

Learning to Resolve Alliance Dilemmas in Many-Player Zero-Sum Games

Edward Hughes
DeepMind
edwardhughes@google.com

Thomas W. Anthony
DeepMind
twa@google.com

Tom Eccles
DeepMind
eccles@google.com

Joel Z. Leibo
DeepMind
jzl@google.com

David Balduzzi
DeepMind
dbalduzzi@google.com

Yoram Bachrach
DeepMind
yorambac@google.com

ABSTRACT

Zero-sum games have long guided artificial intelligence research, since they possess both a rich strategy space of best-responses and a clear evaluation metric. What’s more, competition is a vital mechanism in many real-world multi-agent systems capable of generating intelligent innovations: Darwinian evolution, the market economy and the AlphaZero algorithm, to name a few. In two-player zero-sum games, the challenge is usually viewed as finding Nash equilibrium strategies, safeguarding against exploitation regardless of the opponent. While this captures the intricacies of chess or Go, it avoids the notion of cooperation with co-players, a hallmark of the major transitions leading from unicellular organisms to human civilization. Beyond two players, alliance formation often confers an advantage; however this requires trust, namely the promise of mutual cooperation in the face of incentives to defect. Successful play therefore requires adaptation to co-players rather than the pursuit of non-exploitability. Here we argue that a systematic study of many-player zero-sum games is a crucial element of artificial intelligence research. Using symmetric zero-sum matrix games, we demonstrate formally that alliance formation may be seen as a social dilemma, and empirically that naïve multi-agent reinforcement learning therefore fails to form alliances. We introduce a toy model of economic competition, and show how reinforcement learning may be augmented with a peer-to-peer contract mechanism to discover and enforce alliances. Finally, we generalize our agent model to incorporate temporally-extended contracts, presenting opportunities for further work.

KEYWORDS

deep reinforcement learning; multi-agent learning; bargaining and negotiation; coalition formation (strategic)

ACM Reference Format:

Edward Hughes, Thomas W. Anthony, Tom Eccles, Joel Z. Leibo, David Balduzzi, and Yoram Bachrach. 2020. Learning to Resolve Alliance Dilemmas in Many-Player Zero-Sum Games. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 10 pages.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1 INTRODUCTION

Minimax foundations. Zero-sum two-player games have long been a yardstick for progress in artificial intelligence (AI). Ever since the Minimax theorem [25, 81, 86], researchers has striven for human-level play in grand challenge games with combinatorial complexity. Recent years have seen progress in games with increasingly many states: both perfect information (e.g. backgammon [79], checkers [66], chess, [15], Hex [1] and Go [73]) and imperfect information (e.g. Poker [53] and Starcraft [85]).

The *Minimax theorem* [81] states that every finite zero-sum two-player game has an optimal mixed strategy. Formally, if x and y are the sets of possible mixed strategies of the players and A is the payoff matrix of the game, then $\max_x \min_y x^T A y = \min_y \max_x x^T A y = V$, with V referred to as the value of the game. This property makes two-player zero-sum games inherently easier to analyze, as there exist optimal strategies which guarantee each player a certain value, no matter what their opponent does.

Zero-sum two-player games have three appealing features:

- (1) There is an unambiguous measure of algorithm performance, namely performance against a human expert.
- (2) The size of the game tree gives an intuitive measure of complexity, generating a natural difficulty ordering for research.
- (3) The minimax and Nash strategies coincide, so in principle there is no need for a player to adapt to another’s strategy.

Limitations of zero-sum two-player games. The above properties of two-player zero-sum games makes them relatively easy to analyze mathematically, but most real-world interactions extend beyond direct conflict between two individuals. Zero-sum games are rare. Indeed, the chances of a random two-player two-action game being epsilon-constant-sum follow a triangular distribution, as shown empirically in Figure 1. Importantly, this distribution has low density near $\epsilon = 0$: zero-sum games are even rarer than one might naïvely expect. Thus, a research programme based around agent performance in two-player zero-sum games may fail to capture some important social abilities. More explicitly, property (3) is problematic: human intelligence is arguably predicated on our sociality [59], which exactly represents our ability to dynamically respond to a variety of co-players. Relatedly, natural environments are rarely purely adversarial; most environments are mixed-motive, where interactions between individuals comprise a combination of cooperation and competition. In other words, two-player constant-sum is a reasonable mathematical starting point for coarse-grained multi-agent interactions, but to understand the fine detail, we must move beyond this class of games.




Figure 1: Two-player games were randomly generated with payoffs for each player drawn from $U(0, 1)$. A game is defined to be epsilon-constant-sum if the sums of both players’ payoffs are mutually within ℓ_1 distance epsilon of each other. 1000 games were sampled and a histogram plotted.

Approaches to multi-agent cooperation. There are various ways to move beyond zero-sum two-player interactions. One avenue is provided by zero-sum games between two teams of agents, as done in environments such as soccer [48] or Dota 2 [56]. Here, each team of agents faces an internal coordination problem. However, this does not address the limitation of (3) at the level of the team’s strategy. Moreover, as the teams are fixed, the individuals face no problem of team formation [3, 31, 63], an important feature of socio-economic interactions, evinced by the substantial literature on cooperative game theory; see [23] for a review.

Alternatively, one might consider lifting the zero-sum restriction. Much recent work in multi-agent reinforcement learning (MARL) has focussed on social dilemmas with two or more players, e.g. [26, 29, 34, 44, 45, 58, 87]. These domains challenge learning agents to adapt and shape the learning of their co-players to achieve coordinated, cooperative outcomes without being exploited. In this sense, they directly address the limitations of (3). However, the general-sum case lacks the benefits conferred by properties (1) and (2). Indeed, there is no canonical metric of multi-agent performance, which makes comparisons between algorithms particularly difficult.

A third strand of work investigates scenarios in which cooperation is a given. Examples of such multi-agent learning problems include negotiation [16], coalition formation [2] and communication, both through a “cheap talk” channel [28, 54] and grounded in environment actions. These tasks certainly probe the abilities of agents to co-adapt, for they require intricate coordination. However, this is achieved at the expense of removing some of the adversarial dynamics. For example, in Hanabi any subset of agents are by definition incentivized to work together [6].

Finally, back in the two-player zero-sum setting, it is possible to optimize not for win-rate against an arbitrary opponent, but rather for maximum winning margin against some fixed background of exploitable opponents. For instance, in rock-paper-scissors against an opponent who always plays rock, you should always play paper, eschewing the mixed Nash. By definition one must adapt to one’s

co-players. Indeed, the identification of learning of robust best-responses has been extensively studied; see [5, 43, 84, 84] for a range of applications. Nevertheless, the necessity of choosing background opponents removes some of the advantages conferred by properties (1) and (2).

Our contributions. We have three goals. First, to mathematically define the challenge of forming alliances. Second, to demonstrate that state-of-the-art reinforcement learning, used to great effect in two-player zero-sum games, fails to form alliances. Third, to provide a parsimonious and well-motivated mechanism for the formation of alliances by RL agents, namely contracts.

We focus on alliance formation in *many-player* (> 2 -player) zero-sum games as a relatively understudied arena for AI research. We examine symmetric zero-sum many-player games, and provide empirical results showing that alliance formation in such games often yields a social dilemma, thus requiring online adaptation between co-players. Viewing these games through the prism of reinforcement learning, we empirically show that standard independent reinforcement learning agents fail to learn to form alliances. We then propose a simple protocol that can be used to augment reinforcement learning agents, allowing them to learn peer-to-peer contracts, thus enabling them to make alliances. We study the impact of this protocol via a toy model of economic competition, demonstrating that contract-augment agents outperform the model-free baseline. Finally, we extend our model to incorporate enforcement by punishment, demonstrating the potential for our method to scale to spatially and temporally extended scenarios.

