of h'_{j+1}, and I'_1 is equal to I'_k. The next theorem assumes the abstract game is acyclic. This enables induction over information sets.

Theorem 1. *Given an* ϵ*-Nash equilibrium* σ' *for an acyclic abstract game, any lifted strategy profile* $\sigma^{\uparrow\sigma'}$ *is an* ϵ'*-Nash equilibrium in the original game where*

$\epsilon' = \max_{i\in N} \epsilon_i$ and

$$\epsilon'_i = \epsilon + \text{Mdiff}(r, \sigma^*, \sigma'_{-i}) - \text{Mdiff}(r, \sigma^{\uparrow\sigma'}, \sigma'_{-i}) + \sum_{I\in\mathcal{I}_i} \pi^{\sigma^*}(I)\left[\text{Pdiff}(I'_I, \sigma^{*\prime}_{I'\to I}) - \text{Pdiff}(I'_I, \sigma^{*\prime})\right]$$

here $\sigma^* = (\sigma^*_i, \sigma^{\uparrow\sigma'}_{-i})$ is $\sigma^{\uparrow\sigma'}$ except Player i plays any best response strategy for the original game, $\sigma^{*\prime} = (\sigma^{*\prime}_i, \sigma'_{-i})$ is such that $\sigma^{*\prime}(I', a') = \sum_{g^{-1}(a')} \sigma^*(I, a)$ where $I \in f^{-1}(I')$ is chosen for each I' in order to maximize $W_i^{\sigma^{*\prime}}(r)$, and $\sigma^{*\prime}_{I'\to I}$ is $\sigma^{*\prime}$ except that at I' we set the strategy according to I, i.e. $\sigma^{*\prime}(I', a') = \sum_{g^{-1}(a')} \sigma^*(I, a)$.

This theorem is the first to show results for mapping an ϵ'-Nash equilibrium in the abstract game to an ϵ-Nash equilibrium in the original game. Prior results have been for abstract strategies that are either exact Nash equilibria [25] or with bounded counterfactual regret [30, 27]. That is because all prior proofs were based on applying a worst-case counterfactual regret bound as part of the inductive step (which works for exact Nash equilibrium or strategies with bounded counterfactual regret but not ϵ-Nash equilibrium); our proof instead constructs an expression for $W_i^{\sigma^{*\prime}}(r')$ (i.e., for the value of the whole abstract game) before using the fact that σ' is an ϵ-Nash equilibrium. We next show that our framework can also measure differences for strategies with bounded counterfactual regret.

Theorem 2. *For an abstract strategy profile σ' with bounded counterfactual regret $r(I')$ at every information set $I' \in \mathcal{I}'$, any lifted strategy profile $\sigma^{\uparrow\sigma'}$ is an ϵ-Nash equilibrium where*

$$\epsilon = \max_{i\in N} \epsilon_i, \quad \epsilon_i \leq \sum_{I\in\mathcal{I}_i} \pi^{\sigma^*}(I)\left[\delta_{f(I)_I} r(f(I)) + \text{Pdiff}(I'_I, \sigma'_{I\to\sigma^{*\prime}}) - \text{Pdiff}(I'_I, \sigma')\right] + \text{Mdiff}(r, \sigma^*, \sigma'_{-i}) - \text{Mdiff}(r, \sigma^{\uparrow\sigma'}, \sigma'_{-i})$$

where $\sigma^* = (\sigma^*_i, \sigma^{\uparrow\sigma'}{}_{-i})$ is $\sigma^{\uparrow\sigma'}$ except for Player i best responding, and each $\sigma'_{I\to\sigma^*}$ is equal to σ' except that $\sigma'_{I\to\sigma^*}(f(I), a') = \sum_{a\in g^{-1}(a')} \sigma^*(I, a)$ for all $a' \in A_{f(I)}$.

We will show in the next sections that our two main theorems generalize prior results. In addition, our theorems are the first to give an exact expression for the abstraction error; the inequalities arise only from inexactly solving the abstract game.

6 Generalizing prior results

We now show that if the reach of leaf nodes and child information sets in the original and abstract game are the same (without considering Chance moves), the exact results from the previous section subsume all prior solution quality bounds for EFGs [30, 25, 27] (which also make that assumption or stronger assumptions). In order to state our result, we let χ_i be the set of pure strategies belonging to Player i. We will often use $\boldsymbol{a}$ as an index where a single action a would go according to our definitions; in such cases $\boldsymbol{a}$ should be interpreted as the specific action in $\boldsymbol{a}$ that pertains to the definition, usually the action at a given information set I prescribed by $\boldsymbol{a}$. In a slight abuse of notation, we let $g(\boldsymbol{a})$ denote the pure strategy in the abstract game corresponding to $\boldsymbol{a}$ when applying g.

Proposition 1. *If an abstract strategy profile σ' and a lifted strategy profile $\sigma^{\uparrow\sigma'}$ are such that for all $i, I \in \mathcal{I}$, $\Delta^P_{-i,-0}(z'|I,a,\sigma,\sigma') = 0$, $\Delta^P_{-i,-0}(z|I,\sigma,\sigma') = 0$, and $\Delta^P_{-i,-0}(\hat{I}'_{\hat{I}}|I,a,\sigma,\sigma') = 0$ then for all players i and $\sigma = (\sigma_i, \sigma^{\uparrow\sigma'}_{-i})$ we have*

$$\mathrm{Mdiff}(r,\sigma,\sigma'_{-i}) - \mathrm{Mdiff}(r,\sigma^{\uparrow\sigma'},\sigma'_{-i}) \leq 2\max_{\boldsymbol{a}\in\chi_i} \sum_{I\in\mathcal{I}_i} \pi^{\sigma^{\uparrow\sigma'}}(I|\boldsymbol{a}) \Bigg[\sum_{z\in Z_I^{\boldsymbol{a},\Rightarrow}} \pi^{\sigma^{\uparrow\sigma'}}_{-i}(z|I)\Delta^R(z,\phi_I(z))$$

