
Off-Belief Learning

Hengyuan Hu¹ Adam Lerer¹ Brandon Cui¹ Luis Pineda¹ Noam Brown¹ Jakob Foerster¹

Abstract

The standard problem setting in Dec-POMDPs is self-play, where the goal is to find a set of policies that play optimally together. Policies learned through self-play may adopt arbitrary conventions and rely on multi-step counterfactual reasoning based on assumptions about other agents’ actions and thus fail when paired with humans or independently trained agents. In contrast, no current methods can learn optimal grounded policies that do not rely on counterfactual information from observing other agents’ actions. To address this, we present *off-belief learning* (OBL): at each time step OBL agents interpret all *past actions* as if they were taken by a given, fixed policy, but assume that *future actions* will be taken by an optimal policy. When π_0 is uniform random, OBL learns the optimal grounded policy. OBL can be iterated in a hierarchy, where the optimal policy from one level becomes the input to the next. This introduces counterfactual reasoning in a controlled manner. Unlike traditional RL which may converge to any equilibrium policy, OBL converges to a unique policy, making it more suitable for zero-shot coordination. OBL can be scaled to high-dimensional settings with a *fictitious transition* mechanism and shows strong performance in both a toy-setting and the benchmark human-AI/zero-shot coordination problem Hanabi.

1. Introduction

An important goal of multi-agent reinforcement learning (MARL) research is to develop AI systems that can coordinate with novel partners, such as humans or other artificial agents, unseen at training time. However, the standard problem setting in cooperative MARL is *self-play* (SP), where a team of agents is trained to find optimal policies that achieve the highest scores when all parties follow such policies faith-

fully, without considering performance with other agents. As a result, a strong joint policy often relies on efficient, yet arbitrary conventions, and the inflexible assumption that other agents will adhere to their side of the policy (Foerster et al., 2019).

Although SP agents achieve high scores among themselves, they often fail when paired with humans because they cannot infer human intentions, and vice versa (Carroll et al., 2019; Bard et al., 2020; Hu et al., 2020). In imperfect information settings, agents rely on the actions of others to infer things not observed or known based on their beliefs (implicit or otherwise) about how other agents would act in different states, a phenomenon as known as *counterfactual reasoning* (Foerster et al., 2019).

For example, suppose in the board game Hanabi (detailed game rules in Section 6) our partner reveals that a particular card is green. We now know the fact that the card is green, which is often referred to as *grounded* or *factual* information. But we may also infer information counterfactually. For example if our partner and us have agreed on a convention that we would only reveal the color of a card if it were playable, we can infer that the green card must be playable. We might then say there is a shared *convention* where in that situation “green” also means “playable”, not simply “green”.

Conventions may of course vary. Perhaps “green” means to discard, rather than play. Perhaps it means the adjacent card to it is playable. Perhaps something entirely different. Agents trained via self-play tend to learn an unpredictable set of conventions where the meaning of each action in each possible situation can be completely different in different training runs, and implicitly assume arbitrary cognitive depth, each recursively expecting the other to understand perfectly. When paired with a novel agent at test time, this behavior is disastrously fragile.

By contrast, humans are often observed to limit themselves to a cognitive depth of one or two levels (Camerer et al., 2003), and fewer when other agents are anticipated to be less experienced (Agranov et al., 2012) across a wide variety of competitive and collaborative games. When collaborating with unknown partners for the first time, humans may still successfully communicate by means of simple factual information sharing, without making fragile inferences based on

¹Facebook AI Research. Correspondence to: Hengyuan Hu <hengyuan@fb.com>, Jakob Foerster <jnf@fb.com>.

their partner’s unknown policy, even while further planning their own actions effectively and assuming that other players will do the same. This is a behavior that current methods cannot replicate.

In this paper, we present *off-belief learning* (OBL), a MARL algorithm that addresses this problem by controlling the counterfactual reasoning depth, thereby preventing the use of arbitrary conventions while allowing for optimal planning under such constraint. Given a policy π_0 , OBL agents follow a policy π_1 that at each step, that acts optimally assuming that all past actions were taken by π_0 and all future actions will be taken by π_1 . OBL also accounts for the fact that at each future step from π_1 will again be interpreted as if they were chosen by π_0 . OBL can be applied iteratively to obtain a hierarchy of policies that converges to a unique solution, making it a natural candidate for solving zero-shot coordination problems where the goal is to maximize the cross-play performance between policies from independent training runs of the same algorithm.

Intuitively, OBL address a fundamental problem in zero-shot coordination (ZSC), i.e. “Can an agent act intelligently/optimally, and expect other agents to do the same, *without* assuming conventions?”. If π_0 is a random policy, then π_1 is able learn a grounded policy that plays the game relying only on the facts they know about the states without assuming any conventions since a random agent sends any signal with equal chance regardless of state. By assuming π_1 recursively in the future, OBL assumes partners will act intelligently. Since π_1 will take gainful actions based on factual information, π_1 will also signal its partner with helpful factual information, despite lacking pre-agreed conventions.

We first characterize OBL theoretically. We start with proving that OBL converges to a unique policy, i.e. given the same π_0 , OBL produces the the same π_1 in independent training runs. We then prove that OBL is a policy improvement operator under certain condition. We further show that when π_0 doesn’t depend on private information, OBL will produce a fully *grounded* policy whose belief over the current state does not depend on partner’s policy. We then evaluate OBL in both a toy setting and Hanabi. In the toy setting, we demonstrate that OBL learns the optimal grounded policy while other existing methods such as SP and k-level reasoning fail to do so. In Hanabi, OBL finds fully-grounded policies that reach a score of 20.92 in SP without relying on conventions, a important data point that tells us how well we can perform in this benchmark without conventions. OBL performs strongly in zero-shot coordination settings, achieving significantly higher cross-play scores than previous methods. It also plays well in ad-hoc team play with an agent trained by imitation learning from human data and two distinct agents trained with reinforcement learning, all of which are unseen at training time.

2. Related Work

This work is inspired by the notion of Cognitive Hierarchies (CH) (Camerer et al., 2004) and k-level reasoning. CH is a framework for decision-making in multi-agent settings. Like OBL, CH starts with an ‘uninformative’ level-0 policy and constructs a hierarchy of policies that are optimal assuming that other agents play the policies at lower levels. However, in contrast to OBL, CH does not differentiate between reasoning about other agents’ *past behavior* and expected *future behavior*. In other words, assuming that π_0 is a uniform random policy, π_1 is simply a *best response* to a random player. While this seems like a reasonable default, it has some rather drastic consequences. In particular, unlike OBL, it is impossible for CH to learn optimal grounded policies in many settings. The intuition here is simple: Only the first CH level, the best response to a random agent, interprets actions as entirely grounded, i.e., without any counterfactual reasoning. However, the first level also assumes that the future moves of the partner is entirely random, which means it cannot learn strategies that require the partner to cooperate. One class of behavior that is immediately affected is *grounded signaling*, where an agent can reveal information to another agent through a costly action, as we show in Section 5.

The striking difference between OBL and CH can also be illustrated in fully observable turn-based settings. While OBL learns the optimal policy at the first level, CH will in general need as many levels as the length of the episode to converge to an optimal policy. Another difference is the role of π_0 : While in CH the exact π_0 can change the final policy, in OBL any policy π_0 that ignores the private observation will result in the same, perfectly uninformative belief and lead to the same OBL policy π_1 .

OBL is conceptually related to the Rational Speech Acts (RSA) framework (Frank & Goodman, 2012). RSA assumes a speaker-listener setting with grounded hints and starts out with a literal listener (LL) that only considers the grounded information provided. Like CH, RSA then introduces a hierarchy of speakers and listeners, each level defined via Bayesian reasoning (i.e., a best response) given the level below. OBL allows us to train an analogous hierarchy and thus introduce pragmatic reasoning in a controlled fashion. However, while RSA is mostly a conceptual framework that is focused on simple communication tasks, OBL is designed to deal with high-dimensional, complex Dec-POMDPs in which agents have to both act and communicate through their actions. The fact that RSA has been able to explain a large amount of human communication behavior supports the importance of OBL.

There is a large body of work on coordination in behavioral game theory. One of the most well-known frameworks is focal points (Schelling, 1980), i.e., common knowledge

labels in the environment that allow agents to coordinate on a joint action. In contrast to this line of work we address the *zero-shot coordination* setting (Hu et al., 2020), where the structure of the Dec-POMDP itself needs to be used for coordination and players specifically cannot coordinate on labels.

In this setting *Other-play* (OP) (Hu et al., 2020) has been proposed as a method to prevent agents from learning an arbitrary joint policy out of a set of equivalent but mutually incompatible ones. OP focuses on enforcing equivariance to symmetries in the Dec-POMDP, which must be provided as an input to the algorithm. While OP prevents agents from breaking symmetries of the environment to encode conventions, it does not prevent all arbitrary convention formation. On the other hand, OBL, by starting from grounded information, does avoid both symmetry-breaking and arbitrary convention formation when run with an appropriate π_0 , and unlike OP also does not require those symmetries to be manually identified. Additionally, OBL specifically addresses how agents *communicate* information through their actions, and is tailored for turn-based interactions rather than those with simultaneous moves. The differences between OP and OBL are highlighted by the different playing strategies that the OBL and OP bot develop in Hanabi, see Section 6. OP learns policies that use rank-hints for signalling, while OBL agents learn more human-compatible, grounded behavior.

