

# Mixture of Public and Private Distributions in Imperfect Information Games.

Jérôme Arjonilla  
Paris Dauphine University - PSL  
Paris, France  
jerome.arjonilla@hotmail.fr

Abdallah Saffidine  
University of New South Wales  
Sydney, Australia  
abdallah.saffidine@gmail.com

Tristan Cazenave  
Paris Dauphine University - PSL  
Paris, France  
Tristan.Cazenave@dauphine.psl.eu

**Abstract**—In imperfect information games (e.g. Bridge, Skat, Poker), one of the fundamental considerations is to infer the missing information while at the same time avoiding the disclosure of private information. Disregarding the issue of protecting private information can lead to a highly exploitable performance. Yet, excessive attention to it leads to hesitations that are no longer consistent with our private information. In our work, we show that to improve performance, one must choose whether to use a player’s private information. We extend our work by proposing a new belief distribution depending on the amount of private and public information desired. We empirically demonstrate an increase in performance and, with the aim of further improving performance, the new distribution should be used according to the position in the game. Our experiments have been done on multiple benchmarks and in multiple determinization-based algorithms (PIMC and IS-MCTS).

**Index Terms**—Imperfect Information Games, Search Algorithms, Belief Distributions

## I. INTRODUCTION

Search in artificial intelligence has been constantly evolving over the last few decades, and game-oriented research has always been a cornerstone of this success. Chess, Go [1]–[3], Poker [4], Skat, Contract Bridge or Dota [5] are among the most famous ones.

Perfect information games (Chess, Go) — where all information is available for each player — have been the most studied, and many algorithms have been able to achieve a level far beyond the level of a human professional player. On the other hand, Imperfect Information Games (IIGs) (Poker, Skat, Bridge) — where some information is hidden — have been less studied, and only a few algorithms are capable of beating professional human player [4], [6].

In IIG, the complexity is heightened by the missing information, as one must try to infer the missing information of the opponents and, at the same time, be wary to not reveal private hidden information to opponents. Among the methods used in IIG, determinization-based algorithms — where the hidden information is fixed according to a belief distribution — such as Perfect Information Monte Carlo (PIMC) [7], Recursive PIMC [8], Information Set MCTS [9] or AlphaMu [10] achieve state-of-the-art performance in many trick-taking card games (Contract-Bridge, Skat).

In the work cited above, the determinization operates by sampling the hidden information according to the private information of a given player, *i.e.* what has happened since the beginning, from the point of view of a given agent. However, by doing so, one can indirectly reveal private information to opponents, which can lead to a highly exploitable performance.

Recently, the concept of public knowledge [11] — where a distinction is made between observations accessible to everyone and those accessible individually — has emerged in IIGs. This concept has resulted in many breakthroughs thanks to the decomposition, which made the calculations feasible [12], [13]. Despite this large benefit, there are limitations to its use, especially in the context of belief distribution. By doing so, we completely remove the knowledge observed by the acting player, and one might wonder whether not using the private information was useful.

In this work, we analyze the impact of using one method rather than another and present a new belief distribution, which is a mixture of both public and private belief distribution. We extend the study by analyzing different mixtures, depending on the position within the game. Our experiments are carried out on determinization-based algorithms, which use the belief distribution to fix the incertitude.

The paper is organized as follows: Section II presents notation and current determinization-based algorithms; Section III explains the different belief distributions used with their advantages and drawbacks, and presents our new belief distribution; Section IV empirically shows that using the new belief distribution allows us to improve past performance; and the last section summarizes our work and future work.

## II. NOTATION AND BACKGROUND

### A. Notation

We use the notation based on factored-observation stochastic games (FOSGs [11]). This formalism distinguishes between private and public observations.

A Game  $G$  is composed of the following elements. The set of agents  $\mathcal{N} = \{1, 2, \dots, N\}$  agents, the set of **world state** possible  $\mathcal{W}$ . In each world state  $w \in \mathcal{W}$ , the acting player  $i$  chooses an action  $a \in \mathcal{A}(w)$ , where  $\mathcal{A}(w)$  denotes the legal actions at  $w$ . After an action  $a$ , we reach the next world state  $w'$  from the probability distribution of playing  $a$  in  $w$ .

During the transition from  $w$  to  $w'$  by playing  $a$ , two observations are received: a **public observation** and a **private observation**. Public observation is the observation visible by every player noted  $o_{pub} \in \mathcal{O}_{pub}(w, a, w')$  where  $\mathcal{O}_{pub}(w, a, w')$  refers to all the possible public observations. Private observation is the observation visible by a precise player  $i$ , noted  $o_{priv(i)} \in \mathcal{O}_{priv(i)}(w, a, w')$  where  $\mathcal{O}_{priv(i)}(w, a, w')$  refers to all the possible private observations.

A *history* is a finite sequence of legal actions and world states, denoted  $h^t = (w^0, a^0, w^1, a^1, \dots, w^t)$ . For describing the point of view of an agent  $i$  of a history  $h$ , we introduce an **infostate**  $s_i(h)$ . An **infostate** for agent  $i$  is a sequence of an agent's observations and actions  $s_i^t = (o_i^0, a_i^0, o_i^1, a_i^1, \dots, o_i^t)$  where  $o_i^k = (o_{pub}^k, o_{priv(i)}^k)$ . A **public infostate** is a sequence of public observations  $s_{pub}^t = (o_{pub}^0, o_{pub}^1, \dots, o_{pub}^t)$ .