2 PRELIMINARIES

2.1 Related work

We are far from the first to study many-player zero-sum games; see for instance [13, 41, 55, 57]. Bonnet et al. [11] have recently studied coalitions in Piglet, although they consider the coalition to be fixed from the beginning of the game. Our novel contributions come from studying alliance formation rigorously as a social dilemma, demonstrating the failure modes of MARL in alliance formation, and providing a mechanism for the learning of contracts. Despite much progress in this field, the problem of learning to form lasting, adaptable alliances without human data remains open.

The cooperative game theory literature provides several methods for coordinating agents in competitive settings [12, 17]. That line of work typically focuses on the question of how to share the joint gains of a team between the team members, in a fair way [4, 10, 52] or in a way that fosters the stability of the coalition [17, 21, 24, 60]. Such work has also been used as a foundation for building agents that cooperate and negotiate with humans [36, 39, 51, 61]. In particular, researchers have proposed algorithms for robust team formation [31, 38, 70, 89] and sharing resources [19, 62, 68]. Our work builds on these strong foundations, being the first to apply MARL in this setting, and demonstrating the efficacy of a concrete protocol that allows MARL agents to form alliances.

Our MARL approach augments agents with the ability to negotiate and form contracts regarding future actions. Many protocols have been suggested for multi-agent negotiation, as discussed in various surveys on the topic [39, 40, 62, 68]. Some propose agent interaction mechanisms [30, 42, 64, 65, 75], whereas others aim to

characterize the stability of fair outcomes for agents [20, 35, 38, 80]. However, the focus here lies on solution concepts, rather than the online learning of policies for both acting in the world and developing social ties with others.

As MARL has become more ubiquitous, so has interest in dynamic team formation increased. In two-player settings, the Coco-Q algorithm achieves social welfare maximizing outcomes on a range of gridworld games when RL agents are augmented with the ability to sign binding contracts and make side payments [76]. For many-agent games, there has been much progress on achieving adhoc teamwork [77] in fully cooperative tasks, including by modelling other agents [8], invoking a centralized manager [72] or learning from past teammates [7]. To our knowledge, no previous work has studied how alliances may be formed by learning agents equipped with a contract channel in many-agent zero-sum games.

During review, two complementary works appeared. [37] equips agents with the ability to negotiate and form contracts as part of a search algorithm, aligned with yet distinct from our reinforcement learning approach. [71] defines a framework for reinforcement learning in N -player games, that can be used to scale up our results, but does not explicitly tackle alliance formation.

2.2 Multi-agent reinforcement learning

We consider multi-agent reinforcement learning (MARL) in Markov games [46, 69]. In each game state, agents take actions based on a (possibly partial) observation of the state space and receive an individual reward. Agents must learn an appropriate behavior policy through experience while interacting with one another. We formalize this as follows. Let \mathcal{M} be an n -player partially observable Markov game defined on a finite set of states \mathcal{S} . The observation function $O : \mathcal{S} \times \{1, \dots, n\} \rightarrow \mathbb{R}^d$ specifies each player’s d -dimensional view on the state space. From each state, each player may take actions from each of the sets $\mathcal{A}^1, \dots, \mathcal{A}^n$ respectively. As a result of the joint action $(a^1, \dots, a^n) \in \mathcal{A}^1 \times \dots \times \mathcal{A}^n$ the state changes, following the stochastic transition function $T : \mathcal{S} \times \mathcal{A}^1 \times \dots \times \mathcal{A}^n \rightarrow \Delta(\mathcal{S})$, where $\Delta(\mathcal{S})$ denotes the set of discrete probability distributions over \mathcal{S} . We denote the observation space of player i by $O^i = \{o^i \mid s \in \mathcal{S}, o^i = O(s, i)\}$. Each player receives an individual reward denoted by $r^i : \mathcal{S} \times \mathcal{A}^1 \times \dots \times \mathcal{A}^n \rightarrow \mathbb{R}$ for player i . We denote the joint actions, observations and policies for all players by \vec{a}, \vec{o} and $\vec{\pi}$ respectively.

In MARL, each agent learns, independently through its own experience of the environment, a behavior policy $\pi^i : O^i \rightarrow \Delta(\mathcal{A}^i)$, which we denote $\pi(a^i | o^i)$. Each agent’s goal is to maximize a long term γ -discounted payoff,¹ namely

$$V_{\pi}^i(s_0) = \mathbb{E} \left(\sum_{t=0}^{\infty} \gamma^t r^i(s_t, \vec{a}_t) \right), \quad (1)$$

where \vec{a}_t is a sampled according to the distribution $\vec{\pi}_t$ and s_{t+1} is sampled according to the distribution $T(s_t, \vec{a}_t)$. In our setup, each agent comprises a feedforward neural network and a recurrent module, with individual observations as input, and policy logits and value estimate as output. We train our agents with a policy gradient method known as advantage actor-critic (A2C) [78], scaled up using the IMPALA framework [27]. Details of neural network

¹For all of our experiments, $\gamma = 0.99$.

configuration and learning hyperparameters are provided alongside each experiment to aid reproducibility.

3 ALLIANCE DILEMMAS

In this section we formally introduce the notion of an alliance dilemma, a definition which captures the additional complexity present in n -player zero-sum games when $n > 2$.

Intuition. The simplest class of multi-player zero-sum games are the symmetric 2-action, 3-player repeated matrix games. Particularly interesting games in this class are those which reduce to two-player social dilemmas if one of the participants adheres to a fixed policy. Here, one agent can make the team formation problem hard for the other two, and may be able to unreasonably generate winnings, since independent reinforcement learning algorithms in social dilemmas converge to defecting strategies. In other words, a appropriately configured “stubborn” agent can force the two dynamic agents into a social dilemma, whence MARL finds the worst collective outcome for them: the Nash.²

Social dilemmas. A two-player matrix game,

R, R	S, T
T, S	P, P

is called a *social dilemma* if and only if [18, 49]:

- (1) $R > P$: mutual cooperation is preferred to mutual defection.
- (2) $R > S$: mutual cooperation is preferred to being exploited by a defector.
- (3) either $T > R$: exploiting a cooperator is preferred over mutual cooperation (*greed*), or $P > S$: mutual defection is preferred over being exploited by a defector (*fear*).

We say that the social dilemma is *strict* iff in addition

- (4) $2R > T + S$: mutual cooperation is preferred over an equal mix of defection and cooperation.

There is some debate as to the importance of this final condition: for example, several papers in the experimental literature violate it [9]. We shall demonstrate that pathological behavior arises in gradient-based learning even in non-strict social dilemmas. Thus from the perspective of alliances, the first three conditions should be considered to be the important ones.

Many-player zero-sum games. An n -player 2-action game is called *zero-sum* if and only if for each vector of simultaneous payoffs (U_1, \dots, U_n) , the following holds: $\sum_{i=1}^n U_i = 0$. In other words, a gain in utility for any one player is exactly balanced by a loss in utility for the rest. A well-known corollary of this is that every outcome is Pareto efficient. However, from the perspective a k -player alliance ($k < n$) within the zero-sum game, Pareto improvements are possible. Note that many-player zero-sum is strictly more general than the class of pairwise zero-sum games recently studied in the context of learning [5]. Importantly, we make no restriction on the payoffs for interactions between strict subsets of players.

Alliance dilemmas. An *alliance dilemma* in an n -player zero-sum 2-action matrix game is a social dilemma which arises for a

²If each agent is running a no-regret algorithm then the average of their strategies converges to a coarse-correlated equilibrium [32]. Empirically, there is much evidence in the literature to suggest that reinforcement learning does reliably find the Nash in social dilemma settings, but this is not theoretically guaranteed.

