

Courtesy as a Means to Coordinate

Panayiotis Danassis

École Polytechnique Fédérale de Lausanne (EPFL)
Artificial Intelligence Laboratory
Lausanne, Switzerland
panayiotis.danassis@epfl.ch

Boi Faltings

École Polytechnique Fédérale de Lausanne (EPFL)
Artificial Intelligence Laboratory
Lausanne, Switzerland
boi.faltings@epfl.ch

ABSTRACT

We investigate the problem of multi-agent coordination under rationality constraints. Specifically, role allocation, task assignment, resource allocation, etc. Inspired by human behavior, we propose a framework (CA³NONY) that enables fast convergence to efficient and fair allocations based on a simple convention of courtesy. We prove that following such convention induces a strategy which constitutes an ϵ -subgame-perfect equilibrium of the repeated allocation game with discounting. Simulation results highlight the effectiveness of CA³NONY as compared to state-of-the-art bandit algorithms, since it achieves more than two orders of magnitude faster convergence, higher efficiency, fairness, and average payoff.

KEYWORDS

[Learning and Adaptation] Multiagent learning; [Economic Paradigms] Noncooperative games: theory & analysis

ACM Reference Format:

Panayiotis Danassis and Boi Faltings. 2019. Courtesy as a Means to Coordinate. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 9 pages.

1 INTRODUCTION

In multi-agent systems (MAS), agents are often called upon to implement a joint plan in order to maximize their rewards. Typically, coordination in a joint plan incorporates (possibly a combination of) two distinct elements: agents may be required to take the same action [14, 31], or agents may be required to take distinct actions [13, 19]. The latter is ubiquitous in everyday life, e.g. managing common-pool resources, or deciding on car/ship/airplane routes in traffic management. This paper studies coordination in repeated allocation games. Consider for example the problem of channel allocation in wireless networks [36]. In such a problem, N wireless devices contend for R transmission slots (i.e. particular frequency band in a certain period of time), where $N \gg R$. If more than one agent access a slot simultaneously, a collision occurs and the colliding parties incur a cost $\zeta < 0$. The goal then for the agents is to transmit on different slots to minimize collisions over time. Other scenarios include role allocation (e.g. teammates during a game), task assignment (e.g. employees of a factory), resource allocation (e.g. parking spaces and/or charging stations for autonomous vehicles), etc. What follows is applicable to *any such scenario*. For simplicity hereafter we will refer only to indivisible resources (i.e.

resources which can not be shared). Beyond the scope of repeated games, the proposed framework can also be applied as a negotiation protocol in one-shot interactions. E.g. self-driving vehicles attempting to get a parking space can utilize such a protocol in a simulated environment with message exchange. Conforming to the relevant literature, in what follows we refer to the coordination problem of selecting distinct actions as the anti-coordination problem [9, 13, 16, 19, 23].

The most straightforward solution would be to have a central coordinator with complete information recommend an action to each agent. This, though, would require the agents to communicate their preferences/plans, which creates high overhead and raises incentive compatibility and truthfulness issues. Moreover, in real-world applications with partial observability, agents might not be willing to trust such recommendations. On the other hand, agents can learn to anti-coordinate their actions in a completely decentralized manner. However, in a fully decentralized scheme, agents have an incentive to 'bully' others (stick to your action until you drive out the competition) [25], which makes it impossible to converge to solutions that are both efficient and fair. As such, a mechanism is needed that employs some sort of external authority. In this paper, we employ simple monitoring authorities (MA), with the primary goal to keep track of successful accesses. We do not require any planning capabilities by the MAs, nor any knowledge of the plans and preferences of the participants; only the ability to monitor successful accesses on their respective resource. Such MAs already exist naturally in many domains (e.g. the port authority in maritime traffic management [1], or the access point in wireless networks, etc.), and can be easily introduced in the rest (e.g. an agent monitoring each charging / parking station, etc.). At the end, agents learn to anti-coordinate their actions in a decentralized manner, while the introduced MAs allow us to impose quotas to the use of resources and punishments upon violating them to align individual incentives with an efficient and fair correlated equilibrium.

A second aspect of anti-coordination problems involves convergence speed. Inter-agent interactions often need to take place in an ad-hoc fashion. Typical approaches (e.g. Monte Carlo algorithms, Bayesian learning, bandit algorithms, etc.) tend to require too many rounds to converge to be feasible in dynamic environments and real-life applications. Yet, humans are able to routinely coordinate in such an ad-hoc fashion. One key concept that facilitates human ad-hoc (anti-)coordination is the use of *conventions* [24]. Behavioral conventions are a fundamental part of human societies, yet they have not appeared meaningfully in empirical modeling of MAS. Inspired by human behavior, we propose the adoption of a simple convention of *courtesy*. Courtesy arises by the social nature of humans. Society demands that an individual should conduct himself in consideration of others. This allows for fast convergence, albeit

it is not game theoretically sound; people adhere to it due to social pressure. Such problems become even more severe in situations with scarcity of resources. Under such conditions, courtesy breaks down and in the name of self-preservation people exhibit urgency and competitive behavior [20]. Thus, to satisfy our rationality constraint (no incentive to deviate for self-interested agents) in an artificial system, we need a deterrent mechanism.

In this paper we present a framework (CA³NONY: Contextual Anti-coordination in Ad-hoc **A**nonymous games) for repeated allocation games with discounting. CA³NONY reproduces courtesy based on the allocation algorithm of [13], and uses monitoring and punishments as a deterrent mechanism. It exhibits *fast convergence*, and high *efficiency* and *fairness*, while relying only on occupancy feedback which facilitates scalability and does not require any inter-agent communication. The main contributions are:

- Introduction of an anti-coordination framework (CA³NONY) which consists of a *courteous convention* and a *monitoring scheme*.
- Proof that under such a framework, the use of the courteous convention induces strategies that constitute an approximate *subgame-perfect equilibrium*.
- Comparison to state-of-the-art bandit algorithms.

2 RELATED WORK

In ad-hoc multi-agent (anti-)coordination the goal is to design autonomous agents that achieve high flexibility and efficiency in a setting that admits no prior coordination between the participants [33]. Typical scenarios include the use of Monte Carlo algorithms [6], Bayesian learning [2], or bandit algorithms [7, 11]. Traditionally, pure ad-hoc approaches suffer from slow learning [15], which makes pure ad-hoc coordination a very ambitious goal for real-life applications. In this paper we propose a middle-ground approach. Inspired by human ad-hoc coordination, we incorporate prior knowledge in the form of simple *conventions*. The coordination can still be considered ad-hoc as it is not pre-programmed, rather it involves learning. This allows for faster convergence compared to pure ad-hoc approaches.

A convention is defined as a customary, expected and self-enforcing behavioral pattern [24, 37]. In MAS, there are two scopes through which we study conventions. First, a convention can be considered as a behavioral rule, designed and agreed upon ahead of time or decided by a central authority [32, 35]. Second, a convention may emerge from within the system itself through repeated interactions [26, 35]. The proposed courteous convention falls on the first category. It is incorporated as prior knowledge, and it is self-enforcing since the induced strategies constitute an approximate subgame-perfect equilibrium of the repeated allocation game.

