

The Value of Recall in Extensive-Form Games

Ratip Emin Berker^{1, 2}, Emanuel Tewolde^{1, 2}, Ioannis Anagnostides¹,
Tuomas Sandholm^{1, 4, 5, 6}, Vincent Conitzer^{1, 2, 3}

¹Carnegie Mellon University ²Foundations of Cooperative AI Lab (FOCAL) ³University of Oxford
⁴Strategic Machine, Inc. ⁵Strategy Robot, Inc. ⁶Optimized Markets, Inc.
{rberker, etewolde, ianagnos, sandholm, conitzer}@cs.cmu.edu

Abstract

Imperfect-recall games—in which players may forget previously acquired information—have found many practical applications, ranging from game abstractions to team games and testing AI agents. In this paper, we quantify the utility gain by endowing a player with perfect recall, which we call the *value of recall* (VoR). While VoR can be unbounded in general, we parameterize it in terms of various game properties, namely the structure of chance nodes and the *degree of absentmindedness* (the number of successive times a player enters the same information set). Further, we identify several pathologies that arise with VoR, and show how to circumvent them. We also study the complexity of computing VoR, and how to optimally apportion *partial recall*. Finally, we connect VoR to other previously studied concepts in game theory, including the price of anarchy. We use that connection in conjunction with the celebrated *smoothness* framework to characterize VoR in a broad class of games.

Introduction

Game theory offers a principled framework for reasoning about complex interactions involving multiple strategic players. It continues to propel landmark results in long-standing challenges in artificial intelligence (AI), ranging from poker (Bowling et al. 2015; Moravčík et al. 2017; Brown and Sandholm 2018) to diplomacy (Bakhtin et al. 2022). A common premise in game-theoretic modeling is *perfect recall*—players never forget information once acquired. The perfect-recall assumption is often called into question for games involving human players; however, it is difficult to come up with a faithful model in such cases due to the unpredictability of when and what human players will forget. In contrast, AI agents can be specifically designed to relinquish certain information, thereby making the imperfect-recall framework directly applicable. But why should one consider AI agents with imperfect recall?

An early, influential application of imperfect-recall games revolves around *abstraction*: games encountered in practice are typically too large to represent exactly, and so one resorts to abstraction to compress its description. In particular, one way of doing so consists of allowing players to carefully

forget less important aspects of previously held information. Indeed, imperfect-recall abstractions have been a crucial component of state-of-the-art algorithms in poker solving (Brown, Ganzfried, and Sandholm 2015; Johanson et al. 2013; Waugh et al. 2009; Ganzfried and Sandholm 2014; Čermák, Bosanský, and Lisý 2017). Imperfect recall also naturally arises in so-called *adversarial team games* (Celli and Gatti 2018; Zhang, Farina, and Sandholm 2023; Zhang, An, and Subrahmanian 2022; Emmons et al. 2022; von Stengel and Koller 1997), wherein a team of players—which can be construed as a single player with imperfect recall—faces an adversary. The benefit of reinforcing the communication capacity of the team in such settings—corresponding to boosting recall—is an active area of research, prominently featured in a recent NeurIPS competition (Meisheri and Khadilkar 2020; Resnick et al. 2020). Relatedly, natural notions of *correlated equilibria* can be modeled via an imperfect-recall mediator, endowed with the ability to provide recommendations (Zhang and Sandholm 2022); in that context, imperfect recall can serve to safeguard players’ private information, a consideration that also arises in other settings (Conitzer 2019). Finally, another possible application revolves around simulating and testing AI agents before their deployment in the real world (Kovarič, Oesterheld, and Conitzer 2023, 2024; Chen, Ghersengorin, and Petersen 2024). As a result, it is becoming increasingly pressing to expand our scope beyond the assumption of perfect recall.

In this paper, we examine a question at the heart of this research agenda: *how does perfect recall affect players’ utilities under various natural solution concepts?* More specifically, we contrast the utilities obtained by a player in an initial imperfect-recall game (in extensive form) to those in a perfect recall refinement thereof; we refer to the corresponding ratio as the *value of recall* (VoR). Here, our main contribution is to provide a broad characterization of VoR for different solution concepts in terms of natural game properties.

Many strategic interactions demonstrate that perfect recall offers a significant advantage. In the popular card game blackjack, the house is expected to prevail in the long run against a player with poor recall, but certain memorization strategies tip—at least under the earlier rules followed by casinos—the balance in the player’s favor (Thorp 2016), as pop-culture has hyperbolically portrayed. The role of memory is even more pronounced in other card games such as




Figure 1: A game with imperfect recall. Giving Bobble (■) perfect recall hurts both players. Terminals show utilities for Bobble and Alice (●). Infosets are joined by dotted lines.

solitaire (Kirkpatrick 1954; Foerster and Wattenhofer 2013), where remembering the previously dealt cards drastically increases the odds of winning. We are interested in quantifying how much players benefit from perfect recall.

A plot twist: perfect recall can hurt

The previous examples notwithstanding, surprisingly, endowing a player with perfect recall can end up diminishing every player’s utility! Consider Figure 1: Alice has a small amount of money ($\epsilon > 0$) and interacts with an investment bot Bobble, starting from a free trial to see if the bot is defective (*i.e.*, Bobble plays **d**, in which case the game is over and Alice receives a small compensation of ϵ). If Bobble cooperates (**c**), the game continues and it gains access to Alice’s money, which it multiplies through investments. If Bobble defects (**d**) now, it gets to run away with all the money. However, if it has imperfect recall (cannot remember if the free trial is over), then it has the incentive to cooperate (**c**) with Alice on both counts, as attempting **d** has a greater chance of causing it to get caught during the free trial. Knowing this, Alice is incentivized to trust (**t**) Bobble, leading to the cooperative outcome. On the other hand, if Bobble is given perfect recall, it has every incentive to cooperate in the free trial and then defect after getting the money; anticipating this, Alice walks out (**w**) without interacting with Bobble (Proposition 13 formalizes this example).

Intuitively, this demonstrates that a player gaining perfect recall can result in the other players trusting it less, eliminating a cooperative outcome that is arbitrarily better for everyone. This is in line with prior work showing that the ability of a player to be simulated by others can benefit everyone in trust-based games (Kovarík, Oesterheld, and Conitzer 2023).

In spirit, this phenomenon is similar to the famous Braess paradox (Braess 1968), which predicts that augmenting a network with more links can result in worse equilibria. We formalize this type of hurtful recall in later sections, and also provide necessary conditions under which it does not arise.

Overview of our results

We formally introduce the value of recall (Definition 12) in (imperfect-recall) extensive-form games for a broad set of solution concepts. In particular, building on prior work, our definition is based on the coarsest information refinement

of a game that attains perfect recall (Definition 9). In the remainder of the paper, we investigate a number of questions relating to the value of recall.

We first formalize the observation made earlier regarding hurtful recall (Figure 1) by showing the existence of games in which a single player getting perfect recall can arbitrarily hurt all players, including themselves, for all the solution concepts considered in this paper (Proposition 13). Even more surprisingly, this type of behavior can also arise in single-player games under certain solution concepts (Example 15); we argue that this is a pathology as the single player can always choose to ignore information. We show that this issue can be circumvented by replacing each of these solution concepts with an appropriate refinement thereof, one of which is a novel definition (Definition 17).