2-player subset, on the basis of the current policies of their co-players. Mathematically, denote the action space for player i by $\{a_i^0, a_i^1\}$ with $1 \leq i \leq n$. An alliance dilemma is a social dilemma for players i and j which arises when the game is reduced according to some fixed policies for the rest of the players $\pi_k(a_k^0, a_k^1)$ for $k \neq i, j$.

It is not immediately clear whether such situations are ubiquitous or rare. In this section we shall show that at least in the atomic setting described above, alliance dilemmas arise often in randomly generated games. Moreover, we shall justify the informal statement that gradient-based learning fails to achieve reasonable outcomes in the presence of alliance dilemmas, by considering the learning dynamics in two easily interpretable examples.

3.1 Counting alliance dilemmas

How often can a stubborn agent in symmetric three-player two-action matrix games generate a social dilemmas for their two opponents? To answer this question we randomly sample 1000 such games, parameterized by $0 \leq p \leq 1$ and $0 \leq q \leq 1$ as follows:

Actions	Payoffs
(0, 0, 0)	$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$
(0, 0, 1)	$(\frac{1-p}{2}, \frac{1-p}{2}, p)$
(0, 1, 0)	$(\frac{1-p}{2}, p, \frac{1-p}{2})$
(0, 1, 1)	$(q, \frac{1-q}{2}, \frac{1-q}{2})$
(1, 0, 0)	$(p, \frac{1-p}{2}, \frac{1-p}{2})$
(1, 0, 1)	$(\frac{1-q}{2}, q, \frac{1-q}{2})$
(1, 1, 0)	$(\frac{1-q}{2}, \frac{1-q}{2}, q)$
(1, 1, 1)	$(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

For each game we consider 11 possibilities for the stubborn player’s policy, by discretizing the probability of taking action 0 with step size 0.1. For each of these we compute the resulting matrix game for the remaining players and check the social dilemma conditions. If these yield a social dilemma, then the original game is an alliance dilemma, by definition. We find that 54% of games contain an alliance dilemma, with 12% of these being strict alliance dilemmas. Figure 2(A) depicts the distribution of these alliance dilemmas in p - q space. The structure is seen to be the effect of the linear inequalities in the definition of a social dilemma. Figure 2(B) demonstrates that alliance dilemmas arise for stubborn policies that are more deterministic. This stands to reason: the stubborn player must leach enough of the zero-sum nature of the game to leave a scenario that calls for cooperation rather than pure competition.

3.2 The failure of gradient-based learning

There are two values of p and q which result in easily playable games. When $p = 1$ and $q = 0$ we obtain the *Odd One Out* game. Here, you win outright if and only if your action is distinct from that of the other two players. Similarly, the combination $p = 0$ and $q = 1$ defines the *Matching* game. Here, you lose outright if and only if your action is distinct from that of the other two players. It is easy to verify that both of these games contain an alliance dilemma. Odd One Out has a non-strict greed-type dilemma, while Matching has a strict fear-type dilemma. This classification is intuitively sensible: the dilemma in Odd One Out is generated by the possibility of realizing gains, while the dilemma in Matching arises because of the desire to avoid losses.

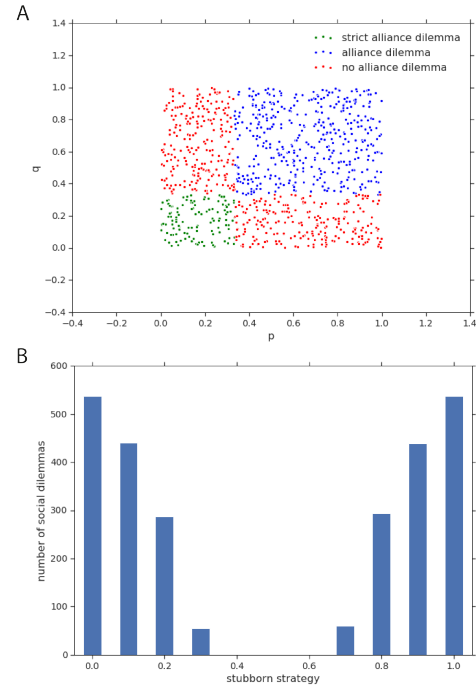


Figure 2: The presence of absence of alliance dilemmas in 1000 simulated three-player zero-sum matrix games. (A) Alliance dilemmas arise for characteristic values of the parameters p and q of the three-player zero-sum game. (B) Alliance dilemmas become more common as the policy of the stubborn agent becomes more deterministic.

Despite the simplicity of these games, gradient-based learning fails to find alliances when pitched against a stubborn agent with a deterministic policy (see Figure 3). More precisely, in Matching the optimal solution for agents in an alliance is to match actions, taking the opposite action from the stubborn agent to gain reward $\frac{1}{2}$ each. However, policies initialised anywhere other than these optimal policies converge to both taking the same action as the stubborn agent, each getting reward $\frac{1}{3}$. In Odd One Out, the optimal symmetric alliance solution is to each match the stubborn agent $\frac{3}{4}$ of the time, for an average reward of $\frac{3}{8}$.³ The learning dynamics of the system tend towards a fixed point in which both agents never match the stubborn agent, gaining no reward at all. In this game, we also see that the system is highly sensitive to the initial conditions or learning rates of the two agents; small differences in the starting point lead to very different trajectories. We provide a mathematical derivation in Appendix ?? which illuminates our empirical findings.

Despite the fact that Matching yields a strict social dilemma while Odd One Out gives a non-strict social dilemma, the learning dynamics in both cases fail to find the optimal alliances. As we anticipated above, this suggests that condition (4) in the definition of social dilemmas is of limited relevance for alliance dilemmas.

³If non-stubborn agents both match the stubborn agent with probability p , their expected joint return is $\frac{2}{3}p^2 + 2p(1-p)$.

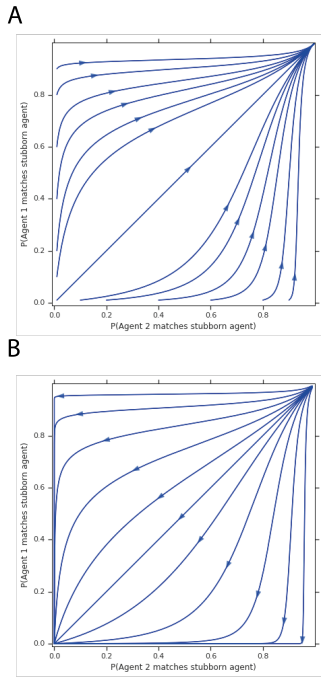


Figure 3: The learning dynamics of two example games. The lines shown are the trajectories of agents under softmax policy gradients with the same learning rate. In Matching (A), both agents learn to match the stubborn agent. In Odd One Out (B), neither agent learns to match the stubborn agent.

3.3 Gifting: an alliance dilemma

Previously we saw that alliance dilemmas are a ubiquitous feature of simultaneous-move games, at least in simple examples. We now demonstrate that the same problems arise in a more complex environment. As our definition of alliance dilemma relies on considering the policy of a given agent as fixed at each point in time, we shift our focus to sequential-move games, where such analysis is more natural, by dint of only one player moving at a time.

In the *Gifting* game (see Figure 4), each player starts with a pile of m chips of their own colour (for our experiments, $m = 5$). On a player’s turn they must take a chip of their own colour and either gift it to another player or discard it from the game. The game ends when no player has any chips of their own color left; that is to say, after $n \times m$ turns. The winner is the player with the most chips (of any colour), with k -way draws possible if k players have the same number of chips. The winners share a payoff of 1 equally, and all other players receive a payoff of 0. Hence Gifting is a constant-sum game, which is strategically equivalent to a zero-sum game, assuming no externalities. Moreover, it admits an interpretation as a toy model of economic competition based on the distribution of scarce goods, as we shall see shortly.