3. Background

Dec-POMDPs and Grounded Policy. We consider Dec-POMDP environments G with state s , action a , observation function $\Omega^i(s)$ for each player i and transition function $\mathcal{T}(s, a)$. Players receive a common reward $R(s, a)$. We denote the historical trajectory as $\tau = (s_1, a_1, \dots, a_{t-1}, s_t)$ and *action-observation history* (AOH) for player i as $\tau^i = (\Omega^i(s_1), a_1, \dots, a_{t-1}, \Omega^i(s_t))$, which encodes the trajectory from player i 's point of view. A policy $\pi^i(a|\tau_t^i)$ for player i takes as input an AOH and outputs a distribution over actions. We denote the joint policy as π . We remark that while the value $V^\pi(\tau)$ of policy π from any state s is well-defined, the value of playing π from an AOH τ^i is not well-defined without specifying what policy was played *leading* to τ^i , because it affects the distribution of world states at τ^i . The distribution of world states is often referred to as belief $\mathcal{B}_\pi(\tau|\tau^i) = P(\tau|\tau^i, \pi)$. Normally the belief depends on the policy π through Bayes' rule. However, *grounded belief* is a special case where it only conditions on the observations but not the partners' policies, i.e.

$$\mathcal{B}_G(\tau|\tau^i) = \frac{P(\tau) \prod_t P(o_t^i|\tau)}{\sum_{\tau'} P(\tau') \prod_t P(o_t^i|\tau')}. \quad (1)$$

We can then call the policy π *grounded policy* if $\mathcal{B}_\pi = \mathcal{B}_G$. For example, a random policy that uniformly picks an action

at each time step is a grounded policy since we cannot infer any additional information from their actions through counterfactual reasoning.

In this work, we restrict ourselves to *bounded-length, turn-based Dec-POMDPs*. Formally, we assume that G reaches a terminal state after at most t_{max} steps, and that only a single agent acts at each state s (other agents have a null action). See Section 4.5 for a discussion of turn-based vs. simultaneous-action environments.

Deep MARL has been successfully applied in many Dec-POMDP settings (de Witt et al., 2019; Foerster et al., 2019). The specific OBL algorithm to be introduced later uses Recurrent Replay Distributed Deep Q-Network (R2D2) (Kapturowski et al., 2019) as backbone. In deep Q-learning (Mnih et al., 2015) the agent learns to predict the expected total return for each action given the AOH, $Q(\tau_t^i, a) = E_{\tau \sim P(\tau|\tau_t^i)} R_t(\tau)$. $R_t(\tau) = \sum_{t'=t}^{\infty} \gamma^{(t'-t)} r_{t'}$ is the forward looking return from time t where $r_{t'}$ is reward at step t' with $a_t = a$ and $a_{t'} = \arg\max_{a'} Q(\tau_{t'}, a')$ for $t' > t$ and γ is an optional discount factor. The state-of-the-art R2D2 algorithm incorporates many modern best practices such as double-DQN (van Hasselt et al., 2016), dueling network architecture (Wang et al., 2016), prioritized experience replay (Schaul et al., 2016), distributed training setup with parallel running environments (Horgan et al., 2018) and recurrent neural network for handling partial observability.

The straightforward way to apply deep Q-learning to Dec-POMDP settings is Independent Q-Learning (IQL) (Tan, 1993) where each agent treats other agents as part of the environment and learns an independent estimate of the expected return without taking other agents' actions into account in the bootstrap process. Many methods (Sunehag et al., 2017; Rashid et al., 2018) have been proposed to learn joint Q-functions to take advantage of the centralized training and decentralized control structure in Dec-POMDPs. In this work, however, we use the IQL setup with shared neural network weights θ for simplicity and note that any potential improvements to the RL algorithm used are orthogonal to our contribution.

Zero-Shot Coordination. The most common problem setting for learning in Dec-POMDPs is *self-play* (SP) where a team of agents is trained and tested together. Optimal SP policies typically rely on *arbitrary conventions*, which the entire team can jointly coordinate on during training. However, many real-world problems require agents to coordinate with other, unknown AI agents and humans at test time. This desiderata was formalized as the *Zero-Shot Coordination* (ZSC) setting by Hu et al. (2020), where the goal is stated as finding algorithms that allow agents to coordinate with *independently trained* agents at test time. ZSC immediately rules out *arbitrary conventions* as optimal solutions and in-

stead requires learning algorithms that produce robust and, ideally, unique solutions across multiple independent runs.

4. Off-Belief Learning

One of the big challenges in ZSC under partially observable settings is to determine how to *interpret* the actions of other agents and how to select actions that will be interpretable to other agents. In the following we introduce OBL, a method that allows us to address this problem by learning optimal *grounded* policies, which do not interpret actions of other agents. We will discuss OBL both as a formal method with corresponding proofs and as a scalable algorithm based on a *fictitious transition* mechanism applicable to the deep RL setting.

4.1. The OBL Operator

We first introduce the OBL operator that computes π_1 given any π_0 . If a common knowledge policy π_0 is played by all agents up to τ^i , then agent i can compute a belief distribution $\mathcal{B}_{\pi_0}(\tau|\tau^i) = P(\tau|\tau^i, \pi_0)$ conditional on its AOH. This belief distribution fully characterizes the effect of the history on the current state. Notice that the return for (all players) playing a ‘counterfactual’ policy π_0 to τ^i and π_1 thereafter, which we denote $V^{\pi_0 \rightarrow \pi_1}(\tau^i)$, is precisely the expected return for all players playing according to π_1 , starting from a trajectory $\tau \sim \mathcal{B}_{\pi_0}(\tau^i)$. Therefore we define a counterfactual value function as follows:

$$V^{\pi_0 \rightarrow \pi_1}(\tau^i) = \mathbb{E}_{\tau \sim \mathcal{B}_{\pi_0}(\tau^i)} [V^{\pi_1}(\tau)]. \quad (2)$$

We can similarly define counterfactual Q values as

$$Q^{\pi_0 \rightarrow \pi_1}(a|\tau_t^i) = \sum_{\tau_t, \tau_{t+1}} \mathcal{B}_{\pi_0}(\tau_t|\tau_t^i) [R(s_t, a) + \mathcal{T}(\tau_{t+1}|\tau_t) V^{\pi_1}(\tau_{t+1})]. \quad (3)$$

Then *OBL* operator is defined to be the operator that maps an initial policy π_0 to a new policy π_1 as follows:

$$\pi_1(a|\tau^i) = \frac{\exp(Q^{\pi_0 \rightarrow \pi_1}(a|\tau^i)/T)}{\sum_{a'} \exp(Q^{\pi_0 \rightarrow \pi_1}(a'|\tau^i)/T)} \quad (4)$$

for some temperature hyperparameter T .

4.2. Properties of Off-Belief Learning

Theorem 1. *For any $T > 0$ and starting policy π_0 , OBL computes a unique policy π_1 .*

Proof in the Appendix. “Unique” means that that, once T is fixed, π_1 is a well-defined deterministic function of π_0 . This means that, unlike most MARL methods, OBL always converges to the same answer in principle, regardless of random initialization or other hyperparameters.

Corollary 1. *For any $T > 0$ and starting policy π_0 , N agents independently computing OBL joint policy $\pi_1 = \{\pi_1^i\}$ and playing their parts of it (π_1^i) will achieve the same return as if they had computed a centralized π_1 policy (zero-shot coordination).*

Proof. Since all π_1 are identical, this follows trivially. \square

Theorem 1 is trivial in a single-agent context, but illustrates a substantial departure from traditional multi-agent learning rules, under which independently computed policies for each agent are typically *not* unique or compatible due to the presence of multiple equilibria, a particularly severe example of which is the formation of ‘arbitrary’ conventions for communication in a game like Hanabi.

Theorem 2. *For every policy π_1 generated by OBL from π_0 , $J(\pi_1) \geq J(\pi_0) - t_{max}T/e$, i.e. OBL is a policy improvement operator except for a term that vanishes as $T \rightarrow 0$.*

Proof in the Appendix. Theorem 2 implies that self-play performance will gradually increase if we apply OBL iteratively. Together with Corollary 1, it also guarantees that zero-shot coordination performance, i.e. cross-play between independent training runs, will also improve at the same pace.

Notably, unlike standard MARL, the fixed points of the OBL learning rule *are not* guaranteed to be equilibria of the game. However, we have the theorem:

Theorem 3. *If repeated application of the OBL policy improvement operator converges to a fixed point policy π , then π is an ϵ -subgame perfect equilibrium of the Dec-POMDP, where $\epsilon = t_{max}T/e$.*

Proof in the Appendix.

4.3. Optimal Grounded Policies

In multi-agent cooperative settings with imperfect information, the optimal policy for an agent depends not only on assumptions about the future policy, but also on counterfactual assumptions about other agents’ policies – i.e. what do those agents’ prior actions say about their private information? Making wrong assumptions is a source of coordination failure and we thus may wish to consider grounded policies that avoid counterfactual reasoning about partners’ actions altogether.

Given the grounded belief defined in 1, we can define an *optimal grounded policy* π_G to be any policy that at each AOH τ^i plays an action that maximizes expected reward assuming the state distribution at τ^i is drawn from the grounded belief, and assuming that π_G is played thereafter.

Theorem 4. *Application of OBL to any constant policy $\pi_0(a|\tau^i) = f(a)$ – or in fact any policy that only conditions on public state – yields an optimal grounded policy in the limit as $T \rightarrow 0$.*

4.4. Iterated Application of OBL

While an optimal grounded policy can be an effective baseline in some settings, clearly there are many settings where counterfactual reasoning improves performance. In humans this ability is commonly described as *theory of mind* and a variety of experimental studies (Camerer et al., 2003; Agronov et al., 2012; Kleiman-Weiner et al., 2016) have shown that humans typically carry out a limited number of counterfactual reasoning steps, depending on the exact situation, even in zero-shot settings.

