Journal of Artificial Intelligence Research 49 (2014) 323-361 Submitted 08/13; published 02/14

Symmetric Subgame-Perfect Equilibria in Resource Allocation

Ludek Cigler LUDEK.CIGLER@EPFL.CH
Boi Faltings BOI.FALTINGS@EPFL.CH
Artificial Intelligence Laboratory
Ecole Polytechnique Fédérale de Lausanne
CH-1015 Lausanne, Switzerland

Abstract

We analyze symmetric protocols to rationally coordinate on an asymmetric, efficient allocation in an infinitely repeated N-agent, C-resource allocation problems, where the resources are all homogeneous. Bhaskar proposed one way to achieve this in 2-agent, 1-resource games: Agents start by symmetrically randomizing their actions, and as soon as they each choose different actions, they start to follow a potentially asymmetric "convention" that prescribes their actions from then on. We extend the concept of convention to the general case of infinitely repeated resource allocation games with N agents and C resources. We show that for any convention, there exists a symmetric subgame-perfect equilibrium which implements it. We present two conventions: *bourgeois*, where agents stick to the first allocation; and *market*, where agents pay for the use of resources, and observe a global coordination signal which allows them to alternate between different allocations. We define *price of anonymity* of a convention as a ratio between the maximum social payoff of any (asymmetric) strategy profile and the expected social payoff of the subgame-perfect equilibrium which implements the convention. We show that while the price of anonymity of the bourgeois convention is infinite, the market convention decreases this price by reducing the conflict between the agents.

1. Introduction

In many situations, agents have to coordinate their use of some resource. One wireless channel can only be used by one device, one parking slot may only be occupied by one vehicle, etc. The problem is that often, the agents have identical preferences: Everyone prefers to access rather than yield. Similarly, everyone prefers to have a parking slot rather than leave their car at home. However, if multiple agents try to use one resource simultaneously, they collide and everyone loses.

Consider a simple example: two agents want to access a single resource. We can describe the problem as a game. Both agents have two actions: yield (Y) and access (A). If agent α yields, it gets a payoff of 0. When agent α accesses the resource while the other agent yields, it gets a payoff of 1. But if both agents access the resource at the same time, they both incur a cost $\gamma > 0$.

The normal form of such a game looks as follows:

	Y	A
Y	0, 0	0, 1
A	1, 0	$-\gamma$, $-\gamma$

This is a symmetric game, but the two efficient Nash equilibria (NE) are asymmetric: either one agent yields and the other one accesses the resource, or vice versa. The only symmetric equilibrium outcome is the mixed NE where both agents access the resource with probability $\Pr(A) := \frac{1}{|\gamma|+1}$. However, this mixed equilibrium is not efficient, because the expected payoff of both agents is 0.

Asymmetric equilibria of symmetric games are undesirable for two reasons: First, they are not fair. In our example, only one agent can access the resource. Second, coordinating on an asymmetric equilibrium is difficult. Imagine that the agents are all identical and anonymous, i.e. they cannot observe their own identity, nor the identity of any other agent. We cannot prescribe a different strategy for each of the agents. Agents in some peer-to-peer file-sharing networks are assumed to be anonymous (Chothia & Chatzikokolakis, 2005), as well as agents in some wireless sensor networks (Durresi, Paruchuri, Durresi, & Barolli, 2005).

Consider the following example: Millions of wireless sensors are produced all by the same pipeline. We take two of them randomly, and put them in a room. There is only one frequency on which the sensors can transmit their measurements. How can each sensor know when to transmit and when to stay quiet? The factory could program half of the sensors to transmit in odd slots, and the other half to transmit in the even slots. Nevertheless, it would be just as likely to have an odd-even pair of sensors, as it would be to have a pair where the sensors transmit at the same time.

Aumann (1974) proposed the notion of *correlated equilibria* which fixes some of our issues with the Nash equilibria in the resource allocation game above. A correlated equilibrium (CE) is a probability distribution over the joint strategy profiles in the game. A correlation device samples this distribution and recommends an action for each agent to play. The probability distribution is a CE if agents do not have an incentive to deviate from the recommended action. The correlation device takes away the burden of coordination from the anonymous agents. They can all follow the same strategy: *"do what the correlation device has told me"*.

What if such "smart" correlation device, which can send each agent a different private signal, is not available? Can we still reach a correlated equilibrium outcome, one in which anonymous agents can play identical strategies, and yet achieve an efficient and fair allocation? In our previous work (Cigler & Faltings, 2011), we have proposed an algorithm that allows agents to learn a correlated equilibrium outcome through repeated play. We considered a special case of a resource allocation problem. We proposed to use a global coordination signal and multi-agent learning to reach a symmetric, fair and efficient outcome (Wang et al. (2011) later implemented this approach in an actual wireless network and achieved throughput $3\times$ higher than standard ALOHA protocols).

How does the coordination signal from our previous work (Cigler & Faltings, 2011) differ from the "smart" correlation device assumed by Aumann (1974)? Firstly, it is public and cannot send private signals to the agents. Such private signals are necessary for anonymous agents to implement the desirable correlated equilibrium in a single stage resource allocation game. The anonymous agents all have to follow the same strategy for each given public signal value. Secondly, the coordination signal is not specific to the game. The only requirement is that it is ergodic, i.e. it regularly sends each of its possible values. An example of such signal is the day of the week, the decimal value of a price of a certain stock, or even a noise on some frequency.

However, our previous solution had a major limitation: The learning algorithm itself was not an equilibrium of the repeated game. A selfish agent could force everyone else to yield by accessing all the time, securing the resource for herself. Therefore, in this paper, we focus on learning algorithms which are themselves equilibria of the repeated game. We propose a distributed algorithm to find an allocation of a set of resources which is not only symmetric and fair, but also an equilibrium. We draw inspiration from the works of Bhaskar (2000) and Kuzmics, Palfrey, and Rogers (2010) on symmetric equilibria for symmetric repeated games.

Assume that agents play an infinitely repeated game, and they discount future payoffs with a common discount factor $0 < \delta < 1$. A strategy for an agent is a mapping from any history of the play to a probability distribution over the actions. Our goal is to find a symmetric *subgame perfect equilibrium*. A subgame perfect equilibrium is a strategy profile (vector of strategies for every agent) which is a NE in any history, including those that cannot occur on the equilibrium path.

The symmetric subgame perfect equilibria that we study have the following form: The agents start by choosing their actions randomly, all according to a given probability distribution. As soon as they play an (asymmetric) pure-strategy Nash equilibrium of the game, they adopt a *convention*, that prescribes their actions deterministically from then on. Bhaskar (2000) gives two examples of conventions for symmetric 2-agent, 2-action games:

Bourgeois Agents keep using the action they played in the last round;

Egalitarian Agents play the action of their opponent from the last round.

In this paper, we extend the notion of convention to arbitrary resource allocation problems with N agents and C homogeneous resources, and we show that for any convention, there exists a symmetric subgame-perfect equilibrium that reaches this convention. We give a closed form expression to calculate the subgame-perfect equilibrium for the bourgeois convention, and show that for a small number of resources C, this convention leads to zero expected payoff. This means that the price of anonymity of the bourgeois convention is ∞.

We present the market convention as a generalization of the egalitarian convention of Bhaskar (2000). The main idea is that 1) agents pay a price for each successful access of a resource, and 2) before each round of the game, they observe a global coordination signal $k \in \{1, \ldots, K\}$, based on which they decide whether and which resource they access. The agents have a decreasing marginal utility from accessing more often. The price helps to decrease the demand for the resources, while the global coordination signal effectively increases the capacity K-times. We show that compared to the bourgeois convention, the market convention improves the expected payoff. Its price of anonymity is therefore finite.

This paper is structured as follows: In Section 2, we review some basic notions from game theory, and we present the general definitions of conventions and their implementations. In Section 3, we formally define the resource allocation game of N players and C resources, and show that for any convention, there exists a symmetric subgame-perfect equilibrium which implements it. In Section 4 we present two concrete examples of a convention: *bourgeois* and *market* conventions and discuss their properties. In Section 5 we discuss the relationship of this work to the work on folk theorems in game theory. Finally, Section 6 concludes.

2. Preliminaries

In this section, we will first introduce some basic concepts of game theory that we are going to use throughout the paper. Then, we will define the notion of price of anonymity. Finally, we will give the general definition of a convention and its implementation.

