

# Continual Depth-limited Responses for Computing Counter-strategies in Sequential Games

Anonymous Author(s)

Submission Id: 8

## ABSTRACT

In zero-sum games, the optimal strategy is well-defined by the Nash equilibrium. However, it is overly conservative when playing against suboptimal opponents and it can not exploit their weaknesses. Limited look-ahead game solving in imperfect-information games allows defeating human experts in massive real-world games such as Poker, Liar's Dice, and Scotland Yard. However, since they approximate Nash equilibrium, they tend to only win slightly against weak opponents. We propose methods combining limited look-ahead solving with an opponent model, in order to 1) approximate a best response in large games or 2) compute a robust response with control over the robustness of the response. Both methods can compute the response in real time to previously unseen strategies. We present theoretical guarantees of our methods. We show that existing robust response methods do not work combined with limited look-ahead solving of the shelf, and we propose a novel solution for the issue. Our algorithm performs significantly better than multiple baselines in smaller games. and outperform state-of-the-art methods against SlumBot.

## KEYWORDS

large games, approximating best response, robust response, opponent exploitation, imperfect information, depth limited solving, gadgets

### ACM Reference Format:

Anonymous Author(s). 2024. Continual Depth-limited Responses for Computing Counter-strategies in Sequential Games. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 16 pages.

## 1 INTRODUCTION

We can not enumerate all the decision points in large games, which makes computing optimal strategy, a Nash equilibrium (NE) in two-player zero-sum games, infeasible. A breakthrough that allowed approximating the NE and defeating human experts in several large imperfect-information games is limited look-ahead solving or search, which adapts the well-known approach from perfect-information games to games with imperfect-information [3, 19, 20]. Limited look-ahead solving takes advantage of decomposition. It iteratively builds the game to some depth and solves a small part of the game while summarising the required values from the rest of the game by a value function. The value function is commonly learned using neural networks. When the algorithms solve the game step by step, it is called continual depth-limited solving or continual resolving.

The vast majority of theoretically sound, continual depth-limited solving algorithms assume perfect rationality of the opponent and do not allow explicit modeling of an opponent and exploitation of the opponent's mistakes. As a result, even very weak opponents exploitable by the heuristic local best response (LBR) [12] can tie or lose very slowly against these methods [25]. Therefore, there has been a significant amount of work towards computing strategies to use against imperfect opponents to create AI systems that would perform well in the real world, for example, against humans [1, 8, 10, 16, 17, 21, 23].

The opponent modeling and exploitation process consists of two steps: opponent modeling and model exploitation. Opponent modeling requires building a model from previous data or actions observed during an online play. Model exploitation is finding a good strategy against the given model and is the main focus of this paper. In smaller games, we can trivially compute a best response to exploit the opponent maximally, or we can use methods to compute robust responses [8, 9] if there is uncertainty in the model and we want to be safer, meaning we want to limit the possible loss when facing the worst-case adversary. However, even the best response (BR) computation in large games is non-trivial, and currently, no approach can compute it while interacting in real-time.

This work explores the full model exploitation and proposes continual depth-limited best response (CDBR). CDBR relies on the value function used in the standard limited look-ahead solving, and we prove theoretical guarantees on the performance. A drawback of using the same value function is decreased performance, and we could improve CDBR by training a specific value function for a particular opponent model. However, it would be impractical since the training is expensive. Furthermore, in cases where we learn the opponent model in real-time interaction and update it after each step, it would be impossible.

The best response and CDBR are useful, e.g., for evaluating the quality of strategies, but they are brittle in game play. We can lose significantly when facing an opponent different from the expected model. In the real world, we will never have exact models, which makes BR and CDBR impractical for game play. To address the issue, robust responses are used [6, 8, 9]. They introduce a notion of safety, and the safety criterion requires the response to stay close to the NE. In other words, only to lose a limited amount to the worst-case adversary. Trivially, we can compute both BR and NE and create a linear combination where we can control the safety by a parameter. However, previous work shows that we can perform significantly better and recover the whole Pareto set of maximally exploiting strategies with maximal safety [9]. We adapt the method to limited look-ahead solving, creating a continual depth-limited restricted Nash response (CDRNR). Similarly to the full robust response, CDRNR significantly outperforms the linear combination. However, it comes with drawbacks in the limited

look-ahead solving. Namely, we need to keep the previously solved subgames as a path to the root to ensure theoretical soundness, which linearly increases the size of the game solved each step, making it scalable to games with low depth like Poker or Goofspiel but impractical in games with high depth.

Our contributions are: **1)** We formulate the algorithms to find the responses given the opponent strategy and an evaluation function. This results in the best performing theoretically sound robust response applicable to large games. **2)** We prove the soundness of the proposed algorithms. **3)** We provide an analysis of problems that arise when using opponent models in limited look-ahead solving and propose a solution we call a full gadget. **4)** We empirically evaluate the algorithms on poker and goofspiel variants and compare them to multiple baselines. We show that our responses exploit the opponents, and CDBR outperforms domain-specific local best response [12] on poker. We also compare CDBR with the approximate best response (ABR) on smaller games and on full Heads-up No-Limit Texas Hold'em (HUNL), where we exploit SlumBot significantly more than ABR.

## 2 BACKGROUND

A two-player extensive-form game (EFG) consists of a set of players  $N = \{1, 2, c\}$ , where  $c$  denotes the chance, 1 is the maximizer and 2 is the minimizer, a finite set  $A$  of all actions available in the game, a set  $H \subset \{a_1 a_2 \dots a_n \mid a_j \in A, n \in \mathbb{N}\}$  of histories in the game. We assume that  $H$  forms a non-empty finite prefix tree. We use  $g \sqsubset h$  to denote that  $h$  extends  $g$ . The root of  $H$  is the empty sequence  $\emptyset$ . The set of leaves of  $H$  is denoted  $Z$ , and its elements  $z$  are called *terminal histories*. The histories not in  $Z$  are *non-terminal histories*. By  $A(h) = \{a \in A \mid ha \in H\}$ , we denote the set of actions available at  $h$ .  $P : H \setminus Z \rightarrow N$  is the *player function* which returns who acts in a given history. Denoting  $H_i = \{h \in H \setminus Z \mid P(h) = i\}$ , we partition the histories as  $H = H_1 \cup H_2 \cup H_c \cup Z$ .  $\sigma_c$  is the *chance strategy* defined on  $H_c$ . For each  $h \in H_c$ ,  $\sigma_c(h)$  is a fixed probability distribution over  $A(h)$ . Utility functions assign each player utility for each leaf node,  $u_i : Z \rightarrow \mathbb{R}$ . The game is zero-sum if  $\forall z \in Z : u_1(z) + u_2(z) = 0$ . In the paper, we assume all the games are zero-sum. The game is of *imperfect information* if all players do not fully observe some actions or chance events. The information structure is described by *information sets* for each player  $i$ , which forms a partition  $\mathcal{I}_i$  of  $H_i$ . For any information set  $I_i \in \mathcal{I}_i$ , any two histories  $h, h' \in I_i$  are indistinguishable to player  $i$ . Therefore  $A(h) = A(h')$  whenever  $h, h' \in I_i$ . For  $I_i \in \mathcal{I}_i$  we denote by  $A(I_i)$  the set  $A(h)$  and by  $P(I_i)$  the player  $P(h)$  for any  $h \in I_i$ .

A *strategy*  $\sigma_i \in \Sigma_i$  of player  $i$  is a function that assigns a distribution over  $A(I_i)$  to each  $I_i \in \mathcal{I}_i$ . A *strategy profile*  $\sigma = (\sigma_1, \sigma_2)$  consists of strategies for both players.  $\pi^\sigma(h)$  is the probability of reaching  $h$  if all players play according to  $\sigma$ . We can decompose  $\pi^\sigma(h) = \prod_{i \in N} \pi_i^\sigma(h)$  into each player's contribution. Let  $\pi_{-i}^\sigma$  be the product of all players' contributions except that of player  $i$  (including chance). For  $I_i \in \mathcal{I}_i$  define  $\pi^\sigma(I_i) = \sum_{h \in I_i} \pi^\sigma(h)$ , as the probability of reaching information set  $I_i$  given all players play according to  $\sigma$ .  $\pi_i^\sigma(I_i)$  and  $\pi_{-i}^\sigma(I_i)$  are defined similarly. Finally, let  $\pi^\sigma(h, z) = \frac{\pi^\sigma(hz)}{\pi^\sigma(h)}$  if  $h \sqsubset z$ , and zero otherwise.  $\pi_i^\sigma(h, z)$  and  $\pi_{-i}^\sigma(h, z)$  are defined similarly. Using this notation, *expected payoff* for player  $i$  is  $u_i(\sigma) = \sum_{z \in Z} u_i(z) \pi^\sigma(z)$ . A *best response* (BR) of




Figure 1: Illustration of the depth-limited solving.


player  $i$  to the opponent's strategy  $\sigma_{-i}$  is a strategy  $\sigma_i^{BR} \in BR_i(\sigma_{-i})$ , where  $u_i(\sigma_i^{BR}, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$  for all  $\sigma'_i \in \Sigma_i$ . A tuple of strategies  $(\sigma_i^{NE}, \sigma_{-i}^{NE})$ ,  $\sigma_i^{NE} \in \Sigma_i, \sigma_{-i}^{NE} \in \Sigma_{-i}$  is a *Nash Equilibrium* (NE) if  $\sigma_i^{NE}$  is an optimal strategy of player  $i$  against strategy  $\sigma_{-i}^{NE}$ . Formally:  $\sigma_i^{NE} \in BR(\sigma_{-i}^{NE}) \quad \forall i \in \{1, 2\}$ .

In a two-player zero-sum game, the **exploitability** of a strategy is the expected utility a fully rational opponent can achieve above the value of the game. Formally, exploitability  $\mathcal{E}(\sigma_i)$  of strategy  $\sigma_i \in \Sigma_i$  is  $\mathcal{E}(\sigma_i) = u_{-i}(\sigma_i, \sigma_{-i}) - u_{-i}(\sigma^{NE})$ ,  $\sigma_{-i} \in BR_{-i}(\sigma_i)$ .

**Safety** is defined based on exploitability and  $\epsilon$ -safe strategy is a strategy which has exploitability at most  $\epsilon$ .

We define **gain** of a strategy against a model as the expected utility we receive above the value of the game. We formally define the gain  $\mathcal{G}(\sigma_i, \sigma_{-i})$  of the strategy  $\sigma_i$  against a strategy  $\sigma_{-i}$  as  $\mathcal{G}(\sigma_i, \sigma_{-i}) = u_i(\sigma_i, \sigma_{-i}) - u_i(\sigma^{NE})$ .