Next, we turn our attention to the computational aspects of the value of recall. We show that VoR is NP-hard to compute, and to approximate, for all solution concepts considered in this paper, even in single-player games (Theorem 1). While this mostly follows from existing hardness results for solving imperfect-recall games (Tewolde et al. 2023), we prove new hardness results for some solution concepts, which even rule out any multiplicative approximation factor.

Those hardness results notwithstanding, we characterize VoR under optimal play in single-player games based on certain natural properties of the game tree. In particular, we show in Proposition 21 that value degradation due to imperfect recall can be fully explained by two sources: *absent-mindedness* (an infoset being entered multiple times in a path of play) and external stochasticity. In Propositions 23 and 26, we provide tight upper bounds for VoR for each of these sources separately. Finally, as our main characterization result, we show that those two bounds compose for games that exhibit both absentmindedness and external stochasticity (Theorem 2).

The aforescribed characterization applies only to optimal play. To extend it to more permissive solution concepts, we make a connection with the price of anarchy literature. Namely, inspired by the homonymous class of games introduced by Roughgarden (2015), we introduce the notion of a *smooth* (imperfect-recall) single-player game (Definition 29), and show VoR can be bounded in terms of the smoothness parameters of the game, in conjunction with our previous bound concerning optimal strategies. Besides this connection with the price of anarchy, we further observe that VoR captures some previously studied concepts, such as the *price of uncorrelation* in adversarial team games (Celli and Gatti 2018) and the *price of miscoordination* in security games (Jiang et al. 2013), which enables interpreting their results as bounds on VoR in those games.

Finally, we examine the value of recall with respect to *partial recall*—instead of perfect recall—refinements. In particular, we focus on the natural problem of refining an imperfect-recall game so as to maximize the utility gain, subject to constraining the amount of new recall. We show that, even with oracle access to optimal strategies, that problem is NP-hard even in single-player games (Theorem 3). We conclude with a number of interesting future directions stemming from our work.

Preliminaries

Before we proceed, we provide some necessary background on imperfect-recall games and solution concepts for them.

Games with imperfect recall

We start by introducing extensive-form games, following the formalism introduced by Fudenberg and Tirole (1991).

Definition 1. An *extensive-form game* Γ specifies

1. A rooted tree with node set \mathcal{H} and edges that represent *actions*. The game starts at the root, and actions are taken to traverse down the tree, until the game finishes at a leaf node, called *terminal node*. The set of terminal nodes is denoted by $\mathcal{Z} \subset \mathcal{H}$, and the set of actions available at any nonterminal node $h \in \mathcal{H} \setminus \mathcal{Z}$ is denoted by A_h .
2. A finite set $\mathcal{N} \cup \{c\}$ of $N+1$ players where $N \geq 1$. Set \mathcal{N} contains the *strategic players*, and c stands for a *chance* “player” that models exogenous stochasticity. Each non-terminal node h is assigned to a particular player $i \in \mathcal{N} \cup \{c\}$, who chooses an action to take from A_h . Set \mathcal{H}_i denotes all nodes assigned to Player i .
3. For each chance node $h \in \mathcal{H}_c$, a probability distribution $\mathbb{P}_c(\cdot | h)$ on A_h with which chance elects an action at h .
4. For each strategic player $i \in \mathcal{N}$, a (without loss of generality) nonnegative *utility (payoff)* function $u_i : \mathcal{Z} \rightarrow \mathbb{R}_{\geq 0}$ which returns what i receives when the game finishes at a terminal node. Player i aims to maximize that utility.¹
5. For each strategic player $i \in \mathcal{N}$, a partition $\mathcal{H}_i = \sqcup_{I \in \mathcal{I}_i} I$ of the nodes of i into information sets (*infosets*). Nodes of the same infoset are considered indistinguishable to the player at that infoset. For that, we also require $A_h = A_{h'}$ for $h, h' \in I$. This also makes action set A_I well-defined.

The *game tree* of Γ refers to \mathcal{H} , $\{A_h\}_{h \in \mathcal{H} \setminus \mathcal{Z}}$, and $\{\mathbb{P}_c(\cdot | h)\}_{h \in \mathcal{H}_c}$ (but not its infoset partitioning or utilities). We now formalize games where players may *forget* previously available information.

Definition 2 ((Im)perfect recall). For a decision node h of a game Γ , let $\text{hist}(h) = (h_k)_{k=0}^{\text{depth}(h)-1}$ be the ordered sequence of nodes from the root node h_0 to h (excluding h) and let $\text{seq}(h) = (i_k, I_k, a_k)_{k=0}^{\text{depth}(h)-1}$ be the corresponding sequence of tuples showing which player i_k acts at h_k , the infoset I_k of node h_k , and what action a_k was taken at h_k . For a player $i \in \mathcal{N}$, let $\text{seq}_i(h)$ be the ordered subsequence of tuples from $\text{seq}(h)$ for which $i_k = i$. We say player i has *perfect recall* in Γ if for all of i ’s infosets $I \in \mathcal{I}_i$, and all pairs of nodes $h, h' \in I$, we have $\text{seq}_i(h) = \text{seq}_i(h')$. Otherwise, we say Player i has *imperfect recall*. We say that Γ is a perfect-recall game if all players $i \in \mathcal{N}$ have perfect recall in Γ . Otherwise, we say Γ is an imperfect-recall game.

Strategies and utilities Players can select a probability distribution—a *randomized action*—over the actions at an infoset. A (behavioral) *strategy* π_i of a player $i \in \mathcal{N}$ specifies a randomized action $\pi_i(\cdot | I) \in \Delta(A_I)$ at each infoset

¹Whenever relevant for computational results, we assume all numbers to be rationals represented in binary.

$I \in \mathcal{I}_i$. We say π_i is *pure* if it assigns probability 1 to a single action for each infoset. A (strategy) *profile* $\pi = (\pi_i)_{i \in \mathcal{N}}$ specifies a strategy for each player. We use the common notation $\pi_{-i} = (\pi_1, \dots, \pi_{i-1}, \pi_{i+1}, \dots, \pi_n)$. We denote the strategy set of Player i with S_i , and $S = \times_{i \in \mathcal{N}} S_i$.

We denote the reach probability of a node h' from another node h under a profile π as $\mathbb{P}(h' | \pi, h)$. It evaluates to 0 if $h \notin \text{hist}(h')$, and otherwise to the product of probabilities with which the actions on the path from h to h' are taken under π and chance. We denote with $u_i(\pi | h) := \sum_{z \in \mathcal{Z}} \mathbb{P}(z | \pi, h) \cdot u_i(z)$ the expected utility of Player i given that the game is at node h and the players are following profile π . We overload notation for the special case the game starts at root node h_0 by defining $\mathbb{P}(h | \pi) := \mathbb{P}(h | \pi, h_0)$ and $u_i(\pi) := u_i(\pi | h_0)$. Finally, let I^{1st} refer to the nodes $h \in I$ for which I does not appear in $\text{seq}(h)$. Then the reach probability of I (from h_0) is $\mathbb{P}(I | \pi) := \sum_{h \in I^{\text{1st}}} \mathbb{P}(h | \pi)$.