Iterating OBL is simple. When the OBL operator is applied to a random policy π_0 , it computes an optimal grounded policy π_1 , which we also refer to as OBL “level 1”. We may compute OBL level 2 similarly by using OBL level 1 as π_0 and compute OBL level 3 by replacing π_0 with OBL level 2, and so on.

Level 2 acts optimally while using beliefs induced by level 1 to interpret partners’ actions. Different from OBL level 1, a level 2 agent will assume that other agents will take positive-utility actions instead of purely random actions and share useful factual information to others. For example, if it sees another agent avoid an action that would likely be beneficial, it will infer that the other agent probably has private information that the action is not so beneficial.

Since all independently trained agents use the same belief (albeit small learning errors), this training procedure is helpful in ZSC, unlike self-play where agents may learn arbitrary beliefs and signaling. Indeed we show in Section 6 that while OBL level 1 already cooperates better with a human-like policy than any prior method, level 2 and higher cooperate even better with it and other tested agents.

The iterated application of OBL allows us to directly control for the number of counterfactual reasoning loops to be carried out by our agent. Different from the cognitive hierarchies framework, at each level the agent plays optimally assuming optimal game play in the future, rather than assuming the partners are fixed and play one level below. Therefore it is possible to vary the overall *quality* of the policy independently of the depth of counterfactual reasoning in OBL but not in cognitive hierarchy because we cannot reach the optimal policy at each level.

4.5. Simultaneous-action vs turn-based environments

So far we have restricted ourselves to *turn-based* games in which a single agent acts at each state. If two players acted simultaneously, then OBL couldn’t guarantee convergence

to a unique policy (Thm 1); for example, consider a one-step coordination game with payoff matrix $\begin{bmatrix} 1, 0 \\ 0, 1 \end{bmatrix}$ (which is a single-state, single-step Dec-POMDP). OBL is identical to self-play in a one step game, and could thus converge to either equilibrium.

This may seem like a strong restriction of OBL to turn-based games. However, any simultaneous-action game can be converted to an equivalent stage game by ordering the simultaneous actions into separate timesteps and only letting players observe the actions of others once the timesteps corresponding to the simultaneous actions have completed. An OBL policy can be generated for this equivalent turn-based game. The ordering of players’ actions solves the equilibrium selection problem, because the player who acts second assumes that the other player acted according to π_0 . The player order we use to convert the simultaneous-action game to a turn-based one may change the OBL policy.

4.6. Algorithms for Off-Belief Learning

Equation 4 immediately suggests a simple algorithm for computing an OBL policy in small tabular environments: compute $\mathcal{B}_{\pi_0}(\tau^i)$ for each AOH, and then compute $Q^{\pi_0 \rightarrow \pi_1}(\tau^i)$ for each AOH in ‘backwards’ topological order. However, such approaches for POMDPs are intractable for all but the smallest size problems.

In order to apply value iteration methods, we can write the Bellman equation for $Q^{\pi_0 \rightarrow \pi_1}$ for each agent i as follows:

$$Q^{\pi_0 \rightarrow \pi_1}(a_t|\tau_t^i) = \mathbb{E}_{\tau_t, \tau_{t+k}} \left[\sum_{t'=t}^{t+k-1} R(\tau_{t'}, a_{t'}) + \sum_{a_{t+k}} \pi_1(a_{t+k}|\tau_{t+k}^i) Q^{\pi_0 \rightarrow \pi_1}(a_{t+k}|\tau_{t+k}^i) \right] \quad (5)$$

where τ_{t+k} is the next history at which player i acts, $\tau_t \sim \mathcal{B}_{\pi_0}(\tau_t^i)$, and $\tau_{t+k} \sim (\mathcal{T}, \pi_1)$ denotes that τ_{t+k} is sampled from the distribution of successor states where all (other) players play according to π_1 .

Now, how can we perform the Bellman iteration in Eq. 5? One approach, which we will denote *Q-OBL*, is to perform independent Q -learning where each player uses π_0 as the exploration policy to play until time t and use π_1 afterwards. We can then train π_1 at time t but not other timesteps. This would guarantee that $\tau_t \sim \mathcal{B}_{\pi_0}(\tau_t^i)$. However in order to compute the value for playing π_1 in the future, we must simulate the transition $\tau_{t+k} \sim (\mathcal{T}, \pi_1)$ by having the other players playing according to their current π_1 . This is guaranteed to converge to the OBL policy (in the tabular case) as long as $\forall \tau^i \forall a \pi_0(a|\tau^i) > 0$, by the same inductive argument used in Theorems 1 and 2, since Q and π_1 at each AOH only depend on Q and π_1 at successor AOHs.

In practice, this approach has a downside: the states reached




Figure 1. Comparison of independent Q learning (left) with LB-OBL (right). Superscripts i and $-i$ mean player i and its partner $-i$. In LB-OBL, the target value for an action a_t^i is computed by first sampling a state from a belief model \hat{B}_{π_0} that assumes π_0 has been played, and then simulating play (blue) from that state to i 's next turn using the current policy π_1 .




Figure 2. Cross-play matrix for each of the different algorithms in the toy communication game. 10 agents were independently trained with each method; each cell (i, j) of the matrix represents the average score of the i th and j th agent. OBL with a depth of 1 (b) achieves a higher score than CH (e) by learning optimal grounded communication, while avoiding arbitrary handshakes that lead to the poor XP performance of SP (a) or deep counterfactual reasoning (d, g).

by π_1 may be reached with very low probability under π_0 ,¹ so this procedure will have low sample efficiency. Instead, we first learn an *approximate belief* \hat{B}_{π_0} that takes an AOH τ^i and can sample a trajectory from an approximation of $P(\tau|\tau^i, \pi_0)$. We compute this belief model following the procedure described in Hu et al. (2021). We then perform Q-learning with a modified rule for computing the target value for $Q(a|\tau_t^i)$: we re-sample a new τ' from $\hat{B}_{\pi_0}(\tau_t^i)$, simulate a transition to τ'_{t+1} with other agents play their policy π_1 , and use $\max_a Q(a|\tau'_{t+1})$ as the target value. We refer to this variant as *Learned-belief OBL* (LB-OBL).

An illustration comparing the LB-OBL algorithm with independent Q-learning (IQL) is shown in Figure 1. In both cases, we assume a two player setting with shared Q-function and 2-step Q-learning, i.e. $k = 2$ in Equation 5. The active player at time t is i and the partner is $-i$. In IQL, each player simply observes the world at each time-step and takes turns to act. The target is computed on the actual trajectory with r_t, r_{t+1} and $\max_a Q(a|\tau_{t+2})$ using target network Q^i . By contrast, LB-OBL involves only *fictitious transitions*. For the active player i , we first decide the action a_t^i given AOH τ_t^i . However, in addition

to applying a_t^i to the actual environment state S_t , we also apply a_t^i to a fictitious state S'_t sampled from the learned belief model B_i and forward the fictitious environment to S'_{t+1} . Note that the action a_t^i can be directly applied to the fictitious state because the observation from the fictitious state is the same as that of the real state. Next, we need to evaluate the partner's policy to produce a fictitious action a_{t+1}^{-i} on $\tau_{t+1}^{-i} = [\tau^{-i}; a_t^i, \Omega^{-i}(S'_{t+1})]$. The learning target in LB-OBL is the sum of *fictitious* rewards r'_t, r'_{t+1} and the *fictitious* bootstrapped value $\max_a Q(a|\tau'_{t+2})$. We also note that it is not possible to simply train a Q-learning agent that takes the *grounded beliefs* as an input. The agents would still develop conventions and the *grounded belief* would simply lose its semantic meaning.

5. Experiments in a Toy Environment

We first test OBL in a simple fully cooperative toy environment that includes both a cheap-talk channel and a grounded communication channel: As shown in Figure 3, Alice observes a random binary variable, $\text{pet} \in \{\text{cat}, \text{dog}\}$, and has four possible actions. She can either turn on the light bulb, bail out of the game with fixed reward of 1 or remove a barrier to let Bob see the pet, which will cost them 5. After

¹The bound is $P(\tau|\pi_0)/P(\tau|\pi_1) \geq (\min_a P(a))^t$




Figure 3. Toy cooperative communication game: Alice (left) observes the pet and can either signal to Bob by turning on the light-bulb, pay 5 to remove the barrier, so that Bob can see the pet, or bail out for a reward of 1. Bob then needs to guess the identity of the pet or can bail out for a reward of 0.5.

Alice takes her action, Bob observes the outcome and has three options. He can either bail for a fixed reward of 0.5 or guess the identity of the pet. If the guess is correct, they obtain an additional 10, otherwise they lose 10.

Results. We run three different algorithms on this simple toy game: SP, OBL and cognitive hierarchies (CH). Figure 2 shows the cross-play (XP) results for 10 independent runs for each of the methods as a matrix, where the entry i, j is the reward obtained from pairing the first agent (Alice) from run i with the second agent (Bob) from run j . As can be seen from the checkerboard-like pattern, the SP agents learn to use the cheap-talk channel, obtaining +10 in about half the cases and -10 in the others. Clearly, these agents have developed an arbitrary handshake where “on” / “off” correspond to dog / cat or vice versa, each incompatible with the other half of the agents that made the opposite choice. In contrast, OBL learns the optimal *grounded policy* at the first level, choosing to remove the barrier which results in a reward of +5, both in XP and SP. The CH agents at level 1 learn to simply bail out of the game, since there is no way to win when playing with a random agent. They also fail to discover the grounded solution in between and proceed straight to an arbitrary handshake for higher levels, as is indicated by the +10 / -10 rewards.