**Determinization** refers to the fact that we sample a world state according to a belief distribution of the world states possible. Determinizing the belief distribution is not new and a similar concept exists in other formalisms such as belief state in Partially Observable Markov Decision Process (POMDP) problems [14] or occupancy-state in Decentralised-POMDPs problems [15].

### B. Determinization-based algorithms

Each determinization-based algorithm has its own characteristics. Nevertheless, they share some common features such as (i) sampling a world state according to a belief distribution over the possible world states; and (ii) using a perfect information algorithm for estimating the value of the sampled world state.

The algorithms are simple and, in practice, they achieve great results, mainly due to the use of perfect information algorithms (AlphaBeta [16], MCTS [17] or Value Network) that are fast and efficient.

In the following, we present two determinization-based algorithms that are baseline and will, at a later stage, be used in our experiments.

1) *PIMC*: Perfect Information Monte Carlo (PIMC) is the state of the art of many IIG problems such as Contract-Bridge, Skat, and many others.

The algorithm is defined in Algorithm 1 and works as follows: (i) samples a world state by using the player's private information; (ii) plays each action of the sampled world state; (iii) estimates the reward of the new world state by using an algorithm available in perfect information setting; (iv) repeats until the budget is over; and, (v) selects the action that produces the best result in average. In practice, PIMC often uses AlphaBeta as the perfect information evaluator.

2) *IS-MCTS*: Information Set Monte Carlo Tree Search [9] uses Monte Carlo Tree Search (MCTS) [17] according to a sampled world state.

MCTS is a state-of-the-art tree search algorithm in perfect information games. It works as follows (i) selection — selects a path of nodes based on an exploitation policy; (ii) expansion — expands the tree by adding a new child node; (iii)

---

### Algorithm 1: PIMC

---

```

Function PIMC (s) :
  for m ∈ Moves (s) do
    | score[m] ← 0;
  end
  while budget do
    | w ← InfoSampling(s);
    | for m ∈ Moves (w) do
      | | score [m] ← score[m] + PerfectAlgo (w,
      | | m);
    | end
  end
  return Best action on average

```

---

playout — estimates the child node by using an exploration policy; and, (iv) backpropagation — backpropagates the result obtained from the playout through the nodes chosen during the selection phase. In practice, MCTS often uses random playout as the perfect information evaluator, and UCB1 in the selection phase.

IS-MCTS works by using MCTS according to a sampled world state, *i.e.* the selection and playout are done on the sampled world state.

---

### Algorithm 2: IS-MCTS

---

```

Function IS-MCTS (s) :
  while budget do
    | w ← InfoSampling(s);
    | MCTS conditioned on w.;
  end
  return Normalise visit count for each action
Function MCTS (w) :
  u ← Selection(w);
  u ← Expansion (u,w);
  u ← Simulation (u,w);
  Backpropagation(u);

```

---

## III. BELIEF DISTRIBUTIONS

To present the different belief distributions, with their advantages and drawbacks, we use the following example throughout the section to facilitate understanding.

The example is based on the famous game 'Liar's Dice' (an explanation of the game is given in Subsection IV-A2). In our case, two players, each with 1 die of 2 sides. We denote  $\{P_1 : X; P_2 : Y\}$  for player 1 has  $X$  and player 2 has  $Y$ . There are four world states possible ( $w_1 = \{P_1 : 1; P_2 : 1\}$ ,  $w_2 = \{P_1 : 1; P_2 : 2\}$ ;  $w_3 = \{P_1 : 2; P_2 : 2\}$ ,  $w_4 = \{P_1 : 2; P_2 : 1\}$ ).

For each player, there are two infostates possible and one public infostate  $s_{pub} = \{o_{pub}^1 = \emptyset, o_{pub}^2 = \emptyset\}$  (no observation). For the player 1 we have  $s_1 = \{o_{priv(1)}^1 = 1, o_{priv(1)}^2 = \emptyset\}$  or  $s'_1 = \{o_{priv(1)}^1 = 2, o_{priv(1)}^2 = \emptyset\}$  (*i.e.* Player 1 observes the die rolled but not the die rolled by the

other player), and for the player 2, we have  $s_2 = \{o_{priv(2)}^1 = \emptyset, o_{priv(2)}^2 = 1\}$  or  $s'_2 = \{o_{priv(2)}^1 = \emptyset, o_{priv(2)}^2 = 2\}$  (i.e. Player 2 observes the die rolled but not the die rolled by the other player).

In the following, we suppose that the world state of this example is  $w_2$ . Therefore, for the player 1, the infostate is  $s_1$  with two world states possible ( $\{w_1, w_2\}$ ) and for the player 2, the infostate is  $s'_2$  with two world states possible ( $\{w_2, w_3\}$ ). Fig. 1 represents the different belief distributions presented throughout the section.

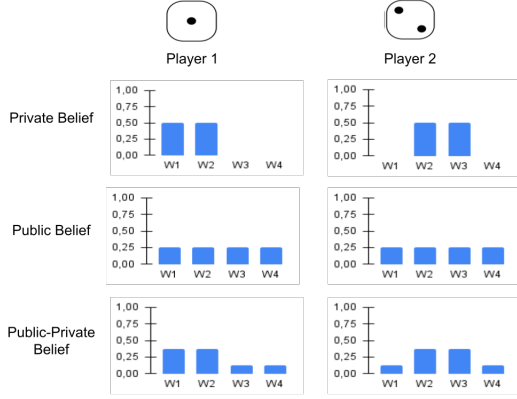


Fig. 1: Multiple belief distributions for the game Liar's Dice with 1 dice of 2 sides each. Four world states possible  $w_1$ ,  $w_2$ ,  $w_3$  and  $w_4$ . The Public-Private belief uses the mixture distribution with  $\lambda = 0.5$ .