$$+ \sum_{z'\in Z^{g(\boldsymbol{a}),\Rightarrow}_{I'_I}} \left[\Delta_0(z'[I]|I)\pi^{\sigma'}_{-i}(z'|z'[I]) + \sum_{h\in I_{h'}} \sum_{h_0\in\mathcal{H}_0:h\sqsubseteq h_0} \pi^{\sigma^{\uparrow\sigma'}}_{-i}(h_0|I)\Delta^A_0(h_0)\pi^{\sigma'}_{-i}(z'|t^{h'_0}_{a'})u_i(z') \right]$$

$$\sum_{\hat{I}\in\mathcal{C}^{\boldsymbol{a}}_I} \sum_{\hat{h}'\in\hat{I}'_{\hat{I}}} \left[\Delta_0(\hat{h}'[I'_I]|I)\pi^{\sigma'}_{-i}(\hat{h}'|\hat{h}'[I'_I]) + \sum_{h\in I_{h'}} \sum_{h_0\in\mathcal{H}_0:h\sqsubseteq h_0} \pi^{\sigma^{\uparrow\sigma'}}_{-i}(h_0|I)\Delta^A_0(h_0)\pi^{\sigma'}_{-i}(\hat{h}'|t^{h'_0}_{a'}) \right] W^{\sigma'}_i(\hat{I}'_{\hat{I}}) \Bigg]$$

$$\stackrel{\text{def}}{=} \overline{\mathrm{Mdiff}_{\mathrm{i}}}(\sigma^{\uparrow\sigma'},\sigma')$$

We can combine Proposition 1 with Theorem 1 to get a bound that is independent of the best-response strategy:

Corollary 1. *If σ' is an abstract ϵ'-Nash equilibrium, satisfies the condition of Proposition 1, and* Pdiff *is zero everywhere, then any lifted strategy profile $\sigma^{\uparrow\sigma'}$ is an ϵ-Nash equilibrium where ϵ is less than* $\max_{i\in N} \overline{\mathrm{Mdiff}_{\mathrm{i}}}(\sigma^{\uparrow\sigma'},\sigma') + \epsilon'$

The game class discussed by Kroer and Sandholm [25] is easily shown to satisfy the assumptions in Proposition 1. Thus this shows a more general bound similar to that of Kroer and Sandholm [25], where we leave in several expectations rather than taking maxima everywhere (the result by Kroer and Sandholm [25] required taking several maxima where we leave in the expectation because their proof is based on upper-bounding as part of the inductive step). Therefore, Corollary 1 yields tighter results despite also being more general.

Corollary 1 shows a result for ϵ-Nash equilibrium computed in the abstraction. An analogous corollary for abstract strategies with bounded immediate regret can easily be obtained by combining Proposition 1 with Theorem 2.

We now show that, similar to mapping error, if the reach of leaf nodes in the original and abstract game are the same without considering Chance moves, we can bound partitioning error with an expression that does not depend on the best response σ^*_i of Player i.

Proposition 2. *If σ' is such that $\pi^{\sigma'}_{-0}(z'|I'_I, a') = \pi^{\sigma'}_{-0}(\hat{z}'|I', a')$ for all $I'_I, a', z', \hat{z}' \in \phi_{I'_I}(z')$, then*

$$\text{Pdiff}(I'_I, \sigma'_{I\to\sigma^*}) - \text{Pdiff}(I'_I, \sigma') \leq 2 \max_{a'\in A_{I'}} \Big[\sum_{z'\in \mathcal{Z}^{a'}_{I'}} \pi^{\sigma'}(z'|I', a') \Delta^R_i(z'|I'_I)$$
$$+ \pi^{\sigma'}_{-0}(z'|I'_I, a') \sum_{\hat{z}'\in\phi_{I'_I}(z')} \left[\pi^{\sigma'}_0(\hat{z}'|I', a')) - \pi^{\sigma'}_0(z'|I'_I, a'\right] \Big] \stackrel{\text{def}}{=} \overline{\text{Pdiff}}(I'_I, \sigma'_{I\to\sigma^*}, \sigma'), \qquad \forall \sigma'_{I\to\sigma^*}$$

This can be combined with our main theorems in order to get results for ϵ-Nash equilibrium or strategies with bounded regret where the partition error does not depend on the best response.

Corollary 2. *If σ' has bounded counterfactual regret $r(I')$ at every information set $I' \in \mathcal{I}'$, satisfies the condition of Proposition 2, and* Mdiff *is zero everywhere, then any lifted strategy $\sigma^{\uparrow\sigma'}$ is an ϵ-Nash equilibrium where $\epsilon = \max_{i\in N} \epsilon_i$ and $\epsilon_i \leq \sum_{I\in\mathcal{I}_i} \pi^{\sigma^*}(I) \left[\delta_{f(I)_I} r(f(I)) + \overline{\text{Pdiff}}(I'_I, \sigma'_{I\to\sigma^*}, \sigma')\right]$*

Kroer and Sandholm [27] took maxima in several places where we left in the expectation: they take a maximum over the decisions of Player i in $\pi^{\sigma^*}(I)$, and they maximize over the partitions in I'. Taking these maxima avoids dependence on σ^*. Taking these maxima could easily be done in Corollary 2 as well. Kroer and Sandholm [27] also separate $\Delta^P_0(z'|I'_I)$ into separate terms for Chance error that occurs before and after reaching I'; this potentially leads to a looser bound than ours (and never tighter since we could combine our Corollary 2 with their separation). An analogue to Corollary 2 but for ϵ-Nash equilibrium can be obtained by combining Theorem 1 with Proposition 2.

6.1 Neccessity of distributional similarity of reach probabilities

We now show that the style of bound given by Lanctot et al. [30] as well as our corrolaries 1 and 2 cannot generalize to games where opponents do not have the same sequence of information-set-action pairs, or in our case the slightly weaker requirements in Propositions 1 and 2, for game nodes that map to each other in the abstraction. The two games that we will use as counterexamples are shown in Figure 2. From the perspective of our results, the usefulness of assuming the same sequence of information-set-action pairs is that it implies the condition used in Propositions 1 and 2; the following counterexamples thus also show that this assumption is a useful way to disallow bad abstractions such as the ones presented here (although overly restrictive from a practical perspective). Contrary to the prior results, our Theorems 1 and 2 still apply to the games below. Our two theorems would give weak bounds commensurate with the large error in the abstract equilibrium; this error is contained in the terms that depend on Δ^P.