An alternative way to model the anti-coordination problem is as a multi-armed bandit (MAB) problem [3]. In MAB problems an agent is given a number of arms and at each time-step has to decide which arm to pull to get the maximum expected reward. Bandit (or no-regret) algorithms typically minimize the total regret of each agent, which is the difference between the expected received payoff and the payoff of the best strategy in hindsight. As such, they satisfy our rationality constraint since they constitute an approximate correlated or coarse correlated equilibrium [21, 27, 30]. However,

the studied problem presents many challenges: there is no stationary distribution (adversarial rewards), all agents are able to learn (similar to recursive modeling) which results to a moving-target problem, and yielding gives a reward of 0 (desirable option for minimizing regret, but not in respect to fairness). Moreover, regret minimization does not necessarily lead to payoff maximization [15]. Nevertheless, due to their ability to learn from partial feedback, bandit algorithms constitute the natural choice for a pure ad-hoc approach. The latter motivates our choice to use them as a baseline, since our agents only receive binary feedback of success or failure upon taking their action.

Game theoretic equilibria are desirable (since they satisfy the rationality constraint), but hard to obtain. Deciding whether an anti-coordination (anonymous) game has a pure Nash equilibrium (NE) is NP-complete [10]. Furthermore, allocation games often admit undesirable equilibria: pure NE which are efficient but not fair, or mixed-strategy NE which are fair but not efficient [13]. Hence, iterative best-response algorithms are not satisfactory. On the other hand, correlated equilibria (CE) [5] can be both efficient & fair, while from a practical perspective they constitute perhaps the most relevant non-cooperative solution concept [21]. An optimal CE of an anonymous game may be found in polynomial time [29]. Moreover, it is possible to achieve a CE without a correlation device (central agent) [18, 21], using the history of the actions taken by the opponents. However, we are interested in information-restrictive learning rules (i.e. completely uncoupled [34]), where each agent is only aware of his own history of action/reward pairs. Such an approach was applied in [13] to design a decentralized algorithm for reaching efficient & fair CE in wireless channel allocation games. Yet, while the algorithm reaches an equilibrium in a polynomial number of steps, cooperation to achieve this state is not rational. A self-interested agent could keep accessing a resource forever, until everyone else backs off (also known as ‘bully’ strategy [25]). In this paper, we build upon the ideas of Cigler and Faltings and develop an anti-coordination strategy that constitutes an approximate subgame-perfect equilibrium, i.e. cooperation with the algorithm is a best-response strategy at each sub-game of the original stage game, given any history of the play.

A generalization of anti-coordination games, called dispersion games, was described in [19]. In a dispersion game, agents are able to choose from several actions, favoring the one that was chosen by the smallest number of agents (analogous to minority games [12]). The authors in [19] define a maximal dispersion outcome as an outcome where no agent can switch to an action chosen by fewer agents. The agents themselves do not have any particular preference for the attained equilibrium. Contrary to that, we are interested in achieving an efficient and fair outcome. Besides, in many real world applications, the agents are indifferent to which role/task/resource they attain, as long as they receive one (e.g. wireless frequencies). Tackling dispersion games, and therefore non-binary utilities, remains open for future research.

A similar approach to ours was introduced for weighted matching in [17]. The authors propose a heuristic which is decentralized, requires only partial feedback, and has *constant* in the total problem size running time, under reasonable assumptions on the preference domain of the agents. However, it is applicable in cooperative scenarios, and not in the presence of strategic agents.

3 THE CA³NONY FRAMEWORK

3.1 The Repeated Allocation Game

Let a ‘resource’ be any element that can be successfully assigned to only one agent at a time. At each time-step, $\mathcal{N} = \{1, \dots, N\}$ agents try to access $\mathcal{R} = \{1, \dots, R\}$ identical and indivisible resources, where possibly $N \gg R$. The set of available actions is denoted as $\mathcal{A} = \{Y, A_1, \dots, A_R\}$, where Y refers to yielding and A_r refers to accessing resource r . We assume that access to a resource is slotted and of equal duration¹. A successful access yields a positive payoff, while no access has a payoff of 0. If more than one agent access a resource simultaneously, a collision occurs and the colliding parties incur a cost $\zeta < 0$. Thus, the agents only receive a binary feedback of success or failure. Let a_n denote agent n ’s action, and $a_{-n} = \times_{n' \in \mathcal{N} \setminus \{n\}} a_{n'}$ the joint action for the rest of the agents. The payoff function is defined as:

$$u_n(a_n, a_{-n}) = \begin{cases} 0, & \text{if } a_n = Y \\ 1, & \text{if } a_n \neq Y \wedge a_i \neq a_n, \forall i \neq n \\ \zeta, & \text{otherwise} \end{cases} \quad (1)$$

We assume that rewards are discounted by $\delta \in (0, 1)$, and, conforming to real-world scenarios, that each agent n is only aware of his own history of action/reward pairs.

Finally, we assume that the agents can observe side information from their environment at each time-step t . We call this side information context (e.g. time, date etc.). The agents utilize this context as a common signal in their decision-making process, a means to learn and anti-coordinate their actions. Let $\mathcal{K} = \{1, \dots, K\}$ denote the context space. The rationale behind the introduction of the common context is that, under completely uncoupled learning rules, having positive probability mass on undesirable actions (e.g. collisions) is unavoidable. Moreover, from a practical perspective, common environmental signals are amply available to the agents [21]. We do not assume any a priori relation between the context space and the problem. The only constraints are that the values should repeat periodically, and satisfy $K = \lceil N/R \rceil$.

The context signals could be produced either on the resource side (e.g. by the port authority in maritime traffic management, or the access point in wireless networks), or in a decentralized manner (e.g. in distributed networks with no authorities the senders can attach identifier signals to data traffic [36]). Finally, in situations where communication is possible, the agents can agree upon the signal themselves by solving the distributed consensus problem.

3.2 Adopted Convention

The adopted convention is based on the cooperative allocation algorithm of [13]. Each agent n has a strategy $g_n : \mathcal{K} \rightarrow \mathcal{A}$ that determines a resource to access at time-step t after having observed context k_t . The strategy is initialized uniformly at random in \mathcal{A} . If $g_n(k_t) = A_r$, then agent n accesses resource r . Otherwise, if $g_n(k_t) = Y$, the agent does not access a resource but instead chooses uniformly at random a resource r to monitor for activity. If it is free, then the agent updates $g_n(k_t) \leftarrow A_r$ (see Alg. 1).

¹This is done to facilitate the proofs. Real world problems have this property anyway (e.g. access to radio channels, role allocation in a plan, etc.), and the algorithms we compare to all work for slotted resources.

Algorithm 1 Learning rule.