2.1 Game Theory

Game theory is the study of interactions among independent, self-interested agents. An agent who participates in a game is called a *player*. Each player has a utility function associated with each state of the world. Self-interested players take actions so as to achieve a state of the world that maximizes their utility. Game theory studies and attempts to predict the behaviour, as well as the final outcome of such interactions. Leyton-Brown and Shoham (2008) give a more complete introduction to game theory.

The basic way to represent a strategic interaction (*game*) is using the so-called *normal form*.

Definition 1. (Normal form game) A finite, N-person *normal-form game* is a tuple $G = (\mathbf{N}, \mathcal{A}, u)$, where

- $\mathbf{N}$ is a set of N players;
- $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \ldots \times \mathcal{A}_N$, where $\mathcal{A}_i$ is a set of actions available to player i. Each vector $\mathbf{a} = (a_1, a_2, \ldots, a_N) \in \mathcal{A}$ is called an *action profile*;
- $u = (u_1, u_2, \ldots, u_N)$, where $u_i : \mathcal{A} \to \mathbb{R}$ is a *utility function* for player i that assigns each action vector a certain utility (payoff).

In this paper, we will be studying *symmetric* games. In such games, the players are anonymous, and the only thing that influences the outcome is the number of agents who took a certain action.

Definition 2. (Symmetric game) We say that a normal-form game $G = (\mathbf{N}, \mathcal{A}, u)$ is a *symmetric game*, if for any permutation of the vector of players $\eta : \mathbf{N} \leftrightarrow \mathbf{N}$, it holds that for any strategy vector $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ and any $i \in \mathbf{N}$,

$$u_i(\sigma_1, \sigma_2, \ldots, \sigma_N) = u_{\eta(i)}(\sigma_{\eta(1)}, \sigma_{\eta(2)}, \ldots, \sigma_{\eta(N)}).$$

Besides playing a single deterministic action, the player can also choose her action randomly from a certain probability distribution.

Definition 3. (Mixed strategy) A *mixed* strategy selects a probability distribution over the entire action space, i.e. $\sigma_i \in \Delta(\mathcal{A}_i)$. A *mixed strategy profile* is a vector of mixed strategies for each player. For a mixed strategy σ_i, we define its support $supp(\sigma_i)$ as

$$supp(\sigma_i) = \{a_i \in \mathcal{A}_i : \sigma_i(a_i) > 0\}.$$

Given a game specified using its normal form, how should the players choose their strategy? When players know the strategies of the others, they can choose their action quite easily: just pick the strategy that maximizes the payoff *given* what everyone else is playing:

Definition 4. (Best response) We say that a mixed strategy σ_i^* of player i is a *best response* to the strategy profile of the opponents σ_{-i} if for any strategy σ_i',

$$u_i(\sigma_i^*, \sigma_{-i}) \geq u_i(\sigma_i', \sigma_{-i}).$$

As we mentioned earlier, one of the basic goals of game theory is to predict an outcome of a strategic interaction. Such outcome should be stable – therefore, it is usually called an equilibrium. One requirement for an outcome to be an equilibrium is that none of the players has an incentive to change their strategy, i.e. all players play their best-response to the strategies of the others. This defines perhaps the most important equilibrium concept, the Nash equilibrium:

Definition 5. (Nash equilibrium) A strategy profile $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$ is a *Nash equilibrium* (NE) if for every player i, her strategy σ_{-i} is a best response to the strategies of the others σ_{-i}.

Correlated equilibrium extends the notion of Nash equilibrium. In the canonical interpretation, it assumes that there is a central correlation device which samples the space of possible outcomes of the game according to some probability distribution, and then recommends an action to play to each player. No player has an incentive to deviate from the recommended action. The formal definition is as follows:

Definition 6. (Correlated equilibrium) Given an N-player game $G = (\mathbf{N}, \mathcal{A}, u)$, a *correlated equilibrium* is a tuple (v, π, μ), where v is a tuple of random variables $v = (v_1, v_2, \dots, v_N)$ with domains $D = (D_1, D_2, \dots, D_N)$, π is a joint probability distribution over v, $\mu = (\mu_1, \mu_2, \dots, \mu_N)$ is a vector of mappings $\mu_i : D_i \mapsto \mathcal{A}_i$, and for each player i and every mapping $\mu_i' : D_i \mapsto \mathcal{A}_i$ it is the case that

$$\sum_{d \in D} \pi(d) u_i \left(\mu_1(d_1), \mu_2(d_2), \dots, \mu_N(d_N)\right) \geq \sum_{d \in D} \pi(d) u_i \left(\mu_1'(d_1), \mu_2'(d_2), \dots, \mu_N'(d_N)\right).$$

In an equilibrium, each agent chooses the best strategy for himself. Oftentimes, the end result is not the best for the agents as a whole. To analyze the overall utility of a game outcome to all of the agents, we define its *social payoff*:

Definition 7. (Social payoff) For a (mixed) strategy vector $(\sigma_1, \sigma_2, \dots, \sigma_N)$, we define its *social payoff* as the sum of utilities of all the players, $\sum_{i=1}^{N} u_i(\sigma_1, \sigma_2, \dots, \sigma_N)$.

2.2 Repeated Game

In a *repeated game*, the same players play a given game (for example specified by its normal form) repeatedly. We call the normal form game that is being played in each round the *stage game*.

Definition 8. (Future discounted payoff) Given an infinite sequence of payoffs $r_i^{(1)}, r_i^{(2)}, \ldots$ for player i and a discount factor δ, $0 < \delta < 1$, the *future discounted payoff* of player i is

$$E_i := \sum_{j=1}^{\infty} \delta^j r_i^{(j)}.$$

Definition 9. (Infinitely repeated game) Let $G = (\mathbf{N}, \mathcal{A}, u)$ be a normal form game. An *infinitely repeated* version $\mathcal{G}$ of the game G *with discounting* is a game where the players play the normal form game G for an infinite number of rounds. The payoff of player i in game $\mathcal{G}$ is defined as its future discounted reward $r_i(\delta)$.

In this paper, we will study symmetric equilibria of an extended version of the repeated game, so-called *augmented game*. We assume that in every round of the game, the players can observe a common coordination signal, on which they can condition what strategy they will use. In general, this coordination signal is just a random integer taken from set $\{0, 1, \ldots, K-1\}$. In practice, it can be any piece of information observable by everyone: price of a certain stock at a given time, temperature in the room, day of the week, etc. Such a signal will allow agents to coordinate more efficiently, while at the same time it is more realistic than a general correlation device which recommends actions to the agents, as is assumed in the definition of correlated equilibria.

Definition 10. (Augmented repeated game) Let $G = (\mathbf{N}, \mathcal{A}, u)$ be a normal form game, let $\mathcal{K} := \{0, 1, \ldots, K-1\}$ be a set of *coordination signals*. An *augmented infinitely repeated* version $\mathcal{G}$ of the game G with discounting is a game where players play the normal form game G for an infinite number of rounds. In each round t, the players observe a coordination signal $k_t \in \mathcal{K}$. The coordination signal is chosen from a uniform distribution over $\mathcal{K}$. The players discount future payoff with a discount factor δ.

W.l.o.g., we always assume that repeated games are *augmented*, since in an ordinary repeated game, we can just assume that there is only one coordination signal. Therefore, in the rest of the paper, whenever we refer to a repeated game or its strategy etc., we always assume that the game is augmented with a coordination signal.

Definition 11. (History of a repeated game) Let $\mathcal{G}$ be an infinitely repeated game with discounting. We define the *history* h^t of the play in round $t \geq 0$ as

$$h^t := \big(((a_1^0, a_2^0, \ldots, a_N^0), k_0), \ldots, ((a_1^{t-1}, a_2^{t-1}, \ldots, a_N^{t-1}), k_{t-1})\big)$$

where a_i^t is the action taken by player i in round t, and k_t is the signal that the players observe in round t.

Definition 12. (Strategy of a repeated game) A *strategy in the repeated game* of a player i is a function χ_i from the history h^t and a currently observed coordination signal k_t to a probability distribution over the action space,

$$\chi_i : (h^t, k_t) \mapsto \Delta(\mathcal{A}_i).$$

We can define the Nash equilibrium of the repeated game in the same way as for the stage game (we can treat the repeated game as if it was just a normal form game where players commit to their strategy for the entire game up front).