Solution concepts

The value of recall, which we introduce in the next section, does not only depend on the underlying game, but also on our assumptions on what reasoning capabilities each player has. These are formally captured by *solution concepts*.

Nash equilibrium This is the most classic solution concept in game theory (Nash 1950).

Definition 3. A profile $\pi \in S$ is a *Nash equilibrium* of a game Γ if for each player $i \in \mathcal{N}$,

$$\pi_i \in \operatorname{argmax}_{\pi'_i \in S_i} u_i(\pi'_i, \pi_{-i}). \quad (1)$$

In the special case that Γ is a single-player game, we use the term *optimal strategy* instead of Nash equilibrium.

Unfortunately, a Nash equilibrium is hard to compute, even in a single-player game with imperfect recall (Koller and Megiddo 1992; Gimbert, Paul, and Srivathsan 2020; Tewolde et al. 2024). To make matters worse, it may not even exist (Wichardt 2008). This motivates considering two relaxations based on the *multiselves approach* (Kuhn 1953).

Multiselves equilibria The multiselves approach interprets a player with imperfect recall as a team of multiple instantiations of the player (referred to as *agents* to distinguish from the original player) who independently act at distinct infosets on behalf of the original imperfect-recall player.

For strategy $\pi_i \in S_i$ of Player i , infoset $I \in \mathcal{I}_i$, and randomized action $\sigma \in \Delta(A_I)$, we denote by $\pi_i^{I \rightarrow \sigma}$ the strategy that plays according to π_i except at I , where it plays σ .

Definition 4. A profile $\pi \in S$ is an *EDT equilibrium* of a game Γ if for each player $i \in \mathcal{N}$ and each of its infosets $I \in \mathcal{I}_i$, the randomized action $\pi_i(\cdot | I)$ satisfies $\pi_i(\cdot | I) \in \operatorname{argmax}_{\sigma \in \Delta(A_I)} u_i(\pi_i^{I \rightarrow \sigma}, \pi_{-i})$.

EDT abbreviates *evidential decision theory*; we refer to Piccione and Rubinstein (1997); Briggs (2010); Oesterheld and Conitzer (2024) for a detailed treatment and motivation. A third equilibrium concept that arose from the aforementioned literature is based on *causal decision theory (CDT)*. It differentiates from EDT only in games with *absentmindedness*, which is when a single infoset I appears

multiple times in $\text{seq}(h)$ for some $h \in \mathcal{H}$ (Figure 2, left). Its original definition is not central to this work and deferred to the appendix of the full version of the paper. Instead, below we give an equivalent characterization of it (Tewolde et al. 2024) using *Karush-Kuhn-Tucker (KKT)* points (Boyd and Vandenberghe 2004), which generalize the concept of a *stationary point* of a function over an unconstrained domain.

Definition 5. A profile $\pi \in S$ is a *CDT equilibrium* of a game Γ if for each player $i \in \mathcal{N}$, strategy π_i is a KKT point of the utility maximization problem (1).

These solution concepts form a strict inclusion hierarchy.

Lemma 6 (Oosterheld and Conitzer 2024). *Nash equilibria are EDT equilibria, which in turn are CDT equilibria.*

In particular, Nash equilibria are the hardest to compute, but they coincide in the following special cases.

Remark 7. *EDT and CDT equilibria coincide in games without absentmindedness. Nash and EDT equilibria coincide in games with only one infoset per player.*

Value of Recall

To introduce the novel concept of the value of recall, we first formalize an ordering among infoset partitionings:

Definition 8 (Game refinements/coarsenings). Given two extensive-form games Γ and Γ' with the same game tree and utilities but potentially different infosets $\{\mathcal{I}_i\}_{i \in \mathcal{N}}$ and $\{\mathcal{I}'_i\}_{i \in \mathcal{N}}$, and a player $i \in \mathcal{N}$, we denote $\Gamma \succeq_i \Gamma'$ if for each $I' \in \mathcal{I}'_i$, there exists $\mathcal{J}_i \subseteq \mathcal{I}_i$ such that $I' = \bigsqcup_{I \in \mathcal{J}_i} I$. That is, the infosets in \mathcal{I}' are (disjointly) partitioned by the infosets in \mathcal{I} . In this case, we say Γ (resp. Γ') is a *refinement* (*coarsening*) of Γ' (Γ) with respect to player i . We denote $\Gamma \succeq \Gamma'$ if $\Gamma \succeq_i \Gamma'$ for all $i \in \mathcal{N}$ and say Γ (resp. Γ') is an *all-player refinement* (*coarsening*) of Γ' (Γ).

We are now ready to define the perfect recall refinement of an imperfect-recall game.

Definition 9 (Perfect recall refinements). Given imperfect-recall game Γ , for all nodes $h \in \mathcal{H}$ and players $i \in \mathcal{N}$, define $\text{seq}_i(h)$ as in Definition 2. For infoset $I \in \mathcal{I}_i$ and nodes $h, h' \in I$, define the equivalence relation $h \sim h'$ if $\text{seq}_i(h) = \text{seq}_i(h')$. We say that the (*coarsest*) *perfect recall refinement* of Γ with respect to player $i \in \mathcal{N}$ is an extensive-form game $\mathcal{R}_i(\Gamma)$ with the same game tree and utilities as Γ , but an infoset partition where each $I \in \mathcal{I}_i$ is partitioned into infosets defined by the equivalence relation \sim , and the infosets of all other players are unchanged. The *all-player* (*coarsest*) perfect recall refinement of Γ is an extensive-form game $\mathcal{R}(\Gamma)$ with the same game tree as Γ , where the infosets of all players are partitioned as above.

An equivalent definition to $\mathcal{R}(\Gamma)$ was introduced by Čermák et al. (2018). Both $\mathcal{R}_i(\Gamma)$ and $\mathcal{R}(\Gamma)$ are well-defined, easy to compute, with $\mathcal{R}_i(\Gamma) \succeq_i \Gamma$ and $\mathcal{R}(\Gamma) \succeq \Gamma$. As claimed, $\mathcal{R}_i(\Gamma)$ is the coarsest refinement of Γ with respect to i that gives i perfect recall. We formalize this below:

Proposition 10. *Given imperfect-recall game Γ and another game Γ' that has the same game tree as Γ but potentially*




Figure 2: (Left) An imperfect-recall game Γ . Boxes indicate chance nodes. (Middle) $\mathcal{R}_1(\Gamma)$, the perfect recall refinement of Γ with respect to \blacktriangle . (Right) Γ with perfect information.

*different infosets, if $\Gamma' \succeq_i \Gamma$ and i has perfect recall in Γ' , then $\Gamma' \succeq_i \mathcal{R}_i(\Gamma)$. Moreover, i has perfect recall in $\mathcal{R}_i(\Gamma)$.*²

(Most proofs are in the appendix due to space constraints.)