Require: Initialize g_n u.a.r. in \mathcal{A} . Set $accessed_n \leftarrow False$.

```

1: for  $k_t \in \mathcal{K}$  do
2:   Agents observe context  $k_t$ 
3:   if  $g_n(k_t) = A_r$  &  $accessed_n = False$  then
4:     Agent  $n$  accesses resource  $r$ 
5:     if Collision( $r$ ) then
6:       Set  $g_n(k_t) \leftarrow Y$  with probability  $p_{backoff} > 0$ 
7:     else
8:       Set  $accessed_n \leftarrow True$ 
9:     end if
10:  else if  $g_n(k_t) = Y$  then
11:    Agent  $n$  monitors random resource  $r \in \mathcal{R}$ 
12:    if Free( $r$ ) then
13:      Set  $g_n(k_t) \leftarrow A_r$ 
14:    end if
15:  end if
16: end for
17: Set  $accessed_n \leftarrow False$ 

```

In [13], agents back-off probabilistically in case of a collision (set $g_n(k_t) \leftarrow Y$ with probability $p_{backoff}$). In such a setting, it is possible to reach a symmetric subgame-perfect equilibrium. But in order to actually play it, the agents need to be able to calculate it. It is not always possible to obtain the closed form of the back-off probability distribution of each resource. Furthermore, a self-interested agent could stubbornly keep accessing a resource forever, until everyone else backs off (‘bully’ strategy [25]).

Instead, we adopted a simple convention where agents are being *courteous*, i.e. if there is a collision, the colliding agents will back-off with some constant positive probability: $p_{backoff} = p > 0, \forall n \in \mathcal{N}$. Being courteous, though, does not satisfy the rationality constraint. However, a uniform distribution of resources is socially optimal (i.e. fair allocations maximize the social welfare). Hence, if we introduce quotas to the resources and punishments upon violating them, courtesy induces rational strategies. In the following sections we introduce a monitoring scheme and prove that the resulting strategy constitutes an ϵ -subgame-perfect equilibrium.

3.3 Rationality

In order to ensure the proposed convention’s rationality, the agents must be assured that they will eventually be successful, i.e. we must provide safeguards against the monopolization of resources. The proposed framework employs simple Monitoring Authorities (MA). The MAs do not require any planning capabilities, nor any knowledge of the agents’ plans and preferences. Their primary goal is to keep track of successful accesses. Depending on the domain, MAs may already exist naturally, or can easily be introduced. Examples include centralized MAs like the port authority in maritime traffic management, or a set of decentralized MAs (one per resource) like the access points in wireless networks, or agents monitoring a charging / parking station. Their function is twofold. First, they could provide occupancy signals. Agents (e.g. CA³NONY, bandit, or Q-learning agents) must be able to receive some form of feedback from their environment to inform collisions and detect free

resources. This could be achieved (if possible) by the use of various sensors, or by receiving occupancy signals (e.g. 0, 1) from the MAs. Second, (which is their principal purpose) MAs deter agents from monopolizing resources to the point that each agent can access a resource only for one context value out of K . To achieve the latter, MAs must be able to keep track of successful accesses. Upon the violation of the imposed quotas, the framework is responsible to enforce the necessary punishments. Punishments are application specific. They can be individual (e.g. in a wireless scenario, if an agent transmits to some other than the designated channel, then his packets will no longer be relayed), or group punishments (e.g. if quotas are exceeded, access is denied for everyone). The imposed punishments make exceeding the quotas an irrational strategy (simulated in Alg. 1 by the $accessed_n$ flag), which in turn aligns the individual incentives with an efficient & fair correlated equilibrium.

3.3.1 Access Monitoring. The primary function of the employed MAs is the tracking of successful accesses by the agents. The latter can be achieved in various ways, depending on the domain. The simplest one would be to maintain a log with successful accesses per episode (period of context values), whether we are dealing with a centralized MA or set of decentralized MAs that are able to communicate to ensure coherence. For more involved scenarios (e.g. if we require agent anonymity per access) we may employ more complex schemes, e.g. using tokens, or artificial currency (solely as an internal mechanism). Note that the latter does not fence the resources (i.e. punishments are still required). In what follows, we present such a self-regulated monitoring scheme.

Initially all the agents that 'buy-in' are issued the same amount $m \in \mathbb{R}$ of artificial cash (AC). This amount also corresponds to the initial fee for every resource $f_r \leftarrow m$. To allow access to resource r , the MA of r charges f_r units of AC, and monitors the event. If there was a successful access, the MA reimburses the amount of $(1 - \xi)f_r$ AC to the accessing agent, where $\xi \rightarrow 0 \in \mathbb{R}$ is a commission fee. Otherwise, the MA reimburses the full amount of f_r AC to the colliding agents, so that they are able to try again for a different context value. Finally, after each episode, the MAs lower the fee to $f'_r \leftarrow (1 - \xi)f_r, \forall r \in \mathcal{R}$. After every successful access, the amount of AC that an agent possesses drops below the access fee of a resource. Waiting for the fee to drop to the point that $f_r = m/2$ is not rational since, assuming $\xi \rightarrow 0$, the number of iterations required to allow accessing two resources at the same time will reach ∞ . At that point the rest of the agents will have reached a correlated equilibrium and the adversarial agent will not have an incentive to access an additional resource, besides the one that corresponds to him, since it would result in a collision. If, due to implementation constraints, we can not select a small ξ , the MAs can change the artificial currency every I episodes, invalidating the old one and again making such strategy irrational.

3.4 Rate of convergence

THEOREM 3.1. *In a repeated allocation game with N agents and R resources, the expected number of steps before Alg. 1 converges to a correlated equilibrium is bounded by:*

$$O\left(N\left(\log\left\lceil\frac{N}{R}\right\rceil+1\right)(\log N+R)\right) \quad (2)$$

PROOF. The adopted learning rule is based on the allocation algorithm of [13]. For $N, R, K \geq 1$, and back-off probability $0 < p < 1$, the expected number of steps before the algorithm converges is bounded by (3)², which for a constant back-off probability and $K = \lceil N/R \rceil$ gives the required bound.

$$O\left(\left(K \log K + 2K\right) R \frac{2-p}{2(1-p)} \left(\frac{1}{p} \log N + R\right)\right) \quad (3)$$

□

COROLLARY 3.2. *Under a common, constant back-off probability assumption, $p^* = 2 - \sqrt{2}$ minimizes the convergence time of Alg. 1 in high congestion scenarios, i.e. $N \gg R$.*

PROOF. According to bound (3), in high congestion scenarios (i.e. $\frac{N}{R} = K \rightarrow \infty$), the common, constant back-off probability that minimizes the convergence time is:

$$p^* = \arg \min \left(\frac{2-p}{2(1-p)} \frac{1}{p} \right) = 2 - \sqrt{2} \quad (4)$$

□

3.5 Courtesy Pays Off

In this section we prove that if the agents back-off with a constant positive probability $p_{backoff} > 0$, then Alg. 1 induces a strategy that is an ϵ -equilibrium.