Definition 13. (Nash equilibrium of a repeated game) A strategy profile $\chi = (\chi_1, \chi_2, \dots, \chi_N)$ is a *Nash equilibrium of the infinitely repeated game* if for each player i,

$$E_i(\chi_i, \chi_{-i}) \geq E_i(\chi_i', \chi_{-i}) \tag{1}$$

for any alternative strategy of the repeated game χ_i'. Here $E_i((\chi_i, \chi_{-i}), h^t, k_t)$ is the future discounted payoff of player i when she adopts strategy χ_i and the other players adopt a strategy vector χ_{-i}.

In the following text, we will use the notion of *future discounted social payoff*:

Definition 14. (Future discounted social payoff) Given a strategy profile χ of the infinitely repeated game $\mathcal{G}$, the *future discounted social payoff* is defined as

$$E(\chi) := \sum_{i=1}^{N} E_i(\chi). \tag{2}$$

For the repeated games, there exists a stronger notion of equilibria, which is a refinement of the standard Nash equilibrium definition.

Definition 15. (Subgame-perfect equilibrium) Let $\mathcal{G}$ be an infinitely repeated game with a discount factor $0 < \delta < 1$. A strategy vector $\chi = (\chi_1, \chi_2, \dots, \chi_N)$ is a *subgame-perfect equilibrium* of the game $\mathcal{G}$ if for each player i,

$$E_i((\chi_i, \chi_{-i}), h^t, k_t) \geq E_i((\chi_i', \chi_{-i}), h^t, k_t)$$

for any strategy χ_i', history h^t and coordination signal k_t.

In the subgame-perfect equilibrium, players play a best-response strategy given any history of the play, including the histories which cannot occur if they follow the equilibrium strategy from the beginning. The notion of subgame-perfect equilibria eliminates this way non-credible threats, or equilibria in which a player threatens someone else with a strategy which the player might be prefer to avoid if it was supposed to be executed.

2.3 Price of Anonymity

In Section 1, we have seen that in the simple resource-allocation game, the symmetric equilibrium leads to a significantly lower payoff than the asymmetric equilibria. Symmetry of the equilibria is a natural requirement when players are all the same, i.e. anonymous. How much social payoff do we have to sacrifice for the requirement of symmetry? Inspired by the price of anarchy of Koutsoupias and Papadimitriou (1999), we propose the *price of anonymity* as a measure of how efficient a given symmetric strategy vector is (the term "price of anonymity" was used previously in a different context by Bonnet & Raynal, 2011). For a given symmetric strategy vector of the stage game σ, we calculate the ratio between the social payoff of the most efficient (potentially asymmetric) Nash equilibrium of the game, and the social payoff of strategy vector σ. The formal definition is as follows:

Definition 16. (Price of anonymity of a Nash equilibrium) Let G be a symmetric game, let $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_N)$ be a symmetric Nash equilibrium (that is $\forall i, j : \sigma_i = \sigma_j$), and let τ be a (mixed) Nash equilibrium of the game G with the maximum social payoff. We define the *price of anonymity of strategy vector* σ as follows:

$$R_G(\sigma) := \frac{E(\tau)}{E(\sigma)}.$$

Definition 17. (Price of anonymity of a stage game) Let G be a symmetric game, let $\sigma^* = (\sigma_1^*, \sigma_2^*, \ldots, \sigma_N^*)$ be a symmetric Nash equilibrium with minimal social payoff, and let τ be a (mixed) Nash equilibrium of the game G with the maximum social payoff. We define the *price of anonymity of the game* G as follows:

$$R_G := \frac{E(\tau)}{E(\sigma^*)}.$$

For infinitely repeated games, we define the price of anonymity for their subgame-perfect equilibria:

Definition 18. (Price of anonymity of a repeated game) Let $\mathcal{G}$ be a symmetric game, let $\chi^* = (\chi_1^*, \chi_2^*, \ldots, \chi_N^*)$ be a symmetric subgame-perfect equilibrium with minimal social payoff, and let ψ be a subgame-perfect equilibrium of the game $\mathcal{G}$ with the maximum social payoff. We define the *price of anonymity of the game* $\mathcal{G}$ as follows:

$$R_{\mathcal{G}} := \frac{E(\psi)}{E(\chi^*)}.$$

2.4 Conventions and Implementations

As we have shown for the example of the 2-agent, 1-resource allocation game in Section 1, there exist symmetric games that have nevertheless only asymmetric efficient equilibria. If we allow for a central coordination device, the agents can play a symmetric and efficient correlated equilibrium that selects randomly from the set of efficient Nash equilibria. Without such a device, in the stage game, there is no way to reach a symmetric efficient outcome in an equilibrium.

However, if the agents play the game repeatedly, they can use the history of the play to condition their strategy. If two agents have different histories, they can take different actions in the future. In the first round of the game though, the history is empty for everyone. Therefore, a symmetric strategy for the players has to randomize in order to ever reach a point when the histories of the agents are distinct.

Bhaskar (2000) considered the problem of playing asymmetric outcomes of the stage game using a symmetric strategy of the repeated game. His work considers games with 2 players and 2 actions, such as the 1-resource allocation game. The idea is that the two players start by playing randomly, using the same probability distribution over actions. They randomize until they reach a round t where they happen to play some pure-strategy Nash equilibrium (that is, they take a different action each). We call this round the *asynchrony* round. Then, the agents start following a so-called *convention*. A convention maps the asymmetric pure-strategy Nash equilibrium to a (potentially asymmetric) strategy vector that the agents then adopt.

We have already mentioned the two basic conventions proposed by Bhaskar (2000): the bourgeois and egalitarian convention. In the 1-resource allocation game, in the asynchrony round, one agent chooses action A and the other one chooses Y. We will call the agent who chose A in the asynchrony round the *winner*. The other agent is the *loser*. The bourgeois convention guarantees that the agents will keep playing this NE forever after. This way, the winner will be forever guaranteed a higher payoff than the loser. In the egalitarian convention, the players alternate between the two pure-strategy Nash equilibria. That way the payoffs of the winner and a loser will be closer.

In the infinitely repeated game with discounting, the social payoff will depend on two things: the discount factor δ, and the probability of a collision, that is the probability that the players both play action A. When there is a big difference between the winner and loser payoff, the losers will "fight back" harder, so they will play their most preferred action A with higher probability. This will increase the probability of a collision. In the egalitarian convention, the payoffs to the loser are closer to the winner. Therefore, the agents will collide less often, and they will also reach the asynchrony faster.

As another example of a convention, Kuzmics et al. (2010) analyze the *Nash demand game*. The Nash demand game is a game of N players who choose between N actions labeled $1, \ldots, N$. If all the players choose a distinct action, each player receives a payoff equal to the label of her chosen action. If there are any two players who chose the same action, every player (including those who chose an action alone) receives zero payoff. In a pure-strategy Nash equilibrium, all the players choose a different action. Naturally, each player prefers the equilibrium where she is the one who chose action N.

In the Nash demand game, we can also define bourgeois and egalitarian conventions. Kuzmics et al. (2010) define three notions of *payoff symmetry*:

Ex-ante All agents have the same expected payoffs before the game starts.

Ex-post All agents have the same expected payoffs when asynchrony occurs (regardless of who was the winner).

Strong ex-post All agents have the same payoff along any realization of the play.

The bourgeois convention is only ex-ante payoff symmetric, since once asynchrony occurs, the winner gets a higher payoff than the loser. The egalitarian convention is strong ex-post payoff symmetric. In fact, Kuzmics et al. (2010) show that in the Nash demand game, if a convention is socially efficient, it must be strong ex-post payoff symmetric. The intuition is that in order to maximize social efficiency, we want to reach asynchrony as fast as possible. This is only possible if agents choose their actions uniformly at random. They will only do that if they are indifferent between which action they choose at the moment asynchrony occurs.

We will now formally define the convention for an *augmented* repeated game of N agents.