### A. Private Distribution

As previously introduced, current determinization-based algorithms work by sampling world states according to the player's private information distribution, i.e. knowing a player's private and public observation, we sample a world state.

Let  $S_j(s_i)$  be the set of possible infostates for player  $j$  conditioned on the infostate  $s_i$  of the player  $i$ . In our example, the infostate possible for the player 2 when the player 1 has  $s_1$  is  $S_2(s_1) = \{s_2, s'_2\}$ , i.e. having the die 1 for the player 1 does not exclude the player 2 to have a 1 or a 2. Depending on the game  $S_j(s_i)$  can be restrictive, e.g. in trick-taking card games if the player  $i$  has the card 'Queen of Hearts', no opponent can have it.

**Definition** (Private Belief Distribution). *Let  $S_j(s_i)$  be the set of possible infostates for player  $j$  conditioned on the infostate  $s_i$ . Let  $\Delta S_j(s_i)$  denotes the probability distribution over the elements of  $S_j(s_i)$ . We define the private belief distribution as  $\Delta_i(s_i) = (\Delta S_1(s_i), \dots, \Delta S_i(s_i), \dots, \Delta S_N(s_i)) = (\Delta S_1(s_i), \dots, s_i, \dots, \Delta S_N(s_i))$ .*

In Fig. 1, using Player 1's private belief state provides the following belief distribution  $\Delta_1(s_1) = (\{s_1 : 100\%, \{s_2 : 50\%; s'_2 : 50\%\})$ , which results in two equiprobable world states ( $w_1, w_2$ ).

When using the private distribution for determinization, the algorithm samples a world state ( $w_1$  or  $w_2$ ) consistent with the current player's information ( $s_1$ ) and, as the state-of-the-art in trick-taking game shows, great performance is obtained. Yet, by doing so, 3 problems arise.

(i) It is not consistent with the other player's belief, e.g. if we use it with the first player, the algorithm samples  $w_1$  or  $w_2$  but never  $w_3$ , which is nevertheless, a world state possible from the point of view of the player 2.

(ii) It is not able to mislead others. In our example, two actions are possible for the first player, 'I have a one' and 'I have a two'. The action 'I have a two' is a lie, however, one may want to play this action with the aim of deceiving the opponent. However, in our case only  $w_1$  or  $w_2$  can be sampled and, in each world, the action 'I have a two' results in a defeat because the second player will say 'This is a lie'. Therefore, lying is never an option, as it never succeeds.

(iii) It, indirectly, allows the opponents to infer our private information, e.g. after playing multiple matches, the second player understands that, if the first player plays 'I have a two', it is because he really has a two as it can not lie, and therefore, play to counter it.

Trying to infer missing information is one of the key components of IIG, and using the private belief distribution could result in a highly exploitable performance. To remove this problem, one can use public belief distribution, as presented in the next section.

### B. Public Distribution

Recently in IIG, many algorithms [12], [13] have been using the concept of public observation. This concept has resulted in many breakthroughs thanks to decomposition, which made the calculations feasible. One application of public observation is the creation of a public belief distribution over the world states possible according to the public observations observed so far.

**Definition** (Public Belief Distribution [13]). *Let  $S_j(s_{pub})$  be the set of possible infostates for player  $j$  conditioned on the public infostate  $s_{pub}$ . Let  $\Delta S_j(s_{pub})$  denote the probability distribution over the elements of  $S_j(s_{pub})$ . We define the public belief distribution as  $\Delta_{pub}(s_{pub}) = (\Delta S_1(s_{pub}), \dots, \Delta S_N(s_{pub}))$ .*

In our example, using the public belief distribution from the point of view of the player 1 would result in the following belief distribution  $\Delta_{pub} = (\{s_1 : 50\%; s'_1 : 50\%\}, \{s_2 : 50\%; s'_2 : 50\%\})$ . In other words, every world state are equiprobable, this is due to the public infostate that does not contain any information.

Using a public belief distribution instead of a private belief distribution removes the problem defined in Section III-A.

(i) It is consistent with the other player's doubts, e.g. it samples the world  $w_3$  which is a world state possible of the second player.

(ii) It is capable of misleading others, *e.g.* when sampling  $w_3$  or  $w_4$  the action ‘I have a two’ does not result in a defeat for the first player, therefore, allows the first player to play the action ‘I have a two’.

(iii) It no longer reveals private information, *i.e.* as the reasoning is no longer biased toward the private information, it can not be used against it.

Nevertheless, using public distribution has a significant drawback as it does not consider a player’s private information, and one might wonder whether it is useful to not use private information. It is straightforward to consider that the extent to which private information should be kept hidden depends on the game being played and, in certain games, it is not necessary to keep the information concealed.

In addition, by using public distribution, one must be aware as there are more world states possible (*e.g.* by using private distribution, we have two world states possible and by using public distribution, we have four world states possible), which can be intractable in large games.

### C. Mixture between public and private distribution

To solve both of the problems defined in Section III-A and in Section III-B, we propose to use a mixture of private and public distribution.