Corollary 11. *If $\Gamma' \succeq \Gamma$ and Γ' is a perfect-recall game, then $\Gamma' \succeq \mathcal{R}(\Gamma)$. Moreover, $\mathcal{R}(\Gamma)$ is a perfect-recall game.*

By using the coarsest refinement, we seek to isolate the impact of recall on the utility, while filtering out other factors. For instance, the optimal strategy for P1 (\blacktriangle) in game Γ in Figure 2(left) is to play **L** with probability 1/3, bringing an expected utility of 2/3. If we give the player perfect information, and hence perfect recall in the process, the player can achieve a utility of 2 (Figure 2, right). However, we argue this refinement misrepresents the “value of recall” of this game, since P1 now learns the outcome of the chance node, unlike in Γ . Instead, using the coarsest perfect recall refinement, $\mathcal{R}(\Gamma)$ per Def. 9, leads to utility 3/2 (Figure 2, middle) and properly captures what P1 can gain if its only advantage is to remember everything it once knew.

The previous example notwithstanding, we should caution that distinguishing perfect recall and perfect information can become blurry: any imperfect information game can be turned into a strategically-equivalent one with only imperfect recall by adding dummy nodes, as we demonstrate in the appendix.

Now, given an imperfect-recall game Γ , a player of interest (always labelled Player 1), and a solution concept **SC**, let $u_1(\text{SC}(\Gamma))$ be the utility that Player 1 receives under that solution concept in game Γ , assuming it exists. In order to ensure that the utility under **SC** is uniquely defined (since, for example, there might be multiple Nash equilibria of Γ with different utilities for Player 1), we also require **SC** to specify whether it is the best or worst possible outcome of that solution concept from Player 1’s perspective; this is similar to the definition of solution concepts in the value of commitment (Letchford, Korzhyk, and Conitzer 2014). In particular, $\underline{\text{Nash}}$, $\underline{\text{EDT}}$, $\underline{\text{CDT}}$ (resp. $\overline{\text{Nash}}$, $\overline{\text{EDT}}$, $\overline{\text{CDT}}$) refer to the best (worst) possible outcome for Player 1 under the corresponding solution concept.

²While intuitive, this last statement is not just definitional: even though nodes h, h' are placed in the same infoset of $\mathcal{R}_i(\Gamma)$ only if $\text{seq}_i(h) = \text{seq}_i(h')$, the infosets in these sequences are also potentially partitioned, causing the sequences to change too.

Definition 12. Given solution concept SC and Γ , the value of recall (VoR) in Γ under SC is

$$\text{VoR}^{\text{SC}}(\Gamma) = \frac{u_1(\text{SC}(\mathcal{R}_1(\Gamma)))}{u_1(\text{SC}(\Gamma))}.$$

If we are instead given a game class \mathcal{C} , we say that the value of recall (VoR) in \mathcal{C} under SC is

$$\text{VoR}^{\text{SC}}(\mathcal{C}) = \sup_{\Gamma \in \mathcal{C}} \frac{u_1(\text{SC}(\mathcal{R}_1(\Gamma)))}{u_1(\text{SC}(\Gamma))}.$$

We can now formalize the situation that arises in Figure 1 and was discussed earlier in the introduction. To do so, we note that strategies π and π' are *realization-equivalent* if they induce the same reach probability $\mathbb{P}(h \mid \pi) = \mathbb{P}(h \mid \pi')$ for all $h \in \mathcal{H}$ (thus achieving the same utility).

Proposition 13. For any $\varepsilon > 0$, there exists a two-player game Γ such that $\frac{u_i(\text{SC}(\mathcal{R}_1(\Gamma)))}{u_i(\text{SC}(\Gamma))} \leq \varepsilon$ for all $i \in \mathcal{N}$, where SC is the only CDT equilibrium of Γ , up to realization equivalence. In particular, $\text{VoR}^{\text{SC}}(\Gamma) = 0$ for $\text{SC} \in \{\underline{\text{CDT}}, \overline{\text{CDT}}, \underline{\text{EDT}}, \overline{\text{EDT}}, \underline{\text{Nash}}, \overline{\text{Nash}}\}$.

Computational complexity of value of recall

We now show that computing the value of recall is hard. For this theorem alone, we assume (WLOG) for all $z \in \mathcal{Z}$ that $u_1(z) \geq \eta$ for some $\eta > 0$, to ensure VoR is bounded.

Theorem 1. Given a game Γ , computing $\text{VoR}^{\text{SC}}(\Gamma)$ is NP-hard for $\{\underline{\text{CDT}}, \overline{\text{CDT}}, \underline{\text{EDT}}, \overline{\text{EDT}}, \underline{\text{Nash}}, \overline{\text{Nash}}\}$. Moreover,

1. Unless $\text{NP} = \text{ZPP}$, none of them admits an FPTAS. In particular, if $\text{SC} \in \{\underline{\text{CDT}}, \underline{\text{EDT}}\}$, then approximation to any multiplicative factor is NP-hard.
2. NP-hardness and conditional inapproximability holds even if Γ is a single-player game.

A fully polynomial-time approximation scheme (FPTAS) takes as input a game Γ , a solution concept SC , and an $\varepsilon > 0$ and outputs a number in the interval $(1 \pm \varepsilon)\text{VoR}^{\text{SC}}(\Gamma)$. Further, ZPP contains the class of problems solvable by randomized algorithms that always return the correct answer, and whose expected running time is polynomial (Gill 1977).

Most of the proof of Theorem 1 relies on existing hardness results for equilibrium computation in (single-player) imperfect-recall games (Koller and Megiddo 1992; Tewolde et al. 2023; Gimbert, Paul, and Srivathsan 2020). The results for $\underline{\text{CDT}}$ and $\underline{\text{EDT}}$ are new, further establishing stronger inapproximability; both proofs proceed by reducing from 3SAT, as we elaborate in the appendix.

VoR pathologies and how to fix them

While Proposition 13 shows that getting recall can hurt in general, one would expect this to not be the case in single-player games. Indeed, without any opponents, we would expect giving recall to only benefit the player, since it can always ignore the information it can now recall. This is the case if SC represents the optimal strategy (Opt), as getting perfect recall expands the strategy set of a player. Further, since the optimal strategy of a game is also its best CDT and EDT equilibrium (Lemma 6), we have the following:




Figure 3: Perfect recall can lead to worse CDT/EDT eq.

Proposition 14. For any single-player game Γ ,

$$\text{VoR}^{\text{Opt}}(\Gamma) = \text{VoR}^{\overline{\text{EDT}}}(\Gamma) = \text{VoR}^{\overline{\text{CDT}}}(\Gamma) \geq 1.$$

Surprisingly, it turns out that this result in fact does not hold for worst EDT and CDT equilibria of the game:

Example 15. Consider the game Γ in Figure 3a. The only CDT/EDT equilibrium of Γ is the optimal strategy: always play **L**, bringing a utility of 1. In $\mathcal{R}_1(\Gamma)$, however, while the same is still the optimal strategy (and hence a CDT and EDT equilibrium), there is now a second CDT and EDT equilibrium: always play **R** on I_1 and I_{21} , and always play **L** on I_{22} , bringing a utility of ε . Hence, $\text{VoR}^{\underline{\text{EDT}}}(\Gamma) = \text{VoR}^{\underline{\text{CDT}}}(\Gamma) = \varepsilon$, which can be arbitrarily close to 0.