Definition 19. (Convention) Let $G = (\mathcal{N}, \mathcal{A}, u)$ be a symmetric normal form game and let $\mathcal{G}$ be the repeated version of game G. We define a *convention* as a function ξ that maps a vector of pure-strategy Nash equilibria of the game G for each signal value $\mathbf{a} = (\mathbf{a}^1, \mathbf{a}^2, \ldots, \mathbf{a}^K)$ to a vector of strategies of the repeated game $\mathcal{G}$, such that for any permutation $\eta : \mathbf{N} \leftrightarrow \mathbf{N}$ of the set of players,

$$\xi((\eta(\mathbf{a}^1), \ldots, \eta(\mathbf{a}^K))) = \eta(\xi(\mathbf{a}^1, \mathbf{a}^2, \ldots, \mathbf{a}^K)) \tag{3}$$

that is, *"the convention of a permutation is a permutation of a convention"* (here ξ_i denotes the strategy for player i). The strategies can be different for each coordination signal value.

We use the notation $\eta(\mathbf{a}) := \left(a_{\eta(1)}, \ldots, a_{\eta(N)}\right)$, and $\eta(\xi(\mathbf{a})) := \left(\xi_{\eta(1)}(\mathbf{a}), \ldots, \xi_{\eta(N)}(\mathbf{a})\right)$ to denote the permutation of the history vector using η, and the permutation of the strategy vectors respectively.

Our definition of convention generalizes the definition Bhaskar (2000) gave for symmetric games of 2 players and two actions α, β. Bhaskar defined a convention as a mapping from a set of Nash equilibrium action profiles $\{(\alpha, \beta), (\beta, \alpha)\}$ to a set of strategies in which the players alternate the strategy profiles (α, β) and (β, α) in some order. In our definition, a convention maps *any* Nash equilibrium of the stage game to *any* strategy profile, provided that it satisfies the permutation condition.

Intuitively, a convention prescribes each agent a potentially different *role*. The problem for anonymous agents is to learn their role. We will call the learning algorithm they will use an *implementation* of a convention.

Definition 20. (Implementation) Let $\mathcal{G}$ be an infinitely repeated game, and let ξ be a convention defined for this game. An *implementation* π_ξ of a convention ξ is a strategy vector of the infinitely repeated game, that is a function that assigns

$$\pi_\xi : (h^t, k_t) \mapsto \Delta(\mathcal{A}_1) \times \ldots \times \Delta(\mathcal{A}_N),$$

and that satisfies the following conditions:

Let h^t be the history of the game at time t.

1. If the players have already played some pure-strategy Nash equilibrium for all coordination signals $k \in \mathcal{K}$ in some round $t_0 < t$ (t_0 is the round in which they played the NE for the last signal), follow the strategy prescribed by the convention ξ for the history $h^t \setminus h^{t_0}$ (that is, the history from round $t_0 + 1$ onwards).

2. Otherwise, let k_t be the signal observed in the current round, and let vector $\mathbf{a} = (\mathbf{a}^1, \mathbf{a}^2, \ldots, \mathbf{a}^K)$ such that $\mathbf{a}^k$ is the action vector from the last round when the signal k was observed (if signal k was not observed yet, we define $\mathbf{a}^k = \emptyset$). Then, the actions of the players in the current round t only depend on vector $\mathbf{a}$ (abusing the notation, we can write $\pi_\xi(h^t, k_t) = \pi_\xi(\mathbf{a}, k_t)$), and for any permutation $\eta : \{1, 2, \ldots, N\} \leftrightarrow \{1, 2, \ldots, N\}$,

$$(\pi_{\xi,1}(\eta(\mathbf{a}), k_t), \ldots, \pi_{\xi,N}(\eta(\mathbf{a}), k_t)) = \left(\pi_{\xi,\eta(1)}(\mathbf{a}, k_t), \ldots, \pi_{\xi,\eta(N)}(\mathbf{a}, k_t)\right),$$

   that is the strategy for the current round only depends on the actions $\mathbf{a}$ played in the last round each signal was observed, and on the current coordination signal.

In Section 3, we will be concerned with equilibrium strategies for the resource allocation game. That is, we will look at its symmetric subgame-perfect equilibria. To construct such equilibria, we define the concepts of an equilibrium convention, and its equilibrium implementation.

Definition 21. (Equilibrium convention) Let $\mathcal{G}$ be an infinitely repeated game, and let ξ be some convention. We say that the convention ξ is an *equilibrium convention* when for every vector of pure-strategy Nash equilibria $\mathbf{a} = (\mathbf{a}^1, \dots, \mathbf{a}^K)$, $\xi(\mathbf{a})$ is a vector of subgame-perfect equilibria of the game $\mathcal{G}$.

Definition 22. (Equilibrium implementation) Let $\mathcal{G}$ be an infinitely repeated game, ξ some equilibrium convention, and π_ξ an implementation of convention ξ. We say that π_ξ is an *equilibrium implementation* if it is a subgame-perfect equilibrium.

3. Resource Allocation Game

In this section, we will first formally define the resource allocation game, and discuss its Nash equilibria. We will then show that for any equilibrium convention of the resource allocation game, there exists an equilibrium implementation.

3.1 Definitions

We will first define the resource allocation game, and restricted notions of uniform convention and uniform implementation.

Definition 23. (Resource allocation game) A *resource allocation game* $G_{N,C}$ is a game of N agents. Each agent i can access one of C identical resources. The agent chooses its action a_i from $\mathcal{A}_i = \{Y, A_1, A_2, \dots, A_C\}$, where action $a_i = Y$ means to yield, and action $a_i = A_c$ means to access resource c. Because all resources are identical, we can define a special meta-action $a_i = A$. To take action A means to choose to access, and then to choose the resource uniformly at random from the set of available resources.

The payoff function for agent i is defined as follows:

$$u_i(a_1, \dots, a_i, \dots, a_N) := \begin{cases} 0 & \text{if } a_i = Y \\ 1 & \text{if } a_i \neq Y, \\ & \forall j \neq i, a_j \neq a_i \\ -\gamma < 0 & \text{otherwise} \end{cases} \tag{4}$$

This game has a set of pure strategy NEs where C agents each access a resource c_i and $N - C$ agents do not. There is also a symmetric mixed strategy NE in which each agent decides to play action A with probability

$$\Pr(a_i > 0) := \min \left\{ C \cdot \left(1 - \sqrt[N-1]{\frac{|\gamma|}{1 + |\gamma|}} \right), 1 \right\}. \tag{5}$$

Note that for high enough values of C, all agents will always choose to access. [1]

Since we assume that the resources are identical, when the agents start following a convention, their expected future payoff shouldn't depend on which resource they have

1. The resource allocation game as defined here is an instance of a class of games known as *potential games* (Monderer & Shapley, 1996). In an (exact) potential game, there exists a *potential function* $\Phi : \mathcal{A} \to \mathbb{R}$ such that $\forall a_{-i} \in \mathcal{A}_{-i}$, $\forall a_i', a_i'' \in \mathcal{A}_i$,

$$\Phi(a_i', a_{-i}) - \Phi(a_i'', a_{-i}) = u_i(a_i', a_{-i}) - u_i(a_i'', a_{-i}).$$

accessed in the Nash equilibrium. We will therefore restrict ourselves to so-called *uniform* conventions:

Definition 24. (Uniform convention) Let $G_{N,C}$ be a resource allocation game, and $\mathcal{G}_{N,C}$ its infinitely repeated version. Let ξ be a convention. We say that the convention ξ is a *uniform convention*, if the following holds: Let $\mathbf{a} = (\mathbf{a}^1, \mathbf{a}^2, \ldots, \mathbf{a}^K)$ be a vector of pure-strategy Nash equilibria for each coordination signal. Let for each player i, c_i the number of signals for which player i accesses some resource in action vector $\mathbf{a}$. Then

$$\forall i, j : c_i = c_j \implies E_i(\xi(\mathbf{a})) = E_j(\xi(\mathbf{a})).$$

That is, if the number of signals for which the two players access some resource is the same, their expected payoff in the remainder of the game has to be the same too.

Definition 25. (Losers, winners, claimed and unclaimed resources) Let $\mathcal{G}_{N,C}$ be an infinitely repeated resource allocation game, let h_t be the history of play in round t, and let $\mathbf{a}^k = (a_1^k, a_2^k, \ldots, a_N^k)$ be the action vector played in the last round when signal k was observed.

- Player i is a *winner* for signal k if $a_i^k = A_i$ and for all other players $j \neq i$, $a_j^k \neq A_i$;
- Player i is a *loser* for signal k otherwise;
- Resource c is *claimed* for signal k, if there exists exactly one player i such that $a_i^k = A_c$;
- Resource c is *unclaimed* for signal k otherwise.