On the left in Figure 2 is a general-sum game where the two nodes belonging to Player 1 are abstracted into a single information set. If we map ℓ onto ℓ and r onto r we get an abstraction with low payoff error: ϵ at every node. Let $\epsilon > 0$.

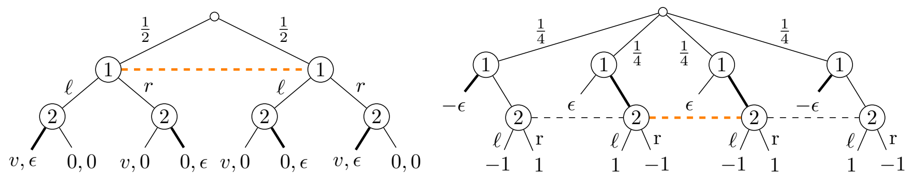

Fig. 2: Left: General-sum EFG with abstraction. Right: zero-sum EFG with abstraction where Player 1 wants to minimize. Orange dashed lines denote information sets joined in the abstraction. Bold edges denote actions taken with probability 1 in the abstracted equilibrium.

Player 2 plays the bolded edges at nodes with non-zero probability of being reached. In the abstraction, Player 1 gets $\frac{v}{2}$ for every strategy. In the full game, Player 1 can choose ℓ in the left subtree and r in the right subtree for a payoff of v. Thus in every equilibrium where Player 2 plays according to the bolded edges (which includes all equilibrium refinements) Player 1 loses $\frac{v}{2}$ from abstracting, despite the payoff error being arbitrarily small. If we set $\epsilon = 0$, equilibria where Player 2 plays the bolded edges still have high loss—despite zero payoff error.

On the right in Figure 2 is a zero-sum game where the two bottom information sets belonging to Player 2 have been abstracted. Consider the following abstract equilibrium: Player 1 plays the bolded edges with probability 1, and Player 2 plays ℓ, r with equal probability. Player 1 gets expected utility $-\frac{\epsilon}{2}$, but in the full game Player 1 can choose ℓ (r) in the left (right) information set to get utility $\frac{1-\epsilon}{2}$. Thus Player 1 has a utility loss of $\frac{1}{2}$ despite a payoff error of 0.

Bibliography

[1] Basak, A., Fang, F., Nguyen, T.H., Kiekintveld, C.: Abstraction methods for solving graph-based security games. In: International Conference on Autonomous Agents and Multiagent Systems. pp. 13–33. Springer (2016)

[2] Basak, A., Fang, F., Nguyen, T.H., Kiekintveld, C.: Combining graph contraction and strategy generation for green security games. In: International Conference on Decision and Game Theory for Security. pp. 251–271. Springer (2016)

[3] Basilico, N., Gatti, N.: Automated abstractions for patrolling security games. In: AAAI Conference on Artificial Intelligence (AAAI) (2011)

[4] Billings, D., Burch, N., Davidson, A., Holte, R., Schaeffer, J., Schauenberg, T., Szafron, D.: Approximating game-theoretic optimal strategies for full-scale poker. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) (2003)

[5] Bowling, M., Burch, N., Johanson, M., Tammelin, O.: Heads-up limit hold'em poker is solved. Science 347(6218) (Jan 2015)

[6] Brown, N., Ganzfried, S., Sandholm, T.: Hierarchical abstraction, distributed equilibrium computation, and post-processing, with application to a champion no-limit Texas Hold'em agent. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) (2015)

[7] Brown, N., Sandholm, T.: Regret transfer and parameter optimization. In: AAAI Conference on Artificial Intelligence (AAAI) (2014)

[8] Brown, N., Sandholm, T.: Simultaneous abstraction and equilibrium finding in games. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) (2015)

[9] Brown, N., Sandholm, T.: Safe and nested subgame solving for imperfect-information games. In: Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS). pp. 689–699 (2017)

[10] Brown, N., Sandholm, T.: Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. Science p. eaao1733 (Dec 2017)

[11] Burch, N., Lanctot, M., Szafron, D., Gibson, R.G.: Efficient Monte Carlo counterfactual regret minimization in games with many player actions. In: Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS). pp. 1880–1888 (2012)

[12] Čermák, J., Bošansky, B., Lisý, V.: An algorithm for constructing and solving imperfect recall abstractions of large extensive-form games. In: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI). pp. 936–942 (2017)

[13] Farina, G., Kroer, C., Sandholm, T.: Regret minimization in behaviorally-constrained zero-sum games. In: International Conference on Machine Learning (ICML) (2017)

[14] Ganzfried, S., Sandholm, T.: Potential-aware imperfect-recall abstraction with earth mover's distance in imperfect-information games. In: AAAI Conference on Artificial Intelligence (AAAI) (2014)

[15] Gilpin, A., Sandholm, T.: A competitive Texas Hold'em poker player via automated abstraction and real-time equilibrium computation. In: Proceedings of the National Conference on Artificial Intelligence (AAAI). pp. 1007–1013 (2006)

[16] Gilpin, A., Sandholm, T.: Better automated abstraction techniques for imperfect information games, with application to Texas Hold'em poker. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). pp. 1168–1175 (2007)

[17] Gilpin, A., Sandholm, T.: Lossless abstraction of imperfect information games. Journal of the ACM 54(5) (2007)

[18] Gilpin, A., Sandholm, T.: Expectation-based versus potential-aware automated abstraction in imperfect information games: An experimental comparison using poker. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI) (2008), short paper

[19] Gilpin, A., Sandholm, T., Sørensen, T.B.: Potential-aware automated abstraction of sequential games, and holistic equilibrium analysis of Texas Hold'em poker. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI) (2007)