*Depth-limited Solving - Figure 1.* We denote  $H_i(h)$  the sequence of player  $i$ 's information sets and actions on the path to a history  $h$ . Two histories  $h, h'$  where player  $i$  does not act are in the same *augmented information set*  $I_i$  if  $H_i(h) = H_i(h')$ . We partition the game histories into **public states**  $PS \subset H$ , which are closed under the membership within the augmented information sets of all players. **Trunk** is a set of histories  $T \subset H$ , closed under prefixes and public states. **Subgame**  $S \subset H$  is a forest of trees with all the roots starting in one public state. It is closed under public states, and the trees can end in terminal public states or often end after a number of moves or rounds in the game. **Range** of a player  $i$  is a probability distribution over his information sets in some public state  $PS_i$ , given we reached the  $PS_i$ . **Value function** is a function that takes the public state and both players' ranges as input and outputs values for each information set in the public state for both players. We assume using an approximation of an **optimal value function**, which is a value function returning the values of using some NE after the depth-limit. **Subgame partitioning**  $\mathcal{P}$  is a partitioning that splits the game into trunk and subgames into multiple different levels based on some depth-limit or other factors (domain knowledge). Subgame partitioning can be naturally created using the formalism of factored-observation stochastic games [11]. By  $u_i(\sigma)_V^T$ , we denote the utility for player  $i$  if we use strategy profile  $\sigma$  in trunk  $T$  and compute values at the depth-limit using value function  $V$ . When resolving a subgame with just the ranges, there are no guarantees on the resulting exploitability of the strategy in the full game, and the exploitability can rise significantly [4]. To address the issue, artificially constructed games called **gadgets** are used to limit the increase in exploitability. They do it by adding



**Figure 2: Simple zero-sum imperfect-information game. Nodes denote the decisions of the players, dotted lines mark information sets, and the leaf shows for player 1**

nodes to the top of the subgame, which simulates that the opponent is allowed to deviate from its strategy in an already solved game.

Figure 2 shows a simple game illustrating depth-limited solving. The game starts with player 2 choosing to either play standard biased matching pennies ( $p$ ) or playing his own version of the game ( $q$ ). In the next round, player 1 does not know which game player 2 chose, and he chooses head ( $H$ ) or tail ( $T$ ). Then player 2 guesses head ( $h_i$ ) or tail ( $t_i$ ), and if he chooses to play the standard version, he receives 2 when correctly guessing head and 1 when correctly guessing tails. Otherwise, the reward is 0. In the modified version, guessing incorrectly gives 2 to the player 2, and guessing correctly gives 1 for heads and -1 for tails.

In the Nash equilibrium of this game, player 1 plays heads with probability  $\frac{2}{3}$  and player 2 chooses his own version of the game with probability  $\frac{2}{3}$  and follows with only heads. Public states in this game are always the whole levels (rows) since the actions are never observable by both players. When we start depth-limited solving, we create a trunk, which we select as just the root with the choice to play  $q$  or  $p$ . We start solving the trunk using an iterative algorithm, e.g., counterfactual regret minimization (CFR) [26].

We initialize strategy to uniform, which gives us *range* in the next public state  $(\frac{1}{2}, \frac{1}{2})$ . We give the *range* to the value function, which returns values as if we played equilibrium in the rest of the game. Value function gives us values in the information sets, which translates to the utility of  $-\frac{5}{3}$  for playing heads and  $-\frac{2}{3}$  for playing tails. We use the values to update regrets in the CFR and perform the next iteration similarly. When we solve the trunk and recover the equilibrium strategy for the first node, we move to a subgame, for example, a game starting in the information set of player 1 and ending after his action. We need to reconstruct what happened earlier. If we replace the already computed strategy with a chance node, which is called unsafe resolving, we are not guaranteed to recover the equilibrium for player 1. Unsafe resolving can produce solutions ranging from heads with probability  $\frac{3}{4}$  to  $\frac{1}{3}$  but the only equilibrium is heads with probability  $\frac{2}{3}$ . The situation is fixed using the mentioned gadgets, which allow the opponent to modify their range above the subgame, forcing the other player to play robustly against all the possible ranges and recover the equilibrium.

### 3 FULLY EXPLOITING THE OPPONENT

Fully exploiting opponent models in small games boils down to computing a best response. This is infeasible in games with an intractable number of information sets for which we use the continual depth-limited solving algorithms. The depth-limited setting does not allow computing BR in one pass anymore. The game we already

saw in Figure 2 can be an example of that. Suppose we know the player 2 always makes a mistake in the first move and plays only to the standard biased matching pennies. If we knew his strategy of guessing heads or tails, we could compute a best response. However, our trunk will end before the choice, and we need to use a value function. Since the value function in this simple case is just a best response of the opponent, the problem is reduced to finding the optimal strategy against a best response, which corresponds to finding NE, and it can not be solved in one pass. In this section, we propose an algorithm for continual depth-limited best response (CDBR), which generalizes a best response to be used with a value function for depth-limited solving.

### Continual Depth-limited Best Response

Given any extensive-form game  $G$  with perfect recall, opponent's fixed strategy  $\sigma_2^F$  and some subgame partitioning  $\mathcal{P}$ , we define continual depth-limited best response (CDBR) recursively from the top, see Figure 1. First, we have trunk  $T_1 = T$  and value function  $V$ . CDBR in the trunk  $T_1$  for player 1 with value function  $V$  is defined as  $\mathcal{B}(\sigma_2^F)_{V}^{T_1} = \arg \max_{\sigma_1} u_1(\sigma_1, \sigma_2^F)_{V}^{T_1}$ . In other words, we maximize the utility over the strategy in the trunk, where we return values from the value function after the depth limit. In each step afterward, we create a new subgame  $S_i$  and create new trunk by joining the old one with the subgame, creating  $T_i = T_{i-1} \cup S_i$ . We fix the strategy of player 1 in the  $T_{i-1}$  and maximize over the strategy in the subgame.  $\mathcal{B}(\sigma_2^F)_{V}^{T_i} = \arg \max_{(\sigma_1^{S_i})} u_1(\sigma_1^{S_i} \cup \sigma_1^{T_{i-1}}, \sigma_2^F)_{V}^{T_i}$ . We continue like that for each step, and we always create a new trunk  $T_i$  using the strategy from step  $T_{i-1}$  until we reach the end of the game. We denote the full CDBR strategy created by joining strategies from all possible branches  $\mathcal{B}(\sigma_2^F)_{V}^{\mathcal{P}}$ .

Intuitively, we always solve the game until the depth limit. The opponent is fixed everywhere above the depth limit, and the rational player is fixed in the already solved parts, and she can play in the part that was added last. Looking at Figure 1 CDBR in  $S_2$  would allow player 1 to play in  $S_2$ , it would replace anything below  $S_2$  with a value function, player 1 would be fixed in  $T$  and  $S_1$  and player 2 would be fixed in  $T, S_1$  and  $S_2$ .

*Computing CDBR and the complexity.* In practice, we will compute CDBR similarly to depth-limited solving with a few key changes. First, we fix the opponent's strategy in the currently resolved part of the game to allow the player to respond to it, which corresponds to the argmax from the definition. Another key change that simplifies the algorithm is that we no longer need a gadget since the opponent is fixed in the parts we already played through, so we do not need to be robust against different ranges than the one taken from the opponent model.

The difference from the standard depth-limited solving is that we fix the opponent's strategy in the resolved part of the game, and we do not use a gadget. Hence, there is less computation required compared to the standard depth-limited solving.

*Convergence in current iterations.* CFR is an algorithm that needs to track average strategies since the current strategy does not converge to an equilibrium. CFR against best response or a fixed strategy is known to converge in the current strategy [5, 14]. The next lemma

says that CDBR also converges in the current strategy even when a value function is used after the depth-limit.

**LEMMA 3.1.** *Let  $G$  be a zero-sum imperfect-information extensive-form game. Let  $\sigma_2^F$  be the fixed opponent's strategy, and let  $T$  be some trunk of the game. If we perform CFR with  $t$  iterations in the trunk for player 1, then for the strategy  $\hat{\sigma}_1$  from the iteration with highest expected utility  $\max_{\sigma_1^* \in \Sigma_1} u_1(\sigma_1^*, \sigma_2^F)_V^T - u_1(\hat{\sigma}_1, \sigma_2^F)_V^T \leq \Delta \sqrt{\frac{A}{t}} |I_{TR}| + t N_S \epsilon_S$  where  $\Delta$  is a span of leaf utilities,  $\Delta = \max_{z \in Z} u_i(z) - \min_{z \in Z} u_i(z)$ ,  $A$  is an upper bound on the number of actions,  $|I_{TR}|$  is a number of information sets in the trunk,  $N_S$  is the number of information sets at the root of any subgame, and value function error is at most  $\epsilon_S$ .*

#### 4 SAFE MODEL EXPLOITATION

While CDBR maximizes the exploitation of the fixed opponent model, it allows a player to be exploited. When we face an opponent unsure if our model is perfect we must limit our exploitability. For example, when we gradually build a model during play, we must limit our exploitability in the initial game rounds when the model is still very inaccurate.

##### Combination of CDBR and Nash Equilibrium

The combination of CDBR and Nash equilibrium (CDBR-NE) is the first approach to limit exploitability. We can simultaneously compute both strategies using depth-limited solving and do a linear combination in every decision node. Let  $p$  be the linear combination parameter and  $\sigma_2^F$  be the opponent model. The gain and exploitability are limited accordingly.

$$\begin{aligned}\sigma_1^{LC} &= p\sigma_1^{NE} + (1-p)\mathcal{B}(\sigma_2^F)_V^{\mathcal{P}} \\ \mathcal{E}(\sigma_1^{LC}) &= p\mathcal{E}(\sigma_1^{NE}) + (1-p)\mathcal{E}(\mathcal{B}(\sigma_2^F)_V^{\mathcal{P}}) \\ \mathcal{G}(\sigma_1^{LC}, \sigma_2^F) &= p\mathcal{G}(\sigma_1^{LC}, \sigma_2^F) + (1-p)\mathcal{G}(\mathcal{B}(\sigma_2^F)_V^{\mathcal{P}}, \sigma_2^F)\end{aligned}$$

Desired exploitability or gain may be achieved by tuning the parameter  $p$  while being only two times slower than the CDBR since we need to find the Nash equilibrium separately and perform CDBR. The required value function is the same for both parts and is still the same as in standard depth-limited solving.

Required computation is exactly running standard depth-limited solving and CDBR in parallel. Since CDBR computation has standard depth-limited solving as an upper bound, the required computation is at most twice as much as standard depth-limited solving.

##### Continual Depth-limited RNR

CDBR-NE is safe, but [9] shows we can get a much better trade-off between gain and exploitability using RNR as it recovers the optimal Pareto set of  $\epsilon$ -safe best responses [15]. It also gives us better control of safety as it links the allowed exploitability to the achieved gain. We combine depth-limited solving with RNR to create CDRNR.

*Description of Restricted Nash Response.* For CDRNR, we first need to explain the RNR method briefly [9]. RNR is solved by computing a modified game, adding an initial chance node with two outcomes that player 1 does not observe. We copy the whole game tree under both chance node outcomes, and in one tree, the opponent plays the fixed strategy, and we denote it  $G^F$ . In the other tree, the opponent

can play as he wants, resulting in a best response to the strategy of player 1. We denote the other tree  $G'$ . Since player 1 does not observe the initial chance node, his information sets span over  $G'$  and  $G^F$ , and we denote the full modified game with both trees  $G^M$ . Parameter  $p$  is the method to control the safety and is the initial probability of picking  $G^F$ .