If signal k was never observed before, all the players are losers and all the resources are unclaimed for signal k.

Definition 26. (Uniform implementation) Let $\mathcal{G}_{N,C}$ be an infinitely repeated resource allocation game, let ξ be some uniform convention. A *uniform implementation* π_ξ is defined as follows: Let h^t be the history of the game at time t, let k_t be the signal observed in the current round.

1. If the players have already played some pure-strategy Nash equilibrium for all coordination signals follow the strategy prescribed by the convention ξ.
2. Otherwise, let n be the number of losers for signal k_t, and let c be the number of unclaimed resources for signal k. The strategy prescribed by implementation π_ξ in round t is then the following:

For an action vector $\mathbf{a}$ such that there are c_o occupied resources, and n_A agents who access some resource, the exact potential function of the resource allocation game is

$$\Phi(\mathbf{a}) := c_o + \gamma \cdot (n_A - c_o).$$

Exact potential games are also referred to as *congestion games* (Rosenthal, 1973). Finite versions of such games are always guaranteed to have a pure-strategy Nash equilibrium. Moreover, agents can reach a pure-strategy Nash equilibrium by starting from an arbitrary action vector $\mathbf{a}_0$ and iteratively playing best-response action, one by one. When the players are anonymous and update their strategies simultaneously, as we study in this paper, this doesn't hold. Hence, the theory of potential games cannot be applied to the scenario we study in this paper.

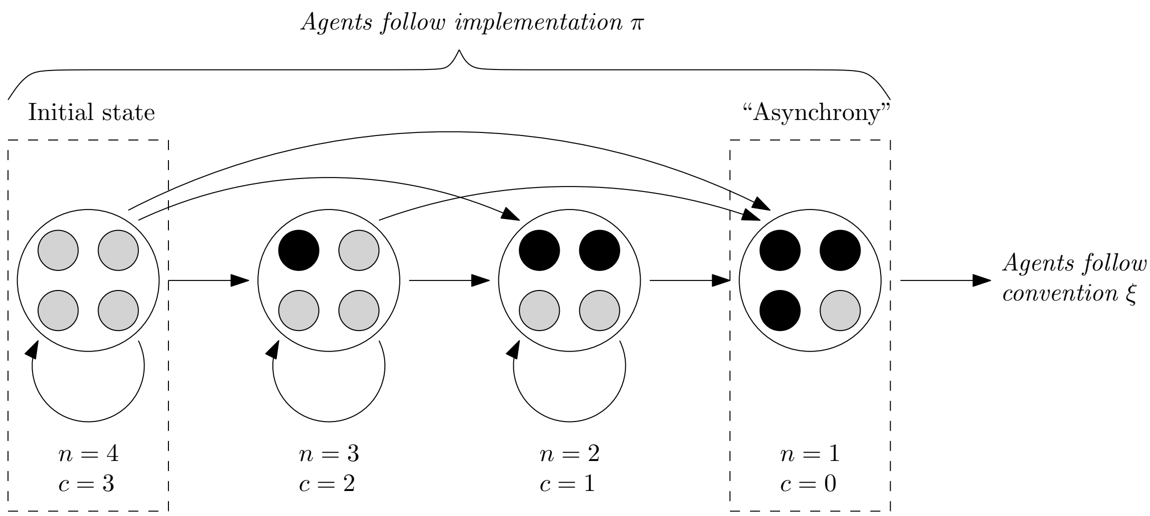

Figure 1: Learning to play a convention in a resource allocation game with $N = 4$ agents and $C = 3$ resources. Under each state, we denote the number of losers in the current state n, and the number of unclaimed resources c. Winners are denoted as black circles, losers as light grey circles. In the asynchrony state, there are 3 winners and one loser. Arrows indicate the possible transitions between the states. Once the players reach the asynchrony state, they start following the convention from the next round on.

- If player i is a winner for signal k_t, she will access the same resource as she did the last time signal k_t was observed.
- If player i is a loser, she will access choose to access an unclaimed resource r with probability $0 \leq \frac{p_{(n,c)}}{c} \leq 1$. The probability of accessing a claimed resource is zero.

In the remainder of this section, instead of studying general strategies for the repeated game, we will limit ourselves to strategies which are a uniform implementation.

Figure 1 shows how the agents learn to follow a convention when $N = 4$ and $C = 3$. Assume that the players adopt a convention ξ, and they use its implementation π. Initially, they are all "losers", and the implementation prescribes the same strategy to all of them. Once an agent accesses some resource alone, she becomes a winner and will access the same resource until the agents reach an asynchrony round (a state where each resource is accessed by exactly one agent).

Definition 27. (Expected payoff functions E_A and E_Y) Let $\mathcal{G}_{N,C}$ be an infinitely repeated resource allocation game, let ξ be some uniform convention and π_ξ its equilibrium implementation. Let h^t be the history of the game in round t, such that some $k \in \mathcal{K}$, the Nash equilibrium has not been reached (and so the convention has not been activated yet). Let n_k the number of losers for signal $k \in \mathcal{K}$ and c_k the number of unclaimed resources for signal $k \in \mathcal{K}$. Let $\mathbf{p} = (p_{n_1,c_1}, \ldots, p_{n_K,c_K})$ be the access probability of the losers for each signal $k \in \mathcal{K}$. Let k_t be the currently observed coordination signal. Let $w_\xi(n_w)$ be the expected payoff of a new winner (player who was a loser in previous rounds and becomes winner in round t) given that there are n_w new winners in round t. Let $l_\xi(n_w)$ be the expected payoff of a player who stays a loser, when there are n_w new winners in round t.

Assume that player α is a loser for signal k_t. We define her *expected payoff functions* E_A *and* E_Y when she takes action A (or Y) for signal k_t, and adopts the strategy prescribed by the implementation π_ξ for other signals:

$$\begin{aligned} &E_A(\mathbf{p}, k_t) := \\ &\sum_{n_w=1}^{\min(n,c)} [\Pr(\alpha \text{ wins } \& \ n_w \text{ winners}|A) w_\xi(n_w) + \Pr(\alpha \text{ loses } \& \ n_w \text{ winners}|A)(-\gamma + l_\xi(n_w))] \\ &+ \Pr(0 \text{ winners}|A) \cdot \left[-\gamma + \frac{\delta}{K} \left(E_A(\mathbf{p}, k) + \sum_{\substack{l=1 \\ l \neq k}}^{K} (p_{n_l,c_l} E_A(\mathbf{p}, l) + (1 - p_{n_l,c_l}) E_Y(\mathbf{p}, l)) \right) \right] \end{aligned} \tag{6}$$

$$\begin{aligned} E_Y(\mathbf{p}, k) := &\sum_{n_w=1}^{\min(n,c)} \Pr(n_w \text{ winners}|Y) \cdot l_\xi(n_w) \\ &+ \Pr(0 \text{ winners}|Y) \cdot \frac{\delta}{K} \left(E_Y(\mathbf{p}, k) + \sum_{\substack{l=1 \\ l \neq k}}^{K} (p_{n_l,c_l} E_A(\mathbf{p}, l) + (1 - p_{n_l,c_l}) E_Y(\mathbf{p}, l)) \right) \end{aligned} \tag{7}$$

3.2 Existence of an Equilibrium Implementation

We are now ready to prove that for any uniform equilibrium convention, there exists its equilibrium implementation.

Lemma 1. *For any signal $k \in \mathcal{K}$, the functions E_A and E_Y are continuous in $\mathbf{p} \in \langle 0, 1\rangle^K$.*

Proof. The probabilities $\Pr(n_w \text{ winners}|A)$ and $\Pr(n_w \text{ winners}|Y)$ are continuous. The functions E_A and E_Y are sums of products of continuous functions, so they must be themselves continuous. □

Lemma 2. *Functions E_A and E_Y are well-defined for any $k \in \mathcal{K}$ and $\mathbf{p} \in \langle 0, 1\rangle^K$.*

Proof. For fixed $\mathbf{p}$, γ and δ, the functions E_A and E_Y define a system of $2K$ linear equations. We can write this system as $(\mathbf{I} - \mathbf{A})\mathbf{E} = \mathbf{b}$, where $\mathbf{E} = (E_{A,1}, \ldots, E_{A,K}, E_{Y,1}, \ldots, E_{Y,K})$ is a vector of variables corresponding to the payoff functions, $\mathbf{b} \in \mathbb{R}^{2K}$ and $\mathbf{I}$ is $2K \times 2K$ unit matrix. The matrix $\mathbf{A}$ is defined as follows: The first K rows correspond to variables $E_{A,k}$ and the second K rows correspond to variables $E_{Y,k}$.