[20] Gilpin, A., Sandholm, T., Sørensen, T.B.: A heads-up no-limit Texas Hold'em poker player: Discretized betting models and automatically generated equilibrium-finding programs. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) (2008)

[21] Hawkin, J., Holte, R., Szafron, D.: Automated action abstraction of imperfect information extensive-form games. In: AAAI Conference on Artificial Intelligence (AAAI) (2011)

[22] Hawkin, J., Holte, R., Szafron, D.: Using sliding windows to generate action abstractions in extensive-form games. In: AAAI Conference on Artificial Intelligence (AAAI) (2012)

[23] Hoda, S., Gilpin, A., Peña, J., Sandholm, T.: Smoothing techniques for computing Nash equilibria of sequential games. Mathematics of Operations Research 35(2) (2010)

[24] Johanson, M., Burch, N., Valenzano, R., Bowling, M.: Evaluating state-space abstractions in extensive-form games. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) (2013)

[25] Kroer, C., Sandholm, T.: Extensive-form game abstraction with bounds. In: Proceedings of the ACM Conference on Economics and Computation (EC) (2014)

[26] Kroer, C., Sandholm, T.: Discretization of continuous action spaces in extensive-form games. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) (2015)

[27] Kroer, C., Sandholm, T.: Imperfect-recall abstractions with bounds in games. In: Proceedings of the ACM Conference on Economics and Computation (EC) (2016)

[28] Kroer, C., Waugh, K., Kılınç-Karzan, F., Sandholm, T.: Faster first-order methods for extensive-form game solving. In: Proceedings of the ACM Conference on Economics and Computation (EC) (2015)

[29] Kroer, C., Waugh, K., Kılınç-Karzan, F., Sandholm, T.: Theoretical and practical advances on smoothing for extensive-form games. In: Proceedings of the ACM Conference on Economics and Computation (EC) (2017)

[30] Lanctot, M., Gibson, R., Burch, N., Zinkevich, M., Bowling, M.: No-regret learning in extensive-form games with imperfect recall. In: International Conference on Machine Learning (ICML) (2012)

[31] Lanctot, M., Waugh, K., Zinkevich, M., Bowling, M.: Monte Carlo sampling for regret minimization in extensive games. In: Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS) (2009)

[32] Moravčík, M., Schmid, M., Burch, N., Lisý, V., Morrill, D., Bard, N., Davis, T., Waugh, K., Johanson, M., Bowling, M.: Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. Science 356(6337) (May 2017)

[33] Nash, J.: Equilibrium points in n-person games. Proceedings of the National Academy of Sciences 36, 48–49 (1950)

[34] Sandholm, T.: Abstraction for solving large incomplete-information games. In: AAAI Conference on Artificial Intelligence (AAAI) (2015), senior Member Track

[35] Sandholm, T., Singh, S.: Lossy stochastic game abstraction with bounds. In: Proceedings of the ACM Conference on Electronic Commerce (EC) (2012)

[36] Shi, J., Littman, M.: Abstraction methods for game theoretic poker. In: CG '00: Revised Papers from the Second International Conference on Computers and Games. pp. 333–345. Springer-Verlag, London, UK (2000)

[37] Waugh, K.: Abstraction in Large Extensive Games. Master's thesis, University of Alberta (2009)

[38] Waugh, K., Morrill, D., Bagnell, D., Bowling, M.: Solving games with functional regret estimation. In: AAAI Conference on Artificial Intelligence (AAAI) (2015)

[39] Wellman, M.P., Reeves, D.M., Lochner, K.M., Cheng, S.F., Suri, R.: Approximate strategic reasoning through hierarchical reduction of large symmetric games. In: Proceedings of the National Conference on Artificial Intelligence (AAAI) (2005)

[40] Zinkevich, M., Bowling, M., Johanson, M., Piccione, C.: Regret minimization in games with incomplete information. In: Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS) (2007)

A Proof of Lemma 1

Proof. We show the statement for $V_i^\sigma(I)$, the result for $W_i^{\sigma'}(I'_I)$ follows by viewing the partitioning as a perfect-recall game. We have

$$V_i^\sigma(I) = \sum_{z \in \mathcal{Z}_I} \pi^\sigma(z|I) u_i(z) = \sum_{\hat{I} \in \mathcal{C}_I} \sum_{z \in \mathcal{Z}_{\hat{I}}} \pi^\sigma(z|I) u_i(z) + \sum_{z \in Z_I} \pi^\sigma(z|I) u_i(z) \qquad (4)$$

Now note that for any $\hat{I} \in \mathcal{C}_I$ we have

$$\sum_{z \in \mathcal{Z}_{\hat{I}}} \pi^\sigma(z|I) u_i(z) = \sum_{h \in \hat{I}} \pi^\sigma(h|I) \sum_{z \in \mathcal{Z}_{\hat{I}}} \pi^\sigma(z|h) u_i(z) = \pi^\sigma(\hat{I}|I) \sum_{h \in \hat{I}} \pi^\sigma(h|\hat{I}) \sum_{z \in \mathcal{Z}_{\hat{I}}} \pi^\sigma(z|h) u_i(z)$$

$$= \pi^\sigma(\hat{I}|I) \sum_{z \in \mathcal{Z}_{\hat{I}}} \pi^\sigma(z|\hat{I}) u_i(z) = \pi^\sigma(\hat{I}|I) V_i^\sigma(\hat{\imath}),$$

where the second equality follows from $\pi^\sigma(h|I) = \pi^\sigma(h|\hat{I})\pi^\sigma(\hat{I}|I)$ and the third equality follows from $\pi^\sigma(z|\hat{I}) = \pi^\sigma(h|\hat{I})\pi^\sigma(z|h)$. Plugging this into (4) gives the result.