**Definition** (Mixture Belief Distribution). *Let  $s_{pub}$  be the public infostate associated with the infostate  $s_i$ . We define the mixture belief distribution as  $\Delta_\lambda(s_i) = (1 - \lambda)\Delta_i(s_i) + \lambda\Delta_{pub}(s_{pub})$*

The mixture belief distribution allows us to be consistent with the problem encountered. When care must be taken not to reveal information, one can increase  $\lambda$ . In contrast, when it is not appropriate to withhold information, one can decrease  $\lambda$ . The private belief distribution is obtained when  $\lambda = 0$  and the public belief distribution is obtained when  $\lambda = 1$ .

In our example, when using the mixture with  $\lambda = 0.5$  for the player 1, we obtain the following belief distribution  $\Delta_{0.5}(s_1) = (\{s_1 : 75\%; s'_1 : 25\%\}, \{s_2 : 50\%; s'_2 : 50\%\})$ .  $w_1$  and  $w_2$  are more probable (37.5% each) than  $w_3$  and  $w_4$  (12.5% each). Nevertheless, their probabilities are not zero, which makes it consistent with the other player’s belief.

It is possible to expand this concept by considering that  $\lambda$  depends on the progress of the game. As an example, in trick-taking card games, it may be important to keep the private information hidden at the beginning of the game (so as not to reveal information) but, as the game progresses, the focus shifts to accumulating points before the end, where the importance of concealing this information may decrease.

### D. Adaptation of algorithms

PIMC and IS-MCTS have been created with private belief distribution in mind. Therefore, it is necessary to modify the algorithms to use the public or a mixture belief distribution. Instead of starting at an infostate  $s_i$ , the algorithms must be

adapted to start at  $s_{pub}$ , where  $s_{pub}$  is the public infostate associated with  $s_i$ .

1) *PIMC*: In the case of PIMC, one must use a distinct PIMC for each infostate possible ( $S_i(s_{pub})$ ), and combine the final result by aggregating the scores using the distribution of possible infostates ( $\Delta S_i(s_{pub})$ ).

In our example, when using the mixture belief distribution, two infostates are possible for the first player ( $s_1$  and  $s'_1$ ). If  $w_2$  or  $w_1$  are sampled, the algorithm used is the one defined for  $s_1$ , on the other hand, if  $w_3$  or  $w_4$  are sampled, the algorithm used is the one defined for  $s'_1$ . In the end, if  $s_1$  has been visited 75% (corresponding to the mixture belief distribution with  $\lambda = 0.5$ ), the action chosen in  $s_1$  will have more impact than the action chosen in  $s'_1$ .

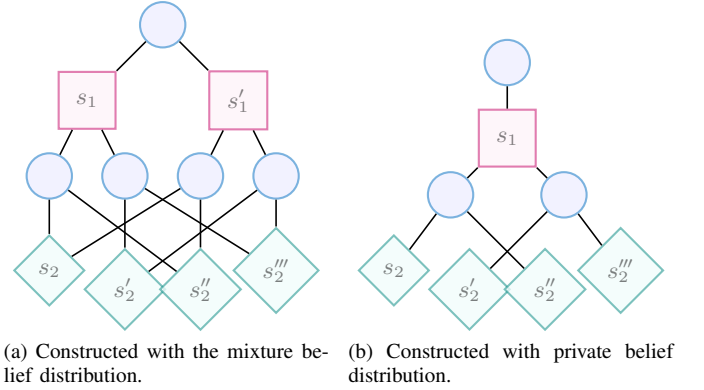


Fig. 2: Example of the tree constructed by IS-MCTS. The first player is acting in the red square, the second player is acting in the green diamond and the blue circle refers to the chance node.

2) *IS-MCTS*: With IS-MCTS, a singular algorithm is feasible as IS-MCTS creates a tree where the nodes represent infostates, and an infostate for player  $j$  may come from several infostates of player  $i$ .

An example is provided in Fig. 2. For the first player, two infostates are possible ( $s_1$  and  $s'_1$ ) and four infostates are possible for the second player after the first player’s action ( $s_2 = \{o_{priv(2)}^1 = \emptyset, o_{priv(2)}^2 = 1, o_{priv(2)}^3 = a_1\}$ ,  $s'_2 = \{o_{priv(2)}^1 = \emptyset, o_{priv(2)}^2 = 1, o_{priv(2)}^3 = a_2\}$ ,  $s''_2 = \{o_{priv(2)}^1 = \emptyset, o_{priv(2)}^2 = 2, o_{priv(2)}^3 = a_1\}$  or  $s'''_2 = \{o_{priv(2)}^1 = \emptyset, o_{priv(2)}^2 = 2, o_{priv(2)}^3 = a_2\}$ ).

For the second player, all infostates are achievable through any infostate of the first player. For example,  $s_2$  is achievable when sampling  $w_1$  (from  $s_1$ ) or when sampling  $w_2$  (from  $s'_1$ ) and playing the action  $a_1$ .

## IV. EXPERIMENTATION

### A. Benchmarks

For our experiments, the following benchmarks are tested ‘Liar’s Dice’ (LD), ‘Card Games (CG)’, and ‘Leduc Poker’ (LP). Each of them is described below.

1) *Card game*: For the purpose of the experimentation, we use a smaller version of classic trick-taking card games. The game is played with two players, 10/20 cards known by all, 2/6 are hidden and the rest is distributed to each player.