The elements in row k corresponding to $E_{A,k}$ are defined as:

$$\mathbf{A}_{k,l} := \begin{cases} \Pr(0 \text{ winners}|A, p_{n_k,c_k}) \cdot \frac{\delta}{K} & \text{for } l = k \\ 0 & \text{for } l = K + k \\ \Pr(0 \text{ winners}|A, p_{n_k,c_k}) \cdot \frac{\delta}{K} \cdot p_{n_l,c_l} & \text{for } l \leq K, l \neq k \\ \Pr(0 \text{ winners}|A, p_{n_k,c_k}) \cdot \frac{\delta}{K} \cdot (1 - p_{n_l,c_l}) & \text{for } K < l \leq 2K, l \neq K + k \end{cases}$$

The elements in row $K+k$ corresponding to $E_{Y,k}$ are defined as:

$$\mathbf{A}_{K+k,l} := \begin{cases} \Pr(0 \text{ winners}|Y, p_{n_k,c_k}) \cdot \frac{\delta}{K} & \text{for } l = K+k \\ 0 & \text{for } l = k \\ \Pr(0 \text{ winners}|Y, p_{n_k,c_k}) \cdot \frac{\delta}{K} \cdot p_{n_l,c_l} & \text{for } l \leq K, l \neq k \\ \Pr(0 \text{ winners}|Y, p_{n_k,c_k}) \cdot \frac{\delta}{K} \cdot (1 - p_{n_l,c_l}) & \text{for } K < l \leq 2K, l \neq K+k \end{cases}$$

This system of equations has a unique solution if the matrix $\mathbf{A}$ is non-singular. This is equivalent to saying that $\det(\mathbf{A}) \neq 0$.

The matrix $\mathbf{A}$ is diagonally dominant, that is $a_{ii} > \sum_{j=1,j\neq i}^{K} |a_{ij}|$. This is because $0 < \delta < 1$, and the rows of the matrix $\mathbf{A}$ sum to $\sum_{l=1}^{2K} \mathbf{A}_{k,l} = \delta \cdot \Pr(0 \text{ winners}|A, p_{n_k,c_k})$ for $1 < k \leq K$, and $\sum_{l=1}^{2K} \mathbf{A}_{K+k,l} = \delta \cdot \Pr(0 \text{ winners}|Y, p_{n_k,c_k})$ for $K < K+k \leq 2K$.

It is known that diagonally dominant matrices are non-singular (Taussky, 1949). Therefore, a unique solution $\mathbf{E}$ of the system exists and the functions E_A, E_Y are well-defined for a fixed $\mathbf{p}$, δ and γ. □

Lemma 3. *There exists $\mathbf{p}^*$ such that for any $k \in \mathcal{K}$, one of the following is true:*

1. $p_k^* = 0$ *and* $E_Y(\mathbf{p}^*, k) > E_A(\mathbf{p}^*, k)$*;*
2. $p_k^* = 1$ *and* $E_A(\mathbf{p}^*, k) > E_Y(\mathbf{p}^*, k)$*;*
3. $0 < p_k^* < 1$ *and* $E_A(\mathbf{p}^*, k) = E_Y(\mathbf{p}^*, k)$*.*

Such $\mathbf{p}^*$ defines a symmetric best-response strategy for the losers.

Proof. Fix γ and δ. We will show that for an arbitrary $\mathbf{p}$ and every signal $k \in \mathcal{K}$, there exists p'_k which satisfies one of the three conditions of the Lemma 3 above.

For contradiction, assume that for $p'_k = 0$, $E_Y(\mathbf{p}', k) \leq E_A(\mathbf{p}', k)$ and for $p'_k = 1$, $E_A(\mathbf{p}', k) \leq E_Y(\mathbf{p}', k)$. Then from the fact that both functions E_A and E_Y are well-defined and continuous for $0 \leq p_k \leq 1$, they must intersect for some $0 < p'_k < 1$.

From this, there must exist a vector $\mathbf{p}^*$ where for *all* $k \in \mathcal{K}$, the conditions of the Lemma 3 are satisfied. □

Corollary 1. *Let $\mathcal{G}_{N,C}$ be an infinitely repeated resource allocation game. For any uniform equilibrium convention ξ of the game $\mathcal{G}_{N,C}$, there exists an equilibrium implementation π_ξ.*

To illustrate the different equilibrium payoffs agents can get when they adopt different conventions, consider the resource allocation game with $N = 4$ agents and $C = 1$ (to simplify the presentation, assume that $K = 1$). Assume that before round t, the resource has been claimed yet, so there are $n = 4$ losers and $c = 1$ unclaimed resource. If some agent becomes a winner in round t, the agents adopt an extended uniform convention that prescribes their strategies from then on.

For comparison, assume that the agents can adopt either a convention $\hat{\xi}_1$, or a convention $\hat{\xi}_2$. If they adopt convention $\hat{\xi}_1$, the winners have an expected payoff $w_{\hat{\xi}_1} = 4$, and the losers an expected payoff $l_{\hat{\xi}_1} = 0$. On the other hand, if they adopt convention $\hat{\xi}_2$, the winners have an expected payoff $w_{\hat{\xi}_2} = 2$, and the losers an expected payoff $l_{\hat{\xi}_2} = 1$.

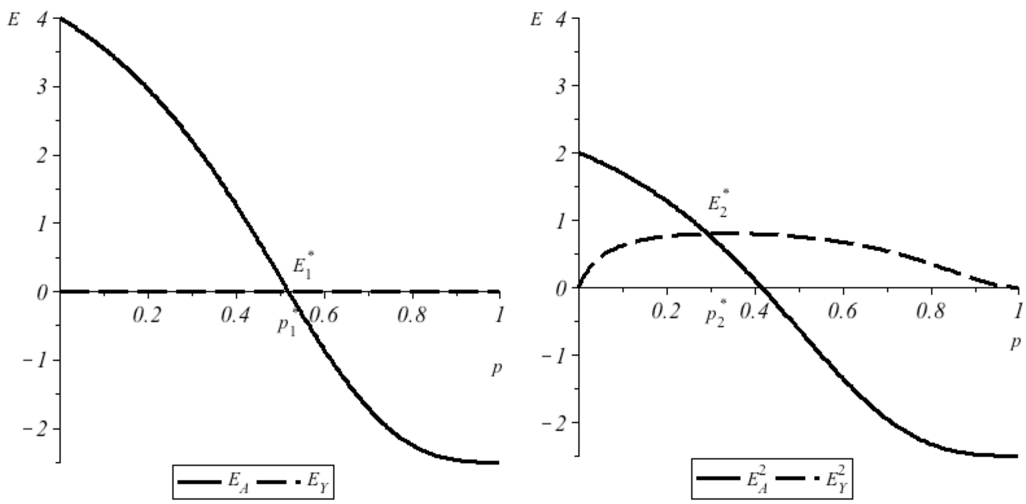

(a) Convention $\hat{\xi}_1$

(b) Convention $\hat{\xi}_2$

Figure 2: Example of expected payoff functions for resource allocation game with $N = 4$ agents, $C = 1$ resources, cost of collision $\gamma = 2$ and discount factor $\delta = 0.8$, given the access probability p. The function E^1_A and E^1_Y are expected payoff functions of accessing and yielding, when the agents use an extended convention $\hat{\xi}_1$. Similarly, E^2_A and E^2_Y are expected payoff functions when the agents use an extended convention $\hat{\xi}_2$. Convention $\hat{\xi}_1$ has an expected winner payoff $w_{\hat{\xi}_1} = 4$, and expected loser payoff $l_{\hat{\xi}_1} = 0$. Convention $\hat{\xi}_2$ has an expected winner payoff $w_{\hat{\xi}_2} = 2$ and expected loser payoff $l_{\hat{\xi}_2} = 1$.

In the equilibrium implementation π_1 of the convention $\hat{\xi}_1$, the agents access the resource with probability p^*_1, and their expected payoff is $E^*_1 = 0$. In the equilibrium implementation π_2 of the convention $\hat{\xi}_2$, the agents access the resource with probability $p^*_2 < p^*_1$, and their expected payoff is $E^*_2 > E^*_1 = 0$.