That “everyone always discards” is a subgame perfect Nash equilibrium is true by inspection; i.e. no player has any incentive to deviate from this policy in any subgame, for doing so would merely




Figure 4: The Gifting game. Each player has taken 1 turn. Player 0 gifted a blue chip to player 2; player 1 discarded a yellow chip; player 2 gifted a red chip to player 0.

advantage another player, to their own cost. On the other hand, if two players can arrange to exchange chips with each other then they will achieve a two-way rather than a three-way draw. This is precisely an alliance dilemma: two players can achieve a better outcome for the alliance should they trust each other, yet each can gain by persuading the other to gift a chip, then renege on the deal. Accordingly, one might expect that MARL fails to converge to policies that demonstrate such trading behavior.

Results.⁴ This expectation is borne out by the results in Figure 5(A–C). Agents start out with different amounts of discarding behavior based on random initialization, but this rapidly increases. Accordingly, gifting behavior to agents decreases to zero during learning. The result is a three-way draw, despite that fact that two agents that agreed to exchange could do better than this. To demonstrate this final point, we replace the second player with a bot which reciprocates the gifting behavior of player 0; see Figure 5(D–F). Under these circumstances, player 0 learns to gift to player 1, leading to a two-way draw. Player 2 initially learns to discard, but soon this behavior confers no additional reward, at which point the entropy regularizer leads to random fluctuation in the actions chosen.

We conclude that reinforcement learning can adapt, but only if an institution supporting cooperative behavior exists. MARL cannot create the reciprocal behavior necessary for alliances ex nihilo. Inspired by the economics literature, we propose a mechanism which solves this problem: learning to sign peer-to-peer contracts.

4 CONTRACTS

4.1 Binding contracts

The origin of alliance dilemmas is the greed or fear motivation that drives MARL towards the Pareto-inefficient Nash equilibrium for any given 2-player subset.⁵ Humans are adept at overcoming this problem, largely on the basis of mutual trust [82]. Much previous

⁴In each episode, agents are randomly assigned to a seat 1, 2 or 3, so must generalize over the order of play. Each agent’s neural network comprises an MLP with two layers of 128 hidden units, followed by an LSTM with 128 hidden units. The policy and value heads are linear layers on top of the LSTM. We train with backpropagation-through-time, using an unroll length equal to the length of the episode. Optimization is carried out using the RMSProp optimizer, with decay 0.99, momentum 0, epsilon 0.001 and learning rate 0.000763. The entropy cost for the policy is 0.001443. We perform training runs initialized with 10 different random seeds and plot the average with 95% confidence intervals.

⁵See the definition of social dilemma in Section 3 for a reminder of these concepts.

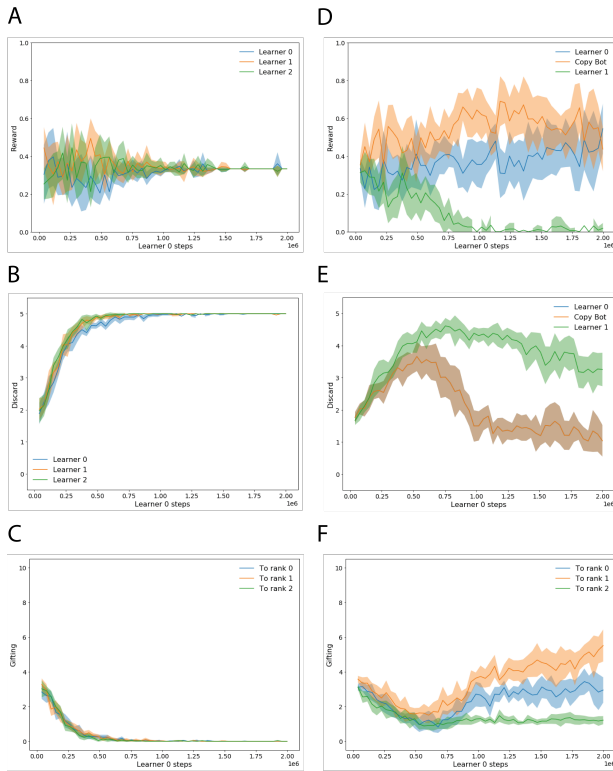


Figure 5: Learning curves for MARL on the Gifting environment. (A–C) Three independent learners quickly converge to the subgame perfect Nash of never gifting, resulting in a three-way draw. (D–F) One of the learners is replaced by a copy bot which always mimics the actions of the first learner. The first learner now learns to gift, since this will surely be reciprocated. The result is a two-way draw. Dependent variables are as follows. (A) Reward. (B) Discard. (C) Gifting. (D) Reward. (E) Discard. (F) Gifting.

work has considered how inductive biases, learning rules and training schemes might generate trust in social dilemmas. Viewing the problem explicitly as an alliance dilemma yields a new perspective: namely, what economic mechanisms enable self-interested agents to form teams? A clear solution comes from contract theory [50]. Binding agreements that cover all eventualities trivialize trust, thus resolving the alliance dilemma.

Contract mechanism. We propose a mechanism for incorporating peer-to-peer pairwise complete contracts into MARL. The core environment is augmented with a contract channel. On each timestep, each player i must submit a contract offer, which comprises a choice of partner j , a suggested action for that partner j , and an action which i promises to take, or no contract. If two players offer contracts which are identical, then these become binding; that is to say, the environment enforces that the promised actions are taken, by providing a mask for the logits of the relevant agents. At each step, agents receive an observation of the contracts which were offered on the last timestep, encoded as a one-hot representation.

Contract-aware MARL. To learn in this contract-augmented environment, we employ a neural network with two policy heads, one for the core environment and another for the contract channel. This network receives the core environment state and the previous timestep contracts as input. Both heads are trained with the A2C algorithm based on rewards from the core environment, similarly to the RIAL algorithm [28]. Therefore agents must learn simultaneously how to behave in the environment and how to leverage binding agreements to coordinate better with peers.

Results.⁶ We run two experiments to ascertain the performance benefits for contract-augmented agents. First, we train two contract-augmented agents and one A2C agent together. Figure 6(A–C) shows that the two contract-augmented agents (Learners 0 and 1) are able to achieve a 2-way draw and eliminate the agent without the ability to sign contracts. We then train three contract-augmented agents together. Figure 6(D) demonstrates the reward dynamics that result as agents vie to make deals that will enable them to do better than a 3-way draw. Figure 7 shows that signing contracts has a significant correlation with accruing more chips in a game, thus demonstrating that contracting is advantageous.

The benefits of contracting have an interesting economic interpretation. Without contracts, and the benefits of mutual trust they confer, there is no exchange of chips. In economic terms, the “gross domestic product” (GDP) is zero. Once agents can sign binding contracts, goods flow around the system — economic productivity is stimulated. Of course, we should take these observations with a large grain of salt, for our model is no more than a toy. Nevertheless, this does hint that ideas from macro-economics may be a valuable source of inspiration for MARL algorithms that can coordinate and cooperate at scale. For instance, our model demonstrates the GDP benefits of forming corporations of a sufficiently large size $k > 1$ but also a sufficiently small size $k < 3$.

4.2 Contracts with temporal extent

As it stands, our contract mechanism requires that contracts are fulfilled immediately, by invoking a legal actions mask on the agent’s policy logits. On the other hand, many real-world contracts involve a promise to undertake some action during a specified future period; that is to say, contracts have temporal extent. Furthermore, directly masking an agent’s action is a heavy-handed and invasive way to enforce that agents obey contractual commitments. By contrast, human contracts are typically enforced according to contract law [67]. Indeed, in the language of [33] “whatever else they do, all legal systems recognize, create, vary and enforce obligations”.

Legal systems stipulate that those who do not fulfil their contractual obligations within a given time period should be punished. Inspired by this, we generalize our contract mechanism. Contracts are offered and signed as before. However, once signed, agents have b timesteps in which to take the promised action. Should an agent fail to do so, they are considered to have broken the contract, and receive a negative reward of r_c . Once a contract is either fulfilled or broken, agents are free to sign new contracts.