The playing phase is decomposed into tricks, the player starting the trick is the one who won the previous trick. The starting player of a trick can play any card in his hand, but the other players must follow the suit of the first player. If they can not, they can play any card they want but, without the possibility of winning the trick. The winner of the trick is the one with the highest-ranking card. At the end of the game, the points of each player are counted (plain version of trick-taking card game). The count is defined by the number of tricks won. A player wins if it has at least half of the points.

2) *Liar’s Dice*: Liar’s dice is a dice game played with two or more players, where each player possesses  $N$  dice of  $K$  sides and in which a player must deceive and be able to detect an opponent’s deception.

In the beginning, each player rolls his dice and observes the values. After that, players take turns guessing the number of dice of a particular type held by everyone. The game continues until a player accuses another of lying. If the player who made the assumption is right, he wins the game, on the opposite, if the challenged player did not lie, the challenged player wins.

During the game, a player can not bid less than previously, *i.e.* he must at least bid more dice than the previous player’s bid, or the same number of dice but with a higher value. Lastly, the highest face is a wild card, *i.e.* the value can be used to count for any other face.

3) *Leduc Poker*: Leduc Poker, as described in [18], is a variation of poker that uses a deck with only two suits, each containing three cards.

The game consists of two rounds. In the first round, each player is dealt a single private card. In the second round, a single board card is revealed. The maximum number of bets allowed is two, with the first round allowing raises of 2 and the second round allowing raises of 4. Both players begin the first round with 1 already in the pot.

## B. Experimentation

In our experiments, our objective is (i) to observe the extent to which an algorithm  $X$  reveals information according to mixture belief distribution; (ii) to analyze how the mixture belief distribution impacts the performance against an opponent that uses the revealed information; and (iii) to analyze how the mixture belief distribution impacts the performance against an opponent that does not use the revealed information.

Our code is based on OpenSpiel [19]. This is a collection of environments and algorithms for research in general reinforcement learning and search/planning in games.

PIMC and IS-MCTS are used with their basic version, *i.e.* PIMC uses AlphaBeta and IS-MCTS uses random rollouts as the perfect information evaluator and an exploration constant of 0.7. For both, 1000 world states are sampled.

To achieve a stable policy (as PIMC and IS-MCTS are online algorithms), we run the algorithm multiple times for every infostate until the policy obtained has less than 1% of variation.

The experiments were conducted according to the player’s playing position (each position reveals more or less information). In the following part, the experiments are carried out for the first player and in the appendix for the second player.

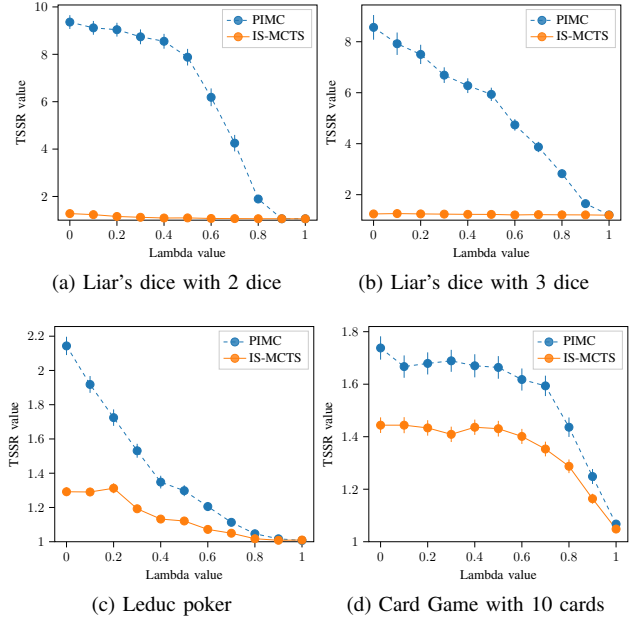


Fig. 3: Average TSSR for IS-MCTS and PIMC on multiple benchmarks according to  $\lambda$  of the mixture distribution.

1) *How much information is revealed*: For analyzing the impact of the revealed information according to the distribution used, we use the formula called True State Sampling Ratio (TSSR) [20]. TSSR measures how much more likely it is for the opponent to guess the current world state when using an algorithm  $X$  rather than using a uniform function.

The formula is  $TSSR(w) = \eta(w | s_i) \cdot |S_i(s_i)|$  where  $s_i$  is the infostate corresponding to  $w$ ,  $\eta(w | s_i)$  is the probability that the true state is guessed given the information set  $s_i$ . The closer the result is to 1, the less likely it is to know the real world state. Fig. 3 presents the TSSR value obtained according to  $\lambda$  of the mixture distribution.

As expected, playing closer to the public belief distribution greatly reduces the probability of knowing the real-world state. In ‘Liar’s Dice’ with 2 dice with PIMC, it is up to 10 fold more likely to guess the real world state when using the private instead of the public belief distribution.

In terms of information revealed, we observe that PIMC reveals more information than IS-MCTS in every benchmark. In ‘Leduc Poker’, it’s up to 2.2 times more likely to deduce the true state with PIMC at  $\lambda = 0.0$  whereas, with IS-MCTS, it is ‘only’ 1.3 times more likely to deduce the true state.

In addition, ‘Liar’s Dice’ is the game that reveals the most information with the algorithm revealing up to 10 times more likely than random, whereas in ‘Leduc Poker’ or ‘Card Game’, it is only up to 2 times more likely than random.