Suppose that in a repeated allocation game with discounting ($\delta \in (0, 1)$) the agents who collide back-off with a constant probability $p_{backoff} > 0$. Let σ^P denote the aforementioned strategy (courteous strategy), and σ^* denote the optimal (best-response) strategy under the monitoring authorities (possibly better than a stage game NE). Moreover, let $U_n(\sigma, \sigma_{-n}, \delta) = \sum_{t=0}^{\infty} \delta^t u_n(a_n^t, a_{-n}^t)$ denote the cumulative payoff of agent n following strategy σ , assuming the rest of the agents follow the strategy σ_{-n} . The following theorem proves that starting at any time-step, agent n does not gain more than ϵ by deviating to the optimal strategy σ^* .

THEOREM 3.3. *Under a high enough discount factor, the courteous strategy σ^P constitutes an approximate subgame-perfect equilibrium, i.e. $\forall \epsilon > 0, \exists \delta_0 \in (0, 1)$ such that $\forall \delta, \delta_0 \leq \delta < 1$:*

$$\mathbb{E}[U_n(\sigma_n^P, \sigma_{-n}^P, \delta)] > (1 - \epsilon) \mathbb{E}[U_n(\sigma_n^*, \sigma_{-n}^P, \delta)]$$

PROOF. Note that $\mathbb{E}[U_n(\sigma_n^*, \sigma_{-n}^*, \delta)] \geq \mathbb{E}[U_n(\sigma_n^*, \sigma_{-n}^P, \delta)]$, since by playing σ_{-n}^P the rest of the agents can potentially introduce additional collisions. Thus, it suffices to prove that $\mathbb{E}[U_n(\sigma_n^P, \sigma_{-n}^P, \delta)] > (1 - \epsilon) \mathbb{E}[U_n(\sigma_n^*, \sigma_{-n}^*, \delta)]$.

The introduced monitoring scheme prohibits the monopolization of resources, i.e. each agent can only access a resource for his corresponding context value. Thus, the best-response strategy's (σ^*) payoff for some δ is bounded by:

$$\mathbb{E}[U_n(\sigma_n^*, \sigma_{-n}^*, \delta)] \leq 1 + \delta^K + \delta^{2K} + \dots = \sum_{i=0}^{\infty} \delta^{iK} = \frac{1}{1 - \delta^K} \quad (5)$$

²Slightly tighter bound on the convergence speed of the adopted learning rule, based on Theorems 12 and 13 of [13]. See the appendix for the proof.

When agents adopt the courteous convention, in each round until the system converges to a correlated equilibrium, the agents receive a payoff between $\zeta < 0$ (collision cost) and 1. Thus, until convergence, the expected payoff is lower bounded by $\zeta \sum_{i=0}^{\tau-1} \delta^{iK} = \zeta \frac{1-\delta^{\tau K}}{1-\delta^K}$, where τ is the number of steps to converge. After convergence, their expected payoff is $\sum_{i=0}^{\infty} \delta^{\tau+iK} = \frac{\delta^\tau}{1-\delta^K}$. Hence, the convention induced strategy's payoff is at least:

$$\mathbb{E}[U_n(\sigma_n^p, \sigma_{-n}^p, \delta)] \geq \sum_{\tau=1}^{\infty} \Pr[\text{conv. in } \tau \text{ steps}] \cdot \left(\frac{\zeta(1-\delta^{\tau K}) + \delta^\tau}{1-\delta^K} \right)$$

We can define a random variable X such that $X = \tau$ if the algorithm converges after exactly τ steps. Since δ^x is a convex function we have that $\mathbb{E}(\delta^x) \geq \delta^{\mathbb{E}(x)}$, therefore:

$$\mathbb{E}[U_n(\sigma_n^p, \sigma_{-n}^p, \delta)] \geq \frac{\zeta(1-\delta^{\mathbb{E}(X)K}) + \delta^{\mathbb{E}(X)}}{1-\delta^K} \quad (6)$$

By dividing (6) by (5) we get:

$$\frac{\mathbb{E}[U_n(\sigma_n^p, \sigma_{-n}^p, \delta)]}{\mathbb{E}[U_n(\sigma_n^*, \sigma_{-n}^*, \delta)]} \geq \zeta(1-\delta^{\mathbb{E}(X)K}) + \delta^{\mathbb{E}(X)} \quad (7)$$

$\mathbb{E}(X)$ does not depend on δ . Moreover, $\delta^{\mathbb{E}(X)}$ is continuous in δ , monotonous, and $\lim_{\delta \rightarrow 1^-} \delta^{\mathbb{E}(X)} = 1$. Thus, we can take the limit of

$$(7) \text{ as } \delta \rightarrow 1^-, \text{ which equals to } \lim_{\delta \rightarrow 1^-} \frac{\mathbb{E}[U_n(\sigma_n^p, \sigma_{-n}^p, \delta)]}{\mathbb{E}[U_n(\sigma_n^*, \sigma_{-n}^*, \delta)]} = 1 \quad \square$$

In order to guarantee rationality, the discount factor δ must be close to 1 since, as δ gets closer to 1, the agents do not care whether they access now or in some future round. Since the proposed monitoring scheme guarantees that every agent will access a resource for his corresponding context value, when $\delta \rightarrow 1$, the expected payoff for agents who are accessing a resource and for those who have not accessed a resource yet will be the same. In other words, the cost (overhead) of learning the correlated equilibrium decreases.

3.6 Indifference Period

In many real world applications, agents are indifferent in claiming a resource in a period of T_{ind} rounds, i.e. $\delta_t = 1, \forall t \leq T_{ind}$. E.g. data of wireless transmitting devices might remain relevant for a specific time-window, during which the agent is indifferent of transmitting. In such cases, we can use the Markov bound to prove that with high probability the proposed algorithm will converge in under T_{ind} time-steps, thus satisfying the rationality constraint. We assume the agents are willing to accept linear 'delay' with regard to the number of resources R , the number of agents N , and the size of the context space K , specifically:

$$T_{ind} = O(RNK) \quad (8)$$

THEOREM 3.4. *Under a linear indifference period T_{ind} (i.e. $\delta_t = 1, \forall t \leq T_{ind} = O(RNK)$) the probability of the system of agents following Alg. 1 not having converged during T_{ind} diminishes as the congestion increases ($\frac{N}{R} \rightarrow \infty$).*

PROOF. Using the Markov bound, it follows that the probability that the system takes more than the accepted number of steps (T_{ind}) to converge is:

$$\Pr[-\text{conv. after } T_{ind}] = O\left(\frac{\left(\log\left\lceil\frac{N}{R}\right\rceil + 1\right)(\log N + R)}{N}\right)$$

Taking the limit: $\lim_{\frac{N}{R} \rightarrow \infty} \Pr[-\text{conv. after } T_{ind}] = 0. \quad \square$

Even though the Markov's inequality generally does not give very good bounds when used directly, Th. 3.4 proves that our algorithm converges in the required time with high probability. The latter holds under high congestion, which constitutes the more interesting scenario since for small number of agents the required time to converge is only a few hundreds of time-steps. Moreover, the higher the indifference period, the higher the probability to converge under such time-constraint. For quasilinear indifference period for example, the system converges in the required time-window with high probability even for a small number of agents. We can further strengthen our rationality hypothesis by using a tighter bound (e.g. Chebyshev's inequality), albeit computing the theoretical variance of the convergence time is an arduous task, thus it remains open for future work.