**Definition.** Given the opponent's fixed strategy  $\sigma_2^F$  and some subgame partitioning  $\mathcal{P}$  of  $G^M$ , we define continual depth-limited restricted Nash response (CDRNR) recursively from the top. First, we have trunk  $T_1^M$  using  $\mathcal{P}$  and value function  $V$ . CDRNR for player 1 in the trunk  $T_1^M$  using value function  $V$  is  $\mathcal{R}(\sigma_2^F, p)_V^{T_1^M} = \arg \max_{\sigma_1} u_1(\sigma_1, \mathcal{B}(\sigma_1))_V^{T_1^M}$ . And then, in every following step, we create the new subgame  $S_i^M$  and enlarge the trunk to incorporate this subgame, creating trunk  $T_i^M = T_{i-1}^M \cup S_i^M$ . Next, we fix strategy  $\sigma_1^{T_{i-1}^M}$  of player 1 in the previous trunk  $T_{i-1}^M$  and the CDRNR is  $\mathcal{R}(\sigma_2^F, p)_V^{T_i} = \arg \max_{\sigma_1'} u_1(\sigma_1', \mathcal{B}(\sigma_1'))_V^{T_i}$  where  $\sigma_1'$  is a combination of the strategy we optimize over and the fixed strategy from the previous step, formally  $\sigma_1' = \sigma_1^{S_i^M} \cup \sigma_1^{T_{i-1}^M}$ .

To summarize, we optimize only over the strategy in the subgame used in the current step while the strategy in the previous parts of the game is fixed for player 1. The strategy of the opponent is fixed in  $G^F$  and free in  $G'$ . We denote the full CDRNR strategy  $\mathcal{R}(\sigma_2^F, p)_V^{\mathcal{P}}$ .

*Computing CDRNR.* In practice, we want to avoid duplicating the tree, and we also want to use the exact same value function as in the standard depth-limited solving. We explain why the RNR does not need the duplicated trees in practice. It only needs the reaches of the fixed strategy injected to the terminal nodes in the ratio defined by the parameter  $p$ . This allows us to precompute the reaches, run CFR as in standard depth-limited solving, and then modify the computed reaches from the iteration using the precomputed fixed reaches. However, we also need to query the value function, which differs from the previous one in the theoretical definition as it spans over the modified public state. However, since the reaches of  $p_1$  are the same for  $G'$  and  $G^F$  we can compute it only once by joining the reaches together as in the previous example and querying the standard value function.

So far, we described exactly the standard depth-limited solving with only one modification: modifying the reaches using the fixed strategy. We also use the gadget since now the opponent can deviate in the  $G'$ . However, standard gadgets will fail due to the addition of imperfect parts of the opponent, and we discuss details along with a solution in the next section.

#### 5 GADGETS AND MODEL EXPLOITATION

When we exploit an opponent model, we need to worsen the strategy in terms of exploitability. We must limit how much the strategy worsens if we want a safe response. Gadgets are used to ensure exploitability does not increase [2, 4, 18], and all the common gadgets work in scenarios where we do not expect our strategy to worsen. However, we need to worsen our strategy to exploit the opponent. We try to gain as much as possible in RNR in  $G^F$ . As

soon as the strategy gets worse and the exploitability increases, the common gadgets fail to quantify this increase, which is crucial in applications doing a delicate trade-off. The requirement for the gadget which would work in CDRNR is in Definition 5.1

*Definition 5.1.* For each information set  $I \in I_1$  we need the value of its part in  $G'$ , formally  $\sum_{h \in I, h \in G', z \in Z, h \sqsubset z} \pi^\sigma(z) u_1(z)$ , to be the same as the value we would get if we let player 2 play BR in full  $G'$ .

The following examples show that the requirement is not satisfied for common resolving gadgets. We tried to construct a gadget that would satisfy the condition, but in the end, we kept the previously resolved parts of the game  $G'$  followed by the value function which the game does not follow. We call the construction the full gadget, it satisfies the condition, and still only increases the size of the solved part linearly. Constant-size gadget fulfilling the Definition 5.1 is an open problem.

### Restricted Nash Response with Gadget

We show that commonly used resolving gadgets are either overestimating or underestimating the values from Definition 5.1 on an example game in Figure 3. In the game, we first randomly pick a red or green coin. Player 2 observes this and decides to place the coin heads up (RH, GH) or tails up (RT, GT). Player 1 cannot observe anything and ultimately chooses whether he wants to play the game (P) or quit (Q).

In equilibrium, player 1 plays action Q, and player 2 can mix actions up to the point where the utility for P is at most 0. This gives the value of the game 0, and counterfactual values in all inner nodes are also 0. Assuming the modified RNR game  $G^M$  with an opponent model playing GT, that makes it worth for player 1 to play (P) in the game, 2 will play (RH, GH) in  $G'$  with utility -3 for player 1 in  $G'$ . We will use gadgets to resolve the game from the player 1 information set.

*Resolving Gadget.* [4] Resolving gadget constructs a game that allows the opponent to choose whether he wants to play in the subgame we created or terminate. It is done by inserting nodes above the roots of the subgame, and the opponent has two actions before each root, either to follow and play the game or to terminate and receive a reward they would get by playing the previously resolved equilibrium. Those nodes are grouped into information sets based on the opponent's augmented information sets at the subgame's roots.

Resolving gadget on the game in Figure 3 has all utilities after *terminate* actions 0. When we resolve the gadget, the utility is 0. However, when player 1 deviates to action P, player 2 plays *follow* action in all but the rightmost node, and the utility of player 1 will be -3.5. Therefore, the common resolving gadget may overestimate the real exploitability of the strategy in the subgame. Overestimating may lead to not exploiting as much as we can and makes it impossible to prove Theorem 5.2 about the minimal gain of our algorithm. Normalization of the chance node might seem to solve the problem, but it would only halve the value to -1.75, which is still incorrect.

*(Reach) Max-margin Gadget.* [2, 18] Both reach max-margin and max-margin gadgets allow the opponent to choose any information

set at the start of the subgame. This is done by inserting a single node to the top, where the opponent has an action for each of his augmented information sets in the root of the subgame. After the action is a chance node to split the information set to the histories, with correct reaches by the resolving player and chance. Furthermore, all the terminal values are adjusted by the same value, which is in the terminate action in the resolving gadget. In the reach max-margin gadget, this value is further modified by an approximation of opponent mistakes.

All the counterfactual best response values are 0, and we assume both players played perfectly before the depth limit. Hence, we do not need to offset any node in the (reach) max-margin gadget. Then, both gadget constructions are identical. We add the initial decision node and the chance nodes (since there is only one state in each information set, the nodes have only one action). When we solve the gadget, player 1 will pick action Q, and the gadget value will be 0. However, when player 1 deviates to action P, player 2 now has a choice between terminal utilities and picks action E to receive the highest one. This will result in utility -2, and we see that (reach) max-margin gadgets can underestimate the real exploitability. It can lead to our algorithm being more exploitable than we want using some  $p$ , and it makes Theorem 5.3 impossible to prove. Similar to the previous gadget, normalizing the chance nodes would lead to double the utility, which is still incorrect.

*Full Gadget.* The only construction fulfilling the requirements we found is to keep all the previously explored parts of the game in a path to the root and use a value function when we leave. Using the optimal value function, the construction simulates the best response, which measures exploitability.

Formally, when we reach subgame  $S_i$  we construct a composite game by joining  $S_i$ , the trunk  $T$ , and all the previous subgames  $S_j, j \in 1, \dots, i-1$ . It corresponds to the illustration in Figure 1, and the value function will evaluate every public state  $PS$  from which the actions lead outside of the game.

We show that other gadgets can overestimate or underestimate exploitability, which could shift the distribution of the parameter  $p$ , and we could still compute the same solutions. However, in Figure 4, we show the results of games created to break the other gadgets. We have two games in which the full gadget behaves correctly, and each breaks the other gadget. Figure 4 shows the results joined together. Both games have five actions for the exploiter, and each one is crucial in reconstructing the full RNR set. The full gadget can recover all five actions using different  $p$ , but other gadgets can only compute two actions regardless of the choice of  $p$ .

*Complexity of CDRNR.* We can use other gadgets in CDRNR to obtain fast algorithms without any theoretical guarantees and with the same bound on computation as we have for the combination of CDBR and Nash equilibrium. The soundness of the algorithm relies on using the full gadget, which requires solving increasingly larger parts of the game as the depth increases. This increase in size is linear with the resolving steps, so the full algorithm complexity is quadratic in the depth of the game compared with vanilla continual resolving or CDBR. It makes the algorithm applicable to shorter games like Poker or Goofspiel but infeasible for long games like Stratego.




Figure 3: (left) A game to show problems with gadgets. (middle) Resolving gadget for the left game. (right) Max-margin and Reach max-margin gadget.




Figure 4: Comparison of Gain and Exploitability of the solutions using Full gadget compared to other gadgets. Max-margin and resolving gadget lines are overlapping.

*Soundness of CDRNR.* For the following theorems, we denote  $\mathcal{S}$  as the set consisting of the trunk and all the subgames explored when computing the response, and  $\mathcal{S}'$  is the same but without the last subgame.  $\mathcal{S}^B$  denotes a border of the subgame. We also denote  $\mathcal{S}^O$ , the set of all the states where we leave the trunk not going in the currently resolving subgame.

**THEOREM 5.2 (GAIN OF CDRNR).** *Let  $G$  be any zero-sum extensive-form game and let  $\sigma_2^F$  be any fixed opponent's strategy in  $G$ . Then we set  $G^M$  as restricted Nash response modification of  $G$  using  $\sigma_2^F$ . Let  $\mathcal{P}$  be any subgame partitioning of the game  $G^M$  and using some  $p \in (0, 1)$ , let  $\sigma_1^R$  be a CDRNR given approximation  $\bar{V}$  of value function  $V$  with error at most  $\epsilon_V$  and opponent strategy  $\sigma_2^F$  approximated in each step with regret at most  $\epsilon_R$ , formally  $\sigma_1^R = \mathcal{R}(\sigma_2^F, p)_V^{\mathcal{P}}$ . Let  $\sigma^{NE}$  be any Nash equilibrium in  $G$ . Then  $u_1(\sigma_1^R, \sigma_2^F) + \sum_{S \in \mathcal{S}'} |I_{SO}|(1-p)\epsilon_V + |\mathcal{S}|\epsilon_R + \sum_{S \in \mathcal{S}'} |I_{SB}|\epsilon_V \geq u_1(\sigma^{NE})$ .*

The previous theorem states that our approaches will receive at least the value of the game when responding to the model. All the proofs are in the appendix.

**THEOREM 5.3 (SAFETY OF CDRNR).** *Let  $G$  be any zero-sum extensive-form game and let  $\sigma_2^F$  be any fixed opponent's strategy in  $G$ . Then we set  $G^M$  as restricted Nash response modification of  $G$  using  $\sigma_2^F$ . Let  $\mathcal{P}$  be any subgame partitioning of the game  $G^M$  and using some  $p \in (0, 1)$ , let  $\sigma_1^R$  be a CDRNR given approximation  $\bar{V}$  of the optimal value function  $V$  with error at most  $\epsilon_V$ , partitioning  $\mathcal{P}$  and opponent strategy  $\sigma_2^F$ , which is approximated in each step with regret at most  $\epsilon_R$ , formally  $\sigma_1^R = \sigma_1^R(\sigma_2^F, p)_V^{\mathcal{P}}$ . Then exploitability has a bound  $\mathcal{E}(\sigma_1^R) \leq \mathcal{G}(\sigma_1^R, \sigma_2^F) \frac{p}{1-p} + \sum_{S \in \mathcal{S}'} |I_{SO}|(1-p)\epsilon_V + |\mathcal{S}|\epsilon_R + \sum_{S \in \mathcal{S}'} |I_{SB}|\epsilon_V$ ,  $\mathcal{E}$  and  $\mathcal{G}$  are defined in Section 2.*

The last theorem is more complex, and it bounds the exploitability by the gain of the strategy against the model. With  $p = 0$ , it

is reduced to the continual resolving, and with  $p = 1$  to CDBR with unbounded exploitability. The theorem shows the parameter  $p$  directly links allowed exploitability to the gain we receive. The same works in RNR without the resolving and value errors; as far as we know, the authors do not explicitly mention it.