B Proposition 3

Proposition 3. *For any player i, abstract strategy σ', real strategy σ, information sets I and $I' = f(I)$, and actions a and $a' = g(a)$*

$$\sum_{z \in Z_I^a} \pi_{-i}^\sigma(z|I) u_i(z) - \sum_{z' \in Z_{I'_I}^{a'}} \pi_{-i}^{\sigma'}(z'|I'_I) u_i(z') = \sum_{z \in Z_I^{a,\Rrightarrow}} \pi_{-i}^\sigma(z|I) \Delta^R(z, \phi_I(z)) + \sum_{z' \in Z_{I'_I}^{a',\Rrightarrow}} \Delta_{-i}^P(z'|I, a, \sigma, \sigma') u_i(z')$$

Proof. We have

$$\sum_{z \in Z_I^{a,\Rrightarrow}} \pi_{-i}^\sigma(z|I) u_i(z) = \sum_{z' \in Z_{I'_I}^{a',\Rrightarrow}} \sum_{z \in \phi_I^{-1}(z') : z \in Z_I^a} \pi_{-i}^\sigma(z|I) u_i(z)$$

$$= \sum_{z' \in Z_{I'_I}^{a',\Rrightarrow}} \sum_{z \in \phi_I^{-1}(z') : z \in Z_I^a} \pi_{-i}^\sigma(z|I)(u_i(z') + \Delta^R(z, \phi_I(z)))$$

$$= \sum_{z' \in Z_{I'_I}^{a',\Rrightarrow}} \left[\pi_{-i}^{\sigma'}(z'|I'_I) + \Delta_{-i}^P(z'|I, a, \sigma, \sigma')\right] u_i(z') + \sum_{z \in Z_I^{a,\Rrightarrow}} \pi_{-i}^\sigma(z|I) \Delta^R(z, \phi_I(z))$$

$$= \sum_{z' \in Z_{I'_I}^{a',\Rrightarrow}} \pi_{-i}^{\sigma'}(z'|I'_I) u_i(z') + \sum_{z \in Z_I^{a,\Rrightarrow}} \pi_{-i}^\sigma(z|I) \Delta^R(z, \phi_I(z)) + \sum_{z' \in Z_{I'_I}^{a',\Rrightarrow}} \Delta_{-i}^P(z'|I, a, \sigma, \sigma') u_i(z').$$

The first equality follows from the fact that every leaf node in $Z_I^{a,\Rrightarrow}$ maps onto some leaf node in $Z_{I'}^{a',\Rrightarrow}$, the second from the definition of Δ^R, the third by rearranging terms and the definition of Δ^P, and the fourth by rearranging terms.

C Proof of Lemma 2

Proof. By Lemma 1 we have

$$V_i^\sigma(I) = \sum_{a \in A_I} \sigma(I, a) \left[\sum_{\hat{I} \in \mathcal{C}_I^a} \pi_{-i}^\sigma(\hat{I}|I) V_i^\sigma(\hat{I}) + \sum_{z \in Z_I^a} \pi_{-i}^\sigma(z|I) u_i(z) \right]$$

We show the result by separately rewriting the two summation terms. For the summation over information sets we have

$$\sum_{\hat{I}\in\mathcal{C}_I^a} \pi_{-i}^{\sigma}(\hat{I}|I)V_i^{\sigma}(\hat{I}) \leq \sum_{\hat{I}\in\mathcal{C}_I^a} \pi_{-i}^{\sigma}(\hat{I}|I)\left[W_i^{\sigma'}(\hat{I}'_{\hat{I}}) + \mathrm{diff}(\hat{I}, f(\hat{I}))\right]$$
$$= \sum_{\hat{I}'\in\mathcal{C}_{I'}^{a'}} \left[\pi_{-i}^{\sigma'}(\hat{I}'_I|I', a') + \Delta_{-i}^{P}(\hat{I}'_I|I, a, \sigma, \sigma')\right] W_i^{\sigma'}(\hat{I}'_{\hat{I}}) + \sum_{\hat{I}\in\mathcal{C}_I^a} \pi_{-i}^{\sigma}(\hat{I}|I)\,\mathrm{diff}(\hat{I}, f(\hat{I}))$$

Where the last step follows by the definition of $\Delta_{-i}^{P}(\hat{I}'_I|I, a, \sigma, \sigma')$. For the summation over $\sum_{z\in Z_I^a}$ we can apply Proposition 3. Adding up terms and using the condition $\sigma'(I', a') = \sum_{a\in g^{-1}(a')} \sigma(I, a)$ then gives the result.

To see why the inequality holds with equality when the bounds on child information sets are equalities, note that the only inequality introduced in the proof comes from applying the bound on the child information sets.

D Proof of Lemma 4

Proof. Using the definition of the value of a partition and applying (2), rearranging, applying (3), rearranging again and using the definition of $W_i^{\sigma'}(I')$ gives

$$W_i^{\sigma'}(I'_I) = \sum_{z'\in\mathcal{Z}_{I'_I}} \pi^{\sigma'}(z'|I'_I)u_i(z') = \sum_{z'\in\mathcal{Z}_{I'_I}} \left[\sum_{\hat{z}'\in\phi_{I'_I}(z')} \pi^{\sigma'}(\hat{z}'|I') + \Delta^{P}(z'|I'_I, \sigma')\right] u_i(z')$$
$$= \sum_{z'\in\mathcal{Z}_{I'_I}} \sum_{\hat{z}'\in\phi_{I'_I}(z')} \pi^{\sigma'}(\hat{z}'|I')u_i(z') + \sum_{z'\in\mathcal{Z}_{I'_I}} \Delta^{P}(z'|I'_I, \sigma')u_i(z'))$$
$$= \sum_{z'\in\mathcal{Z}_{I'_I}} \sum_{\hat{z}'\in\phi_{I'_I}(z')} \pi^{\sigma'}(\hat{z}'|I')\left[\delta_{I'_I}u_i(\hat{z}') + \Delta^{R}(\hat{z}'|I'_I)\right] + \sum_{z'\in\mathcal{Z}_{I'_I}} \Delta^{P}(z'|I'_I, \sigma')u_i(z'))$$
$$=\delta_{I'_I}W_i^{\sigma'}(I') + \sum_{\hat{z}'\in\mathcal{Z}_{I'}} \pi^{\sigma'}(\hat{z}'|I')\Delta^{R}(\hat{z}'|I'_I) + \sum_{z'\in\mathcal{Z}_{I'_I}} \Delta^{P}(z'|I'_I, \sigma')u_i(z'))$$

where the last step follows because $\phi_{I'_I}(\cdot)$ specifies a partitioning of the leaves under I'.