Figure 2 shows the expected payoff functions (E^1_A and E^1_Y for the convention $\hat{\xi}_1$, and E^2_A and E^2_Y for the convention $\hat{\xi}_2$), depending on the access probability p. We can see that the equilibrium implementation payoff E^*_2 of the convention $\hat{\xi}_2$ is higher than the equilibrium payoff E^*_1 of the convention $\hat{\xi}_1$, even though the sum of the winner and loser payoffs is higher for convention $\hat{\xi}_1$. This is because the loser receives a positive payoff when the agents adopt a convention $\hat{\xi}_2$; the agents are less likely to "fight" to become a winner, and they access the resource with a lower probability $p^*_2 < p^*_1$. This way, there will be less collisions, and the agents will receive a higher expected social payoff when they adopt the convention $\hat{\xi}_2$.

3.3 Calculating the Equilibrium

While the symmetric subgame-perfect equilibrium is guaranteed to exist, 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 probability of accessing a resource. Therefore, we will show how to calculate the equilibrium strategy numerically.

Let $\mathbf{p}$ be a probability vector and k a signal. Let $\mathbf{p}_0 := (p_1, p_2, \ldots, p_k = 0, \ldots, p_K)$, i.e. vector $\mathbf{p}$ with p_k set to 0. Let $\mathbf{p}_1 := (p_1, p_2, \ldots, p_k = 1, \ldots, p_K)$. From Lemma 3 we know that either $E_Y(\mathbf{p}_0, k) > E_A(\mathbf{p}_0, k)$, or $E_A(\mathbf{p}_1, k) > E_Y(\mathbf{p}_1, k)$ or the two functions intersect for some $0 \le p_k \le 1$. Furthermore, we know that $E_A(\mathbf{p}_0, k) = w_{\hat{\xi}}(c)$ since the probability of successfully claiming a resource is 1 when everyone else yields, and also $E_Y(\mathbf{p}_0, k) = 0$. Therefore, $E_Y(\mathbf{p}_0, k) > E_A(\mathbf{p}_0, k)$ *iff* $w_{\hat{\xi}}(c) > 0$.

W.l.o.g, we will assume that $w_{\hat{\xi}}(c) > 0$. Algorithm 1 shows then how to calculate the probability vector.

Algorithm 1 Calculating the equilibrium probabilities

---

for Each subset $S \subseteq \{1, 2, \ldots, K\}$ **do**
  Let Σ be a system of equations
  $\forall i \notin S$, Σ contains two equations for $E(\mathbf{p}, i)$. One corresponding to $E_A(\mathbf{p}, i)$, one to $E_Y(\mathbf{p}, i)$ (see Equations 6 and 7).
  $\forall j \in S$, we set $p_j := 1$. Σ contains only one equation for $E(\mathbf{p}, j)$, corresponding to $E_A(\mathbf{p}, j)$.
  So Σ is a system of $2K - |S|$ equations with $2K - |S|$ variables.
  Solve numerically the system of equations Σ.

  if there exists a solution to Σ for which $\forall i \notin S$, $0 \le p_i \le 1$ **then**
    We have found a solution
    break;
  end if
end for

The numerical algorithm has a complexity exponential in K, and is therefore only suitable for small K. In Section 4.2.3, we will show conditions under which the access probabilities are easy to compute and define a ε-equilibrium of the repeated game. That is, no agent can gain more than ε factor more by deviating from the prescribed strategy.

4. Actual Conventions

In the previous section, we have shown that we can find a symmetric way to reach any convention, provided the agents access the resources with a certain probability. We have also shown how to calculate the resource access probability in every stage of the game. In this section, we would like to show specific examples of the conventions that agents can adopt, and discuss their properties.

4.1 Bourgeois Convention

The bourgeois convention is the simplest one. Once an agent has accessed a resource successfully for the first time, he will keep accessing it forever. We say that the agent has *claimed* the resource. We don't need any coordination signal to implement it, so we can set $K := 1$.

We will describe the decision problem from the point of view of agent α. Assume that there are N agents and C resources. At round t, let c_t be the number of resources which have not been claimed yet, and $n_t := N - C + c_t$ the number of players who have not claimed a resource yet. Assume that other players besides α use the following strategy:

- If a player has claimed a resource previously, she will keep accessing it;
- If a player hasn't claimed any resource yet (she is a "loser"), she will choose to access with probability p_{c_t} and then choose the actual resource to access uniformly at random.

Definition 28. (Expected payoff function of the Bourgeois convention) Let $\mathbf{p} := (p_1, p_2, \ldots, p_C)$ be a probability vector, such that p_c is the probability with which any of the losers will access when there are c unclaimed resources. We define the *expected payoff function* to player α should she choose to access (play A, that is choose to access and then choose the resource uniformly at random) or yield (play Y), respectively:

$$E_A(\mathbf{p}, c) := \left(1 - \frac{p}{c}\right)^{n-1} \cdot \frac{1}{1-\delta} + \left[1 - \left(1 - \frac{p}{c}\right)^{n-1}\right] \cdot (-\gamma)$$
$$+ \delta \cdot \sum_{l=0}^{c} \Pr(\alpha \text{ loses and } n_w = l | A) \cdot E(\mathbf{p}, c - l);$$

$$E_Y(\mathbf{p}, c) := \delta \cdot \sum_{l=0}^{c} \Pr(n_w = l | Y) \cdot E(\mathbf{p}, c - l);$$

In both equations, $E(\mathbf{p}, c) = \max\{E_A(\mathbf{p}, c), E_Y(\mathbf{p}, c)\}$.

Lemma 4. *For any $\mathbf{p}$ and $1 \leq c \leq C$, $E(\mathbf{p}, c) \geq 0$.*

Proof. No matter what is the strategy of the opponents, if agent α chooses to always yield, its payoff will be 0. □

Lemma 5. *Let $\mathbf{p}$ be a probability vector which defines the strategies of the other losers, and let c_t be the number of unclaimed resources in round t. If $\forall c \leq c_t$, $E_A(\mathbf{p}, c) = E_Y(\mathbf{p}, c)$, then $\forall c \leq c_t$, $E(\mathbf{p}, c) = 0$.*

Proof. When there are c_t unclaimed resources in round t, in every following round $t' \geq t$ there will be $c \leq c_t$ unclaimed resources (in bourgeois convention, agents never release claimed resources). If the agent α is indifferent between actions Y and A in every round following round t, that means that it is indifferent between a strategy of the subgame that prescribes Y in every round and any other strategy. The (expected) payoff of the strategy that prescribes always Y is 0. Therefore, the expected payoff of any other subgame strategy must be 0 as well. □

For the purpose of our problem, all the unclaimed resources are identical. Therefore the only parameter of the losers' strategy is the probability with which the agents decide to access – the resource itself is then chosen uniformly at random. Lemma 5 shows a necessary condition on $\mathbf{p}$ for agent α to be indifferent. The following lemma shows that such $\mathbf{p}$ exists and is unique.

Lemma 6. *Assume at time t there are c_t unclaimed resources. Let for all $c \leq c_t$ unclaimed resources be $p_c^* = c\left(1 - \sqrt[n-1]{\frac{|\gamma|}{|\gamma|+\frac{1}{1-\delta}}}\right)$ the probability with which the losers play A. Then for all $c \leq c_t$ unclaimed resources, agent α is indifferent between yielding and accessing. For a given c, such probability is unique on the interval $[0, c]$.*

Proof. From Lemma 5 we know that when agent α is indifferent (i.e. $E_A(\mathbf{p}, c) = E_Y(\mathbf{p}, c)$), it must be that $E(\mathbf{p}, c) = 0$ for all $1 \leq c \leq c_t$.

From Definition 28, the expected profit to agent α from playing A and then following best-response strategy (with zero payoff) is

$$\begin{aligned} E_A(\mathbf{p}, c) = \left(1 - \frac{p_c}{c}\right)^{n-1} \cdot \frac{1}{1-\delta} + \left[1 - \left(1 - \frac{p_c}{c}\right)^{n-1}\right] \cdot (-\gamma) \\ + \delta \cdot \Pr(\alpha \text{ loses and } n_w = 0 | A) \cdot E(\mathbf{p}, c). \end{aligned} \tag{8}$$

Here p_c is the probability with which the other losers access. We want $E_A(\mathbf{p}, c) = E_Y(\mathbf{p}, c) = 0$. This holds if p_c^* is defined as in the lemma above.