SES has bound relies on opponent estimation being close to an equilibrium strategy. When the estimation is more different, the bound is infinity for a large portion of the parameter  $\alpha$ . We give a detailed explanation in the appendix.

More intuitively, Theorem 5.2 says that by playing the proposed algorithm, we have at least the same safety guarantees we would get by playing a NE against the opponent we modeled correctly. Theorem 5.3 allows us to choose a trade-off between the exploitation of the opponent behaving according to the model and safety against an opponent who would deviate arbitrarily from the model.

## 6 EXPERIMENTS

We compared CDBR and local best response (LBR) [12]. We empirically show the performance of CDRNR and explore the trade-off between exploitability and gain in CDRNR. The appendix contains hardware setup, domain description, algorithm details, and experiments on more domains. We use two types of opponent strategies: strategies generated by few CFR iterations and random strategies with different seeds.


### SES explanation

Safe exploitation search (SES) [13] is a similar method to the one we propose. However, there are two significant differences. First, the method uses a max-margin gadget without the analysis we did. Hence, the bound of exploitability is very loose, and for a wide range of inputs, the exploitability can be unbounded. Second, the method does not fix the opponent's strategy at all and only uses opponent reaches when resolving the subgame. As a result, SES exploitation is very limited, and as we show in experiments, it is often worse than using the best Nash equilibrium. On top of that, in some games, it fundamentally cannot exploit the opponent, notably in any perfect information game, even with simultaneous moves.

### Exploitability of Robust Responses

We report both gain and exploitability for CDRNR on Leduc Hold'em. Results in Figure 5 show that the proven bound on exploitability works in practice, and we see that the bound is very loose in practice. For example, with  $p = 0.5$ , the bound on the exploitability is the gain itself, but the algorithm rarely reaches even a tenth of the gain in exploitability. This shows that the CDRNR is similar





**Figure 5: Gain and exploitability comparison of BR, RNR, best Nash equilibrium (BNE), CDBR-NE, SES, and CDRNR in Leduc Hold'em against strategies from CFR using a small number of iterations with different  $p$  values. The  $a$  stands for the average of the other values. VF is CDRNR using an imperfect value function.**

to the restricted Nash response because, with a well-set  $p$ , it can significantly exploit the opponent without significantly raising its exploitability. The only exception is the CDRNR with a value function, which shows the added constants from the value function's imprecision. When the opponent is close to optimal, we see that the exploitability can rise above the gain.

In most cases, the gain and exploitability of CDRNR are lower than that of RNR. Gain must be lower because of the different value function, but the exploitability can be higher, as seen with  $p = 0.1$  against low iteration strategies, due to the depth-limited nature. We provide results on other domains and for more parameter values  $p$  in the appendix.

We also compare the algorithm against the best possible Nash equilibrium. We compute the best NE by a linear program, and it serves as the theoretical limit of maximal gain, which does not allow exploitability. It would be impossible to compute for larger games. We can see we can gain more than twice as much, with exploitability still being almost zero.

In our results on smaller games, we use the optimal value function computed by a linear program. To show the performance of imperfect value function we included results where the value function is approximated by CFR with 500 iterations. As expected, the

	ABR	CDBR1	CDBR3	CDBR5	BNE
Leduc	98%	74.4%	96%	97.1%	32.6%
IIGS4	97%	98.5%	100%	100%	77.1%
IIGS5	95%	93.5%	100%	100%	50.4%
IIGS6	97%	90.7%	100%	100%	45.9%

**Table 1: Comparison of ABR and CDBR with BNE baseline on different games against uniform random. The values are the percentage of gain achieved by the best response.**

imperfect value function slightly decreases the performance but is comparable to the algorithms using the optimal value function.

The last comparison is with SES, which performs poorly, and its gain is only slightly above the best Nash equilibrium. Conversely, it is almost not exploitable. Our results are consistent with results in the paper [13] and are a direct consequence of using the information from the opponent model only to set the reaches to the subgame. Only reaches are not enough to do meaningful exploitation, and SES produces strategies that are very close to Nash equilibrium.


## ABR vs CDBR

We compared CDBR with ABR [22] on Leduc and different imperfect information Goofspiel. The results are in Table 1, showing that our method is slightly behind in Leduc, even with the highest search depth. In Goofspiel, CDBR1, which looks only one action into the future, is already pretty good, and as soon as we allow CDBR to look three turns in the future, it fully exploits the opponents. In IIGS4, that is half of the game, but in IIGS6, it is less than  $\frac{1}{3}$  of the game, and it is still enough.

## Local Best Response vs. CDBR

We compare LBR and CDBR in Leduc Hold'em. We also compare CDBR with the BR in imperfect information Goofspiel 5, but without LBR, which is poker-specific. We show that CDBR and LBR are very similar with smaller steps, and as we increase the depth limit of CDBR, it starts outperforming LBR. The behavior differs slightly based on the specific strategy because LBR assumes the player continues the game by only calling till the end, while CDBR uses the perfectly rational extension.

The results in Figure 6 show that both concepts are good at approximating the best response, with CDBR being better against both strategies. LBR looks at one following action, so it is best compared to the CDBR1 in terms of comparability. Next, we observe a lack of monotonicity in step increase, which is explained with an example in the appendix. When we increase the depth limit, the algorithm can exploit some early mistake that causes it to miss a sub-tree where the opponent makes a much bigger mistake in the future. We can clearly see the difference between the algorithm with guarantees and LBR without them. Against strategy from 34 CFR iterations, LBR can no longer achieve positive gain and only worsens with more iterations. In contrast, CDBR can always achieve at least zero gain, assuming we have an optimal value function.



**Figure 6: Gain comparison of best response (BR), local best response (LBR - only poker), and continual depth-limited best response (CDBR) in Leduc Hold'em (top) and IIGoofspiel 5 (bottom) against strategies from CFR using a small number of iterations (left) and random strategies (right). The a stands for the average of the other values in the plot. The number after CDBR stands for the number of actions CDBR was allowed to look at in the future, and CDBRNN is a one-step CDBR with a neural network as a value function.**

	ABR	LBR	CDBR
Win-rate [mbb/h]	1259 $\pm$ ?	1388 $\pm$ 150	1774 $\pm$ 137

**Table 2: Comparison of CDBR with LBR and ABR against SlumBot. Results are reported in milibigblinds per hand (mbb/h) with 95% confidence intervals.**

## Playing SlumBot

We tested our method in HUNL against SlumBot [7], which is a publicly available abstraction-based bot commonly used for benchmarking. We used a fold, call, pot, and all-in (FCPA) abstraction for CDBR, and we also rerun the results of LBR against the new SlumBot (since the author of SlumBot confirmed that the strategy we were provided is different from the one LBR used but the same as the one provided to authors of ABR). CDBR significantly outperforms both ABR and LBR and we report the results in Table 2. We tried to run the LBR restricted to only call in the first two rounds but found it no longer helps against the new SlumBot, and we reported results for LBR with FCPA. Authors in [22] also use FCPA for their method but did not report a confidence interval for the results. However, the difference is large enough to have statistical significance if we assume they played over 50 thousand hands.

## 7 RELATED WORK

This section describes the related work focusing more on distinguishing our novel contributions.

Restricted Nash response (RNR) [9] is an opponent-exploiting scheme. It solves the entire game and allows changing the trade-off

between exploitability and gain. Essentially it always produces  $\epsilon$  safe best response [15]. It accomplishes the goal by copying the whole game and then fixing the opponent in one part while having a chance node at the top decide which game we play.

However, it is impossible to compute RNR in huge games, and we fused the RNR approach with depth-limited solving creating a novel algorithm we call CDRNR. CDRNR is the best performing theoretically sound robust response calculation that can be done in huge games, enabling new opponent exploiting approaches.

Local best response [12] is an evaluation tool for poker. It uses a given abstraction in its action space. It picks the best action in each decision set, looking at the fold probability for the opponent in the next decision node and then assuming the game is called until the end. Our algorithm CDBR is a generalization of the LBR because we can use it on any game solvable by depth-limited solving. In the algorithm, we have explicitly defined value function, which we can exchange for different heuristics.

Approximate best response (ABR) [22] is also a generalization of the LBR and showed promising results in evaluating strategies. However, our approach focuses on model exploitation, which requires crucial differences, such as quick re-computation against unseen models. ABR needs to independently learn the response for every combination of opponent and game, making it unusable in the opponent modeling scenario. Our algorithms learn a single domain-specific value function and can subsequently compute strategies against any opponent in the run-time. Furthermore, ABR and even CDBR are extremely brittle, making it a bad choice if we are unsure about the opponent, which we often are in a game against an unknown opponent. On the other hand, CDRNR tackles exactly this issue and provides powerful exploitation with very limited exploitability.

Another reinforcement-learning (RL) method uses neuroevolution with RL, they show a significant increase in performance over DQN and evaluate their method on HUNL. However, they do not share the details of the baseline opponents they played against. We tried to contact the authors without any response and we could not compare the performance with our methods. [24]

## 8 CONCLUSION

Opponent modeling and exploitation is an essential topic in computational game theory, with many approaches attempting to model and exploit opponents in various games. However, exploiting opponents in very large games is not trivial, and only recently was an algorithm created to exploit models in depth-limited solving. We explain the problem arising from the inability of gadgets to measure exploitability and we propose a full gadget that solves the issue. We propose a new algorithm to quickly compute depth-limited best response and depth-limited restricted Nash response once we have a value function, creating the best performing theoretically sound robust response applicable to large games. Finally, we empirically evaluate the algorithms on multiple games. We show that CDBR outperforms LBR in both Leduc and HUNL and we show that CDBR performs significantly better against SlumBot than any other previous method. Finally, we show that CDRNR outperforms SES in any game and can achieve over half of the possible gain without almost any exploitability.