E Proof of Theorem 1

Proof. The proof consists of showing that $V_i^{\sigma^*}(r)$ can be bounded by $W_i^{\sigma^{*\prime}}(r')$ and some difference terms, where $\sigma^{*\prime}$ is a (intuitively speaking) reversely lifted strategy from σ^*. We can then use the fact that σ' is an ϵ-Nash equilibrium to bound $W_i^{\sigma^{*\prime}}(r')$ in terms of $W_i^{\sigma'}(r')$ and finally show that $W_i^{\sigma'}(r')$ is close to $V_i^{\sigma^{\uparrow\sigma'}}(r)$.

To rewrite $V_i^{\sigma^*}(r)$ in terms of $W_i^{\sigma^{*\prime}}(r')$ we first prove the following inductive statement for $I \in \mathcal{I}_i, I' = f(I)$ and letting $\hat{I}' = f(\hat{I})$ for each $\hat{I}$:

$$V_i^{\sigma^*}(I) \leq W_i^{\sigma^{*\prime}}(I'_I) + \mathrm{Mdiff}(I, \sigma^*, \sigma'_{-i}) + \sum_{\hat{I}\in\mathcal{D}_I\cup\{I\}} \pi^{\sigma^*}(\hat{I}|I)\left[\mathrm{Pdiff}(\hat{I}'_{\hat{I}}, \sigma^{*\prime}_{\hat{I}'\to\hat{I}}) - \mathrm{Pdiff}(\hat{I}'_{\hat{I}}, \sigma^{*\prime})\right]$$

We will start by showing the inductive step, as the base case is the special case of information sets that only have leaves beneath them (i.e. information sets I such that for all a, $\mathcal{C}_I^a = \emptyset$).

Let $I, I' = f(I)$ be such that the inductive statement holds for all $a \in A_I$ and $\hat{I} \in \mathcal{C}_I^a$. The inductive assumption then gives a bound for each $\hat{I} \in \mathcal{D}_I$ as required by Lemma 2. We have

$$V_i^{\sigma^*}(I) \leq W_i^{\sigma^{*\prime}_{I'\to I}}(I'_I) + \mathrm{Mdiff}(I, \sigma^*, \sigma'_{-i}) + \sum_{\hat{I}\in\mathcal{D}_I} \pi^{\sigma^*}(\hat{I}|I)\left[\mathrm{Pdiff}(\hat{I}'_{\hat{I}}, \sigma^{*\prime}_{\hat{I}'\to\hat{I}}) - \mathrm{Pdiff}(\hat{I}'_{\hat{I}}, \sigma^{*\prime})\right], \tag{5}$$

where the result follows by collecting terms that arise from Lemma 2 into the three separate terms above.

Now we can bound $W_i^{\sigma^{*\prime}_{I'\to I}}(I'_I)$

$$\begin{aligned} W_i^{\sigma^{*\prime}_{I'\to I}}(I'_I) =& \delta_{I'_I} W_i^{\sigma^{*\prime}_{I'\to I}}(I') + \mathrm{Pdiff}(I'_I, \sigma^{*\prime}_{I'\to I}); \text{ by Lemma 4} \\ \leq& \delta_{I'_I} W_i^{\sigma^{*\prime}}(I') + \mathrm{Pdiff}(I'_I, \sigma^{*\prime}_{I'\to I}); \text{ because } \sigma^{*\prime} \text{ maximizes } W_i^{\sigma^{*\prime}}(r) \\ =& W_i^{\sigma^{*\prime}}(I'_I) + \mathrm{Pdiff}(I'_I, \sigma^{*\prime}_{I'\to I}) - \mathrm{Pdiff}(I'_I, \sigma^{*\prime}); \text{ by Lemma 4.} \end{aligned} \tag{6}$$

Putting (5) and (6) together gives the inductive statement.

For the base case, note that it follows from the exact same logic but where there are no descendant information sets. Applying the induction to the whole game we get

$$V_i^{\sigma^*}(r) \leq W_i^{\sigma^{*\prime}}(r') + \mathrm{Mdiff}(r, \sigma^*, \sigma'_{-i}) + \sum_{I\in\mathcal{I}_i} \pi^{\sigma^*}(I)\left[\mathrm{Pdiff}(I'_I, \sigma^{*\prime}_{I'\to I}) - \mathrm{Pdiff}(I'_I, \sigma^{*\prime})\right] \tag{7}$$

Now we can bound $W_i^{\sigma^{*\prime}}(r')$ by using the fact that σ' is an ϵ-Nash equilibrium in the abstract game:

$$(7) \leq W_i^{\sigma'}(I') + \epsilon + \mathrm{Mdiff}(r, \sigma^*, \sigma'_{-i}) + \sum_{I\in\mathcal{I}_i} \pi^{\sigma^*}(I)\left[\mathrm{Pdiff}(I'_I, \sigma^{*\prime}_{I'\to I}) - \mathrm{Pdiff}(I'_I, \sigma^{*\prime})\right] \tag{8}$$

Finally we can apply Lemma 3 to get the theorem statement.

F Proof of Theorem 2

Proof. Consider some best-response profile $\sigma^* = (\sigma_i^*, \sigma^{\uparrow\sigma'}{}_{-i})$. We will show that it satisfies this bound. First we show that the following holds by induction for every $I, I' = f(I)$:

$$\begin{aligned} V_i^{\sigma^*}(I) - W_i^{\sigma'}(I'_I) \leq& \sum_{\hat{I}\in\mathcal{D}_I\cup\{I\}} \pi^{\sigma^*}(\hat{I}|I)\left[\delta_{\hat{I}'_{\hat{I}}} r(f(\hat{I})) + \mathrm{Pdiff}(\hat{I}'_{\hat{I}}, \sigma'_{I\to\sigma^*}) - \mathrm{Pdiff}(\hat{I}'_{\hat{I}}, \sigma')\right] \\ &+ \mathrm{Mdiff}(I, \sigma^*, \sigma'_{-i}) \end{aligned}$$

We start by proving the inductive case. The base case follows by being a special case with no descendant information sets.