For the following experiments, it is expected to observe  $\lambda$  closer to 1 for PIMC in ‘Liar’s Dice’, as it reveals more information, and therefore, could be exploited by the opponent.

2) *How does the mixture impact the performance:* To measure how the mixture impacts the performance, we compute the expected utility against the best responder. The best responder is the worst possible enemy of all algorithms, *i.e.* it knows exactly the policy our algorithm will execute, and therefore, can infer the true infostate and plays the best action against it.

The results are available in Table I where the values represent the expected utility of the best responder and must be minimized. The results obtained are exact utility (without variation), as the best responder computes the best strategy knowing all the distributions in every infostate of the game.

We observe that the private belief distribution performs better than the public belief distribution, *i.e.* for all benchmarks and algorithms (better results are obtained when  $\lambda = 0.0$  than when  $\lambda = 1.0$ ).

In ‘Liar’s Dice’ with PIMC, the best performances are obtained when  $\lambda$  is close to 0.5 (with 2 dice, we obtain the best value when  $\lambda = 0.6$ ). These results were expected, as PIMC reveals a lot of information with Liar’s Dice which is then exploited by the best responder.

On the other hand, when the algorithm reveals less information (as observed in ‘Leduc Poker’ or IS-MCTS), it is preferable to use the private belief distribution or very close, as it is not sufficient for the best responder to exploit the revealed information.

3) *Can the use of multiple mixture belief distributions throughout the game improve performance:* In this experiment, we analyze the use of multiple mixtures throughout the game to improve performance. For this purpose, we compute multiple mixture distributions against the best responder.

Fig. 4 represents heatmaps for ‘Leduc Poker’ and ‘Liar’s Dice’ according to the position throughout the game when using PIMC (resp. IS-MCTS). For both games, we have a mixture distribution for the first action and another for the second action.

In all experiments, we observe that using multiple mixtures throughout the game has an impact on the performance. In ‘Leduc Poker’ for both algorithms, not using our private belief distribution is more punished in the second round than in the first round (*e.g.*  $\{0.0, 1.0\}$  has a value of 1.17 whereas  $\{1.0, 0.0\}$  has a value of 1.88 for IS-MCTS). On the other hand, for ‘Liar’s Dice’, we observe that the first round is the most important one.

In addition, we observe that playing multiple mixtures improve performance. In ‘Liar’s Dice’, the best value for IS

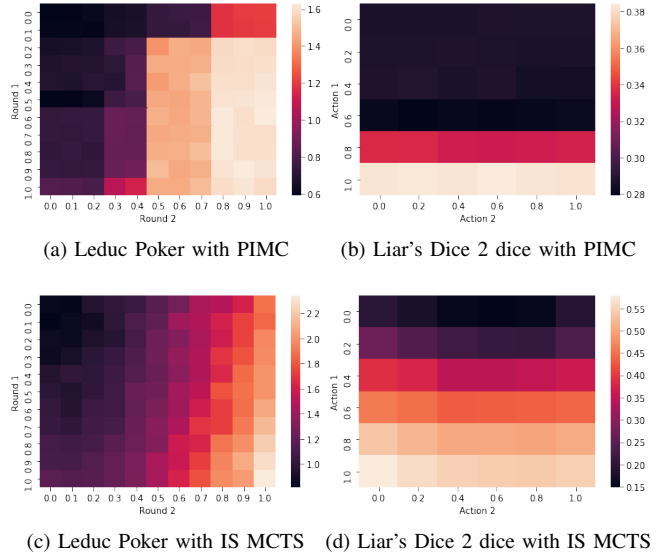


Fig. 4: Heatmap of the expected utility against the best response when playing at the first position.

MCTS is obtained when we have  $\{0.0, 0.6\}$  and for PIMC when we have  $\{0.6, 0.2\}$ .

4) *How does the mixture impact the winning rate:* As observed in the previous experiments, when using a  $\lambda$  closer to the public belief distribution, we obtain a distribution of action less relevant but with the advantage of disclosing less information. Therefore, when faced with an opponent who does not infer on our private information, it is expected to lose the benefit of using a  $\lambda$  closer to the public belief distribution. Nevertheless, using a  $\lambda$  closer to the public belief distribution not only reveals less information but allows it to be more consistent with the other player’s doubts.

To measure the impact of being more consistent with the other player’s doubts, we evaluate the performance against an algorithm that does not try to infer our private information. To do this, we compute the winning rate against ‘PIMC’ over 1000 games which results in 3.1% variation (95% of confidence interval). The scores are available in Table II.

As before, we observe that it is preferable to use private belief distribution instead of public belief distribution. In ‘Liar’s Dice’ with 3 dice with PIMC, we observe a drop of 20.8 in the winning rate between the private and public belief distribution. In addition, we observe that in every benchmark tested and for both algorithms, using a  $\lambda$  between 0.0 to 0.5 does not produce a drop in performance, but provides equivalent results.

These results are surprising, as we could have expected a drop in performance as the actions are less relevant to the current infostate (as we have sampled less often the true infostate). This implies that being more consistent with the doubts of the other players compensates for the loss of the

TABLE I: Expected utility against best responder when playing at the first player position.