## REFERENCES

- [1] Nolan Bard, Michael Johanson, Neil Burch, and Michael Bowling. 2013. Online implicit agent modelling. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. 255–262.
- [2] Noam Brown and Tuomas Sandholm. 2017. Safe and nested endgame solving for imperfect-information games. In *Workshops at the thirty-first AAAI conference on artificial intelligence*.
- [3] Noam Brown and Tuomas Sandholm. 2019. Superhuman AI for multiplayer poker. *Science* 365, 6456 (2019), 885–890.
- [4] Neil Burch, Michael Johanson, and Michael Bowling. 2014. Solving imperfect information games using decomposition. In *Twenty-eighth AAAI conference on artificial intelligence*.
- [5] Trevor Davis, Neil Burch, and Michael Bowling. 2014. Using response functions to measure strategy strength. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- [6] Sam Ganzfried and Tuomas Sandholm. 2015. Safe opponent exploitation. *ACM Transactions on Economics and Computation (TEAC)* 3, 2 (2015), 1–28.
- [7] Eric Griffin Jackson. 2017. Targeted cfr. In *Workshops at the thirty-first AAAI conference on artificial intelligence*.
- [8] Michael Johanson and Michael Bowling. 2009. Data biased robust counter strategies. In *Artificial Intelligence and Statistics*. 264–271.
- [9] Michael Johanson, Martin Zinkevich, and Michael Bowling. 2008. Computing robust counter-strategies. In *Advances in neural information processing systems*. 721–728.
- [10] Kevin B Korb, Ann Nicholson, and Nathalie Jitnah. 2013. Bayesian poker. *arXiv preprint arXiv:1301.6711* (2013).
- [11] Vojtěch Kovářík, Dominik Seitz, Viliam Lisý, Jan Rudolf, Shuo Sun, and Karel Ha. 2019. Value Functions for Depth-Limited Solving in Imperfect-Information Games. *arXiv preprint arXiv:1906.06412* (2019).
- [12] Viliam Lisý and Michael Bowling. 2017. Equilibrium approximation quality of current no-limit poker bots. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*.
- [13] Mingyang Liu, Chengjie Wu, Qihan Liu, Yansen Jing, Jun Yang, Pingzhong Tang, and Chongjie Zhang. 2022. Safe Opponent-Exploitation Subgame Refinement. In *Advances in Neural Information Processing Systems*. <https://openreview.net/forum?id=YpHb0IVJu92>
- [14] Edward Lockhart, Marc Lanctot, Julien Pérolat, Jean-Baptiste Lespiau, Dustin Morrill, Finbarr Timbers, and Karl Tuyls. 2019. Computing Approximate Equilibria in Sequential Adversarial Games by Exploitability Descent. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. 464–470.
- [15] Peter McCracken and Michael Bowling. 2004. Safe strategies for agent modelling in games. In *AAAI Fall Symposium on Artificial Multi-agent Learning*. 103–110.
- [16] Richard Mealing and Jonathan L Shapiro. 2015. Opponent modeling by expectation–maximization and sequence prediction in simplified poker. *IEEE Transactions on Computational Intelligence and AI in Games* 9, 1 (2015), 11–24.
- [17] David Milec, Jakub Černý, Viliam Lisý, and Bo An. 2021. Complexity and Algorithms for Exploiting Quantal Opponents in Large Two-Player Games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5575–5583.
- [18] Matej Moravčík, Martin Schmid, Karel Ha, Milan Hladík, and Stephen Gaukrodger. 2016. Refining subgames in large imperfect information games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30.
- [19] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. 2017. Deepstack: Expert-level artificial intelligence in Heads-Up No-Limit Poker. *Science* 356, 6337 (2017), 508–513.
- [20] Martin Schmid, Matej Moravčík, Neil Burch, Rudolf Kadlec, Josh Davidson, Kevin Waugh, Nolan Bard, Finbarr Timbers, Marc Lanctot, Zach Holland, et al. 2021. Player of Games. *arXiv preprint arXiv:2112.03178* (2021).
- [21] Finnegan Southey, Michael P Bowling, Bryce Larson, Carmelo Piccione, Neil Burch, Darse Billings, and Chris Rayner. 2012. Bayes’ bluff: Opponent modelling in poker. *arXiv preprint arXiv:1207.1411* (2012).
- [22] Finbarr Timbers, Edward Lockhart, Marc Lanctot, Martin Schmid, Julian Schrittwieser, Thomas Hubert, and Michael Bowling. 2020. Approximate exploitability: Learning a best response in large games. *arXiv preprint arXiv:2004.09677* (2020).
- [23] Zhe Wu, Kai Li, Enmin Zhao, Hang Xu, Meng Zhang, Haobo Fu, Bo An, and Junliang Xing. 2021. L2E: Learning to Exploit Your Opponent. *arXiv preprint arXiv:2102.09381* (2021).
- [24] Jiahui Xu, Jing Chen, and Shaofei Chen. 2021. Efficient opponent exploitation in no-limit Texas hold’em poker: A neuroevolutionary method combined with reinforcement learning. *Electronics* 10, 17 (2021), 2087.
- [25] Ryan Zarick, Bryan Pellegrino, Noam Brown, and Caleb Banister. 2020. Unlocking the potential of deep counterfactual value networks. *arXiv preprint arXiv:2007.10442* (2020).
- [26] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. 2008. Regret minimization in games with incomplete information. In *Advances in Neural Information Processing Systems*. 1729–1736.

## A PSEUDOCODE

In this section, we present the pseudocode of both CDBR and CDBRNN.

---

### Algorithm 1 Computing CDBRNN (CDBR)

---

**Require:** game  $G$ , model strategy  $\sigma_2^F$ , value function  $V$   
 create (virtually) modified game  $G^M$  (only CDBRNN)  
 create subgame partitioning  $\mathcal{P}$  from  $G^M$  ( $G$ )  
 $\sigma_1^B$  = empty strategy ready to be filled  
 $I$  = initial information set in which we act  
 $S$  = current constructed subgame  
**while**  $I$  not null **do**  
   **if**  $I$  not in  $S$  **then**  
 $S$  = construct new  $S$  from  $\mathcal{P}$  using previous  $S$  (CDBRNN does not delete trunk)  
 $\sigma_1^B$  += solution of  $S$  using CFR+ with  $V$   
**else**  
   pick action  $A$  according to  $\sigma_1^B$  in  $I$   
   get new  $I$  using  $A$  (or null if the game ends)  
**end if**  
**end while**

---


## B ADDITIONAL CDBR RESULTS

We compare LBR and CDBR in Leduc Holde'm. We also compare CDBR with just the BR in imperfect information Goofspiel, but without LBR, which is poker specific. We show that CDBR and LBR are very similar with smaller steps, and as we increase the depth-limit of CDBR, it starts outperforming LBR. The behavior differs in every strategy because LBR assumes the player continues the game by only calling till the end, while CDBR uses the perfectly rational extension. Furthermore, it is possible to exchange the value function of CDBR, and both concepts would be very similar. However, we would lose the guarantee that CDBR will never perform worse than the value of the game.

Looking at the results in Figure 7, we can see that both concepts are good at approximating the best response, with CDBR being better against both strategies. LBR looks at one following action, so in terms of comparability, it is best compared to the CDBR1. Next, we observe a lack of monotonicity in step increase, which is linked to the counterexample in Figure 10. When we increase the depth-limit, the algorithm can exploit some early mistake that causes it to miss a sub-tree where the opponent makes a much bigger mistake in the future. We can clearly see the difference between the algorithm with guarantees and LBR without them. Against strategy from 34 CFR iterations, LBR can no longer achieve positive gain and only worsens with more iterations. In contrast, CDBR can always achieve at least zero gain (assuming we have an optimal value function).

## C COUNTEREXAMPLE GADGET GAME

Examples in Figure 8 are the games used to generate the Figure 4. The plot in the figure combines two games that have pure actions with the same gain and exploitability. The full gadget reconstructs the Pareto set using all the actions in both games. Other gadgets fail



**Figure 7: Gain comparison of best response (BR), local best response (LBR - only poker), and continual depth-limited best response (CDBR) in Leduc Hold'em (top), Imperfect information Goofspiel (middle) and Small Liar's dice (bottom) against strategies from CFR using a small number of iterations (left) and random strategies (right). The a stands for the average of the other values in the plot. The number after CDBR stands for the number of actions CDBR was allowed to look in the future, and CDBRNN is one step CDBR with a neural network as a value function.**

in one of the presented games in Figure 8. In Figure 9, we show the expected utility of all the actions in the CDBRNN version of the game, showing that for the resolving gadget and max-margin gadget, two actions dominate all the others, and we can not select any  $p$  which would resolve any of the remaining actions.

## D EXPERIMENT DETAILS

### Experimental Setup

For all experiments, we use Python 3.8 and C++17. We solved linear programs using Gurobi 9.0.3, and experiments were done on an Intel i7 1.8GHz CPU with 8GB RAM. We used Leduc Hold'em, imperfect information Goofspiel 5, and Liar's dice for the smaller detailed experiments. We used Goofspiel 6 for the large experiment, and we only ran it against the strategy generated by the CFR with three iterations. We used the torch library for the neural network experiment. For most of the experiments, we wanted to solve the concepts perfectly with perfect value function, so we used LP and fixed the parts of the game that needed to be fixed. For the neural