At higher levels, our OBL implementation starts to diverge between different runs. Although theoretically the OBL policy is unique, some environment states do not occur very often and the noise is amplified higher up. This could be addressed with higher temperature or increased regularization if desired. All results can be reproduced in our notebook implementation <https://bit.ly/3cIia6z>.

6. Hanabi Experiments

We now test our methods in the more complex domain of Hanabi. Hanabi is a fully-cooperative multiplayer partially-observable board game. It has become a popular benchmark environment (Bard et al., 2020) for MARL, theory of mind,

and zero-shot coordination research.

Hanabi is a 2-5 player card game. The deck is composed of 50 cards, split among five different colors (suits) and ranks, with each color having three 1s, two 2s and 3s and 4s, and one 5. In a 2-player game, each player maintains a 5-card hand. Players can see their partner’s hand but not their own. The goal of the team is to play one card of each rank in each color in order from 1 to 5. The team shares 8 hint tokens and 3 life tokens. Taking turns, each turn a player can play or discard a card in their hand, or spend a hint token to give a hint to their partner. Playing a card succeeds if it is the lowest-rank card in its color not yet played, otherwise it fails and loses a life token. Giving a hint consists of choosing a rank or a color that a partner’s hand contains and indicating *all* cards in the partner’s hand sharing that color or rank. Discarding a card or successfully playing a 5 regains one hint token. The team’s score is zero if all life tokens are lost, otherwise it is equal to the number of cards successfully played, giving a maximum possible score of 25.

Experimental Setup. In Hanabi, we implement LB-OBL using distributed deep recurrent Q -learning with prioritized experience replay (Kapturowski et al., 2019). Training consists of a large number of parallel environments that invoke a deep neural network to approximate the Q function at each time step to generate trajectories, a prioritized experience replay buffer to store the trajectories, and a training loop that updates the neural network using samples from the replay buffer. The queries from different environments are dynamically batched together to run efficiently on GPUs (Espeholt et al., 2020). Notably, we introduce several novel engineering designs to make our training infrastructure significantly more efficient than previous ones in Hanabi, allowing us to reach new state-of-the-art results in self-play. Please refer to appendix for more details on the implementation and benchmarks.

Our policy (Q -function) is parameterized by a deep recurrent neural network θ . As illustrated in Figure 1, we should sample trajectory τ from replay buffer and update the neural network with TD-error:

$$\mathcal{L}(\theta|\tau) = \frac{1}{2} \sum_{t=1}^T [r'_t + r'_{t+1} + \max_a Q_{\hat{\theta}}(a|\tau'_{t+2}) - Q_{\theta}(a_t|\tau_t)]^2 \quad (6)$$

where $Q_{\hat{\theta}}$ is the target network that uses a slightly outdated version of θ for stability reasons. In normal Q -learning with RNNs, the target $Q_{\hat{\theta}}(a_t|\tau_t)$ for $t = 1, 2, \dots, T$ can be re-computed efficiently by passing the sequence τ to the RNN at once. However, this is no longer feasible in OBL as each τ'_t contains unique fictitious transitions. To solve this problem, we pre-compute the target $G'_t = r'_t + r'_{t+1} + \max_a Q_{\hat{\theta}}(a|\tau'_{t+2})$ during rollouts and store the sequence of $\{G'_t\}$ along with the real trajectory τ into the replay buffer. Although this might cause stability issues if the stored target

Method	Self-Play	Cross-Play	w/ Other-Play (Rank Bot)	w/ Other-Play (Color Bot)	w/ Clone Bot
SAD(*)	23.97 \pm 0.04	2.52 \pm 0.34	7.29 \pm 1.40	0.18 \pm 0.06	0.83 \pm 0.31
Other-Play	24.14 \pm 0.03	21.77 \pm 0.68	22.81 \pm 0.87	4.05 \pm 0.37	8.55 \pm 0.48
K-Level	16.97 \pm 1.19	17.17 \pm 0.98	14.80 \pm 1.77	12.36 \pm 1.44	13.03 \pm 1.91
OBL (level 1)	20.92 \pm 0.07	20.85 \pm 0.03	10.83 \pm 0.26	13.21 \pm 0.34	13.80 \pm 0.19
OBL (level 2)	23.41 \pm 0.03	23.24 \pm 0.03	15.99 \pm 0.30	18.74 \pm 0.46	15.61 \pm 0.10
OBL (level 3)	23.93 \pm 0.01	23.68 \pm 0.05	15.61 \pm 0.34	20.68 \pm 0.44	16.02 \pm 0.26
OBL (level 4)	24.10 \pm 0.01	23.76 \pm 0.06	14.46 \pm 0.59	21.78 \pm 0.42	16.00 \pm 0.13

Table 1. Average scores in 2-player Hanabi for different pairings of agents. SP indicates play between an agent and itself. XP indicates play between agents from different independently-trained runs of the same algorithm. The other columns indicate play with agents from different algorithms. We evaluate each model on 5000 games and aggregate results of 5 independent training runs (seeds). (*) We use the 12 SAD agents obtained from (Hu et al., 2020) which uses different network architectures. Thus the numbers are not directly comparable.

was computed using a very old target network, in practice we find it to be fairly stable because the replay buffer gets refreshed quickly thanks to the fast trajectory generation speed of our training infrastructure. On average, a trajectory is used 4-5 times before it gets evicted from the buffer.

We set up the belief training following Hu et al. (2021). The belief model is trained to predict the player’s hand given τ^i , which is the only missing information needed to produce fictitious game states in Hanabi. It takes as input τ^i and predicts each of the player’s cards auto-regressively from the oldest to the newest. The belief model is an RNN trained via supervised learning:

$$\mathcal{L}(\mathbf{h}|\tau_t^i) = - \sum_{k=1}^n \log p(h_k|\tau_t^i, h_{1:k-1}), \quad (7)$$

where h_k is the k th card in hand and n is the hand size.

To compare the zero-shot coordination performance of OBL against previous methods, we train three categories of policies with existing methods to serve as the unseen partner for OBL models. “Rank-Bot” is trained with Other-Play. Empirically, Other-Play in Hanabi consistently converges to using rank-based conventions to communicate to a partner to play a card, hence the name. “Color Bot” is roughly the color equivalent of “Rank Bot”, induced by adding extra reward for hinting color at the early stage of training. “Clone Bot” is trained with supervised learning on human game data collected from an online board game platform. When we pair Clone Bot with the Other-Play Bot under the scoring system used by Hu et al. (2020), it achieves an average score of 14.98 ± 0.34 , similar to the average score of 15.75 that the authors reported for an experiment of humans partnering with Other-Play. This suggests that Clone Bot is a reasonable proxy for human play.

Similar to the toy example, we implemented a variant of Cognitive Hierarchies (CH) called k-level reasoning as one of our baselines. The details of neural network design,




Figure 4. Grounded knowledge of cards played. OBL-1 mostly plays cards of known rank and color. Iterated OBL converges to an even balance of color and rank, similar to humans.

hyper-parameters and computation cost of OBL policies and belief models, as well as those of Rank, Color, Clone Bot and K-Level can be found in the appendix. We will open source our code and all models.

For all the methods and OBL levels shown in Table 1, we carry out 5 independent runs with different seeds, except for the 12 SAD models directly downloaded from the open-sourced repository of Hu et al. (2020).

Grounded Policies in Hanabi. In this experiment, we empirically show that OBL can indeed learn surprisingly strong grounded policies in Hanabi.

Although Theorem 4 states that applying OBL on any constant policy π_0 should yield an optimal grounded policy π_1 , in practice we want the trajectories generated by π_0 to be as diverse as possible to help train a good belief model. Hence a simple choice is to make π_0 a uniform-random policy. We first train a belief model B_{π_0} till convergence and then use B_{π_0} to train the OBL policy. We repeat this training procedure with five different seeds. The numerical results are shown in the *OBL (level 1)* row of Table 1. First of all, we learn that a policy with no conventions can achieve nearly 21 points in SP, which by itself is an important milestone

for this popular benchmark environment. The XP score, which is computed by pairing agents with different seeds, is almost as high as the SP score, indicating that they indeed converge to nearly identical policies. The slight gap between XP and SP scores may come from the noise of belief and policy training. It worth noting that despite being a grounded policy, OBL (level 1) already cooperates with human-like policies in Hanabi (proxied by CloneBot) better than any prior method.

One way to quantitatively show whether a policy is grounded in Hanabi is to look at how much the active player knows about a card before playing it. Playing a card unsuccessfully is costly, and therefore playing it without high confidence that it will succeed is usually suboptimal. Cards played can be categorized into four groups. 1) Both color and rank of the card are known. 2) Only color is known. 3) Only rank is known. 4) Neither is known. Here the known information can be revealed by hints or deduced from card counts. The result is shown in the *OBL LI* column on Figure 4. More than 60% of the cards played by this policy are completely known to the player, significantly higher than for other agents, which may infer playability based on counterfactuals rather than grounded knowledge. Although roughly 28% of the time the policy plays cards only knowing the number, inspection showed that in those cases the cards are safe to play regardless of color. These observations are consistent with the theoretical prediction that OBL should converge to the optimal grounded policy in Hanabi.

OBL Hierarchy in Hanabi. As described in Section 4.4, we can repeat the OBL training process to get OBL level $i + 1$ using level i as the new π_0 . In practice, we can speed up the training of next level belief and policy by loading weights from the previous level as a start point.