4 EVALUATION

In this section we model the resource allocation problem as a multi-armed bandit problem and provide simulation results of CA³NONY's performance in comparison to state-of-the-art, well established bandit algorithms, namely the EXP4 [4], EXP4.P [8], and EXP3 [4]. In every case we report the average value over 128 runs of the same simulation. For the EXP family of algorithms, the input parameters are set to their optimal values, as prescribed in the aforementioned publications. We assume a reward of 1 for a successful access, -1 if there is a collision, and 0 if the agent yielded.

4.0.1 Level of Courtesy. We evaluated different back-off probabilities ($p_{backoff} \in \{0.1, 0.25, 2 - \sqrt{2}, 0.75, 0.9\}$) for $R = K \in \{2, 4, 8, 16\}$ and $N = R \times K$. There was no significant difference in convergence time, but since $p_{backoff}$ is directly correlated with the number of collisions, it can have significant impact on the average payoff (see Table 1). It is important to note, though, that agents do not have global nor local knowledge of the level of congestion of the system (they only receive binary occupancy feedback for one resource per time-step). Thus, agents can not select the optimal back-off probability depending on congestion. In what follows, the back-off probability of CA³NONY is set to $p_{backoff} = 2 - \sqrt{2}$ of Eq. 4, which is only optimal under high congestion settings (i.e. $\frac{N}{R} = K \rightarrow \infty$), yet constitutes a safe option. Note that the provided theoretical analysis holds for any constant back-off probability.

4.0.2 Bandits & Monitoring. CA³NONY has three meta actions (Access, Yield, and Monitor $\{A, Y, M\}$), while bandit algorithms have only two ($\{A, Y\}$), and CA³NONY assumes the existence of monitoring authorities (MAs). For fairness' sake in the reported results, we made the following two modifications. First, we include

Table 1: Avg. payoff depending on the level of courtesy (back-off probability), $K = R, N = R \times K$.

$p_{backoff}$	$R = 2$	$R = 4$	$R = 8$	$R = 16$
0.1	37.6	-4.7	-39.5	-63.1
0.25	43.6	10.0	-16.9	-40.0
$2 - \sqrt{2}$	45.7	15.6	-5.2	-23.0
0.75	45.5	15.7	-3.7	-20.6
0.9	44.2	12.5	-5.8	-19.4

a variation of CA³NONY, denoted as CA³NONY*, where we assume agents incur a cost (equal to collision cost ζ) every time they monitor a resource (reflected in Table 3). Second, all the employed bandit algorithms make use of the MAs, which indirectly grants them the ability to monitor resources as well. More specifically, the accumulated payoff is updated only if they are allowed to access a resource. If not, we consider it a monitoring action, which means the agents still receive occupancy feedback. The latter is important, otherwise the bandit algorithms would require significantly longer time to converge. Note that, contrary to CA³NONY*, monitoring is free for bandit algorithms with respect to the accumulated payoff, i.e. they do not incur a collision cost.

4.1 Employed Bandit Algorithms

The reward of each arm does not follow a fixed probability distribution (adversarial setting). Moreover, the agents are able to observe side-information (context) at each time-step t . The arm that yields the highest expected reward can be different depending on the context. Hence we focus on adversarial contextual bandit algorithms (see [38] for a survey). A typical approach is to use expert advice. In this method we assume a set of experts $\mathcal{M} = \{1, \dots, M\}$ who generate a probability distribution on which arm to pull depending on the context. A no-regret algorithm performs asymptotically as well as the best expert. Such algorithms are the EXP4 and EXP4.P, the difference being that EXP4 exhibits high variance [38], while EXP4.P achieves the same regret with high probability by combining the confidence bounds of UCB1 [3] and EXP4. The computational complexity and memory requirements of the above algorithms are linear in M , making them intractable for large number of experts. In order to deal with the increased complexity in larger simulations, we gave an edge to these algorithms by restricting the set of experts \mathcal{M} to the same uniform correlated equilibria (CE) that CA³NONY converges to. The latter enabled us to perform larger simulations (the unrestricted versions could not handle $N \geq 64$), while resulting in the same order of magnitude convergence time, slightly higher average payoff, and significantly higher fairness ($> 35\%$ since by design the experts proposed fair CE). Alternatively, we can use non-contextual adversarial bandit algorithms, such as the EXP3. Moreover, we can convert EXP3 to a contextual algorithm by setting up a separate instance for all $k \in \mathcal{K}$. We call this ‘CEXP3’. This results in a contextual bandit algorithm which has the edge over EXP4 from an implementation viewpoint since its running time at each time-step is $O(R)$ and its memory requirement is $O(KR)$ (CA³NONY’s is $O(1)$ and $O(K)$ respectively).




Figure 1: Utilization: $R = 4, N = 16$ (x -axis in log scale).

4.2 Simulation Results

4.2.1 Convergence Speed & Efficiency. We know that CA³NONY converges to a CE which is efficient. If all agents follow the courteous convention of Alg. 1, the system converges to a state where no resources remain un-utilized and there are no collisions (Th. 13 of [13]). Furthermore, Th. 3.1 argues for fast convergence. The former are both corroborated by Fig. 1 (similar results ($> \times 10^2$ faster convergence) were acquired for $R \in \{2, 8, 16\}$ as well.). Fig. 1 depicts the total utilization of resources for a simulation period of $T = 10^6$ time-steps. Note that the x -axis is in logarithmic scale. CA³NONY converges significantly ($> \times 10^2$) faster than the bandit algorithms to a state of 100% efficiency. On the other hand, the bandit algorithms exhibit high variance, never achieve 100% efficiency, and are not able to handle efficiently the increase in context space size and number of resources.

4.2.2 Fairness. The usual predicament of efficient equilibria for allocation games is that they assign the resources only to a fixed subset of agents, which leads to an unfair result (e.g. an efficient pure NE (PNE) is for R agents to access and $N - R$ agents to yield). This is not the case for CA³NONY, which converges to an equilibrium that is not just efficient but fair as well. Due to the enforced monitoring scheme, all agents acquire the same amount of resources. As a measure of fairness, we will use the Jain index [22]. The Jain index is widely used by economists, and exhibits a lot of desirable properties such as: population size independence, continuity, scale and metric independence, and boundedness. For an allocation game of N agents, such that the n^{th} agent receives an allocation of x_n , the Jain index is given by $\mathbb{J}(\mathbf{x}) = \frac{|\sum_{n \in \mathcal{N}} x_n|^2}{N \sum_{n \in \mathcal{N}} x_n^2}$. An allocation is considered fair, iff $\mathbb{J}(\mathbf{x}) = 1$, $\mathbf{x} = (x_1, \dots, x_N)^T$.