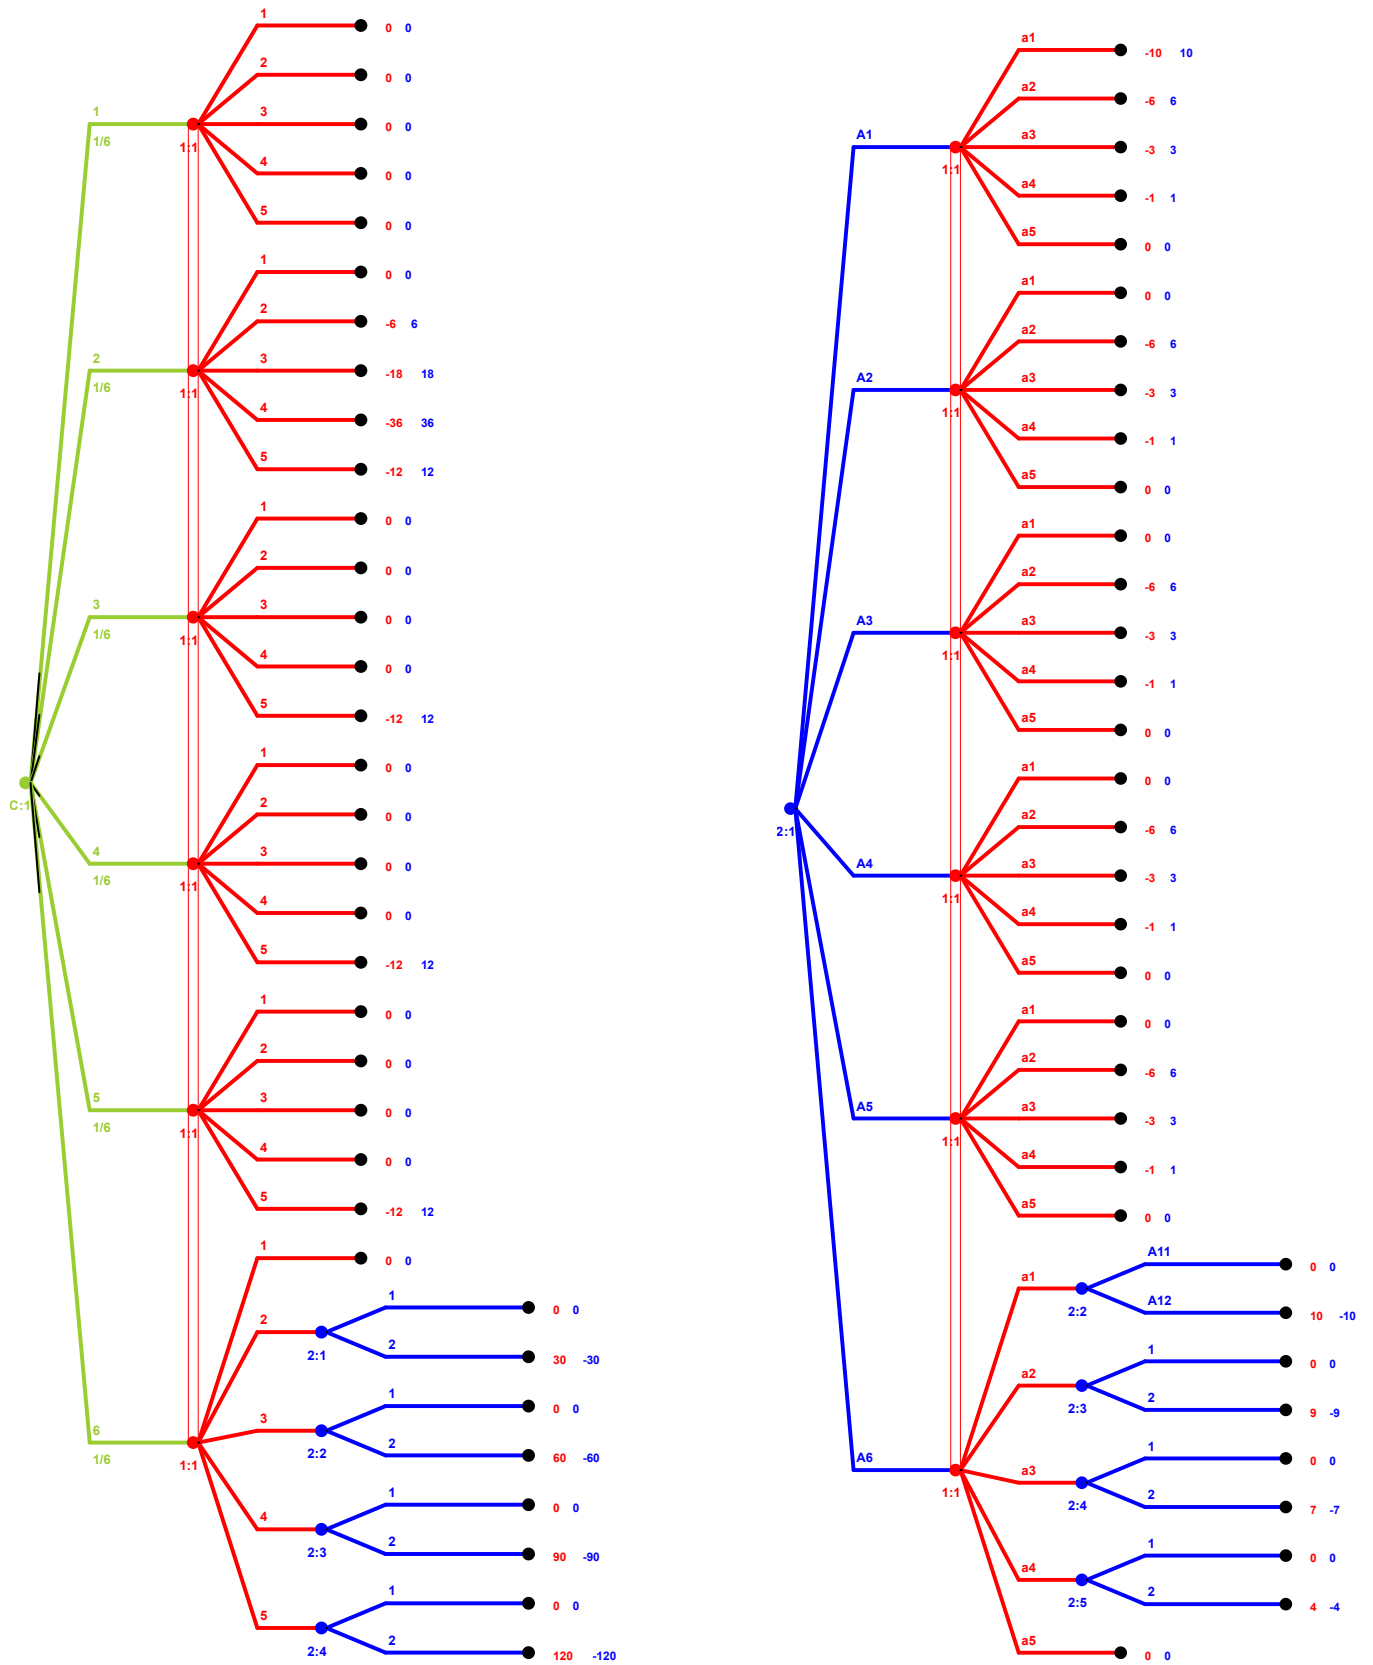


Figure 8: Example games to show the inability of common gadgets to reconstruct the whole Pareto set in the CDRNR setting. Left: Game 1 to break the max-margin gadget. Right: Game 2 to break the resolving gadget.




Figure 9: Value of each action based on  $p$  in the games shown in Figure 8, max-margin gadget in Game 1 (top), full gadget on both Game 1 and 2 (middle) and resolving gadget on Game 2 (bottom)

network experiment with CDBR, we used CFR+ to solve the subgame and the neural network as a value function. For the value function experiment in CDRNR, we used CFR+ with 1000 iterations to solve the game and CFR+ with 500 iterations as a value function in the subgames.

## Domain Definition

**Leduc Hold'em** is a medium-sized poker game. Both players give one chip at the beginning of the match and receive one card from a deck with six cards of 2 suits and three ranks. Then players switch and can call or bet. After a bet, the opponent can also fold, which finishes the game, and he forfeits all the staked money. After both players call or after at most two bets public card is revealed, and another betting round begins. In the first round, the bet size is two, and in the second, it is 4. If the game ends without anyone, folding cards are compared, and the player with pair always wins, and if there is no pair, the player with the higher card wins. If both have the same card, the money is split. **Goofspiel** is a bidding card game where players are trying to obtain the most points. Cards are shuffled and set face-down. Both players have  $K$  cards with values from 1 to  $K$ . These cards may be used as a bid. After bidding with that card, the player cannot play it again. Each turn, the top point card is revealed, and players simultaneously play a bid card; the point card is given to the highest bidder or discarded if the bids are equal. In this implementation, we use a fixed deck with  $K = 5$  and  $K = 6$ . **Liar's Dice** is a game where players have some number of dice and secretly roll. The first player bids rolled numbers, and the other player can either bid more or disbelieve the first player. When bidding ends with disbelief action, both players show dice. If the bid is right, the caller loses one die, and if the bidder is wrong, the bidder loses one die. Then the game continues, but for our computation, we use a version that ends with the loss of a die, and we use only a single die with four sides for each player.

## E PROOFS

LEMMA E.1. Let  $G$  be zero-sum imperfect-information extensive-form game with perfect recall. Let  $\sigma_2^F$  be fixed opponent's strategy, let  $T$  be some trunk of the game. If we perform CFR iterations in the trunk for player 1 then for the best iterate  $\hat{\sigma}_1 \max_{\sigma_1^* \in \Sigma_1} u_1(\sigma_1^*, \sigma_2^F)_V^T - u_1(\hat{\sigma}_1, \sigma_2^F)_V^T \leq \Delta \sqrt{\frac{A}{T}} |I_{TR}| + N_S \epsilon_S$  where  $\Delta$  is variance in leaf utility,  $A$  is an upper bound on the number of actions,  $|I_{TR}|$  is number of information sets in the trunk,  $N_S$  is the number of information sets at the root of any subgame and value function error is at most  $\epsilon_S$ .

PROOF. Using Theorem 2 from [4] we know that regret for player 1 is bounded  $R_1^T = \frac{1}{T} \max_{\sigma_1^* \in \Sigma_1} \sum_{t=1}^T (u_1(\sigma_1^*, \sigma_2^F)_V^T - u_1(\sigma_1^t, \sigma_2^F)_V^T) \leq \Delta \sqrt{AT} |I_{TR}| + TN_S \epsilon_S$ . Then we can directly map the regret to a different regret, that uses time-independent loss function  $l(\sigma_1) = -u_1(\sigma_1, \sigma_2^F)_V^T$ . We can then use Lemma 2 from [14] and we get  $l(\hat{\sigma}_1) - \min_{\sigma_1^* \in \Sigma_1} l(\sigma_1^*) \leq \frac{R_1^T}{T}$ . Substituting  $l$  and  $R_1^T$  back we get

$$\max_{\sigma_1^* \in \Sigma_1} u_1(\sigma_1^*, \sigma_2^F)_V^T - u_1(\hat{\sigma}_1, \sigma_2^F)_V^T \leq \Delta \sqrt{\frac{A}{T}} |I_{TR}| + N_S \epsilon_S$$

□

**THEOREM E.2.** Let  $G$  be zero-sum extensive-form game with perfect recall. Let  $\sigma_2^F$  be fixed opponent's strategy, let  $\mathcal{P}$  be any subgame partitioning of the game  $G$  and let  $\sigma_1^B$  be a CDBR given approximation  $\bar{V}$  of the optimal value function  $V$  with error at most  $\epsilon_V$ , partitioning  $\mathcal{P}$  and opponent strategy  $\sigma_2^F$  approximated in each step with regret at most  $\epsilon_R$ , formally  $\sigma_1^B = \sigma_1^B(\sigma_2^F)_V^{\mathcal{P}}$ . Let  $\sigma^{NE}$  be any Nash equilibrium. Then  $u_1(\sigma_1^B, \sigma_2^F) + |S|\epsilon_R + \sum_{S \in \mathcal{S}'} |I_{SB}|\epsilon_V \geq u_1(\sigma^{NE})$ .

**PROOF.** Using subgame partitioning  $\mathcal{P}$ , let  $T_1$  be the trunk of the game. From the properties of a NE  $u_1(\sigma^{NE}) \leq u_1(\sigma_1^{NE}, \sigma_2^F) \leq u_1(\sigma_1^{NE}, \sigma_2^F)_{V}^{T_1}$ . To compute CDBR we are maximizing in the trunk and using error in the value function and non-zero regret of the computed strategy  $u_1(\sigma_1^B, \sigma_2^F)_V^{T_1} + |I_{TB}|\epsilon_V + \epsilon_R \geq u_1(\sigma_1^{NE}, \sigma_2^F)_V^{T_1}$ . We continue using induction over steps with induction assumption that in step  $i$ ,  $u_1(\sigma_1^B, \sigma_2^F)_V^{T_i} + \epsilon_i \geq u_1(\sigma^{NE})$ . We already know it holds for  $T_1$ . Now we assume we have trunk  $T_{i-1}$  for which the induction step holds and trunk  $T_i$  which is  $T_{i-1}$  joined with new subgame  $S_{i-1}$ . Our algorithm recovers approximate equilibrium in  $S_{i-1}$  using  $\bar{V}$  at the boundary  $S_{i-1}^B$ , which means  $u_1(\sigma_1^B, \sigma_2^F)_V^{T_i} + |I_{SB_{i-1}}|\epsilon_V + \epsilon_R \geq u_1(\sigma_1^{NE_{S_{i-1}}} \cup \sigma_1^{B_{T_{i-1}}}, \sigma_2^F)_V^{T_i}$ . If we use equilibrium for the opponent in the subgame  $S_{i-1}$  we can replace equilibrium in the subgame by the value function  $V$  and we have  $u_1(\sigma_1^{NE_{S_{i-1}}} \cup \sigma_1^{B_{T_{i-1}}}, \sigma_2^F)_V^{T_i} \geq u_1(\sigma_1^B, \sigma_2^F)_V^{T_{i-1}}$  and joining it all together we have  $u_1(\sigma_1^B, \sigma_2^F)_V^{T_{i-1}} \leq u_1(\sigma_1^B, \sigma_2^F)_V^{T_i} + |I_{SB_{i-1}}|\epsilon_V + \epsilon_R$  and  $u_1(\sigma^{NE}) \leq u_1(\sigma_1^B, \sigma_2^F)_V^{T_i} + |I_{SB_{i-1}}|\epsilon_V + \epsilon_R + \epsilon_{i-1}$ . Accumulating the errors through the subgames will give the desired result  $u_1(\sigma_1^B, \sigma_2^F) + |S|\epsilon_R + \sum_{S \in \mathcal{S}'} |I_{SB}|\epsilon_V \geq u_1(\sigma^{NE})$ . We omit last subgame from the accumulated value function error because the last step does not use value function.  $\square$


**THEOREM E.3.** Let  $G$  be any zero-sum extensive-form game with perfect recall and let  $\sigma_2^F$  be any fixed opponent's strategy in  $G$ . Then we set  $G^M$  as restricted Nash response modification of  $G$  using  $\sigma_2^F$ . Let  $\mathcal{P}$  be any subgame partitioning of the game  $G^M$  and using some  $p \in (0, 1)$ , let  $\sigma_1^R$  be a CDRNR given approximation  $\bar{V}$  of the optimal value function  $V$  with error at most  $\epsilon_V$  and opponent strategy  $\sigma_2^F$  approximated in each step with regret at most  $\epsilon_R$ , formally  $\sigma_1^R = \sigma_1^R(\sigma_2^F, p)_V^{\mathcal{P}}$ . Let  $\sigma^{NE}$  be any Nash equilibrium in  $G$ . Then  $u_1(\sigma_1^R, \sigma_2^F) + \sum_{S \in \mathcal{S}'} |I_{SO}|(1-p)\epsilon_V + |S|\epsilon_R + \sum_{S \in \mathcal{S}'} |I_{SB}|\epsilon_V \geq u_1(\sigma^{NE})$ .