We train three additional levels and the evaluation results are shown in Table 1. OBL level 2 infers beliefs over unseen cards assuming other agents are more likely to take good actions and to share factual information about useful cards the way OBL level 1 would, and so on for higher levels. Since other agents, even if they have unknown conventions, will also often act this way, these beliefs are sufficiently robust that OBL level 2 achieves even better scores than OBL level 1 with Clone Bot and with all other tested agents.

For comparison, we show results of Other-Play applied on the same network architecture as OBL. The Other-Play agents are also trained with the auxiliary task from the original paper. We see from the table that both SP and XP of OBL gradually increase after each level. The final SP performance is similar to that of Other-Play while the XP score is significantly higher, showing that OBL is a better method for zero-shot coordination. OBL also performs significantly better than Other-Play when playing with a diverse set of unseen partners. Note that Other-Play gets

high scores with Rank Bot because Rank Bot is randomly chosen from the five Other-Play training runs. Based on its performance with Clone Bot, it is probable that OBL would achieve much better performance than Other-Play when paired with real humans. From Figure 4, we can see that as we progress through OBL levels, the agents play fewer grounded cards and start infer playability based on counterfactuals. Unlike SP where conventions are formed implicitly as each agent adapts to the random biases of its partner, OBL exhibits a consistent path of evolution that can be understood by reasoning through the behaviors of the lower level OBL policy. The final level of OBL is the most similar to Clone Bot in how it decides to play cards.

OBL in 3 Player Hanabi. We also apply OBL to 3 player Hanabi and evaluate its performance in XP and playing with Clone Bot. All OBL levels perform better than Other-Play in these two criteria. In XP, OBL (level 4) gets 23.02, which is significantly higher than the 17.36 of Other-Play. When paired with Clone Bot, Other-Play gets 12.18 while the highest score of OBL is 14.09 from level 2. The Clone Bot itself gets 17.46 in self-play. The detailed results are in the Appendix due to space limit.

7. Conclusion and Future Work

We present *off-belief learning*, a new method that can train optimal *grounded* policies, preventing agents from exchanging information through *counterfactual reasoning*. When used in a hierarchy, each level adds one step of counterfactual reasoning, thus providing a controlled means of introducing conventions and pragmatic reasoning into the learning process. Crucially, OBL removes the ‘weirdness’ of learning in Dec-POMDPs, since each level has a unique optimal policy, given the beliefs induced by the level below. As a consequence, OBL at convergence can solve instances of the ZSC problem.

There are many avenues for future work. For example, we may extend OBL to two-player zero-sum and general-sum settings by training *average beliefs* and sampling the fictitious partner move randomly from levels below. Secondly, OBL cannot adapt to its partner at test time. Even in zero-shot, if each episode consists of many turns, agents could potentially adapt to each other’s policies on the fly, or dynamically form protocols to leverage cheap-talk channels within the episode. OBL is currently incapable of learning cheap-talk at test time.

8. Acknowledgement

We would like to thank Andrei Lupu, Samuel Sokota, Michael Dennis and Wendelin Boehmer for helpful discussions and feedback on the ideas of OBL and the manuscript. We appreciate Broad Game Arena for providing us with

anonymized Hanabi game play data for this research.

References

- Agranov, M., Potamites, E., Schotter, A., and Tergiman, C. Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior*, 75(2):449 – 463, 2012. ISSN 0899-8256. doi: <https://doi.org/10.1016/j.geb.2012.02.002>. URL <http://www.sciencedirect.com/science/article/pii/S089982561200022X>.
- Bard, N., Foerster, J. N., Chandar, S., Burch, N., Lanctot, M., Song, H. F., Parisotto, E., Dumoulin, V., Moitra, S., Hughes, E., et al. The hanabi challenge: A new frontier for ai research. *Artificial Intelligence*, 280:103216, 2020.
- Camerer, C., Ho, T., and Chong, J.-K. A cognitive hierarchy theory of one-shot games and experimental analysis. 08 2003. URL <https://ssrn.com/abstract=411061>.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.
- Carroll, M., Shah, R., Ho, M. K., Griffiths, T., Seshia, S. A., Abbeel, P., and Dragan, A. D. On the utility of learning about humans for human-ai coordination. In Wallach, H. M., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E. B., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pp. 5175–5186, 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/f5b1b89d98b7286673128a5fb112cb9a-Abstract.html>.
- de Witt, C., Foerster, J., Farquhar, G., Torr, P., Boehmer, W., and Whiteson, S. Multi-Agent Common Knowledge Reinforcement Learning. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32, pp. 9927–9939. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/f968fdc88852a4a3a27a81fe3f57bfc5-Paper.pdf>.
- Espeholt, L., Marinier, R., Stanczyk, P., Wang, K., and Michalski, M. Seed rl: Scalable and efficient deep-rl with accelerated central inference. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgvXlrKwH>.
- Foerster, J., Song, F., Hughes, E., Burch, N., Dunning, I., Whiteson, S., Botvinick, M., and Bowling, M. Bayesian action decoder for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pp. 1942–1951. PMLR, 2019.
- Frank, M. C. and Goodman, N. D. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998, 2012.
- Hendon, E., Jacobsen, H. J., and Sloth, B. The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior*, 12(2):274–282, 1996. ISSN 0899-8256. doi: <https://doi.org/10.1006/game.1996.0018>. URL <https://www.sciencedirect.com/science/article/pii/S0899825696900184>.
- Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., van Hasselt, H., and Silver, D. Distributed prioritized experience replay. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL <https://openreview.net/forum?id=H1Dy---0Z>.
- Hu, H. and Foerster, J. N. Simplified action decoder for deep multi-agent reinforcement learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=B1xm3RVtwB>.
- Hu, H., Peysakhovich, A., Lerer, A., and Foerster, J. “other-play” for zero-shot coordination. In *Proceedings of Machine Learning and Systems 2020*, pp. 9396–9407. 2020.
- Hu, H., Lerer, A., Brown, N., and Foerster, J. N. Learned belief search: Efficiently improving policies in partially observable settings, 2021. URL https://openreview.net/forum?id=xP37gkVKa_0.
- Kapturowski, S., Ostrovski, G., Quan, J., Munos, R., and Dabney, W. Recurrent experience replay in distributed reinforcement learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=r1lyTjAqYX>.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.

- Kleiman-Weiner, M., Ho, M. K., Austerweil, J. L., Littman, M. L., and Tenenbaum, J. B. Coordinate to cooperate or compete: abstract goals and joint intentions in social interaction. In *CogSci*, 2016.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529–533, 02 2015. URL <http://dx.doi.org/10.1038/nature14236>.
- Rashid, T., Samvelyan, M., de Witt, C. S., Farquhar, G., Foerster, J. N., and Whiteson, S. QMIX: monotonic value function factorisation for deep multi-agent reinforcement learning. In Dy, J. G. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4292–4301. PMLR, 2018. URL <http://proceedings.mlr.press/v80/rashid18a.html>.
- Schaul, T., Quan, J., Antonoglou, I., and Silver, D. Prioritized experience replay. In Bengio, Y. and LeCun, Y. (eds.), *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1511.05952>.
- Schelling, T. C. *The strategy of conflict*. Harvard university press, 1980.
- Sunehag, P., Lever, G., Gruslys, A., Czarnecki, W. M., Zambaldi, V. F., Jaderberg, M., Lanctot, M., Sonnerat, N., Leibo, J. Z., Tuyls, K., and Graepel, T. Value-decomposition networks for cooperative multi-agent learning. *CoRR*, abs/1706.05296, 2017. URL <http://arxiv.org/abs/1706.05296>.
- Tan, M. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pp. 330–337, 1993.
- van Hasselt, H., Guez, A., and Silver, D. Deep reinforcement learning with double q-learning. In Schuurmans, D. and Wellman, M. P. (eds.), *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pp. 2094–2100. AAAI Press, 2016. URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12389>.
- Wang, Z., Schaul, T., Hessel, M., van Hasselt, H., Lanctot, M., and de Freitas, N. Dueling network architectures for deep reinforcement learning. In Balcan, M. and Weinberger, K. Q. (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 1995–2003. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/wangf16.html>.

A. Proofs of Theorems

We first prove two lemmas that will be necessary for the proofs of the main theorems.

Lemma 1. *For any policies (π_0, π_1, π) ,*

$$\begin{aligned} \sum_{\tau_t \in T_t} P(\tau_t | \pi_0) \sum_{a_t} \pi(a_t | \tau_t^i) \left(R(s_t, a_t) + \sum_{s_{t+1}} \mathcal{T}(s_{t+1} | s_t, a_t) V^{\pi_1}(\tau_{t+1}) \right) \\ = \sum_{\tau_t^i \in T_t^i} P(\tau_t^i | \pi_0) \sum_{a_t} \pi(a_t | \tau_t^i) Q^{\pi_0 \rightarrow \pi_1}(a_t | \tau_t^i) \end{aligned} \quad (8)$$

This Lemma shows an equivalence between an expected value integrated over all *trajectories* of length t , and expected value integrated over all *AOHs* of length t .