We will use the inductive assumption to apply Lemma 2 to the strategy pair $\sigma^*, \sigma'_{I\to\sigma^*}$, which satisfies the condition in the Lemma since $\sigma'_{I\to\sigma^*}$ plays according to σ^* at I'. Lemma 2 requires a bound for each $\hat{I} \in \mathcal{C}_I^a$: we let $\mathrm{diff}(\hat{I}, f(\hat{I}))$ be equal to the right-hand side terms in the inductive assumption. Lemma 2 then gives

$$V_i^{\sigma^*}(I) \leq W_i^{\sigma'_{I\to\sigma^*}}(I'_I) + \sum_{a\in A_I} \sigma^*(I,a) \left[\sum_{z\in Z_I^{a,\Rightarrow}} \pi_{-i}^{\sigma^{\uparrow\sigma'}}(z|I) \Delta^R(z,\phi_I(z)) \right. \tag{9}$$

$$+ \sum_{z'\in Z_{I'_I}^{a',\Rightarrow}} \Delta_{-i}^P(z'|I,a,\sigma^{\uparrow\sigma'},\sigma') u_i(z')$$

$$\left. + \sum_{\hat{I}\in\mathcal{C}_I^a} \pi_{-i}^{\sigma^{\uparrow\sigma'}}(\hat{I}|I)\,\mathrm{diff}(\hat{I},f(\hat{I})) + \sum_{\hat{I}'_{\hat{I}}\in\mathcal{C}_{I'_I}^{a'}} \Delta_{-i}^P(\hat{I}'_{\hat{I}}|I,a,\sigma^{\uparrow\sigma'},\sigma') W_i^{\sigma'}(\hat{I}'_{\hat{I}}) \right] \tag{10}$$

Note that for some quantities that do not depend on player i, we have substituted $\sigma^{\uparrow\sigma'}$ for σ^*, which is valid because ${\sigma^{\uparrow\sigma'}}_{-i} = \sigma^*_{-i}$. The same applies to σ' and $\sigma'_{I\to\sigma^*}$ for quantities that do not depend on I'. We now show how to convert $W_i^{\sigma'_{I\to\sigma^*}}(I'_I)$ into $W_i^{\sigma'}(I'_I)$. First we apply Lemma 4 followed by the bound on immediate regret to get

$$W_i^{\sigma'_{I\to\sigma^*}}(I'_I) = \delta_{I'_I} W_i^{\sigma'_{I\to\sigma^*}}(I') + \mathrm{Pdiff}(I'_I,\sigma'_{I\to\sigma^*}) \leq \delta_{I'_I}\left[W_i^{\sigma'}(I') + r(I')\right] + \mathrm{Pdiff}(I'_I,\sigma'_{I\to\sigma^*}) \tag{11}$$

Now we can apply Lemma 4 again to get

$$(11) = W_i^{\sigma'}(I'_I) + \delta_{I'_I} r(I') + \mathrm{Pdiff}(I'_I,\sigma'_{I\to\sigma^*}) - \mathrm{Pdiff}(I'_I,\sigma')$$

This proves the inductive step: expanding the $\mathrm{diff}(\hat{I} f(\hat{I}))$ and $\mathrm{diff}(\hat{I} f(\hat{I}))$ terms in (10), using the above equality, and collecting terms gives the inductive assumption.

Applying the inductive statement to the whole game almost gives the theorem statement, we only need to convert $W_i^{\sigma'}(r)$ into $V_i^{\sigma^{\uparrow\sigma'}}(r)$ and acquire a negative term $\mathrm{Mdiff}(r,\sigma^{\uparrow\sigma'},\sigma'_{-i})$. This is exactly what we get if we apply Lemma 3 to $W_i^{\sigma'}(r)$, and so the proof is done.

G Proof of Proposition 1

Proof. If we unroll the recursive definition of Mdiff and use the fact that $\sigma_{-i} = \sigma_{-i}^{\uparrow\sigma'}$ we get

$$
\begin{aligned}
&\text{Mdiff}(r,\sigma,\sigma'_{-i}) - \text{Mdiff}(r,\sigma^{\uparrow\sigma'},\sigma'_{-i}) \\
=&\sum_{I\in\mathcal{I}_i}\sum_{a\in A_I}\left[\pi_i^{\sigma}(I)\sigma(I,a) - \pi_i^{\sigma^{\uparrow\sigma'}}(I)\sigma^{\uparrow\sigma'}(I,a)\right]\pi_{-i}^{\sigma^{\uparrow\sigma'}}(I)\left[\sum_{z\in Z_I^{a,\Rightarrow}}\pi_{-i}^{\sigma^{\uparrow\sigma'}}(z|I)\Delta^R(z,\phi_I(z))\right. \\
&\left.+\sum_{z'\in Z_{I'_I}^{g(a),\Rightarrow}}\Delta^P_{-i}(z'|I,a,\sigma^{\uparrow\sigma'},\sigma')u_i(z') + \sum_{\hat{I}'_{\hat{I}}\in\mathcal{C}_{I'_I}^{g(a),\beta}}\Delta^P_{-i}(\hat{I}'_{\hat{I}}|I,a,\sigma^{\uparrow\sigma'},\sigma')W_i^{\sigma'}(\hat{I}'_{\hat{I}})\right] \\
\leq& 2\max_{\boldsymbol{a}\in\chi_i}\sum_{I\in\mathcal{I}_i}\pi^{\sigma^{\uparrow\sigma'}}(I|\boldsymbol{a})\left[\sum_{z\in Z_I^{\boldsymbol{a},\Rightarrow}}\pi_{-i}^{\sigma^{\uparrow\sigma'}}(z|I)\Delta^R(z,\phi_I(z)) + \sum_{z'\in Z_{I'_I}^{g(\boldsymbol{a}),\Rightarrow}}\Delta^P_{-i}(z'|I,\boldsymbol{a},\sigma^{\uparrow\sigma'},\sigma')u_i(z')\right. \\
&\left.+\sum_{\hat{I}'_{\hat{I}}\in\mathcal{C}_{I'_I}^{g(\boldsymbol{a})}}\Delta^P_{-i}(\hat{I}'_{\hat{I}}|I,\boldsymbol{a},\sigma^{\uparrow\sigma'},\sigma')W_i^{\sigma'}(\hat{I}'_{\hat{I}})\right] \qquad (12)
\end{aligned}
$$