Table 2 presents the expected Jain Index of the evaluated algorithms at the end of the time horizon T . CA³NONY converges to a fair equilibrium, achieving a Jain index of 1. The EXP4 and EXP4.P were the fairest amongst the bandit algorithms, achieving a Jain index of close to 1. This is to be expected since the set of experts \mathcal{M} is limited to the same set of equilibria that CA³NONY converges to. On the other hand, EXP3 and CEXP3 performed considerably worse, with the EXP3 exhibiting the worst performance in terms of fairness, equal to a PNE’s: $\mathbb{J}_{PNE}(\mathbf{x}) = \frac{R^2}{NR} = \frac{1}{K} = \mathbb{J}_{EXP3}(\mathbf{x})$.

4.2.3 Average Payoff. The average payoff corresponds to the total discounted payoff an agent would receive in the time horizon T . Table 3 presents the average payoff for the studied algorithms. The clustered pairs of bandit algorithms exhibited $< 1\%$ difference, hence we included the average value of the pair. The discount factor was set to $\delta = 0.99$. At the top half we do not assume any

Table 2: Fairness (Jain Index), $K = R, N = R \times K$.

	$R = 2$	$R = 4$	$R = 8$	$R = 16$
CA ³ NONY	1.0000	1.0000	1.0000	1.0000
EXP3	0.5000	0.2500	0.1250	0.0625
CEXP3	0.7018	0.5865	0.5341	0.9638
EXP4	1.0000	0.9999	0.9959	0.9198
EXP4.P	1.0000	0.9999	0.9798	0.7503

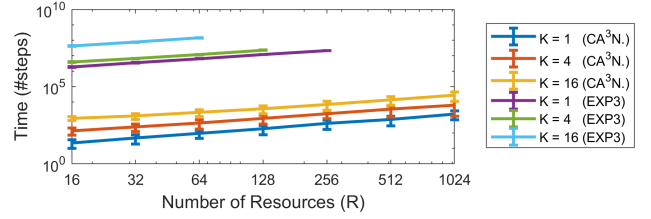

Table 3: Average Payoff, $K = R, N = R \times K$.

	$R = 2$	$R = 4$	$R = 8$	$R = 16$
CA ³ NONY	45.4	15.7	-5.3	-23.0
CA ³ NONY*	20.0	-22.4	-50.6	-72.6
(C)EXP3	-72.4	-93.1	-99.8	-100.0
EXP4(P)	-59.0	-81.0	-90.7	-95.4
CA ³ NONY	53.3	79.0	502.5	4057.5
CA ³ NONY*	23.8	-55.5	-1337.1	-26723.9
(C)EXP3	-84.0	-331.0	-4175.7	-60595.1
EXP4(P)	-68.4	-288.4	-3807.1	-62631.5

indifference period, while at the bottom half we assume linear indifference period ($\delta_t = 1, \forall t \leq T_{ind}$, where T_{ind} is given by Eq. 8). Recall that CA³NONY and the bandit algorithms do not incur any cost for monitoring resources, while CA³NONY* incurs a cost equal to the collision cost ζ .

Once more, CA³NONY (and even CA³NONY*) significantly outperforms all the bandit algorithms. The latter have relatively similar performance, with EXP4 and EXP4.P being the best amongst them. It is worth noting that adding an indifference period has a dramatic effect on the results. Comparing CA³NONY to bandit algorithms we observe that CA³NONY achieves a large increase on average payoff, while the opposite happens for the bandit algorithms. This is because the learning rule of Alg. 1 prohibits from accessing an already claimed resource, thus minimizing collisions. On the contrary, bandit algorithms constantly explore (they assign a positive probability mass to every arm) which leads to collisions. In a multi-agent system where every agent learns this can have a cascading effect. The latter becomes apparent when fixing $\delta = 1$ for T_{ind} steps. The collision cost remains high for longer which, as seen by Table 3, has a significant impact on the bandit algorithms' performance.

4.2.4 Large Scale Systems. The innovation of CA³NONY stems from the adoption of a simple convention, which allows its applicability to large scale MAS. To evaluate the latter, Fig. 2 and 3 depict the convergence time for increasing number of resources R , and increasing system congestion ($\frac{N}{R} = K$) respectively. Both graphs are in a double logarithmic scale, and the error bars represent one standard deviation of uncertainty. The total number of agents is given by $N = R \times K$. Thus, the largest simulations involve 16.384 agents. Along with CA³NONY, we depict the fastest (based on the previous simulations) of the bandit algorithms, namely EXP3. In both cases we acquire $\times 10^3 - \times 10^5$ faster convergence. The above validate CA³NONY's performance in both scenarios with abundance ($N \approx R$ or small K), and scarcity of resources ($N \gg R$ or large K). As depicted, CA³NONY is significantly faster than the

**Figure 2: Convergence time: increasing #resources R , varying context space size $K, N = R \times K$ (double log scale).****Figure 3: Convergence time: increasing congestion ($\frac{N}{R} = K$), varying #resources $R, N = R \times K$ (double log scale).**

EXP3 and can gracefully handle increasing number of resources, and high congestion. Finally note that, in several of the simulations, EXP3 was unable to reach its convergence goal of 90% efficiency (utilization of resources) in a reasonable amount of computation time (1.5×10^8 time-steps), hence the resulting gaps in EXP3's lines in Fig. 2 and 3. Especially in situations with scarcity of resources the utilization was significantly lower.

5 CONCLUSION

In this paper we proposed CA³NONY, an anti-coordination framework under rationality constraints. It is based on a simple, human-inspired convention of courtesy which prescribes a positive back-off probability in case of a collision. Coupled with a monitoring scheme which deters the monopolization of resources, we proved that the induced strategy constitutes an ϵ -subgame-perfect equilibrium. We compared CA³NONY to state-of-the-art bandit algorithms, namely EXP3, EXP4, and EXP4.P. Simulation results demonstrated that CA³NONY outperforms these algorithms by achieving more than two orders of magnitude faster convergence, a fair allocation, and higher average payoff. The aforementioned gains suggest that human-inspired conventions may prove beneficial in other ad-hoc coordination scenarios as well.

A APPENDIX

A.1 Proof of Bound (3)

In this section we provide a formal proof of bound (3). The proof is an adaptation of the convergence proof of [13]. We will prove the following theorem:

THEOREM A.1. For N agents and $R \geq 1, K \geq 1, 0 < p < 1$ the expected number of steps before the learning algorithm converges to an efficient correlated equilibrium of the allocation game for every $k \in \mathcal{K}$ is $O\left((K \log K + 2K) R \frac{2-p}{2(1-p)} \left(\frac{1}{p} \log N + R\right) + 1\right)$.