Algo	Game	$\lambda$										
		0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
PIMC	LD 2D	0.300	0.298	0.297	0.292	0.294	0.288	<b>0.281</b>	0.290	0.336	0.382	0.382
	LD 3D	0.313	0.276	0.265	0.269	<b>0.235</b>	0.283	0.324	0.356	0.359	0.393	0.458
	LP	0.622	<b>0.616</b>	0.660	0.767	0.797	1.481	1.626	1.480	1.532	1.599	1.611
IS-MCTS	LD 2D	0.513	<b>0.512</b>	0.517	0.528	0.539	0.547	0.552	0.554	0.555	0.562	0.562
	LP	<b>0.797</b>	0.890	0.966	0.959	1.158	1.226	1.402	1.673	1.786	2.083	2.326

TABLE II: Winning rate when the opponent uses ‘PIMC’ when playing at the first player position.

Our	Game	$\lambda$										
		0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
PIMC	LD 3D	48.6	<b>50.4</b>	47.9	47.4	44.6	42.6	39.9	37.5	36.1	28.9	27.8
	LD 5D	43.1	<b>43.4</b>	42.2	<b>43.4</b>	42.5	41.4	39.8	36.3	37.1	29.8	23.6
	CG 10C	<b>48.2</b>	47.9	47.7	47.7	47.6	47.4	46.7	46.	45.5	39.8	31.6
	CG 20C	53.7	53.8	54.2	<b>54.5</b>	53.9	53.2	52.8	52.	47.4	36.3	23.5
IS MCTS	LD 3D	23.7	23.7	24.7	<b>27.</b>	23.1	23.1	21.7	20.	19.3	15.4	16.4
	LD 5D	22.	20.9	21.9	<b>22.2</b>	21.9	20.8	21.6	21.	16.9	15.5	13.4
	CG 10C	45.3	<b>46.3</b>	45.4	45.1	43.8	45.1	45.	43.1	42.7	37.6	30.
	CG 20C	36.5	<b>38.5</b>	38.2	36.2	36.4	36.6	35.5	34.9	33.3	33.1	20.8

player’s private information.

## V. CONCLUSION

In this paper, we study the strengths and weaknesses of probability distributions (private and public) in which particular attention has been paid to the revealed information and the impact of this revealed information on performance. Our study has been carried out on determinization-based algorithms and on multiple imperfect information games.

We complete the study by proposing a new probability distribution, a mixture of the two previous ones, which solves problems encountered by other distributions. We show that using the mixture is beneficial to reduce the revealed information and improve performance. We also show that using multiple mixtures throughout the game improves performance. In addition, we observed that using the mixture against an opponent that does not use our private information revealed results in a good performance as we are being more consistent with the other player’s doubt.

An avenue for improvement would be to extend the utilization of using multiple mixtures throughout the game. For example, by using the mixture at each public infostate instead of a fixed time step or using a different lambda for the opponent player. Another area for improvement would be to extend the study of algorithms that do not use determinization or even, without probability distributions but bearing in mind that one should not always use one’s private information at the risk of revealing information and, on the contrary, that one should not always use one’s public information in order to be more consistent to one’s private knowledge. Lastly, it would be interesting to extend the results at a larger scale, either by using more games or by using larger games.

## APPENDIX

### A. Complementary experiments

The following experiments are identical to those in the primary paper, with the exception that they are conducted for the second player position.

Similar results are observed, *i.e.* PIMC reveals more information than IS-MCTS, the private belief distribution obtains better performance than the public belief distribution against the best responder, using multiple mixtures is useful to improve the performance and it is all as well to play the mixture as the private against an opponent that does not try to infer.

Yet, we also observe some differences, especially that less information is revealed when playing in the second position, which results in  $\lambda$  closer to the private belief distribution against the best responder.




Fig. 5: Average TSSR according to  $\lambda$  value of the mixture distribution.

## REFERENCES

- [1] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever,

TABLE III: Expected utility for best responder against our algorithm being the second player.

Algo	Game	$\lambda$										
		0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
PIMC	LD 2D	<b>0.678</b>	0.695	0.703	0.707	0.716	0.718	0.711	0.741	0.779	0.836	0.836
	LP	<b>0.398</b>	0.400	0.459	0.612	0.796	1.461	1.450	1.509	1.593	1.615	1.632
IS-MCTS	LD 2D	0.697	<b>0.687</b>	0.697	0.716	0.727	0.732	0.740	0.751	0.759	0.768	0.787
	LP	<b>0.784</b>	0.784	0.898	0.800	1.017	1.078	1.186	1.324	1.561	1.728	2.002

TABLE IV: Winning rate when the opponent uses ‘PIMC’ when playing at the second player position.

Our	Game	$\lambda$										
		0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
PIMC	LD 3D	51.9	<b>53.</b>	51.3	49.8	49.7	51.3	48.7	48.6	46.9	46.1	42.7
	LD 5D	<b>56.7</b>	55.5	56.	56.2	54.8	56.1	55.3	53.	51.9	44.7	42.3
IS MCTS	LD 3D	48.4	<b>51.3</b>	49.9	49.	50.1	51.	47.4	44.	39.7	36.9	33.3
	LD 5D	<b>48.4</b>	47.1	48.	46.7	47.8	45.	46.5	40.7	34.4	23.2	14.7

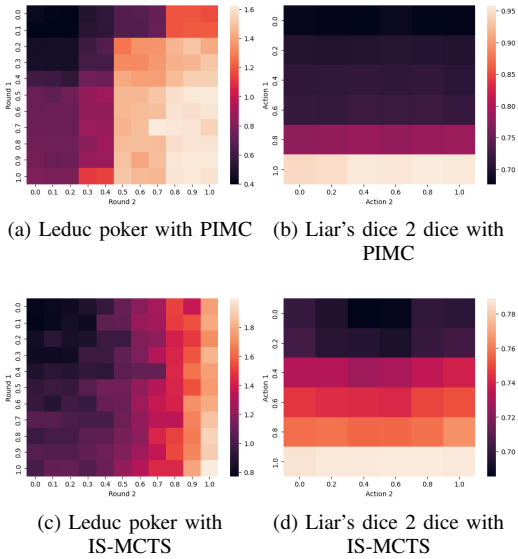


Fig. 6: Heatmap of expected utility against the best response when playing at the second position.