Proof.

$$\sum_{\tau_t \in T_t} P(\tau_t | \pi_0) \sum_{a_t} \pi(a_t | \tau_t^i) \left(R(s_t, a_t) + \sum_{s_{t+1}} \mathcal{T}(s_{t+1} | s_t, a_t) V^{\pi_1}(\tau_{t+1}) \right) \quad (9)$$

$$= \sum_{\tau_t^i \in T_t^i} P(\tau_t^i | \pi_0) \sum_{\tau_t \in \tau_t^i} P(\tau_t | \tau_t^i) \sum_{a_t} \pi(a_t | \tau_t^i) \left(R(s_t, a_t) + \sum_{s_{t+1}} \mathcal{T}(\tau_{t+1} | \tau_t, a_t) V^{\pi_1}(\tau_{t+1}) \right) \quad (10)$$

$$= \sum_{\tau_t^i \in T_t^i} P(\tau_t^i | \pi_0) \sum_{a_t} \pi(a_t | \tau_t^i) \sum_{\tau_t \in \tau_t^i} P(\tau_t | \tau_t^i) \left(R(s_t, a_t) + \sum_{s_{t+1}} \mathcal{T}(s_{t+1} | s_t, a_t) V^{\pi_1}(\tau_{t+1}) \right) \quad (11)$$

$$= \sum_{\tau_t^i \in T_t^i} P(\tau_t^i | \pi_0) \sum_{a_t} \pi(a_t | \tau_t^i) Q^{\pi_0 \rightarrow \pi_1}(a_t | \tau_t^i) \quad (12)$$

□

Lemma 2. *The softmax policy with temperature T is worse than the optimal policy by at most T/e . Formally, for any $x_1 \dots x_N \in \mathbb{R}^N$,*

$$\frac{\sum_{i=1}^N \exp(x_i/T) x_i}{\sum_{j=1}^N \exp(x_j/T)} \geq \max_i x_i - T/e \quad (13)$$

Proof. Let $x_1 \geq x_2 \geq \dots \geq x_N$ w.l.o.g. So $\max_i x_i = x_1$. Let $y_i = x_i - x_1$.

$$\frac{\sum_{i=1}^N \exp(x_i/T) x_i}{\sum_{j=1}^N \exp(x_j/T)} = \frac{\sum_{i=1}^N \exp(x_1/T) \exp(y_i/T) (y_i + x_1)}{\sum_{j=1}^N \exp(x_1/T) \exp(y_j/T)} \quad (14)$$

$$= x_1 + \frac{\sum_{i=1}^N \exp(y_i/T) y_i}{\sum_{j=1}^N \exp(y_j/T)} \quad (15)$$

$$\geq x_1 + \frac{\sum_{i=1}^N \exp(y_i/T) y_i}{N} \quad \text{since } y \leq 0 \quad (16)$$

$$\geq x_1 + \min_{z \leq 0} \exp(z/T) z \quad (17)$$

To compute the minimum of $\exp(z/T) z$ in $(-\infty, 0]$,

$$f(z) = \exp(z/T)z \quad (18)$$

$$\frac{df}{dz} = \exp(z/T)(1 + z/T) = 0 \quad (19)$$

$$z = -T \quad (20)$$

$$f(z) = -\exp(-1)T = -T/e \quad (21)$$

The value of $f(z)$ at both endpoints is 0. Therefore, $\min_{z \leq 0} \exp(z/T)z = -T/e$, which when substituted into Equation 17 proves the theorem. \square

Theorem 1. For any $T > 0$ and starting policy π_0 , OBL computes a unique policy π_1 .

Proof. Since AOHs cannot repeat (a successor AOH is always longer than its predecessor) and the game has bounded length, the AOHs of acting players form a DAG with edges from each AOH to all possible successor AOHs for the next acting player. Ordering AOHs topologically s.t. successor AOHs precede their predecessors, we prove by induction.

In the base case, the first AOH in the topological ordering is a terminal AOH τ_ϵ , so $Q^{\pi_0 \rightarrow \pi_1}(\cdot | \tau_\epsilon) = 0$.

For the inductive case, we must show that if $Q^{\pi_0 \rightarrow \pi_1}(a | \tau_{t+1}^j)$ is unique for every τ_{t+1}^j that is a successor of τ_t^i , then $Q^{\pi_0 \rightarrow \pi_1}(a | \tau_t^i)$ is unique (where j is the player to act at the successor trajectory).

$\pi_1(a | \tau_t^i)$ is a function of $Q^{\pi_0 \rightarrow \pi_1}(\cdot | \tau_t^i)$.

$$Q^{\pi_0 \rightarrow \pi_1}(a | \tau_t^i) = \sum_{\tau_t, \tau_{t+1}} R(s_t, a) + \mathcal{B}_{\pi_0}(\tau_t | \tau_t^i) \mathcal{T}(\tau_{t+1} | \tau_t) V^{\pi_0 \rightarrow \pi_1}(\tau_{t+1}^j) \quad (22)$$

The RHS of Eq. 22 only depends on π_1 at successor AOHs, which are uniquely defined. \square

Theorem 2. For every policy π_1 generated by OBL from π_0 , $J(\pi_1) \geq J(\pi_0) - t_{max}T/e$, i.e. OBL is a policy improvement operator except for a term that vanishes as $T \rightarrow 0$.

Proof. Let $J(\pi_0 \xrightarrow{t} \pi_1)$ be the expected return of playing π_0 for the first $t - 1$ timesteps and π_1 subsequently.

We prove by induction backwards in t that

$$J(\pi_0 \xrightarrow{t} \pi_1) \geq J(\pi_0) - T(t_{max} - t)/e \quad (23)$$

for all t , which implies the theorem for $t = 0$.

The base case of $t = t_{max}$ is trivially true because $(\pi_0 \xrightarrow{t_{max}} \pi_1) \equiv \pi_0$.

For the inductive case, suppose that Eq. (23) holds for all $t > t'$. Let T_t and T_t^i be the set of all trajectories and AOHs of length t , respectively.

$$J(\pi_0) - e(t_{max} - (t' + 1))T \leq J(\pi_0 \xrightarrow{t'+1} \pi_1) \quad (24)$$

$$J(\pi_0 \xrightarrow{t'+1} \pi_1) = \sum_{\tau_{t'+1} \in T_{t'+1}} P(\tau_{t'+1} | \pi_0) V^{\pi_1}(\tau_{t'+1}) \quad (25)$$

$$= \sum_{\tau_{t'} \in T_{t'}} P(\tau_{t'} | \pi_0) \sum_{a_{t'}} \pi_0(a_{t'} | \tau_{t'}^i) \left(R(s_{t'}, a_{t'}) + \sum_{s_{t'+1}} \mathcal{T}(s_{t'+1} | s_{t'}, a_{t'}) V^{\pi_1}(\tau_{t'+1}) \right) \quad (26)$$

$$= \sum_{\tau_{t'}^i \in T_{t'}^i} P(\tau_{t'}^i | \pi_0) \sum_{a_{t'}} \pi_0(a_{t'} | \tau_{t'}^i) Q'_{\pi_0}(a_{t'} | \tau_{t'}^i) \quad (\text{Lemma 1}) \quad (27)$$

$$\leq \sum_{\tau_{t'}^i \in T_{t'}^i} P(\tau_{t'}^i | \pi_0) \max_{a_{t'}} Q'_{\pi_0}(a_{t'} | \tau_{t'}^i) \quad (28)$$

$$\leq \sum_{\tau_{t'}^i \in T_{t'}^i} P(\tau_{t'}^i | \pi_0) \sum_{a_{t'}} \pi_1(a_{t'} | \tau_{t'}^i) Q'_{\pi_0}(a_{t'} | \tau_{t'}^i) + T/e \quad (\text{Lemma 2}) \quad (29)$$

$$= \sum_{\tau_{t'} \in T_{t'}} P(\tau_{t'} | \pi_0) \sum_{a_{t'}} \pi_1(a_{t'} | \tau_{t'}^i) \left(R(s_{t'}, a_{t'}) + \sum_{s_{t'+1}} \mathcal{T}(s_{t'+1} | \tau_{t'}, a_{t'}) V^{\pi_1}(\tau_{t'+1}) \right) + T/e \quad (\text{Lemma 1}) \quad (30)$$

$$= \sum_{\tau_{t'} \in T_{t'}} P(\tau_{t'} | \pi_0) V^{\pi_1}(\tau_{t'}) + T/e \quad (31)$$

$$= \mathbb{E}_{T_{t'}} [V^{\pi_1}(\tau_{t'}) | \pi_0] + T/e \quad (32)$$

$$= J(\pi_0 \xrightarrow{t'} \pi_1) + T/e \quad (33)$$

Lemmas 1 and 2 can be found above. \square

Notably, unlike standard MARL, the fixed points of the OBL learning rule *are not* guaranteed to be equilibria of the game.

Theorem 3. *If repeated application of the OBL policy improvement operator converges to a fixed point policy π , then π is an ϵ -subgame perfect equilibrium of the Dec-POMDP, where $\epsilon = t_{max}T/e$.*

Proof. We'll start with a proof sketch for a "temperature 0" policy $\pi_1(\tau^i) = \arg\max_a Q^{\pi_0 \rightarrow \pi_1}(a | \tau^i)$ and then deal with the softmax policy.

Suppose π is a fixed point of the temperature-0 OBL operator. Then at every AOH τ^i ,

$$\pi(\tau^i) = \arg\max_a Q^{\pi \rightarrow \pi}(a | \tau^i) \quad (34)$$

$$= \arg\max_a Q^{\pi}(a | \tau^i) \quad (35)$$

By the one-shot deviation principle (Hendon et al., 1996), π must be a subgame-perfect equilibrium of G .

Now, we will need to reiterate the proof of the one-shot deviation principle in order to modify it for softmax policies.

Suppose that π is a fixed point of the OBL operator at temperature T . We will show that for any π'_i ,

$$V^{\pi}(\tau^i) \geq V^{\pi'_i, \pi-i}(\tau^i) - T/e(|\tau^i| - t_{max}). \quad (36)$$

which for $\tau^i = \emptyset$ reduces to $J(\pi) \geq J(\pi'_i, \pi_{-i}) - t_{max}T/e$, proving the theorem.

We prove by induction over i 's AOHs, treating successor AOHs before predecessors as before.