⁶The neural network is identical to the baseline experiments, except for the addition of a linear contract head. The optimizer minimizes the combined loss $L_{RL} + \alpha L_{contract}$ where $\alpha = 1.801635$. We include an entropy regularizer for the contract policy, with entropy cost 0.000534. Training and evaluation methodology are identical to the baseline experiments.

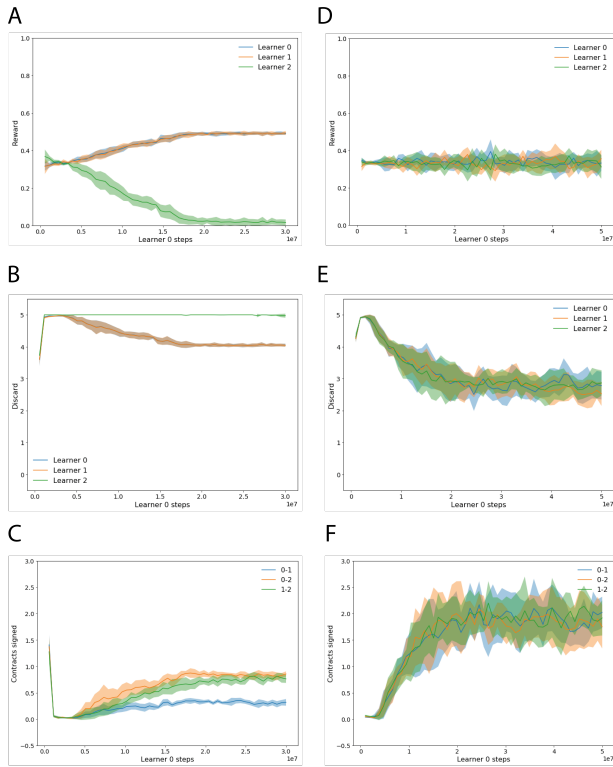


Figure 6: Performance of contract-augmented agents on Gifting. Data comes from evaluation episodes without a trembling hand policy. (A–C) Learners 0 and 1 are endowed with the ability to make contracts, but learner 2 is not. The contract augment agents use this ability to consistently achieve a 2-way draw. (D–F) All learners are endowed with the ability to make contracts. (D) Rewards across training become more chaotic, since agents are vying with each other to sign the contracts to secure a 2-way draw (D). As agents compete to sign contracts guaranteeing them more gifts, the number of discards decreases (E) and the number of contracts signed increases (F). Dependent variables are as follows. (A) Reward. (B) Discard. (C) Contracts signed. (D) Reward. (E) Discard. (F) Contracts signed.

Trembling hand. Learning to sign punishment-enforced contracts from scratch is tricky. This is because there are two ways to avoid punishment: either fulfil a contract, or simply refuse to sign any. Early on in training, when the acting policies are random, agents will learn that signing contracts is bad, for doing so leads to negative reward in the future, on average. By the time that the acting policies are sensibly distinguishing moves in different states of the game, the contract policies will have overfitted to never offering a contract.

To learn successfully, agents must occasionally be forced to sign contracts with others. Now, as the acting policies become more competent, agents learn that contracts can be fulfilled, so do not always lead to negative reward. As agents increasingly fulfil their




Figure 7: 50 games are rolled out in evaluation (after training). For each game, we record the total number of chips accrued (points) and contracts signed by each player. We present a scatter plot of all 150 such values. Linear regression shows a significant positive correlation between the signing of gifting contracts and points accrued during the game.

contractual obligations, they can learn to sign contracts with others. While such contracts are not binding in the sense of legal actions within the game, they are effectively binding since agents have learned to avoid the negative reward associated with breaking a contract.

We operationalize this requirement by forcing agents to follow a *trembling hand* policy for contract offers. Every timestep, the contract mechanism determines whether there are two agents not already participating in a contract. The mechanism chooses a random contract between the two, and applies a mask to the logits of the contract policy of each agent. This mask forces each agent to suggest the chosen contract with a given minimum probability p_c . We also roll out episodes without the trembling hand intervention, which are not used for learning, but plotted for evaluation purposes.

Results.⁷ We ran one experiment to ascertain the ability of agents to learn to form gifting contracts. All agents are contract-aware and have trembling hand policies while learning. We display our results in Figure 8. Initially agents learn to discard, and indeed back this up by signing contract for mutual discarding behavior. However, after some time agents discover the benefits of signing gifting contracts. Thereafter, the agents vie with each other to achieve two-player alliances for greater reward. Interestingly, the agents do not learn to gift as much as in the binding contracts case: compare Figure 8(B) and Figure 6(E). This is likely because contracts are not always perfectly adhered to, so there is a remnant of the fear and greed that drives discarding behavior.

5 CONCLUSION

In this paper, we have made five key contributions. We (1) formalized one of the challenges of many-player zero-sum games by introducing the notion of an alliance dilemma. We (2) demonstrated

⁷The neural network is identical to the binding contracts experiment. For the punishment mechanism, we choose $b = 6$, $r_c = -1$ and $p_c = 0.5$. Training hyperparameters are as follows: learning rate 0.002738, environment policy entropy cost 0.004006, contract loss weight 3.371262 and contract entropy cost 0.002278. All other hyperparameters are unchanged from the previous experiments.

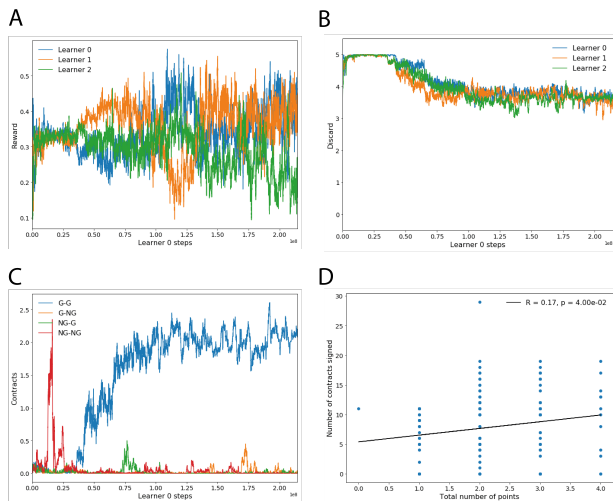


Figure 8: Punishment-enforced contracts lead to alliance formation in Gifting. (A) Initially agents achieve a 3-way draw, receiving $\frac{1}{3}$ reward each. Later agents learn to form alliances in which two agents get reward $\frac{1}{2}$ and the other 0. (B) After an initial period in which all agents discard, gifting behavior is learned. (C) Agents first learn to sign mutually non-gifting contracts (NG-NG), but later discover the benefit of mutually gifting contracts (G-G). (D) There is a significant positive correlation between number of contracts signed and points scored. Dependent variables are as follows. (A) Reward. (B) Discard. (C) Contracts signed. (D) Number of chips accrued vs. number of contracts signed.

that these are ubiquitous, and that (3) gradient-based and reinforcement learning fails to resolve them. We introduced (4) an approach for allowing MARL agents to form dynamic team, by augmenting the agents with a *binding contract* channel. These agents learn to use contracts as a mechanism for trusting a potential alliance partner. Finally, we (5) generalized our contract mechanism beyond the binding case, showing that agents learn to sign temporally-extended contracts enforced through punishment.

Future work. Our model suggests several avenues for further work. Most obviously, we might consider contracts in an environment with a larger state space, such a spatially extended gridworld. Here, agreeing on specific actions to take within an alliance is likely too granular. The alliance dilemma would emerge on the level of the meta-game defined by the policies of different players, defining a *sequential alliance dilemma*. There are at least two promising approaches in such a setting. Firstly, we could incorporate a centralized agent to whose policy agents could defer, should they wish to enter an alliance. Secondly, we could make contracts about an abstraction over the state of the environment, rather than about atomic actions. Further, one might want to include multiple rounds of contract negotiation per time step of the environment, along the lines of [88].