T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, “Mastering the game of go with deep neural networks and tree search,” *Nature*, vol. 529, pp. 484–489, 2016.

[2] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. P. Lillicrap, K. Simonyan, and D. Hassabis, “Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm,” *ArXiv*, vol. abs/1712.01815, 2017.

[3] —, “A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play,” *Science*, vol. 362, pp. 1140 – 1144, 2018.

[4] N. Brown and T. Sandholm, “Superhuman AI for multiplayer poker,” *Science*, vol. 365, pp. 885 – 890, 2019.

[5] C. Berner, G. Brockman, B. Chan, V. Cheung, P. Debiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, R. Józefowicz, S. Gray, C. Olsson, J. W. Pachocki, M. Petrov, H. P. d. O. Pinto, J. Raiman, T. Salimans, J. Schlatter, J. Schneider, S. Sidor, I. Sutskever, J. Tang, F. Wolski, and S. Zhang, “Dota 2 with Large Scale Deep Reinforcement Learning,” *ArXiv*, vol. abs/1912.06680, 2019.

[6] O. Tammelin, N. Burch, M. B. Johanson, and M. Bowling, “Solving

Heads-Up Limit Texas Hold’em,” in *IJCAI*, 2015.

[7] J. R. Long, N. R. Sturtevant, M. Buro, and T. Furtak, “Understanding the Success of Perfect Information Monte Carlo Sampling in Game Tree Search,” in *AAAI*, 2010.

[8] T. Furtak and M. Buro, “Recursive Monte Carlo search for imperfect information games,” *2013 IEEE Conference on Computational Intelligence in Games (CIG)*, pp. 1–8, 2013.

[9] P. I. Cowling, E. J. Powley, and D. Whitehouse, “Information Set Monte Carlo Tree Search,” *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 4, pp. 120–143, 2012.

[10] T. Cazenave and V. Ventos, “The  $\alpha\mu$  Search Algorithm for the Game of Bridge,” in *Monte Carlo Search at IJCAI*, ser. Communications in Computer and Information Science, 2021.

[11] V. Kovařík, M. Schmid, N. Burch, M. H. Bowling, and V. Lisý, “Rethinking Formal Models of Partially Observable Multiagent Decision Making,” *Artif. Intell.*, vol. 303, p. 103645, 2022.

[12] M. Moravčík, M. Schmid, N. Burch, V. Lisý, D. Morrill, N. Bard, T. Davis, K. Waugh, M. B. Johanson, and M. H. Bowling, “DeepStack: Expert-level artificial intelligence in heads-up no-limit poker,” *Science*, vol. 356, pp. 508 – 513, 2017.

[13] N. Brown, A. Bakhtin, A. Lerer, and Q. Gong, “Combining Deep Reinforcement Learning and Search for Imperfect-Information Games,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS’20. Red Hook, NY, USA: Curran Associates Inc., 2020, event-place: Vancouver, BC, Canada.

[14] T. Smith, “Probabilistic planning for robotic exploration,” Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, PA, July 2007.

[15] J. S. Dibangoye, C. Amato, O. Buffet, and F. Charpillet, “Optimally Solving Dec-POMDPs as Continuous-State MDPs,” *Journal of Artificial Intelligence Research*, vol. 55, pp. 443–497, Feb. 2016.

[16] D. E. Knuth and R. W. Moore, “An analysis of alpha-beta pruning,” *Artificial Intelligence*, vol. 6, no. 4, pp. 293–326, 1975.

[17] C. Browne, E. J. Powley, D. Whitehouse, S. M. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. P. Liebana, S. Samothrakis, and S. Colton, “A Survey of Monte Carlo Tree Search Methods,” *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 4, pp. 1–43, 2012.

[18] F. Southey, M. P. Bowling, B. Larson, C. Piccione, N. Burch, D. Billings, and C. Rayner, “Bayes’ bluff: Opponent modelling in poker,” *arXiv preprint arXiv:1207.1411*, 2012.

[19] M. Lanctot, E. Lockhart, J.-B. Lespiau, V. F. Zambaldi, S. Upadhyay, J. Pérolat, S. Srinivasan, F. Timbers, K. Tuyls, S. Omidshafiei, D. Hennes, D. Morrill, P. Muller, T. Ewalds, R. Faulkner, J. Kramár, B. D. Vyllder, B. Saeta, J. Bradbury, D. Ding, S. Borgeaud, M. Lai, J. Schrittwieser, T. W. Anthony, E. Hughes, I. Danihelka, and J. Ryan-Davis, “OpenSpiel: A Framework for Reinforcement Learning in Games,” *ArXiv*, vol. abs/1908.09453, 2019.

[20] C. Solinas, D. Rebstock, and M. Buro, “Improving Search with Supervised Learning in Trick-Based Card Games,” *ArXiv*, vol. abs/1903.09604, 2019.