The base case at terminal AOHs holds trivially. Now suppose that (36) is true for all successors of τ^i . We know from Lemma 2 that

$$\sum_a \pi(\tau^i) Q^{\pi \rightarrow \pi}(a | \tau^i) \geq \max_a Q^{\pi \rightarrow \pi}(a | \tau^i) - T/e \quad (37)$$

$$V^{\pi}(\tau^i) = \sum_a \pi(\tau^i) Q^{\pi}(a | \tau^i) \geq \max_a Q^{\pi}(a | \tau^i) - T/e \quad (38)$$

Let $P(\tau^{i'}|\tau^i, a, \pi_{-i})$ be the probability that player i 's next AOH is $\tau^{i'}$ after playing a at τ^i , and assuming other players play π_{-i} .² Then expanding out Q^π in (38) leads to

$$V^\pi(\tau^i) \geq \max_a \sum_{\tau^{i'}} P(\tau^{i'}|\tau^i, a, \pi_{-i}) V^\pi(\tau^{i'}) - T/e \quad (39)$$

$$\geq \max_a \sum_{\tau^{i'}} P(\tau^{i'}|\tau^i, a, \pi_{-i}) \left(V^{\pi_{i'}, \pi_{-i}}(\tau^{i'}) - T/e - T/e(|\tau^{i'}| - t_{max}) \right) \quad (40)$$

$$\geq \max_a \sum_{\tau^{i'}} P(\tau^{i'}|\tau^i, a, \pi_{-i}) \left(V^{\pi_{i'}, \pi_{-i}}(a|\tau^{i'}) - T/e(|\tau^i| - t_{max}) \right) \quad (41)$$

$$\geq V^{\pi_{i'}, \pi_{-i}}(\tau^i) - T/e(|\tau^i| - t_{max}) \quad (42)$$

Eq. (39) just expands out the expectation for Q^π over all possible next AOHs that player i may reach given the other players' policies π_{-i} . □

Theorem 4. *Application of OBL to any constant policy $\pi_0(a|\tau^i) = f(a)$ – or in fact any policy that only conditions on public state – yields an optimal grounded policy in the limit as $T \rightarrow 0$.*

Proof. We need only show that any policy π_0 that conditions only on public state yields a state distribution at each AOH that matches the grounded beliefs \mathcal{B}_G ; then the optimal grounded policy follows directly from the definition of OBL (Eq. 4)

The true state distribution induced by π_0 at τ^i are

$$P(\tau|\tau^i, \pi_0) = \frac{P(\tau) \prod_t P(o_t^i|\tau) \pi_0(a_t|\tau_t^{-i})}{\sum_{\tau'} P(\tau') \prod_t P(o_t^i|\tau') \pi_0(a_t|\tau_t'^{-i})} \quad (43)$$

If π_0 is constant, i.e. $\pi_0(a_t|\tau^i) = f(a_t)$, then π_0 in the numerator and denominator immediately cancel, yielding the grounded beliefs

$$P(\tau|\tau^i) = \frac{P(\tau) \prod_t P(o_t^i|\tau)}{\sum_{\tau'} P(\tau') \prod_t P(o_t^i|\tau')}. \quad (44)$$

If π only depends on the public state, then the numerator and denominator still cancel for all trajectories τ which contribute to the sum, since any τ that doesn't share a common public state with τ' will lead to different observations, by definition. □

B. Experimental Details for Hanabi

B.1. Reinforcement Learning

We use a highly scalable setup as illustrated in Figure 5 to efficiently train RL models in Hanabi using moderate computational resources (three GPUs). It includes four major components. The first one is a large number of parallel thread workers that handle the interactions between environments and the multiple players in each environment. The player often needs to invoke one or more neural network inferences at each step to compute an action, which is often computationally intensive and thus should be executed on GPUs. Here, we take an asynchronous approach that instead of waiting for the neural network calls to return, the thread worker immediately moves on to execute the next set of environment and players after sending the neural network request to the inference loop. This allows us to run multiple environments in a single thread worker. On top of this, we also run multiple thread workers in parallel. The combination of these two techniques allows us to run a massive amount of environments and therefore generates a large amount of neural network inference requests simultaneously. These requests are then batched together before sending to an everlasting inference loop. The batched inference is executed on GPUs and the inference loop may use multiple GPUs to distribute workloads. Each player collects observation, action and reward at each step and aggregates them into a trajectory at the end of an episode. The trajectory

²Crucially, this transition probability is not dependent on **player i 's** policy to reach $\tau^{i'}$, because AOH τ^i already specifies all of player i 's actions to reach $\tau^{i'}$.

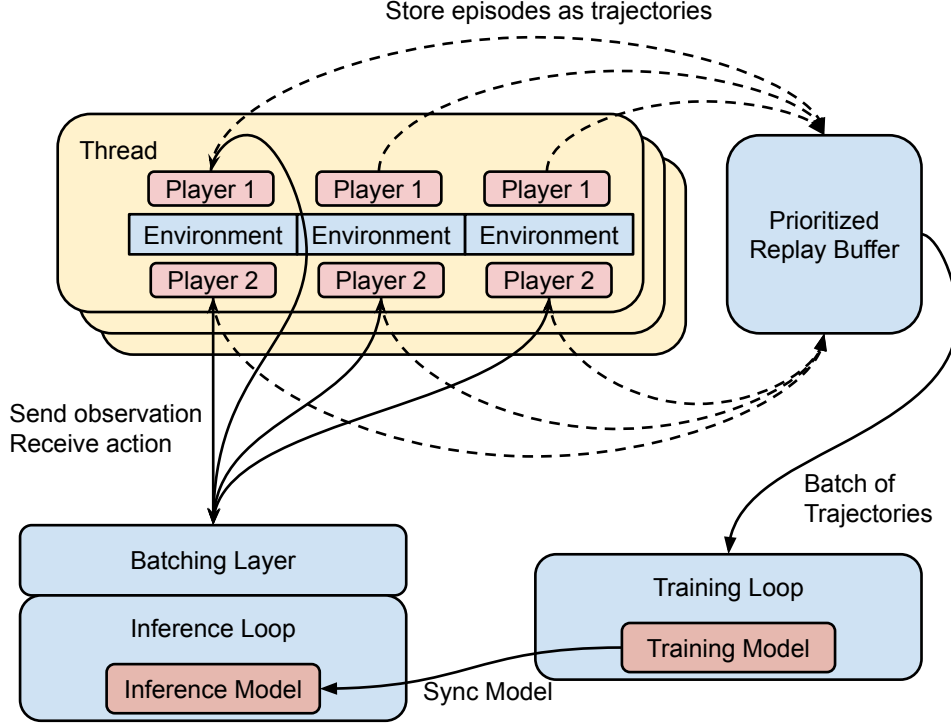


Figure 5. Illustration of RL training setup. Some arrows linking *player 1* and *batching layer* are omitted for legibility.

is padded to a fixed length of 80 time-steps and then stored into the prioritized replay buffer. A training loop, which runs completely in parallel to all the procedures mentioned above, keeps sampling batches from the replay buffer and updates the model with TD-error. The training loop sends a new copy of the model to the inference loop to have the inference model synced every 10 gradient steps. We also follow best Q-learning practices such as dueling architecture (Wang et al., 2016), double DQN (van Hasselt et al., 2016) and prioritized experience replay (Schaul et al., 2016) with priority of the trajectory computed the same way as (Kapturowski et al., 2019). In all of our RL experiments, i.e. both training OBL agents and reproducing Other-Play agents, we run 80 thread workers and 80 environments in each thread worker on 40 CPU cores. We use 2 GPUs for the inference loop and 1 GPU for the training loop. For OBL, we use 1 or 2 additional GPUs for belief model inference. We list out important hyper-parameters in Table 2.

Notably, our implementation runs much faster in terms of both data generation due to a more efficient parallelization and thus can train on more data given a fixed, reasonable amount of time. After training for 40 hours, we are able to significantly outperform the state-of-the-art self-play performance in Hanabi using existing methods. As shown in Table 3, a simple VDN baseline is sufficient to achieve a remarkable 24.27 and 24.32 on average for 2-player and 3-player Hanabi, reducing the distance to perfect (25 points) by 26% and 47% respectively. We also find that with abundant data and fast data generation enabled by our infrastructure, the VDN baseline is strong enough that additional methods such as simplified action decoder (SAD) or auxiliary task (AUX) is no longer helpful for further improving self-play scores. However, we note that AUX still improves cross-play scores significantly.

B.2. Belief Learning

Instead of creating a training/validation dataset, we use a similar setting as the reinforcement learning training, where mini-batches are sampled from an ever-changing replay buffer populated by thread workers. There are four main differences. First, the inference loop uses a pretrained fixed policy instead of syncing with a training policy periodically. Second, we store the true hand of the player alongside the trajectories which will be used as training targets for our belief model. Third, a normal experience replay buffer without priority is used. Finally, the training loop trains an auto-regressive model that predicts cards in hand one-by-one from the oldest to the latest with supervised learning. The network architecture is the same as the one used in (Hu et al., 2021).