More generally, it would be fascinating to discover how a system of contracts might emerge and persist within multi-agent learning

dynamics without directly imposing mechanisms for enforcement. Such a pursuit may eventually lead to a valuable feedback loop from AI to sociology and economics. Relatedly, we might ask how to scale contracts beyond the bilateral case, given the exponential explosion of possible alliances with the number of players in the game. Indeed, real-world institutions sign large numbers of contracts simultaneously, each of which may involve several partners. Finally, we note that our contract mechanism provides a simple grounded “language” for agent interaction. We hope to draw a stronger connection between this work and the emergent communication literature in the future.

Outlook. Many-player zero-sum games are a common feature of the natural world, from the economic competition of Adam Smith [74] to Darwin’s theory of evolution [22] which can be viewed as a zero-sum game for energy [83]. In many-player zero-sum games a single agent cannot necessarily play a policy which is impervious to its opponent’s behavior. Rather, to be successful, an algorithm must influence the joint strategy across many players. In particular, zero-sum multi-player games introduce the problem of dynamic team formation and breakup. This problem is remarkably deep, touching three strands of multi-agent research beyond two-player zero-sum games:

- (1) Emergent teams must coordinate within themselves to effectively compete in the game, just as in team games like soccer.
- (2) The process of team formation may itself be a social dilemma. Intuitively, players should form alliances to defeat others; however, membership of a alliance requires individuals to contribute to a wider good which is not completely aligned with their self-interest.
- (3) Decisions must be made about which teams to join and leave, and how to shape the strategy of these teams. Here communication is vital, analogously to work on negotiation and contractual team formation.

In Section 1 we identified three appealing features of zero-sum two-player games. Many-player zero-sum games inherit the appealing properties (1) and (2) from their two-player counterparts. However, the restrictive property (3) does not hold, since there is no general analogue of the minimax theorem beyond 2-player games.⁸ For example, the minimax- Q algorithm [47] is not applicable, because it fails to account for dynamic teams. We have demonstrated that many-player zero-sum games capture intricate, important and interesting multi-agent dynamics, amenable to formal study with appropriate definitions. We look forward to future developments of AI in this fascinating arena of study.

REFERENCES

- [1] Thomas Anthony, Zheng Tian, and David Barber. 2017. Thinking Fast and Slow with Deep Learning and Tree Search. *CoRR* abs/1705.08439 (2017). arXiv:1705.08439 <http://arxiv.org/abs/1705.08439>
- [2] Yoram Bachrach, Richard Everett, Edward Hughes, Angeliki Lazaridou, Joel Leibo, Marc Lanctot, Mike Johanson, Wojtek Czarnecki, and Thore Graepel. 2019. Negotiating Team Formation Using Deep Reinforcement Learning. (2019). <https://openreview.net/forum?id=HJG0oJCeFm>
- [3] Yoram Bachrach, Pushmeet Kohli, Vladimir Kolmogorov, and Morteza Zadimoghaddam. 2013. Optimal coalition structure generation in cooperative graph games. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*.

⁸Although in restricted classes, a statement can be made [14].