**PROOF.** Let  $T_1^M$  be a trunk of a modified game  $G^M$  using partitioning  $\mathcal{P}$ . We will use  $u^G(\sigma)$  as utility in  $G$ . Utility of player 1 for playing Nash equilibrium of the  $G$  in trunk  $T_1$  will be higher or the same as game value of  $G$ , formally  $u_1^G(\sigma^{NE}) \leq u_1^{G^M}(\sigma^{NE})_{V}^{T_1^M}$ . To compute CDRNR we use the approximate value function. In the fixed part of the game  $G^F$  the value will be worse at most by sum of errors as in the CDBR case. However, in the  $G'$  the situation is more complicated and we use Theorem 2 from [4] to bound the utility increase, resulting in  $u_1^{G^M}(\sigma_1^R, BR(\sigma_1^R))_V^{T_i} + |I_{T^M, B}|\epsilon_V + |I_{T^M}|\epsilon_R \geq u_1^G(\sigma^{NE})$ . We continue using induction over steps with induction assumption that in step  $i$ ,  $u_1^{G^M}(\sigma_1^R, BR(\sigma_1^R))_V^{T_i} + \epsilon_i \geq u_1^G(\sigma^{NE})$  and we already showed it holds for  $T_1$ . Now we assume we have

trunk  $T_{i-1}$  for which the induction step holds and trunk  $T_i$  which is  $T_{i-1}$  joined with new subgame  $S_{i-1}$ . Our algorithm recovers approximate equilibrium in  $S_{i-1}^M$  and we want similar equation as for the CDBR. Part of the game tree  $G^F$  has the errors bounded as in CDBR but because we use gadget in the  $G'$  we need to also consider error in actions ending with value function player 2 can play in the top with error bounded by  $\epsilon_V$ . We have  $|I_{SO}|$  of actions leading out of the tree so the error increase in the  $G'$  going to the next subgame is at most  $|I_{SO}|\epsilon_V + |I_{S_{i-1}^M, B}|\epsilon_V + |I_{S_{i-1}^M}|\epsilon_R$  which together gives us  $u_1^{G^M}(\sigma_1^R, BR(\sigma_1^R))_V^{T_i} + (1-p)|I_{SO}|\epsilon_V + |I_{S_{i-1}^M, B}|\epsilon_V + |I_{S_{i-1}^M}|\epsilon_R \geq u_1^{G^M}(\sigma_1^R, BR(\sigma_1^R))_V^{T_{i-1}}$ , where  $\sigma_1^R$  is a combination of the strategy we approximated in the subgame and the fixed strategy from previous step, formally  $\sigma_1^R = \sigma_1^{S_{i-1}} \cup \sigma_1^{R, T_{i-1}}$ . Joining it with the induction assumption we have  $u_1^{G^M}(\sigma_1^R, BR(\sigma_1^R))_V^{T_i} + (1-p)|I_{SO}|\epsilon_V + |I_{S_{i-1}^M, B}|\epsilon_V + |I_{S_{i-1}^M}|\epsilon_R + \epsilon_{i-1} \geq u_1^G(\sigma^{NE})$ . Accumulating the errors in the last subgame we have  $u_1^{G^M}(\sigma_1^R, BR(\sigma_1^R)) + \sum_{S \in \mathcal{S}'} |I_{SO}|(1-p)\epsilon_V + |S|\epsilon_R + \sum_{S \in \mathcal{S}'} |I_{SB}|\epsilon_V$ . However, we still need to show it works for  $u_1^G(\sigma_1^R, \sigma_2^F)$ . We can do it by replacing strategy of player 2 in the  $G'$  by  $\sigma_2^F$  which will effectively transform  $G^M$  game back to  $G$  with player 2 playing  $\sigma_2^F$ . Since we did this transformation by changing the strategy that was a best response the utility can only increase and  $u_1^G(\sigma_1^R, \sigma_2^F) \geq u_1^{G^M}(\sigma_1^R, BR(\sigma_1^R))$  which concludes the proof.  $\square$

**THEOREM E.4.** Let  $G$  be any zero-sum extensive-form game with perfect recall and let  $\sigma_2^F$  be any fixed opponent's strategy in  $G$ . Then we set  $G^M$  as restricted Nash response modification of  $G$  using  $\sigma_2^F$ . Let  $\mathcal{P}$  be any subgame partitioning of the game  $G^M$  and using some  $p \in (0, 1)$ , let  $\sigma_1^R$  be a CDRNR given approximation  $\bar{V}$  of the optimal value function  $V$  with error at most  $\epsilon_V$ , partitioning  $\mathcal{P}$  and opponent strategy  $\sigma_2^F$ , which is approximated in each step with regret at most  $\epsilon_R$ , formally  $\sigma_1^R = \sigma_1^R(\sigma_2^F, p)_V^{\mathcal{P}}$ . Then exploitability has a bound  $\mathcal{E}(\sigma_1^R) \leq \mathcal{G}(\sigma_1^R, \sigma_2^F) \frac{p}{1-p} \sum_{S \in \mathcal{S}'} |I_{SO}|(1-p)\epsilon_V + |S|\epsilon_R + \sum_{S \in \mathcal{S}'} |I_{SB}|\epsilon_V$ ,  $\mathcal{E}$  and  $\mathcal{G}$  are defined in Section 2.

**PROOF.** We will examine the exploitability increase in each step. First, we define gain in a single step as  $\mathcal{G}(\sigma_1, \sigma_2)_V^{T_i} = u_1(\sigma_1, \sigma_2)_V^{T_i} - u_1(\sigma_1, \sigma_2)_V^{T_{i-1}}$  for  $i > 0$  and  $\mathcal{G}(\sigma_1, \sigma_2)_V^{T_0} = u_1(\sigma_1, \sigma_2)_V^{T_0} - u_1(\sigma^{NE})$ . This is consistent with full definition of gain because sum of gains over all steps will result in  $u_1(\sigma_1, \sigma_2)_V^{T_n} - u_1(\sigma_1, \sigma_2)_V^{T_0} + u_1(\sigma_1, \sigma_2)_V^{T_0} - u_1(\sigma^{NE}) = u_1(\sigma_1, \sigma_2)_V^{T_n} - u_1(\sigma^{NE}) = \mathcal{G}(\sigma_1, \sigma_2)$ . We define exploitability in a single step similarly as  $\mathcal{E}(\sigma_1)_V^{T_i} = u_2(\sigma_1, BR(\sigma_1))_V^{T_i} - u_2(\sigma_1, BR(\sigma_1))_V^{T_{i-1}}$  for  $i > 0$  and  $\mathcal{E}(\sigma_1)_V^{T_0} = u_2(\sigma_1, BR(\sigma_1))_V^{T_0} - u_2(\sigma^{NE})$  and it also sums to full exploitability. In each step we approximate the strategy in the modified game, having full utility in step written as  $\mathcal{G}(\sigma_1^R, \sigma_2^F)_V^{T_i} p - \mathcal{E}(\sigma_1^R)_V^{T_i} (1-p)$ . If we had exact equilibrium in the subgame this would always be at least 0. However, we have  $\bar{V}$  instead of  $V$ , values at the top of the gadget are not exact and the computed strategy has regret  $\epsilon_R$ . As in the previous proof the error is bounded by  $|I_{SO}|(1-p)\epsilon_V + |I_{S_{i-1}^M, B}|\epsilon_V + |I_{S_{i-1}^M}|\epsilon_R$  and we can write



**Figure 10: Example of game where step best response is worse than NE against fixed strategy  $\sigma(h) = \frac{2}{3}, \sigma(x) = 1$ .**

$\mathcal{G}(\sigma_1^R, \sigma_2^F)_{V_i}^{T_i^M} p - \mathcal{E}(\sigma_1^R)_{V_i}^{T_i^M} (1-p) + |I_{SO}|(1-p)\epsilon_V + |I_{S_{M,B}}|\epsilon_V + |I_{S_{M-1}}|\epsilon_R \geq 0$ . We reorganize the equation to get  $\mathcal{G}(\sigma_1^R, \sigma_2^F)_{V_i}^{T_i^M} \frac{p}{1-p} + |I_{SO}|(1-p)\epsilon_V + |I_{S_{M,B}}|\epsilon_V + |I_{S_{M-1}}|\epsilon_R \geq \mathcal{E}(\sigma_1^R)_{V_i}^{T_i^M}$  and summing over all the steps gives us  $\mathcal{G}(\sigma_1^R, \sigma_2^F)_{V_i}^{T_i^M} \frac{p}{1-p} + \sum_{S \in S'} |I_{SO}|(1-p)\epsilon_V + |S|\epsilon_R + \sum_{S \in S'} |I_{S,B}|\epsilon_V \geq \mathcal{E}(\sigma_1^R)$   $\square$

## F CDBR AGAINST NASH STRATEGY

**OBSERVATION 1.** An example in Figure 10 shows that CDBR can perform worse than a Nash equilibrium against the fixed opponent because of the perfect opponent assumption after the depth-limit. An example is a game of matching pennies with a twist. Player 2 can choose in the case of the tails whether he wants to give the opponent 10 instead of only 1. A rational player will never do it, and the equilibrium is a uniform strategy as in normal matching pennies.

Now we have an opponent model that plays  $h$  with probability  $\frac{2}{3}$  and always plays  $x$ . The best response to the model will always play  $T$  and get payoff  $\frac{10}{3}$ . Nash equilibrium strategy will get payoff 2, and CDBR with depth-limit 2 will cut the game before the  $x/y$  choice. Assuming the opponent plays perfectly after the depth-limit and chooses  $y$ , 1 will always play  $H$ . Playing  $H$  will result in receiving payoff  $\frac{2}{3}$ , which is higher than the value of the game ( $\frac{1}{2}$ ) but lower than what Nash equilibrium can get against the model.