Hyper-parameters	Value
# replay buffer related	
burn_in_frames	10,000
replay_buffer_size	100,000
priority_exponent	0.9
priority_weight	0.6
max_trajectory_length	80
# optimization	
optimizer	Adam (Kingma & Ba, 2015)
lr	6.25e-05
eps	1.5e-05
grad_clip	5
batchsize	128
# Q learning	
n_step	1 (OBL), 3 (non-OBL)
discount_factor	0.999
target_network_sync_interval	2500
exploration ϵ	$\epsilon_0 \dots \epsilon_n$, where $\epsilon_i = 0.1^{1+7i/(n-1)}$, $n = 80$

Table 2. Hyper-Parameters for Reinforcement Learning

Agent	2 Players		3 Players	
	Mean	Max	Mean	Max
VDN (Hu & Foerster, 2020)	23.83 \pm 0.03	23.96	23.71 \pm 0.06	23.99
SAD (Hu & Foerster, 2020)	23.87 \pm 0.03	24.01	23.69 \pm 0.05	23.93
SAD+AUX (Hu & Foerster, 2020)	24.02 \pm 0.01	24.08	23.56 \pm 0.07	23.81
VDN (ours)	24.27 \pm 0.01	24.33	24.32 \pm 0.02	24.39
SAD (ours)	24.26 \pm 0.01	24.33	24.24 \pm 0.03	24.42
VDN+AUX (ours)	24.26 \pm 0.01	24.32	24.08 \pm 0.01	24.17

Table 3. Self-play of various methods rerun using our infrastructure, comparing with previous published best results. For each method, we run 10 independent training with different seeds and shown the mean and max of the final model evaluated on 10K games. VDN is value-decomposition network that construct the joint Q-value with the sum of Q-values of each player. SAD is the simplified action decoder that includes additional representations for the greedy action beside the actual action that chosen by the agent via ϵ -greedy or sampling from a softmax function. AUX means the addition of an auxiliary task that predict whether each card is playable/discardable/unknown.

B.3. Other Models

To test the performance of OBL when paired with unseen partners (*ad-hoc* teamplay), we train three types of policies with existing methods.

Rank Bot. The first one is trained with the Other-Play (OP) (Hu et al., 2020) technique. The majority of policies trained with OP use a rank-based convention where they predominantly hint for ranks to indicate a playable card. For example, in a case where “Red 1” and “Red 2” have been played and the partner just draw a new “Red 3”, the other agent will hint 3 and then partner will play that card deeming that 3 being a red card based on the convention and counterfactual reasoning. We refer to this category of agents as “Rank Bot”.

Color Bot. Equivalently, one can also expect a color-based policy that will hint red for that latest card instead. In fact, such color-based policy was also appeared in the original Other-Play paper as the worst partner of the reset of the Other-Play policies. However, they are generally hard to reproduce as only 1 out of the 12 their training runs ended up with such policy. We use a simple reward shaping technique to produce similar policies reliably where we give each “hint color” move an extra reward during the first half of the training process and then disable the extra reward in the second half to wash out any artifacts. However, the reward shaping introduces undesired side effects that lead to arbitrary conventions. In practice,

	2 Players	3 Players
w/ Other-Play	19.93 \pm 0.09	13.07 \pm 0.15
w/o Other-Play	21.01 \pm 0.07	17.46 \pm 0.12

Table 4. Results of Clone Bot. Here Other-Player refer to the color shuffling technique first proposed in the Other-Player paper. We adapt the same technique here but use it for data-augmentation purpose. The trained models are evaluated on 5000 games by greedily selecting the action of the highest probability at test time. Errors shown are standard error of mean (s.e.m.)

we find that hiding the *last action* field of the input observation will make training outcomes more consistent. The reward shaping and feature engineering techniques are only used for producing this specific policy, which we refer to as “Color Bot”.

Clone Bot. One of the goals of the Hanabi challenge (Bard et al., 2020) is to develop artificial agents that can collaborate with humans. To understand how well our models can collaborate with humans without resorting to costly experiments, we train a behavior clone bot mimicking human behaviors to serve as a proxy. We acquire 240,954 2-player games and 113,900 3-player games from the online game platform “Board Game Arena”³ and convert it to a dataset for supervised learning. The model takes in the trajectory of an entire game from the perspective of a single player and predicts its action at each time-step. This is similar to the independent Q-learning setting in the multi-agent RL context. The network consists of one fully layer, followed by ReLU activation, a two-layer LSTM and an output fully connected layer with softmax. Dropout is applied before the output layer to reduce overfitting. The best models use 512 hidden units for each layer in the network and 0.5 dropout rate. We adapt the color shuffling technique of Other-Play (Hu et al., 2020) as a data augmentation tool. At each training step, we randomly shuffle the color space for both observation and action of each trajectory in the mini-batch before feeding them to the network. We find this technique significantly improves the performance of the supervised model, especially for 3-player Hanabi where we have much less data. The results are shown in Table 4. The agent is trained by minimizing the cross-entropy loss. At test time, the agent selects the action with the highest probability at each step. We pick the model that achieves the highest self-play score during training as our final “Clone Bot”.

C. Qualitative Analysis of Learned Hanabi Agents

In this section we share more insights about the policies learned by different methods through a series of qualitative analysis.

As we can see from Figure 4 the OBL level 1 agent plays a significantly higher percentage of grounded cards, *i.e.* cards of which both color and rank are known. This supports our claim that OBL learns a grounded policy given a constant π_0 . However, it is worth mentioning that OBL level 1 still plays a fair number of cards when it only knows their rank. At first this may seem contradictory to the grounded policy claim. However, as we analyze the games, we find that in many cases the best grounded policy is to play those cards even if the color is not known. For example, at an early stage of the game when only a few cards have been played, knowing a card is 1 is sufficient to play it. Moreover, the agent often knows that the card is of a “safe color”. For example, when both red 2 and green 2 have been played, it will be safe to play a card if we know that it is either a red 3 or a green 3. Playing such a card will be classified as “only rank” in the figure. Furthermore, there are multiple lives in Hanabi and sometimes it is worth taking a risk if the odds are good. We observe that the OBL level 1 agent sometimes plays a card blindly without knowing any information about the card, as reflected in the figure by the slightly higher percentage of “none” category than other OBL levels and other agents except for SAD. This is because Hanabi is not designed to be a game that can be mastered by a grounded policy, and therefore the number of the remaining hint tokens are often low due to the aggressive hinting strategy used by the policy. In this case, OBL level 1 resorts to maximizing the utility of the life tokens by playing one or two cards blindly in later stages of the game when the chances that a newly drawn card could be useful is high and remaining life tokens are abundant. Once there is only one life token left, the agent becomes conservative and stops the gambling behavior.

The conventions in the high levels of OBL agents stem from the fact that OBL level 1 agents give hints about a card in certain orders based on the situations. In some cases the OBL level 1 agent first hints at the color of a card it wants the partner to play while in other cases it hints rank first. The order depends on the grounded probability of having a playable card in the partner’s hand after giving the hint assuming only the public information. If the partner knowing the color of some cards leads to a higher chance that they will have a playable card in hand, then we should hint color first and vice

³<https://en.boardgamearena.com/>

Off-Belief Learning

Method	Self-Play	Cross-Play	w/ Clone Bot
Other-Play	23.98 ± 0.03	17.36 ± 0.19	12.18 ± 0.25
OBL (level 1)	20.52 ± 0.05	20.41 ± 0.01	13.10 ± 0.42
OBL (level 2)	22.71 ± 0.04	22.50 ± 0.01	14.09 ± 0.48
OBL (level 3)	23.19 ± 0.03	22.85 ± 0.01	13.96 ± 0.47
OBL (level 4)	23.38 ± 0.04	23.02 ± 0.01	13.88 ± 0.47

Table 5. Performance in 3 player Hanabi. We run 10 independent training with different seeds for each method. Self-Play indicates play between an agent and itself. Cross-Play indicates play between agents from different independently-trained runs of the same algorithm. For 3 player Hanabi, the Cross-Play is computed by pairing agents from 3 different training runs. When evaluating agents with clone bot, we average the results of 2 copies of the same agents with 1 clone bot and 1 agent with 2 copies of the same clone bots. Each combination of agents are evaluated on 5000 games.

versa. This information gets picked up by the belief model and then taken advantage of by the OBL in the next level to form conventions. The conventions gradually get reinforced and lead to more nuanced conventions in the higher level of OBL agents.

For reference, we can also see from the same Figure 4 that the Other-Play Rank Bot and Color Bot have strong preferences to use only rank and only color to exchange information respectively, making it difficult for them to coordinate with each other. The Clone Bot uses both color and rank depending on the situation, which is similar to the highest level OBL agent. The SAD agent, which is trained in the typical self-play setting, seems to play “blindly” a lot. However, the fact that such an agent achieves a high self-play score indicates that it must be using some form of secretive conventions that are not grounded at all, e.g. the conventions such as “hinting red means play the second card”, which completely disregard the grounded meaning of the hint actions. Since these conventions by definition assign arbitrary meaning to actions, they are deemed to fail under zero-shot coordination setting and for human-AI coordination.

D. Off-Belief Learning on 3 Player Hanabi

To better demonstrate the generalizability of off-belief learning, we apply it on 3 player Hanabi. The results are shown in Table 5. The Cross-Play between 3 player can either consists of two different agents with two copies of one the one copy of the other or three different agents. We find the latter to be more challenging and therefore the Cross-Play in the table is computed that way. The score with Clone Bot is the average of two Clone Bot with one other bot and one Clone Bot with two other bot. We run 10 independent training seeds for each method. Despite high Self-Play score, Other-Play agents suffer greatly in the Cross-Play, even more than they do in 2 Player Hanabi. However, the gap between Cross-Play and Self-Play scores for each OBL level are significantly smaller. Therefore, even the OBL level 1 agent that does not use any conventions is able to collaborate better in the zero-shot coordination setting. As we apply OBL for more levels, the Cross-Play score grows on the same pace as Self-Play score. Notably, all OBL levels perform better than Other-Play when collaborating with Clone Bot, indicating OBL as a promising method for human-AI coordination.