Now it remains to note that by Lemmas **??** and **??** we have

$$
\begin{aligned}
&2\max_{\boldsymbol{a}\in\chi_i}\sum_{I\in\mathcal{I}_i}\pi^{\sigma^{\uparrow\sigma'}}(I|\boldsymbol{a})\left[\sum_{z\in Z_I^{\boldsymbol{a},\Rightarrow}}\pi_{-i}^{\sigma^{\uparrow\sigma'}}(z|I)\Delta^R(z,\phi_I(z))\right. \\
&+\sum_{z'\in Z_{I'_I}^{g(\boldsymbol{a}),\Rightarrow}}\left[\Delta_0(z'[I]|I)\pi_{-i}^{\sigma'}(z'|z'[I]) + \sum_{h\in I_{h'}}\sum_{h_0\in\mathcal{H}_0:h\sqsubseteq h_0}\pi_{-i}^{\sigma^{\uparrow\sigma'}}(h_0|I)\Delta_0^A(h_0)\pi_{-i}^{\sigma'}(z'|t_{a'}^{h'_0})u_i(z')\right] \\
&\left.\sum_{\hat{I}\in\mathcal{C}_I^{\boldsymbol{a}}}\sum_{\hat{h}'\in\hat{I}'_{\hat{I}}}\left[\Delta_0(\hat{h}'[I'_I]|I)\pi_{-i}^{\sigma'}(\hat{h}'|\hat{h}'[I'_I]) + \sum_{h\in I_{h'}}\sum_{h_0\in\mathcal{H}_0:h\sqsubseteq h_0}\pi_{-i}^{\sigma^{\uparrow\sigma'}}(h_0|I)\Delta_0^A(h_0)\pi_{-i}^{\sigma'}(\hat{h}'|t_{a'}^{h'_0})\right]W_i^{\sigma'}(\hat{I}'_{\hat{I}})\right]
\end{aligned}
$$

H Proof of Proposition 2

Proof. Since $\pi^{\sigma'_{I\to\sigma}}(z'|I') = \pi^{\sigma'}(z'|I',a')\sigma'_{I\to\sigma}(I',a')$ we can write the difference as

$$
\begin{aligned}
&\text{Pdiff}(I'_I,\sigma'_{I\to\sigma^*}) - \text{Pdiff}(I'_I,\sigma') = \sum_{a'\in A_{I'}}\left[\sigma'_{I\to\sigma^*}(I',a') - \sigma'(I',a')\right]\Bigg[\\
&\sum_{z'\in\mathcal{Z}_{I'}^{a'}}\pi^{\sigma'}(z'|I',a')\Delta^R(z'|I'_I) + \sum_{z'\in\mathcal{Z}_{I'_I}^{a'}}\Delta^P(z'|I'_I,\sigma',a')u_i(z')\Bigg] \qquad (13)
\end{aligned}
$$

We can bound this by two times the maximum value of the expression in the second brackets, where the maximum is over all $a' \in A_{I'}$ to get

$$(13) \leq 2 \max_{a' \in A_{I'}} \left[\sum_{z' \in \mathcal{Z}^{a'}_{I'}} \pi^{\sigma'}(z'|I', a') \Delta^R(z'|I'_I) + \sum_{z' \in \mathcal{Z}^{a'}_{I'_I}} \Delta^P(z'|I'_I, \sigma', a') u_i(z') \right]$$

Now it remains to note that by our condition $\pi^{\sigma'}_{-0}(z'|I'_I, a') = \pi^{\sigma'}_{-0}(\hat{z}'|I', a')$ we have

$$\Delta^P(z'|I'_I, \sigma', a') = \pi^{\sigma'}(z'|I'_I, a') - \sum_{\hat{z}' \in \phi_{I'_I}(z')} \pi^{\sigma'}(\hat{z}'|I', a')$$

$$= \pi^{\sigma'}(z'|I'_I, a') - \sum_{\hat{z}' \in \phi_{I'_I}(z')} \pi^{\sigma'}_{-0}(z'|I'_I, a') \pi^{\sigma'}_0(\hat{z}'|I', a')$$

$$= \pi^{\sigma'}(z'|I'_I, a') - \sum_{\hat{z}' \in \phi_{I'_I}(z')} \pi^{\sigma'}_{-0}(z'|I'_I, a') \left[\pi^{\sigma'}_0(z'|I'_I, a') + \left(\pi^{\sigma'}_0(\hat{z}'|I', a')) - \pi^{\sigma'}_0(z'|I'_I, a' \right) \right]$$

$$= \pi^{\sigma'}_{-0}(z'|I'_I, a') \sum_{\hat{z}' \in \phi_{I'_I}(z')} \left[\pi^{\sigma'}_0(\hat{z}'|I', a')) - \pi^{\sigma'}_0(z'|I'_I, a' \right]$$

which proves the theorem.

I Proof that the game in Figure ?? is a game of ordered signals

Proof. We go through the conditions for games of ordered signals as given by Gilpin and Sandholm [17].

1. The number of players is 2 which is finite.
2. The game gives only a signal tree that can be used to define winners, and thus works with any betting tree.
3. We only give a signal tree so this is not relevant.
4. The set of signals is {**J1,J2,K1,K2**}.
5. $\kappa = \{1\}, \gamma = \{1\}$.