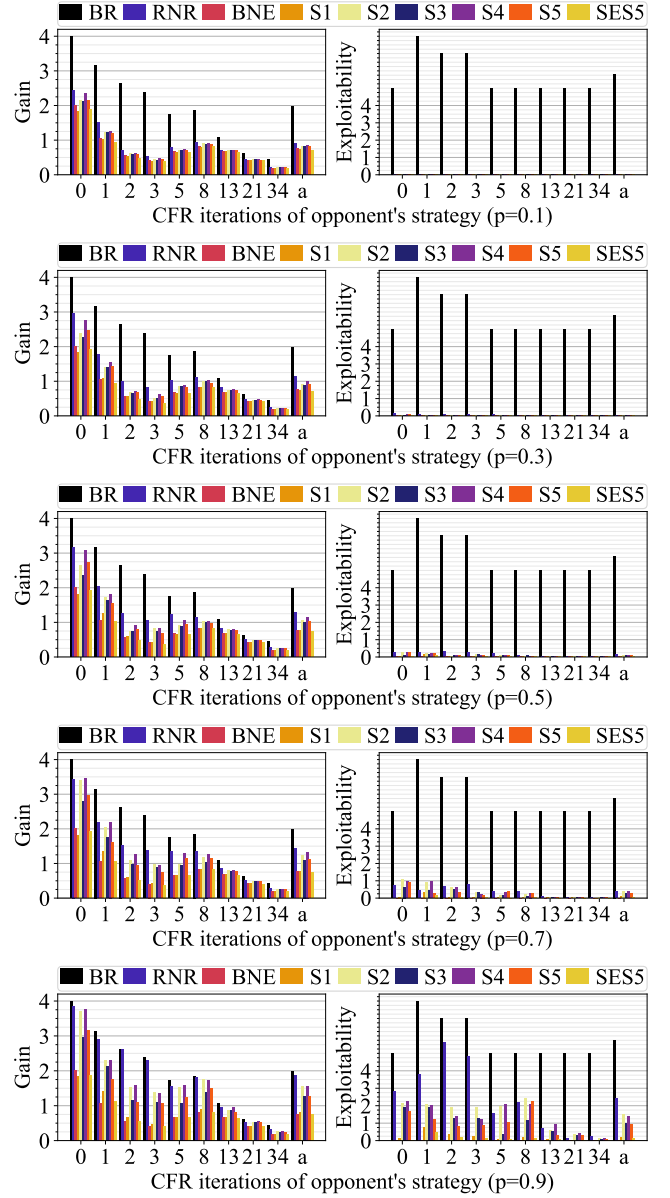
## G ADDITIONAL EMPIRICAL RESULTS

**CDRNR.** We show more results for Goofspiel, Leduc Hold'em, and Liar's dice with different values of  $p$ . SX is CDRNR with step size denoted by  $X$ . We also evaluate SES and only use the highest step value of 5. Next, we show the same setup as in the main text with exactly the same partitioning as they used in SES, and we include more values of  $p$ .

**Repeated RPS.** In Figure 15, we show the strategy sets recovered for all possible  $p$  against one strategy in two round biased RPS where after the round information is revealed. As we explained before, we can see that SES cannot gain anything in a game where only information imperfections are simultaneous moves. Exp-strat can exploit only the second round, and it gains half of the maximum, while the other algorithms can gain the maximum and are more or less successful in achieving the best trade-off. The full gadget is the best, followed by the other gadgets without theoretical guarantees, and then by a combination of Nash and CDBR.

## H SES BOUND

The bound in SES [13] is



**Figure 11: Additional results for CDRNR showing the performance of CDRNR with varying step-size. Generated on Goofspiel 5.**

**THEOREM H.1.** Let  $\mathbb{S}$  be a disjoint set of subgames  $S$ . Let  $\sigma^* = \langle \sigma_1^*, \sigma_2^* \rangle$  be the NE where player 1's strategy is constrained to be the same with  $\sigma_1$  outside  $\mathbb{S}$ . Define  $\Delta = \max_{S \in \mathbb{S}, I_2^i \in S_{top}} |CBV_2^{\sigma_1^*}(I_2^i) - v_2^{\sigma_2^*}(I_2^i)|$ . Let  $\tilde{p}(I_2^i)$  be the reach probability given by  $\sigma_2^*$ . Let  $\hat{p}(I_2^i)$  be the estimation of reach probability  $p(I_2^i)$  given by the real opponent strategy. Define  $\tau = \max_{S \in \mathbb{S}, I_2^i \in S_{top}} \left| \frac{\hat{p}(I_2^i) - \tilde{p}(I_2^i)}{\tilde{p}(I_2^i)} \right|$ . Whenever  $1 - (2\tau + 1)\alpha > 0$ , the exploitability bound is given by:



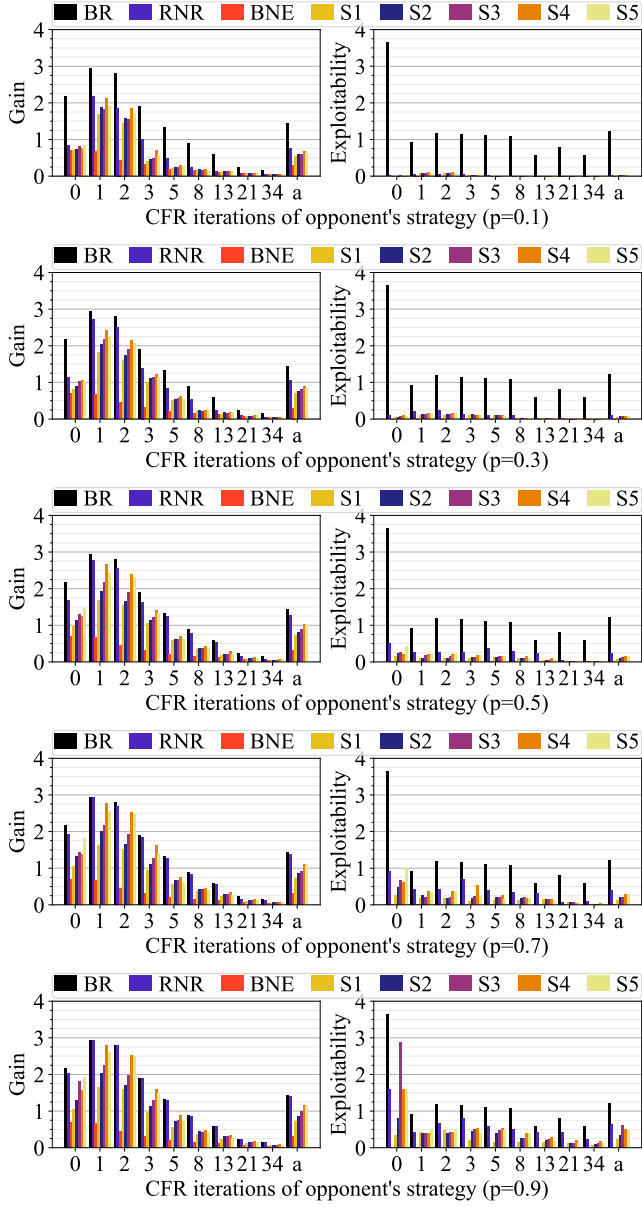


Figure 12: Additional results for CDRNR showing the performance of CDRNR with varying step-size. Generated on Leduc Hold'em.

$$\mathcal{E}(\sigma'_1) \leq \mathcal{E}(\sigma_1^*) + \frac{2}{1 - (2\tau + 1)\alpha} \Delta$$

We switched the players since authors in the previous work use player 2 as the rational player.

We can see that the bound relies on the estimation being close to an equilibrium strategy defined by authors as  $\tau$ . However, it does max over all the differences in reaches to the subgame, and in practice, some of the reaches will be very different, resulting in a large value of  $\tau$ . To demonstrate the difference, we assume

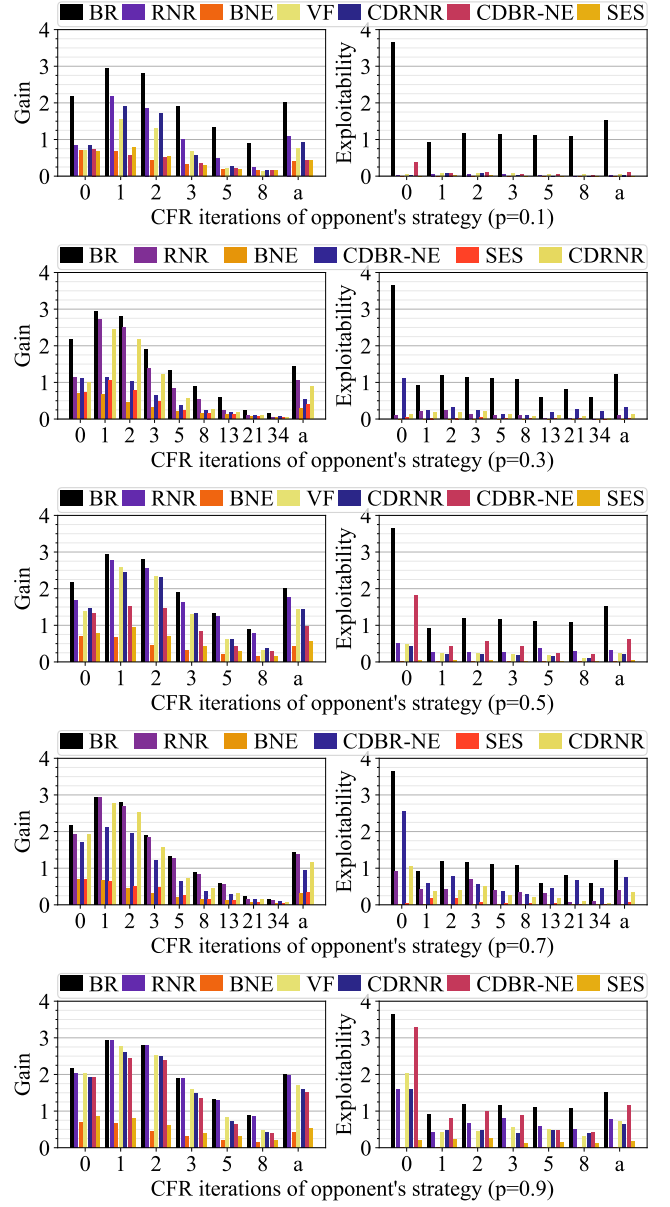




Figure 13: Additional results for CDRNR with different values of  $p$ . Generated on Leduc Hold'em split only by the round.

the opponent model plays such that some action difference from equilibrium is 1, which is the highest it can be, and hence  $\tau = 1$ . Parameter  $\alpha$  in SES directly matches  $p$ . For  $\alpha = 0$ , the bound is the same as in the max-margin gadget, and  $\tau$  is disregarded. However, as  $\alpha$  increases, the bound steeply rises, and as  $\alpha$  goes in the limit to  $\frac{1}{3}$ , the bound goes to infinity, and for any larger  $\alpha$ , the bound says nothing. In comparison, our bound does not have this problem, and in the same setup, with  $p = 0.5$ , our bound still limits the exploitability by exactly the gain achieved. Note that since in SES, they do not account for errors in value function and errors in



**Figure 14: Additional result on Liar's dice. For every  $p$  it exactly mimics the RNR so we only show one value.**



**Figure 15: Results showing gain and exploitability trade-off in two round biased RPS. Max-margin and resolving gadget overlaps**

resolving, for this comparison only, we also omitted error terms caused by those errors.