We begin with the case of having multiple agents but only a single resource ($R = 1$). We will describe the execution of the proposed learning rule as a discrete time Markov chain (DTMC)³. In every time-step, each agent performs a Bernoulli trial with probability of ‘success’ $1 - p$ (remain in the competition), and failure p (back-off). When N agents compete for a single resource, a state of the system is a vector $\{0, 1\}^N$ denoting the individual agents that still compete for that resource. But, since the back-off probability is the same for everyone, we are only interested in how many agents are competing and not which ones. Thus, in the single resource case ($R = 1$), we can describe the execution of the proposed algorithm using the following chain:

Definition A.2. Let $\{X_t\}_{t \geq 0}$ be a DTMC on state space $S = \{0, 1, \dots, N\}$ denoting the number of agents still competing for the resource. The transition probabilities are as follows:

$$\begin{aligned} Pr(X_{t+1} = N | X_t = 0) &= 1 && \text{restart} \\ Pr(X_{t+1} = 1 | X_t = 1) &= 1 && \text{absorbing} \\ Pr(X_{t+1} = j | X_t = i) &= \binom{i}{j} p^{i-j} (1-p)^j && i > 1, j \leq i \end{aligned}$$

(all the other transition probabilities are zero)

Intuitively, this Markov chain describes the number of individuals in a decreasing population, but with two caveats: The goal (absorbing state) is to reach a point where only one individual remains, and if we reach zero, we restart.

LEMMA A.3. [13] Let $A = \{0, 1\}$. The expected hitting time of the set of states A in the Markov chain described in Definition A.2 is $O\left(\frac{1}{p} \log N\right)$.

Definition A.4. Let $\{Y_t\}_{t \geq 0}$ be a DTMC on state space $S = \{0, 1, \dots, N\}$ with the following transition probabilities (two absorbing states, 0 and 1):

$$\begin{aligned} Pr(Y_{t+1} = 0 | Y_t = 0) &= 1 && \text{absorbing} \\ Pr(Y_{t+1} = 1 | Y_t = 1) &= 1 && \text{absorbing} \\ Pr(Y_{t+1} = j | Y_t = i) &= \binom{i}{j} p^{i-j} (1-p)^j && i > 1, j \leq i \end{aligned}$$

(all the other transition probabilities are zero)

Let h_i^A denote the hitting probability of a set of states A , starting from state i . We will prove the following lemma.

LEMMA A.5. The hitting probability of the absorbing state $\{1\}$, starting from any state $i \geq 1$, of the DTMC $\{Y_t\}$ of Definition A.4 is given by Eq. 9. This is a tight lower bound.

$$h_i^{\{1\}} = \Omega\left(\frac{2(1-p)}{2-p}\right), \forall i \geq 1 \quad (9)$$

PROOF. For simplicity we denote $h_i \triangleq h_i^{\{1\}}$. We will show that for $p \in (0, 1)$, $h_i \geq \lambda = \frac{2(1-p)}{2-p}$, $\forall i \geq 1$ using induction. First note

³For an introduction on Markov chains see [28]

that since state $\{0\}$ is an absorbing state, $h_0 = 0$, $h_1 = 1 \geq \lambda$ and that $\lambda \in (0, 1)$.

The vector of hitting probabilities $h^A = (h_i^A : i \in S = \{0, 1, \dots, N\})$ for a set of states A is the minimal non-negative solution to the system of linear equations 10:

$$\begin{cases} h_i^A = 1, & \text{if } i \in A \\ h_i^A = \sum_{j \in S} p_{ij} h_j^A, & \text{if } i \notin A \end{cases} \quad (10)$$

By replacing p_{ij} with the probabilities of Definition A.4, the system of equations 10 becomes:

$$\begin{cases} h_i^A = 1, & \text{if } i \in A \\ h_i^A = \sum_{j=0}^i \binom{i}{j} p^{i-j} (1-p)^j h_j^A, & \text{if } i \notin A \end{cases} \quad (11)$$

Base case:

$$h_2 = (1-p)^2 h_2 + 2p(1-p)h_1 + p^2 h_0 = \frac{2p(1-p)}{1-(1-p)^2} = \frac{2(1-p)}{2-p} \geq \lambda$$

Inductive step: We assume that $\forall j \leq i-1 \Rightarrow h_j \geq \lambda$. We will prove that $h_i \geq \lambda$, $\forall i > 2$.

$$\begin{aligned} h_i &= \sum_{j=0}^i \binom{i}{j} p^{i-j} (1-p)^j h_j \\ &= p^i h_0 + ip^{i-1}(1-p)h_1 + \sum_{j=2}^{i-1} \binom{i}{j} p^{i-j} (1-p)^j h_j + (1-p)^i h_i \\ &\geq p^i h_0 + ip^{i-1}(1-p)h_1 + \sum_{j=2}^{i-1} \binom{i}{j} p^{i-j} (1-p)^j \lambda + (1-p)^i h_i \\ &= ip^{i-1}(1-p) + [1-p^i - (1-p)^i - ip^{i-1}(1-p)]\lambda + (1-p)^i h_i \\ \Rightarrow h_i &= \lambda - \frac{p^i}{1-(1-p)^i} \lambda + \frac{ip^{i-1}(1-p)}{1-(1-p)^i} (1-\lambda) \end{aligned}$$

We want to prove that $h_i \geq \lambda$:

$$\begin{aligned} \frac{ip^{i-1}(1-p)}{1-(1-p)^i} (1-\lambda) &\geq \frac{p^i}{1-(1-p)^i} \lambda \Rightarrow \\ \frac{ip^{i-1}(1-p) + p^i - p^i}{p^i + ip^{i-1}(1-p)} &\geq \lambda \Rightarrow \\ 1 - \frac{p^i}{p^i + ip^{i-1}(1-p)} &\geq \frac{2(1-p)}{2-p} \Rightarrow \\ \frac{p^i}{p^i + ip^{i-1}(1-p)} &\leq \frac{p}{2-p} \Rightarrow \\ p^i(2-p) &\leq p[p^i + ip^{i-1}(1-p)] \Rightarrow \\ 2-2p-i+ip &\leq 0 \Rightarrow 2-i-p(2-i) \leq 0 \Rightarrow \\ (2-i)(1-p) &\leq 0 \Rightarrow 2-i \leq 0 \end{aligned}$$

which holds since $i > 2$. The above bound is also tight since $\exists i \in S : h_i = \lambda$, specifically $h_2 = \lambda$. \square

Plugging the above hitting probability bound to the convergence proof of [13], results on the required bound.