The issue in Example 15 is that of the chicken or the egg: the unreasonable strategy of playing **R** at I_{21} cannot violate CDT/EDT conditions if the player never visits I_{21} , while if the strategy in I_{21} is unreasonable enough then the decision to not visit I_{21} also does not violate them. This shows that CDT/EDT conditions (which, again, are relaxations of Nash equilibrium) are perhaps *too permissive*, accepting strategies that are not reasonable under perfect recall. To rule out such equilibria, we now introduce *equilibrium refinements* for both solution concepts. (It is important to differentiate between *equilibrium refinements*—which narrows the definition of a solution concept—and *information refinements*, per Definition 8—which introduces a new game where players have finer infosets.) The refinements of CDT/EDT that we introduce will force the player to consider its behavior in all infosets it *could have* reached, hence preventing pathologies such as Example 15.

The appropriate equilibrium refinement for CDT has been introduced by Lambert, Marple, and Shoham (2019), which we will refer to as *CDT-Nash*. Due to space constraints, we defer its definition to the appendix. Below, we introduce an analogous, novel refinement called *EDT-Nash*. The relevant properties of both refinements are in Propositions 18 and 19.

Definition 16. A strategy π in a single-player game Γ is *EDT-limit-rational* if there is a sequence $(\pi^{(k)}, \varepsilon^{(k)})_{k \in \mathbb{N}}$ s.t.

1. each $\pi^{(k)}$ is a strategy in Γ such that $\pi^{(k)}(a \mid I) > 0$ for all I and a , and $(\pi^{(k)})_{k \in \mathbb{N}}$ converges to π ;
2. each $\varepsilon^{(k)} > 0$ and $(\varepsilon^{(k)})_{k \in \mathbb{N}}$ converge to 0; and
3. for each k , for all I with $\mathbb{P}(I \mid \pi^{(k)}) > 0$ and $\sigma \in \Delta(A_I)$,

$$\frac{1}{\mathbb{P}(I \mid \pi^{(k)})} \cdot \left(u_1(\pi^{(k)}, I \mapsto \sigma) - u_1(\pi^{(k)}) \right) \leq \varepsilon^{(k)}.$$

Intuitively, the sequence of fully mixed strategies prevents the player from ignoring any infosets it could have reached.

Definition 17. A profile π is an *EDT-Nash equilibrium* of Γ if it is an EDT equilibrium and if for all $i \in \mathcal{N}$, and in the single-player perspective of Γ (where every other player plays fixed π_{-i}), the strategy π_i is realization-equivalent to an EDT-limit-rational strategy π .

The key property of our refinement is that it agrees with the optimal strategy under perfect recall.

Proposition 18. *EDT-Nash equilibria are EDT equilibria. Without absentmindedness, a strategy profile is an EDT-Nash equilibrium iff it is a CDT-Nash equilibrium. Under perfect recall, a strategy profile is an EDT-Nash equilibrium iff it is a Nash equilibrium.*

An analogous result was shown by Lambert, Marple, and Shoham (2019) for CDT-Nash equilibria:

Proposition 19 (Lambert, Marple, and Shoham 2019). *CDT-Nash equilibria are CDT equilibria, and they always exist. Under perfect recall, a strategy profile is a CDT-Nash equilibrium iff it is a Nash equilibrium of Γ .*

The above propositions imply that in a single-player game with perfect recall, the only CDT-Nash and EDT-Nash equilibrium is the optimal strategy. Combined with Proposition 14, this shows that the refinements successfully resolve the pathologies that arose with CDT/EDT.

Corollary 20. *Given single-player game Γ , $\text{VoR}^{\text{SC}}(\Gamma) \geq 1$ for $\text{SC} \in \{\text{CDT-Nash}, \text{CDT-Nash}, \text{EDT-Nash}, \text{EDT-Nash}\}$.*

Bounding the Value of Recall

In this section, we first focus on bounding VoR^{Opt} for single-player games, and show that while it can be arbitrarily large in general, we can still parameterize it using properties of the game tree and the utility functions. Later on, we show how VoR^{SC} for other solution concepts can be bounded in conjunction with these parametrizations.

A key observation is that in single-player games, there are exactly two factors that can lead to a change in the optimal utility when perfect recall is introduced: (1) absentmindedness, and (2) chance nodes. Indeed, if neither is present, the optimal utility remains unchanged.

Proposition 21. *For a single-player game Γ with no chance nodes and with no absentmindedness, $\text{VoR}^{\text{Opt}}(\Gamma) = 1$. Further, for both Γ and $\mathcal{R}_1(\Gamma)$, there is a pure optimal strategy.*

As we will show, either absentmindedness or chance nodes is sufficient to have a game with $\text{VoR}^{\text{Opt}}(\Gamma) > 1$. We first deal with each of these cases separately, before moving on to games that exhibit both.

VoR due to absentmindedness To bound the impact of absentmindedness, we first parameterize the number of times an infoset is visited and an action is taken on the way to a leaf node. Given a single-player game Γ , for any $z \in \mathcal{Z}$ with $\text{seq}(z) = (i_k, I_k, a_k)_{k=0}^{\text{depth}(z)-1}$, for each $I \in \mathcal{I}_1$ and $a \in A_I$ let $n_z(I) = |\{k : I_k = I\}|$, $n_z(a) = |\{k : a_k =$




Figure 4: (Left) Example 24, $n = 4$. (Right) Ex. 28, $n = 2$

$a\}$], and $p_z(a) = \frac{n_z(a)}{n_z(I)}$. Then, we define

$$\alpha(z) = \prod_{\substack{I \in \mathcal{I}_1 : n_z(I) > 1 \\ a \in A_I : n_z(a) > 0}} p_z(a)^{n_z(a)} \in (0, 1]$$

to be the *absentmindedness coefficient* of z . Intuitively, it describes how easy it is to reach z under absentmindedness:

Lemma 22. *Given single-player game Γ with no chance nodes, for all $z \in \mathcal{Z}$, there exists a strategy π_z that reaches z with probability $\alpha(z)$, achieving $u_1(\pi_z) \geq \alpha(z)u_1(z)$.*

We are now ready to introduce our upper bounds for VoR in terms of the absentmindedness coefficients:

Proposition 23. *In a single-player game Γ without chance nodes, we have*

$$\text{VoR}^{\text{Opt}}(\Gamma) \leq \frac{\max_{z \in \mathcal{Z}} u_1(z)}{\max_{z \in \mathcal{Z}} \alpha(z)u_1(z)} \leq \frac{1}{\alpha(z^*)}$$

where $z^* = \text{argmax}_{z \in \mathcal{Z}} u_1(z)$.

As we see next, the inequalities in Proposition 23 are tight.