- [4] Yoram Bachrach, Evangelos Markakis, Ezra Resnick, Ariel D Procaccia, Jeffrey S Rosenschein, and Amin Saberi. 2010. Approximating power indices: theoretical and empirical analysis. *Autonomous Agents and Multi-Agent Systems* 20, 2 (2010), 105–122.
- [5] David Balduzzi, Marta Garnelo, Yoram Bachrach, Wojciech M. Czarnecki, Julien Pérolat, Max Jaderberg, and Thore Graepel. 2019. Open-ended Learning in Symmetric Zero-sum Games. *CoRR* abs/1901.08106 (2019). arXiv:1901.08106 <http://arxiv.org/abs/1901.08106>
- [6] Nolan Bard, Jakob N. Foerster, Sarath Chandar, Neil Burch, Marc Lanctot, H. Francis Song, Emilio Parisotto, Vincent Dumoulin, Subhdeep Moitra, Edward Hughes, Iain Dunning, Shihb Mourad, Hugo Larochelle, Marc G. Bellemare, and Michael Bowling. 2019. The Hanabi Challenge: A New Frontier for AI Research. *CoRR* abs/1902.00506 (2019). arXiv:1902.00506 <http://arxiv.org/abs/1902.00506>
- [7] Samuel Barrett, Avi Rosenfeld, Sarit Kraus, and Peter Stone. 2017. Making friends on the fly: Cooperating with new teammates. *Artificial Intelligence* 242 (2017), 132 – 171. <https://doi.org/10.1016/j.artint.2016.10.005>
- [8] Samuel Barrett, Peter Stone, Sarit Kraus, and Avi Rosenfeld. 2012. Learning Teammate Models for Ad Hoc Teamwork. In *AAMAS 2012*.
- [9] Martin Beckenkamp, Heike Hennig-Schmidt, and Frank Maier-Rigaud. 2007. Cooperation in Symmetric and Asymmetric Prisoner’s Dilemma Games. *Max Planck Institute for Research on Collective Goods, Working Paper Series of the Max Planck Institute for Research on Collective Goods* (03 2007). <https://doi.org/10.2139/ssrn.968942>
- [10] JM Bilbao, JR Fernandez, A Jiménez Losada, and JJ Lopez. 2000. Generating functions for computing power indices efficiently. *Top* 8, 2 (2000), 191–213.
- [11] François Bonnet, Todd W Neller, and Simon Viennot. 2018. Towards Optimal Play of Three-Player Piglet and Fig. (2018).
- [12] Rodica Branzei, Dinko Dimitrov, and Stef Tijs. 2008. *Models in cooperative game theory*. Vol. 556. Springer Science & Business Media.
- [13] Noam Brown and Tuomas Sandholm. 2019. Superhuman AI for multiplayer poker. *Science* 365, 6456 (2019), 885–890.
- [14] Yang Cai and Constantinos Daskalakis. 2011. On Minmax Theorems for Multi-player Games. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA ’11)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 217–234. <http://dl.acm.org/citation.cfm?id=2133036>. 2133056
- [15] Murray Campbell, A Joseph Hoane, and Feng hsiung Hsu. 2002. Deep Blue. *Artificial Intelligence* 134, 1 (2002), 57 – 83. [https://doi.org/10.1016/S0004-3702\(01\)00129-1](https://doi.org/10.1016/S0004-3702(01)00129-1)
- [16] Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z. Leibo, Karl Tuyls, and Stephen Clark. 2018. Emergent Communication through Negotiation. *CoRR* abs/1804.03980 (2018). arXiv:1804.03980 <http://arxiv.org/abs/1804.03980>
- [17] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. 2011. Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5, 6 (2011), 1–168.
- [18] A.R.A.M. Chammah, A. Rapoport, A.M. Chammah, and C.J. Orwant. 1965. *Prisoner’s Dilemma: A Study in Conflict and Cooperation*. University of Michigan Press. <https://books.google.co.uk/books?id=yPtNnKjXaj4C>
- [19] Han-Lim Choi, Luc Brunet, and Jonathan P How. 2009. Consensus-based decentralized auctions for robust task allocation. *IEEE transactions on robotics* 25, 4 (2009), 912–926.
- [20] Vincent Conitzer and Tuomas Sandholm. 2004. Computing Shapley values, manipulating value division schemes, and checking core membership in multi-issue domains. In *AAAI*, Vol. 4. 219–225.
- [21] Vincent Conitzer and Tuomas Sandholm. 2006. Complexity of constructing solutions in the core based on synergies among coalitions. *Artificial Intelligence* 170, 6–7 (2006), 607–619.
- [22] Charles Darwin. 1859. *On the Origin of Species by Means of Natural Selection*. Murray, London. or the Preservation of Favored Races in the Struggle for Life.
- [23] Julia Drechsel. 2010. Selected Topics in Cooperative Game Theory. (07 2010). https://doi.org/10.1007/978-3-642-13725-9_2
- [24] Paul E Dunne, Wiebe van der Hoek, Sarit Kraus, and Michael Wooldridge. 2008. Cooperative boolean games. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 1015–1022.
- [25] Claude E. Shannon. 1950. XXII. Programming a computer for playing chess. *Philos. Mag.* 41 (03 1950), 256–275. <https://doi.org/10.1080/14786445008521796>
- [26] Tom Eccles, Edward Hughes, János Kramár, Steven Wheelwright, and Joel Z. Leibo. 2019. Learning Reciprocity in Complex Sequential Social Dilemmas. *CoRR* abs/1903.08082 (2019). arXiv:1903.08082 <http://arxiv.org/abs/1903.08082>
- [27] Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. 2018. IMPALA: Scalable Distributed Deep-RL with Importance Weighted Actor-Learner Architectures. *CoRR* abs/1802.01561 (2018). arXiv:1802.01561 <http://arxiv.org/abs/1802.01561>
- [28] Jakob N. Foerster, Yannis M. Assael, Nando de Freitas, and Shimon Whiteson. 2016. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. *CoRR* abs/1605.06676 (2016). arXiv:1605.06676 <http://arxiv.org/abs/1605.06676>
- [29] Jakob N. Foerster, Richard Y. Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. 2017. Learning with Opponent-Learning Awareness. *CoRR* abs/1709.04326 (2017). arXiv:1709.04326 <http://arxiv.org/abs/1709.04326>
- [30] Nicoletta Fornara and Marco Colombetti. 2002. Operational specification of a commitment-based agent communication language. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 2*. ACM, 536–542.
- [31] Matthew E Gaston and Marie DesJardins. 2005. Agent-organized networks for dynamic team formation. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*. ACM, 230–237.
- [32] Geoffrey J. Gordon, Amy Greenwald, and Casey Marks. 2008. No-regret Learning in Convex Games. In *Proceedings of the 25th International Conference on Machine Learning (ICML ’08)*. ACM, New York, NY, USA, 360–367. <https://doi.org/10.1145/1390156.1390202>
- [33] Leslie Green. 2012. Legal Obligation and Authority. In *The Stanford Encyclopedia of Philosophy* (winter 2012 ed.), Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [34] Edward Hughes, Joel Z. Leibo, Matthew G. Philips, Karl Tuyls, Edgar A. Duñez-Guzmán, Antonio García Castañeda, Iain Dunning, Tina Zhu, Kevin R. McKee, Raphael Koster, Heather Roff, and Thore Graepel. 2018. Inequity aversion resolves intertemporal social dilemmas. *CoRR* abs/1803.08884 (2018). arXiv:1803.08884 <http://arxiv.org/abs/1803.08884>
- [35] Samuel Jeong and Yoav Shoham. 2005. Marginal contribution nets: a compact representation scheme for coalitional games. In *Proceedings of the 6th ACM conference on Electronic commerce*. ACM, 193–202.
- [36] Nicholas R Jennings, Peyman Faratin, Alessio R Lomuscio, Simon Parsons, Michael J Wooldridge, and Carles Sierra. 2001. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation* 10, 2 (2001), 199–215.
- [37] Dave Jonge and Dongmo Zhang. 2020. Strategic negotiations for extensive-form games. *Autonomous Agents and Multi-Agent Systems* 34 (04 2020). <https://doi.org/10.1007/s10458-019-09424-y>
- [38] Matthias Klusch and Andreas Gerber. 2002. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems* 17, 3 (2002), 42–47.
- [39] Sarit Kraus. 1997. Negotiation and cooperation in multi-agent environments. *Artificial intelligence* 94, 1-2 (1997), 79–97.
- [40] Sarit Kraus and Ronald C Arkin. 2001. *Strategic negotiation in multiagent environments*. MIT press.
- [41] Sarit Kraus, Daniel Lehmann, and E. Ephrati. 1989. An automated Diplomacy player. *Heuristic Programming in Artificial Intelligence: The 1st Computer Olympiad* (1989), 134–153.
- [42] Kazuhiro Kuwabara, Toru Ishida, and Nobuyasu Osato. 1995. AgenTalk: Coordination Protocol Description for Multiagent Systems. In *ICMAS*, Vol. 95. 455–461.
- [43] Marc Lanctot, Vinícius Flores Zambaldi, Audrunas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. 2017. A Unified Game-Theoretic Approach to Multiagent Reinforcement Learning. *CoRR* abs/1711.00832 (2017). arXiv:1711.00832 <http://arxiv.org/abs/1711.00832>
- [44] Joel Z. Leibo, Vinícius Flores Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-agent Reinforcement Learning in Sequential Social Dilemmas. *CoRR* abs/1702.03037 (2017). arXiv:1702.03037 <http://arxiv.org/abs/1702.03037>
- [45] Adam Lerer and Alexander Peysakhovich. 2017. Maintaining cooperation in complex social dilemmas using deep reinforcement learning. *CoRR* abs/1707.01068 (2017). arXiv:1707.01068 <http://arxiv.org/abs/1707.01068>
- [46] M. L. Littman. 1994. Markov Games as a Framework for Multi-Agent Reinforcement Learning. In *Proceedings of the 11th International Conference on Machine Learning (ICML)*. 157–163.
- [47] Michael L. Littman. 1994. Markov Games As a Framework for Multi-agent Reinforcement Learning. In *Proceedings of the Eleventh International Conference on International Conference on Machine Learning (ICML ’94)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 157–163. <http://dl.acm.org/citation.cfm?id=3091574.3091594>
- [48] Siqi Liu, Guy Lever, Nicholas Heess, Josh Merel, Saran Tunyasuvunakool, and Thore Graepel. 2019. Emergent Coordination Through Competition. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BkG8sJR5Km>
- [49] Michael W. Macy and Andreas Flache. 2002. Learning dynamics in social dilemmas. *Proceedings of the National Academy of Sciences* 99 (2002), 7229–7236. <https://doi.org/10.1073/pnas.092080099>
- [50] David Martimort. 2017. *Contract Theory*. Palgrave Macmillan UK, London, 1–11. https://doi.org/10.1057/978-1-349-95121-5_2542-1
- [51] Moshe Mash, Yoram Bachrach, and Yair Zick. 2017. How to form winning coalitions in mixed human-computer settings. In *Proceedings of the 26th international joint conference on artificial intelligence (IJCAI)*. 465–471.
- [52] Tomasz P Michalak, Karthik V Aadithya, Piotr L Szczepanski, Balaraman Ravindran, and Nicholas R Jennings. 2013. Efficient computation of the Shapley value for game-theoretic network centrality. *Journal of Artificial Intelligence Research* 46 (2013), 607–650.