REFERENCES

- [1] Lucas Agussurja, Akshat Kumar, and Hoong Chuin Lau. 2018. Resource-Constrained Scheduling for Maritime Traffic Management. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17116>
- [2] Stefano V Albrecht, Jacob W Crandall, and Subramanian Ramamoorthy. 2016. Belief and truth in hypothesised behaviours. *Artificial Intelligence* 235 (2016), 63–94.
- [3] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning* 47, 2 (2002), 235–256. <https://doi.org/10.1023/A:1013689704352>
- [4] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 2002. The nonstochastic multiarmed bandit problem. *SIAM journal on computing* 32, 1 (2002), 48–77.
- [5] Robert J. Aumann. 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 1 (1974), 67 – 96. [https://doi.org/10.1016/0304-4068\(74\)90037-8](https://doi.org/10.1016/0304-4068(74)90037-8)
- [6] 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.
- [7] Samuel Barrett and Peter Stone. 2011. Ad hoc teamwork modeled with multi-armed bandits: An extension to discounted infinite rewards. In *Proceedings of 2011 AAMAS Workshop on Adaptive and Learning Agents*. 9–14.
- [8] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. 2011. Contextual bandit algorithms with supervised learning guarantees. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. 19–26.
- [9] Yann Bramoullé, Dunia López-Pintado, Sanjeev Goyal, and Fernando Vega-Redondo. 2004. Network formation and anti-coordination games. *International Journal of Game Theory* 33, 1 (2004), 1–19.
- [10] Felix Brandt, Felix Fischer, and Markus Holzer. 2009. Symmetries and the Complexity of Pure Nash Equilibrium. *J. Comput. Syst. Sci.* 75, 3 (May 2009), 163–177. <https://doi.org/10.1016/j.jcss.2008.09.001>
- [11] Mithun Chakraborty, Kai Yee Phoebe Chua, Sanmay Das, and Brendan Juba. 2017. Coordinated Versus Decentralized Exploration In Multi-Agent Multi-Armed Bandits. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 164–170. <https://doi.org/10.24963/ijcai.2017/24>
- [12] Damien Challet, Matteo Marsili, Yi-Cheng Zhang, et al. 2013. Minority games: interacting agents in financial markets. *OUP Catalogue* (2013).
- [13] Ludek Cigler and Boi Faltings. 2013. Decentralized anti-coordination through multi-agent learning. *Journal of Artificial Intelligence Research* 47 (2013), 441–473.
- [14] Russell Cooper. 1999. *Coordination Games*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511609428>
- [15] Jacob W Crandall. 2014. Towards minimizing disappointment in repeated games. *Journal of Artificial Intelligence Research* 49 (2014), 111–142.
- [16] Panayiotis Danassis and Boi Faltings. 2018. Learning in Ad-hoc Anti-coordination Scenarios. *AAAI Spring Symposium Series* (2018). <https://www.aaai.org/ocs/index.php/SSS/SSS18/paper/view/17485>
- [17] Panayiotis Danassis, Aris Filos-Ratsikas, and Boi Faltings. 2019. Anytime Heuristic for Weighted Matching Through Altruism-Inspired Behavior. *arXiv:1902.09359* (2019). <https://arxiv.org/abs/1902.09359>
- [18] Dean P Foster and Rakesh V Vohra. 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior* 21, 1-2 (1997), 40.
- [19] Trond Grenager, Rob Powers, and Yoav Shoham. 2002. Dispersion games: general definitions and some specific learning results. In *AAAI/LAAI*. 398–403.
- [20] Shipra Gupta and James W. Gentry. 2016. The behavioral responses to perceived scarcity àÀš the case of fast fashion. *The International Review of Retail, Distribution and Consumer Research* 26, 3 (2016), 260–271. <https://doi.org/10.1080/09593969.2016.1147476>
- [21] Sergiu Hart and Andreu Mas-Colell. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68, 5 (2000), 1127–1150.
- [22] Raj Jain, Dah-Ming Chiu, and W. Hawe. 1998. A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems. *CoRR cs.NI/9809099* (1998). <http://arxiv.org/abs/cs.NI/9809099>
- [23] Jeremy Kun, Brian Powers, and Lev Reyzin. 2013. Anti-coordination games and stable graph colorings. In *International Symposium on Algorithmic Game Theory*. Springer, 122–133.
- [24] David Lewis. 2008. *Convention: A philosophical study*. John Wiley & Sons.
- [25] Michael L. Littman and Peter Stone. 2002. *Implicit Negotiation in Repeated Games*. Springer Berlin Heidelberg, Berlin, Heidelberg, 393–404. https://doi.org/10.1007/3-540-45448-9_29
- [26] Mihail Mihaylov, Karl Tuyls, and Ann Nowé. 2014. A decentralized approach for convention emergence in multi-agent systems. *Autonomous Agents and Multi-Agent Systems* 28, 5 (2014), 749–778.
- [27] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. 2007. *Algorithmic game theory*. Vol. 1. Cambridge University Press Cambridge.
- [28] James R Norris. 1998. *Markov chains*. Number 2. Cambridge university press.
- [29] Christos H. Papadimitriou and Tim Roughgarden. 2008. Computing Correlated Equilibria in Multi-player Games. *J. ACM* 55, 3, Article 14 (Aug. 2008), 29 pages. <https://doi.org/10.1145/1379759.1379762>
- [30] Tim Roughgarden. 2016. *Twenty Lectures on Algorithmic Game Theory* (1st ed.). Cambridge University Press, New York, NY, USA.
- [31] Thomas C Schelling. 1960. The strategy of conflict. *Cambridge, Mass* (1960).
- [32] Yoav Shoham and Moshe Tennenholtz. 1995. On social laws for artificial agent societies: off-line design. *Artificial Intelligence* 73, 1 (1995), 231 – 252. [https://doi.org/10.1016/0004-3702\(94\)00007-N](https://doi.org/10.1016/0004-3702(94)00007-N) Computational Research on Interaction and Agency, Part 2.
- [33] Peter Stone, Gal A. Kaminka, Sarit Kraus, and Jeffrey S. Rosenschein. 2010. Ad Hoc Autonomous Agent Teams: Collaboration without Pre-Coordination. In *Proceedings of the Twenty-Fourth Conference on Artificial Intelligence*.
- [34] Mohammad Sadegh Talebi. 2013. Uncoupled Learning Rules for Seeking Equilibria in Repeated Plays: An Overview. *CoRR abs/1310.5660* (2013). <http://arxiv.org/abs/1310.5660>
- [35] A. Walker and M. J. Wooldridge. 1995. Understanding the Emergence of Conventions in Multi-Agent Systems. In *ICMAS95*. San Francisco, CA, 384–389. <http://groups.lis.illinois.edu/amag/langev/paper/walker95understandingThe.html>
- [36] L. Wang, K. Wu, M. Hamdi, and L. M. Ni. 2013. Attachment-Learning for Multi-Channel Allocation in Distributed OFDMA-Based Networks. *IEEE Transactions on Wireless Communications* 12, 4 (April 2013), 1712–1721. <https://doi.org/10.1109/TWC.2013.022013.120627>
- [37] H Peyton Young. 1996. The economics of convention. *The Journal of Economic Perspectives* 10, 2 (1996), 105–122.
- [38] Li Zhou. 2015. A survey on contextual multi-armed bandits. *arXiv preprint arXiv:1508.03326* (2015).