Example 24. *Consider a single-player game Γ where Lenny needs to pick between the action **L** and the action **R** for n consecutive rounds for some even n . He gets utility 1 if he first plays **L** exactly $n/2$ times followed by **R** the remaining $n/2$ times. If he does anything else, the game is over and he gets 0 utility. Moreover, his memory is reset each time.*

The game tree of Γ has n nodes, $n+1$ leaves, and a single infoset I . Figure 4(Left) depicts this for $n = 4$. Let z^ be the single leaf node with $u_1(z^*) = 1$. The optimal strategy in $\mathcal{R}_1(\Gamma)$ (where each node is its own infoset) is to arrive at z^* , guaranteeing a utility of 1. In Γ , however, Lenny cannot do anything better than playing uniformly at random, achieving an expected utility of 2^{-n} . Moreover, $\alpha(z^*) = 2^{-n}$. Hence, for Γ , all of the inequalities in Proposition 23 are tight.*

Example 24 is the worst-case scenario with regard to absentmindedness: only one leaf node brings positive utility, and reaching it requires playing each action equally often.

Importantly, $\alpha(z)$ is independent of the utilities of Γ . This allows us to interpret Proposition 23 in two parts: a tighter bound on Γ using its utilities, and another bound that applies to all games that differ from Γ only by their utility functions.

Corollary 25. Given a single-player game Γ without chance nodes, say \mathcal{C} is the class of games that share the same game tree and infoset partition as Γ . Then

$$\text{VoR}^{\text{Opt}}(\mathcal{C}) = \max_{z \in \mathcal{Z}} \frac{1}{\alpha(z)}.$$

VoR due to chance nodes We now do a similar analysis for chance nodes. Given a single-player game Γ , for any $z \in \mathcal{Z}$ with $\text{seq}(z) = (i_k, I_k, a_k)_{k=0}^{\text{depth}(z)-1}$, say k_1, k_2, \dots, k_ℓ are steps that correspond to chance nodes, i.e., $i_{k_j} = c$ for all $j \in [\ell]$. Then, the *chance coefficient* of leaf node z is

$$\chi(z) = \prod_{j=1}^{\ell} \mathbb{P}_c(a_{k_j} | h_{k_j})$$

if $\ell > 0$ and $\chi(z) = 1$ otherwise. $\chi(z)$ is the probability of reaching z in $\mathcal{R}_1(\Gamma)$ (i.e., under perfect recall), given that the player is trying to reach it. For each chance node $h \in \mathcal{H}_c$, and each $a \in A_h$, say $H_{ha} \subset \mathcal{H}_c$ are the chance nodes in the subtree rooted at ha (the node reached when chance plays a at h). Then the *branching factor* of h is $\beta(h) = \sum_{a \in A_h} b_h(a)$, where

$$b_h(a) = \begin{cases} 1 & \text{if } |H_{ha}| = 0 \\ \max_{h \in H_{ha}} \beta(h) & \text{otherwise} \end{cases}.$$

One can compute $\beta(h)$ for each $h \in \mathcal{H}_c$ recursively, starting from the bottom of the tree (that is, the leaf nodes). We now have all the tools we need for characterizing the impact of chance nodes on VoR.

Proposition 26. In a single-player game Γ without absentmindedness, we have³

$$\text{VoR}^{\text{Opt}}(\Gamma) \leq \frac{u_1(\text{Opt}(\mathcal{R}_1(\Gamma)))}{\max_{z \in \mathcal{Z}} \chi(z) u_1(z)} \leq \max_{h \in \mathcal{H}_c} \beta(h).$$

Corollary 27. Given a single-player game Γ without absentmindedness, say \mathcal{C} is the class of games that share the same game tree and infoset partition as Γ . Then

$$\text{VoR}^{\text{Opt}}(\mathcal{C}) \leq \max_{h \in \mathcal{H}_c} \beta(h).$$

The reason we have an inequality for the game class, unlike in Corollary 25, is that while absentmindedness does imply imperfect recall, chance nodes alone do not tell us anything about the information structure of the game. We now show that the bounds in Proposition 26 are also tight.

Example 28. Consider a game Γ that starts with a single chance node h_c with $|A_{h_c}| = n$, each played with equal probability. Under each outcome, Dory needs to act twice, using the same action set as chance A_{h_c} , and gets utility 1 only if she replicates the action of the chance node both times, and 0 otherwise. Each of Dory's nodes immediately following the chance node is in its own information set of size 1, and every other node is in a single information set. Figure 4(Right) shows the game tree for $n = 2$.

In $\mathcal{R}_1(\Gamma)$, Dory has perfect information and can guarantee utility 1. However, with imperfect recall, the best she

can do is select the correct action the first time she acts, and then any strategy she will follow on the large information set will bring her expected utility $1/n$. Moreover, $\beta(h_c) = n$ and $\max_{z \in \mathcal{Z}} \chi(z) u_1(z) = 1/n$, showing that for Γ all the inequalities in Proposition 26 are tight.

We end this section by showing that our results from Propositions 23 and 26 do in fact compose, hence giving a parameterization of VoR^{Opt} for any single-player game.

Theorem 2. For a single-player game Γ , say \mathcal{C} is the class of games that share the same game tree and infosets. Then

$$\text{VoR}^{\text{Opt}}(\mathcal{C}) \leq \max_{z \in \mathcal{Z}, h \in \mathcal{H}_c} \frac{\beta(h)}{\alpha(z)}.$$

Smooth imperfect-recall games

Remaining on single-player games, here we bound the value of recall for a broader set of equilibria. Our approach is driven by a connection with the *price of anarchy* (PoA). In particular, we introduce the notion of a *smooth* (single-player) imperfect-recall game, which is based on the homonymous class of (multi-player) games by Roughgarden (2015). Below, we denote by $(\pi_I)_{I \in \mathcal{I}_1} \in S$ the player's strategy, and use the notation $\pi_{-I} := (\pi_{I'})_{I' \in \mathcal{I}_1 \setminus \{I\}}$.

Definition 29. A single-player game Γ is (λ, μ) -smooth if there exists $\pi^* \in S$ such that for any $\pi \in S$,

$$\frac{1}{|\mathcal{I}_1|} \sum_{I \in \mathcal{I}_1} u_1(\pi_I^*, \pi_{-I}) \geq \lambda u_1(\text{Opt}(\Gamma)) - \mu u_1(\pi). \quad (2)$$

The rationale behind this definition is that it enables disentangling the left-hand side of (2) via a suitable strategy π^* , with the property that if followed by each infoset separately, a non-trivial fraction of the optimal utility can be secured *no matter the strategy in the rest of the infosets*. While this may appear like an overly restrictive property, it manifests itself in many important applications (Roughgarden, Syrgkanis, and Tardos 2017). In Definition 29, infosets play the role of strategic players in Roughgarden's formalism; we provide a concrete example of a smooth imperfect-recall game in the appendix.

Now, by definition, an EDT equilibrium π satisfies $u_1(\pi) \geq \frac{1}{|\mathcal{I}_1|} \sum_{I \in \mathcal{I}_1} u_1(\pi_I^*, \pi_{-I})$ (by applying Definition 4 successively for each infoset). Combining with (2), we immediately arrive at the following conclusion.