- [53] Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael H. Bowling. 2017. DeepStack: Expert-Level Artificial Intelligence in No-Limit Poker. *CoRR* abs/1701.01724 (2017). arXiv:1701.01724 <http://arxiv.org/abs/1701.01724>
- [54] Igor Mordatch and Pieter Abbeel. 2017. Emergence of Grounded Compositional Language in Multi-Agent Populations. *CoRR* abs/1703.04908 (2017). arXiv:1703.04908 <http://arxiv.org/abs/1703.04908>
- [55] Joseph Antonius Maria Nijssen. 2013. *Monte-Carlo tree search for multi-player games*. Maastricht University.
- [56] OpenAI. 2018. OpenAI Five. (2018). <https://blog.openai.com/openai-five/>
- [57] Philip Paquette, Yuchen Lu, Steven Bocco, Max O. Smith, Satya Ortiz-Gagne, Jonathan K. Kummerfeld, Satinder Singh, Joelle Pineau, and Aaron Courville. 2019. No Press Diplomacy: Modeling Multi-Agent Gameplay. (2019). arXiv:cs.AI/1909.02128
- [58] Julien Pérolat, Joel Z. Leibo, Vinícius Flores Zambaldi, Charles Beattie, Karl Tuyls, and Thore Graepel. 2017. A multi-agent reinforcement learning model of common-pool resource appropriation. *CoRR* abs/1707.06600 (2017). arXiv:1707.06600 <http://arxiv.org/abs/1707.06600>
- [59] Simon M. Reader and Kevin N. Laland. 2002. Social intelligence, innovation, and enhanced brain size in primates. *Proceedings of the National Academy of Sciences* 99, 7 (2002), 4436–4441. <https://doi.org/10.1073/pnas.062041299> arXiv:<https://www.pnas.org/content/99/7/4436.full.pdf>
- [60] Ezra Resnick, Yoram Bachrach, Reshef Meir, and Jeffrey S Rosenschein. 2009. The cost of stability in network flow games. In *International Symposium on Mathematical Foundations of Computer Science*. Springer, 636–650.
- [61] Ariel Rosenfeld and Sarit Kraus. 2018. Predicting human decision-making: From prediction to action. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 12, 1 (2018), 1–150.
- [62] Jeffrey S Rosenschein and Gilad Zlotkin. 1994. *Rules of encounter: designing conventions for automated negotiation among computers*. MIT press.
- [63] Tuomas Sandholm, Kate Larson, Martin Andersson, Onn Shehory, and Fernando Tohmé. 1999. Coalition structure generation with worst case guarantees. *Artificial Intelligence* 111, 1-2 (1999), 209–238.
- [64] Tuomas Sandholm, Victor R Lesser, et al. 1995. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *ICMAS*, Vol. 95, 12–14.
- [65] Tuomas W Sandholm and Victor R Lesser. 1996. Advantages of a leveled commitment contracting protocol. In *AAAI/IAAI*, Vol. 1, 126–133.
- [66] Jonathan Schaeffer, Joseph Culberson, Norman Treloar, Brent Knight, Paul Lu, and Duane Szafron. 1992. A world championship caliber checkers program. *Artificial Intelligence* 53, 2 (1992), 273 – 289. [https://doi.org/10.1016/0004-3702\(92\)90074-8](https://doi.org/10.1016/0004-3702(92)90074-8)
- [67] Alan Schwartz and Robert E. Scott. 2003. Contract Theory and the Limits of Contract Law. *The Yale Law Journal* 113, 3 (2003), 541–619. <http://www.jstor.org/stable/3657531>
- [68] Jeff S Shamma. 2007. *Cooperative control of distributed multi-agent systems*. Wiley Online Library.
- [69] L. S. Shapley. 1953. Stochastic Games. In *Proc. of the National Academy of Sciences of the United States of America* (1953).
- [70] Onn Shehory and Sarit Kraus. 1998. Methods for task allocation via agent coalition formation. *Artificial intelligence* 101, 1-2 (1998), 165–200.
- [71] Alexander Shmakov, John Lanier, Stephen McAleer, Rohan Achar, Cristina Lopes, and Pierre Baldi. 2019. ColosseumRL: A Framework for Multiagent Reinforcement Learning in *N*-Player Games. (2019). arXiv:cs.MA/1912.04451
- [72] Tianmin Shu and Yuandong Tian. 2019. M3RL: Mind-aware Multi-agent Management Reinforcement Learning. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=BkzeUiRcY7>
- [73] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the Game of Go with Deep Neural Networks and Tree Search. *Nature* 529, 7587 (Jan. 2016), 484–489. <https://doi.org/10.1038/nature16961>
- [74] Adam Smith. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. McMaster University Archive for the History of Economic Thought. <https://EconPapers.repec.org/RePEc:hay:hetboo:smith1776>
- [75] Reid G Smith. 1980. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on computers* 12 (1980), 1104–1113.
- [76] E. Sodomka, E.M. Hilliard, M.L. Littman, and Amy Greenwald. 2013. Coco-Q: Learning in stochastic games with side payments. *30th International Conference on Machine Learning, ICML 2013* (01 2013), 2521–2529.
- [77] Peter Stone, Gal A. Kaminka, Sarit Kraus, and Jeffrey S. Rosenschein. 2010. Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination. In *AAAI*.
- [78] Richard S. Sutton and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction* (second ed.). The MIT Press. <http://incompleteideas.net/book/the-book-2nd.html>
- [79] Gerald Tesauro. 1995. Temporal Difference Learning and TD-Gammon. *Commun. ACM* 38, 3 (March 1995), 58–68. <https://doi.org/10.1145/203330.203343>
- [80] Maksim Tsvetovat, Katia Sycara, Yian Chen, and James Ying. 2000. Customer coalitions in the electronic marketplace. In *AAAI/IAAI*, 1133–1134.
- [81] J. v. Neumann. 1928. Zur Theorie der Gesellschaftsspiele. *Math. Ann.* 100, 1 (01 Dec 1928), 295–320. <https://doi.org/10.1007/BF01448847>
- [82] P.A.M. van Lange, B. Rockenbach, and T. Yamagishi. 2017. *Trust in Social Dilemmas*. Oxford University Press. <https://books.google.co.uk/books?id=e-fwsWEACAAJ>
- [83] L Van Valen. 1980. Evolution as a zero-sum game for energy. *Evolutionary Theory* 4 (1980), 289–300.
- [84] Alexander Sasha Vezhnevets, Yuhuai Wu, Rémi Leblond, and Joel Z. Leibo. 2019. Options as responses: Grounding behavioural hierarchies in multi-agent RL. *CoRR* abs/1906.01470 (2019). arXiv:1906.01470 <http://arxiv.org/abs/1906.01470>
- [85] Oriol Vinyals, Timo Ewalds, Sergey Bartunov, Petko Georgiev, Alexander Sasha Vezhnevets, Michelle Yeo, Alireza Makhzani, Heinrich Küttler, John Agapiou, Julian Schrittwieser, John Quan, Stephen Gaffney, Stig Petersen, Karen Simonyan, Tom Schaul, Hado van Hasselt, David Silver, Timothy P. Lillicrap, Kevin Calderone, Paul Keet, Anthony Brunasso, David Lawrence, Anders Ekermo, Jacob Repp, and Rodney Tsing. 2017. StarCraft II: A New Challenge for Reinforcement Learning. *CoRR* abs/1708.04782 (2017). arXiv:1708.04782 <http://arxiv.org/abs/1708.04782>
- [86] John VON NEUMANN and Oskar Morgenstern. 1953. Theory of Games and Economic Behavior / J. von Neumann, O. Morgenstern ; introd. de Harold W. Kuhn. *SERBIULA (sistema Librum 2.0)* 26 (01 1953).
- [87] Jane X. Wang, Edward Hughes, Chrisantha Fernando, Wojciech M. Czarnecki, Edgar A. Duéñez-Guzmán, and Joel Z. Leibo. 2018. Evolving intrinsic motivations for altruistic behavior. *CoRR* abs/1811.05931 (2018). arXiv:1811.05931 <http://arxiv.org/abs/1811.05931>
- [88] Colin Williams, Valentin Robu, Enrico Gerding, and Nicholas Jennings. 2012. Negotiating Concurrently with Unknown Opponents in Complex, Real-Time Domains.
- [89] Gilad Zlotkin and Jeffrey S Rosenschein. 1994. *Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains*. Alfred P. Sloan School of Management, Massachusetts Institute of Technology.