Proposition 30. Let Γ be a (λ, μ) -smooth, single-player game. For any EDT equilibrium $\pi \in S$,

$$u_1(\pi) \geq \frac{\lambda}{1 + \mu} u_1(\text{Opt}(\Gamma)).$$

In words, $\rho := \lambda/(1 + \mu)$ measures the degradation incurred in an EDT equilibrium, which is referred to as the *robust price of anarchy* in the parlance of Roughgarden (2015). In light of Proposition 30, bounding $\text{VoR}^{\text{EDT}}(\Gamma)$ reduces to relating $u_1(\text{Opt}(\Gamma))$ in terms of $u_1(\text{Opt}(\mathcal{R}(\Gamma)))$, which was accomplished earlier in Theorem 2.

Corollary 31. Let Γ be a (λ, μ) -smooth, single-player game. Then

$$\text{VoR}^{\text{EDT}}(\Gamma) \leq \frac{1 + \mu}{\lambda} \max_{z \in \mathcal{Z}, h \in \mathcal{H}_c} \frac{\beta(h)}{\alpha(z)}.$$

³By convention, we assume $\max_{h \in \mathcal{H}_c} \beta(h) = 1$ if $\mathcal{H}_c = \emptyset$.




Figure 5: Partial recall gives a worse EDT-Nash equilibrium: In (a), the only EDT-Nash equilibrium is playing **L**; in (b), playing **R** in both infosets is also an EDT-Nash equilibrium.

This also applies to EDT-Nash always and (by Remark 7) to CDT and CDT-Nash when Γ has no absentmindedness.

Further connections

In addition, we note that the value of recall (Definition 12) encompasses several notions from prior literature. First, the *price of uncorrelation* in adversarial team games (Celli and Gatti 2018), which measures how much of a team of players facing a single adversary can gain from communicating, corresponds to $\text{VOR}^{\text{Nash}}(\mathcal{C}^{2p0s})$, where \mathcal{C}^{2p0s} is the class of two-player zero-sum games (based on their construction, one of the players—corresponding to the adversary—has perfect recall).

Second, the *price of miscoordination* in security games (Jiang et al. 2013), which measures the utility loss due to having multiple defenders rather than a single one, corresponds to the VoR in this game class; here, **SC** corresponds to *Stackelberg equilibria*, which involves Player 1 committing to a strategy and its opponent best responding. We expand on the above connections in the appendix.

Partial Recall Refinements

So far, we have defined the value of recall based on the (coarsest) perfect-recall refinement. It is also natural to consider the change in utility due to obtaining *partial* recall.

Definition 32 (Partial recall refinements). Given games Γ and Γ' with the same game tree but possibly different info sets, Γ' is a *partial recall refinement* of Γ with respect to a player $i \in \mathcal{N}$ if $\mathcal{R}_i(\Gamma) \succeq_i \Gamma' \succeq_i \Gamma$. Further, Γ' is an *all-player partial recall refinement* of Γ if $\mathcal{R}(\Gamma) \succeq \Gamma' \succeq \Gamma$.

Partial recall refinements introduces further interesting properties. For example, Figure 5 shows that partial recall can lead to a worse EDT-Nash equilibrium in a single-player game; this stands in contrast to perfect recall refinements (Corollary 20).

In what follows, we study the complexity of perhaps the most natural problem arising from Definition 32: how should one refine an imperfect-recall game so as to maximize the utility gain, subject to constraining the amount of new recall. This problem is well-motivated from the literature on abstraction (e.g., Kroer and Sandholm 2016), but to our knowledge, it has not been studied in this form. To formalize constraints on recall, we first introduce the following notation: consider games Γ, Γ' that differ solely on info sets for $i \in \mathcal{N}$ (\mathcal{I}_i and \mathcal{I}'_i , respectively). We write $\Gamma' \vdash_i \Gamma$ if there is $I \in \mathcal{I}_i$

such that $\mathcal{I}'_i = \mathcal{I}_i \setminus \{I\} \cup \{I_1, I_2\}$, where $I = I_1 \sqcup I_2$; i.e., Γ' results from splitting a single info set of i in Γ .

Definition 33. Fix a player $i \in \mathcal{N}$. We say Γ is its own *0-partial recall refinement*. Γ' is a *k-partial recall refinement* of Γ if it is a partial recall refinement of Γ and $\Gamma' \vdash_i \Gamma''$, where Γ'' is some $(k - 1)$ -partial recall refinement of Γ .

This restriction is motivated by the fact that many practical algorithms scale with the number of info sets, and so one naturally strives to minimize that when abstracting a game (Kroer and Sandholm 2014, 2016).

Then, the computational problem $k\text{-BESTPARTIAL}(\Gamma)$ asks: given a parameter $k \in \mathbb{N}$ and a single-player game Γ , compute its k -partial recall refinement Γ' that maximizes $u_1(\text{Opt}(\Gamma'))$. For this task, we assume to be given access to an oracle \mathcal{O} that outputs the optimal utility of any single-player game; even though such an oracle can only make the problem easier, we show the following hardness result.

Theorem 3. $k\text{-BESTPARTIAL}(\Gamma)$ is NP-hard.

Our proof relies on a reduction from *exact cover by 3-sets* (Garey and Johnson 1979), which asks to exactly cover a set of items using a given family of subsets of size three. Our construction consists of a chance node with an action per item, followed by player nodes with an action per subset.

Conclusions and Future Research

We introduced the value of recall, which measures the utility gain by granting a player perfect recall. Our work opens many interesting avenues for future research. First, the value of recall could be used to guide abstraction techniques. We also observed the interesting phenomenon that perfect recall can be hurtful to all players. It would be interesting to provide a broader characterization of games where this is so—a natural candidate being *simulation games* (Kovarík, Oesterheld, and Conitzer 2024), and quantify the *price* of recall therein. Furthermore, we have focused on the value of recall from the perspective of a single player, but understanding the impact on *social welfare* is a natural next step.

Acknowledgments

We are grateful to the anonymous AAAI reviewers for many helpful comments that improved the exposition of this paper. We also thank Brian Hu Zhang for many discussions. Ratip Emin Berker, Emanuel Tewolde, and Vincent Conitzer thank the Cooperative AI Foundation, Polaris Ventures (formerly the Center for Emerging Risk Research) and Jaan Tallinn’s donor-advised fund at Founders Pledge for financial support. Tuomas Sandholm is supported by the Vannevar Bush Faculty Fellowship ONR N00014-23-1-2876, National Science Foundation grants RI-2312342 and RI-1901403, ARO award W911NF2210266, and NIH award A240108S001.

References

Bakhtin, A.; Brown, N.; Dinan, E.; Farina, G.; Flaherty, C.; Fried, D.; Goff, A.; Gray, J.; Hu, H.; Jacob, A. P.; Komeili, M.; Konath, K.; Kwon, M.; Lerer, A.; Lewis, M.; Miller, A. H.; Mitts, S.; Renduchintala, A.; Roller, S.; Rowe, D.; Shi, W.; Spisak, J.; Wei, A.; Wu, D.; Zhang, H.; and Zijlstra,

- M. 2022. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science*, 378(6624): 1067–1074.
- Bowling, M.; Burch, N.; Johanson, M.; and Tammelin, O. 2015. Heads-up Limit Hold'em Poker is Solved. *Science*, 347(6218): 145–149.
- Boyd, S.; and Vandenberghe, L. 2004. *Convex Optimization*. Cambridge University Press.
- Braess, D. 1968. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12: 258–268.
- Briggs, R. 2010. Putting a Value on Beauty. In Gendler, T. S.; and Hawthorne, J., eds., *Oxford Studies in Epistemology: Volume 3*, 3–34. Oxford University Press.
- Brown, N.; Ganzfried, S.; and Sandholm, T. 2015. 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)*.
- Brown, N.; and Sandholm, T. 2018. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374): 418–424.
- Celli, A.; and Gatti, N. 2018. Computational Results for Extensive-Form Adversarial Team Games. In *Conference on Artificial Intelligence (AAAI)*.
- Čermák, J.; Bošanský, B.; Horák, K.; Lisý, V.; and Pěchouček, M. 2018. Approximating maxmin strategies in imperfect recall games using A-loss recall property. *International Journal of Approximate Reasoning*, 93: 290–326.
- Čermák, J.; Bosanský, B.; and Lisý, V. 2017. 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)*.
- Chen, E. O.; Ghersengorin, A.; and Petersen, S. 2024. Imperfect Recall and AI Delegation.
- Conitzer, V. 2019. Designing Preferences, Beliefs, and Identities for Artificial Intelligence. In *Conference on Artificial Intelligence (AAAI)*.
- Emmons, S.; Oesterheld, C.; Critch, A.; Conitzer, V.; and Russell, S. 2022. For Learning in Symmetric Teams, Local Optima are Global Nash Equilibria. In *International Conference on Machine Learning (ICML)*.
- Foerster, K.-T.; and Wattenhofer, R. 2013. The solitaire memory game. Technical report, ETH Zurich.
- Fudenberg, D.; and Tirole, J. 1991. *Game Theory*. MIT Press.
- Ganzfried, S.; and Sandholm, T. 2014. Potential-Aware Imperfect-Recall Abstraction with Earth Mover's Distance in Imperfect-Information Games. In *Conference on Artificial Intelligence (AAAI)*.
- Garey, M.; and Johnson, D. 1979. *Computers and Intractability*. W. H. Freeman and Company.
- Gill, J. 1977. Computational Complexity of Probabilistic Turing Machines. *SIAM Journal on Computing*, 6(4): 675–695.
- Gimbert, H.; Paul, S.; and Srivathsan, B. 2020. A Bridge between Polynomial Optimization and Games with Imperfect Recall. In *Autonomous Agents and Multi-Agent Systems*.
- Jiang, A. X.; Procaccia, A. D.; Qian, Y.; Shah, N.; and Tambe, M. 2013. Defender (Mis)coordination in Security Games. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Johanson, M.; Burch, N.; Valenzano, R.; and Bowling, M. 2013. Evaluating State-Space Abstractions in Extensive-Form Games. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Kirkpatrick, P. 1954. Probability theory of a simple card game. *The Mathematics Teacher*, 47(4): 245–248.
- Koller, D.; and Megiddo, N. 1992. The Complexity of Two-Person Zero-Sum Games in Extensive Form. *Games and Economic Behavior*, 4(4): 528–552.
- Kovarič, V.; Oesterheld, C.; and Conitzer, V. 2023. Game Theory with Simulation of Other Players. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.
- Kovarič, V.; Oesterheld, C.; and Conitzer, V. 2024. Recursive Joint Simulation in Games. *CoRR*, abs/2402.08128.
- Kroer, C.; and Sandholm, T. 2014. Extensive-Form Game Abstraction With Bounds. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- Kroer, C.; and Sandholm, T. 2016. Imperfect-Recall Abstractions with Bounds in Games. In *Proceedings of the ACM Conference on Economics and Computation (EC)*.
- Kuhn, H. W. 1953. Extensive Games and the Problem of Information. In *Contributions to the Theory of Games*, volume 2 of *Annals of Mathematics Studies*, 28, 193–216. Princeton University Press.
- Lambert, N. S.; Marple, A.; and Shoham, Y. 2019. On equilibria in games with imperfect recall. *Games and Economic Behavior*, 113: 164–185.
- Letchford, J.; Korzhyk, D.; and Conitzer, V. 2014. On the value of commitment. *Autonomous Agents and Multi-Agent Systems*, 28: 986–1016.
- Meisheri, H.; and Khadilkar, H. 2020. Sample Efficient Training in Multi-Agent Adversarial Games with Limited Teammate Communication.
- 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.
- Nash, J. 1950. *Non-cooperative games*. Ph.D. thesis, Princeton University.
- Oesterheld, C.; and Conitzer, V. 2024. Can *de se* choice be *ex ante* reasonable in games of imperfect recall? A complete analysis. <https://www.andrew.cmu.edu/user/coesterh/DeSeVsExAnte.pdf>. Working paper. Accessed: 2024-07-13.
- Piccione, M.; and Rubinstein, A. 1997. On the Interpretation of Decision Problems with Imperfect Recall. *Games and Economic Behavior*, 20: 3–24.

Resnick, C.; Gao, C.; Márton, G.; Osogami, T.; Pang, L.; and Takahashi, T. 2020. Pommerman & NeurIPS 2018. In *The NeurIPS '18 Competition*, 11–36. Springer International Publishing.

Roughgarden, T. 2015. Intrinsic Robustness of the Price of Anarchy. *Journal of the ACM*, 62(5): 32:1–32:42.

Roughgarden, T.; Syrgkanis, V.; and Tardos, É. 2017. The Price of Anarchy in Auctions. *Journal of Artificial Intelligence Research*, 59: 59–101.

Tewolde, E.; Oesterheld, C.; Conitzer, V.; and Goldberg, P. W. 2023. The Computational Complexity of Single-Player Imperfect-Recall Games. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Tewolde, E.; Zhang, B. H.; Oesterheld, C.; Zampetakis, M.; Sandholm, T.; Goldberg, P. W.; and Conitzer, V. 2024. Imperfect-Recall Games: Equilibrium Concepts and Their Complexity. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Thorp, E. O. 2016. *Beat the dealer: A winning strategy for the game of twenty-one*. Vintage.

von Stengel, B.; and Koller, D. 1997. Team-Maxmin Equilibria. *Games and Economic Behavior*, 21(1): 309–321.

Waugh, K.; Zinkevich, M.; Johanson, M.; Kan, M.; Schnizlein, D.; and Bowling, M. 2009. A Practical Use of Imperfect Recall. In *Symposium on Abstraction, Reformulation and Approximation (SARA)*.

Wichardt, P. C. 2008. Existence of Nash equilibria in finite extensive form games with imperfect recall: A counterexample. *Games and Economic Behavior*, 63(1): 366–369.

Zhang, B. H.; Farina, G.; and Sandholm, T. 2023. Team Belief DAG: Generalizing the Sequence Form to Team Games for Fast Computation of Correlated Team Max-Min Equilibria via Regret Minimization. In *International Conference on Machine Learning (ICML)*.

Zhang, B. H.; and Sandholm, T. 2022. Polynomial-Time Optimal Equilibria with a Mediator in Extensive-Form Games. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*.

Zhang, Y.; An, B.; and Subrahmanian, V. S. 2022. Correlation-Based Algorithm for Team-Maxmin Equilibrium in Multiplayer Extensive-Form Games. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.