Function E_A is decreasing in p_c on the interval $[0, c]$, while function E_Y is constantly 0. Therefore, their intersection is unique on an interval $[0, c]$. □

Lemma 7. *Assume that all the opponents who haven't claimed any resource access a resource with probability $p < p_c^*$. Then it is best-response for agent α to access.*

Proof. The probability that agent α claims successfully a resource after playing A is

$$\Pr(\text{claim some resource} | A) := \left(1 - \frac{p}{c}\right)^{n-1} \tag{9}$$

This probability increases as p decreases. Therefore the expected profit of playing A is increasing as p decreases, whereas the profit of playing Y stays 0. □

Theorem 8. *Define an agent's strategy τ as follows: If there are c unclaimed resources, play A with probability $p_c := min\,(1, p_c^*)$ (where p_c^* is defined in Lemma 6). Then a joint strategy profile $\tau = (\tau_1, \tau_2, \ldots, \tau_N)$ where $\forall c, \tau_c = \tau$ is a subgame-perfect equilibrium of the infinitely repeated resource allocation game.*

Proof. From Lemma 6, if $p_c^* < 1$, any agent is indifferent between playing Y and playing A, therefore will happily follow strategy τ. From Lemma 7, if $p_c = 1 < p_c^*$, it is best response for any agent to play A, just as the strategy τ prescribes. □

Theorem 9. *For all $c \in \mathbb{N}$, if $p_c = p_c^*$, $E(\mathbf{p}, c) = 0$.*

Proof. We will proceed by induction.

For $c = 0$, the expected payoff is trivially $E(\mathbf{p}, 0) = 0$, because there are no free resources.

Let $\forall j < c, E(\mathbf{p}, j) = 0$ and $p_c = p_c^*$. If agent α plays Y, the expected payoff is clearly 0 (it will be 0 now and 0 in the future by the induction hypothesis). If agent α plays A, the expected payoff is (by Definition 28):

$$\begin{aligned} E_A(\mathbf{p}, c) := \left(1 - \frac{p_c}{c}\right)^{n-1} \cdot \frac{1}{1-\delta} \\ + \left[1 - \left(1 - \frac{p_c}{c}\right)^{n-1}\right] \cdot (-\gamma) + \delta \cdot \sum_{l=0}^{c} \Pr(\alpha \text{ loses and } n_w = l | A) \cdot E(\mathbf{p}, c - l) \end{aligned} \tag{10}$$

Because of the way the p_c^* is defined, and from the induction hypothesis $E(\mathbf{p}, j) = 0$ for $j < c$, we get

$$\begin{aligned} E_A(\mathbf{p}, c) &:= \Pr(\alpha \text{ loses and } n_w = 0|A) \cdot E(c, \tau_{-\alpha}) \\ &= \delta \Pr(\alpha \text{ loses and } n_w = 0|A) \cdot \max\{E_A(\mathbf{p}, c), E_Y(\mathbf{p}, c)\} \end{aligned} \tag{11}$$

Since $\delta \cdot \Pr(\alpha \text{ loses and } n_w = 0|A) < 1$, it must be that $E_A(\mathbf{p}, c) = 0$. □

Theorem 10. *If $p_c < p_c^*$, $E(\mathbf{p}, c) > 0$.*

Proof. From Lemma 7 we know that when $p_c < p_c^*$, it is a best response to access, so $E(\mathbf{p}, c) = E_A(\mathbf{p}, c)$. From Lemma 4 we know that for all j, $E(\mathbf{p}, j) \geq 0$. If $p_c < p_c^*$, by Definition 28 we see that $E(\mathbf{p}, c) > 0$. □

Theorem 10 shows that if we have enough resources so that $p_c^* \geq 1$, the expected payoff for the agents, even when they access all the time, will be positive.

Given the number of agents N, discount factor δ and collision cost γ, the necessary number of resources c^* for the expected payoff to be positive is:

$$c^* := \frac{1}{1 - \sqrt[n-1]{\frac{|\gamma|}{|\gamma| + \frac{1}{1-\delta}}}} \tag{12}$$

Figure 3 illustrates the value of c^* depending on N, δ, and γ respectively. Figure 3a shows how the number of resources c^* increases as N increases – naturally, more agents need more resources.

Figure 3b shows on the other hand that with an increasing discount factor δ, the necessary number of resources drops. This is because for high δ, the agents are almost indifferent between winning now and winning later. In Section 4.2.3, we will explore this idea in more detail – we will show that for high enough delta, a strategy which prescribes the agents to access with a constant probability until they reach the asynchrony is an ε-equilibrium of the resource allocation game.

Finally, Figure 3c shows an increasing number of resources which are necessary for the bourgeois convention to have positive payoff, as the collision cost γ increases. The increase is almost linear in γ. This is because the higher the cost of collision, the lower the expected payoff of accessing E_A. For the bourgeois convention to have positive expected payoff, we need $E_A > 0$ for all $0 \leq p \leq 1$.

Let us now look at the price of anonymity for the bourgeois convention (as defined in Definition 16).

Theorem 11. *The price of anonymity of the bourgeois convention is infinite.*

Proof. The highest social payoff any strategy profile τ can achieve in an N-agent, C-resource allocation game ($N \geq C$) is

$$\max E(\tau) := \frac{C}{1 - \delta}. \tag{13}$$

This is achieved when in every round, every resource is accessed by exactly one agent. Such strategy profile is obviously asymmetric.

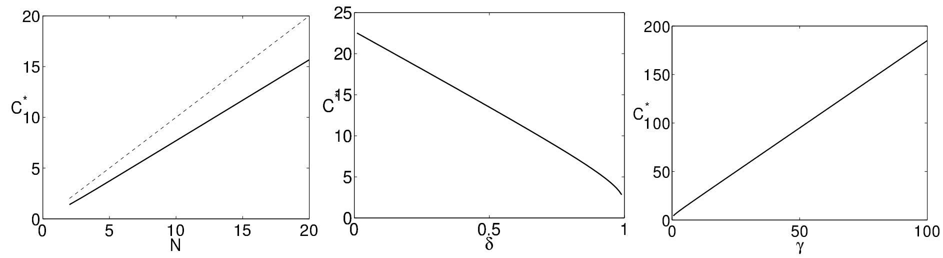

(a) N (b) δ (c) γ

Figure 3: Minimum number of resources c^* needed for the expected payoff of bourgeois convention to be positive, depending on N, δ, and γ. One parameter is varying, the other parameters are set to $N = 10$, $\delta = 0.8$, $\gamma = 2$. For varying N, the dashed line shows when $c = N$.

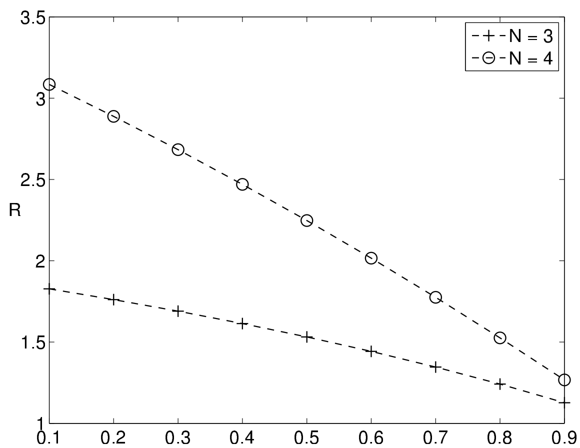

Figure 4: Market convention: Price of anonymity for $C = 1$, $K = N$, $\gamma = 0.5$ and varying δ.

If each agent knew which part of the bourgeois convention to play at the beginning of the game, this convention would be socially efficient. However, when the agents are anonymous, they have to learn which part of the convention they should play through randomization. For the bourgeois convention for small C relative to N, this randomization is such that the agents are indifferent between accessing some resource and yielding, and their expected payoff of both is zero. Therefore, its price of anonymity is infinite. □

4.2 Market Convention

We saw that the bourgeois convention leads to zero expected social payoff for a small number of resources. We would like to improve the expected payoff here. In the bourgeois convention, the agents receive zero expected payoff because the demand for resources is too high compared to the supply. We need to decrease the demand, while increasing the

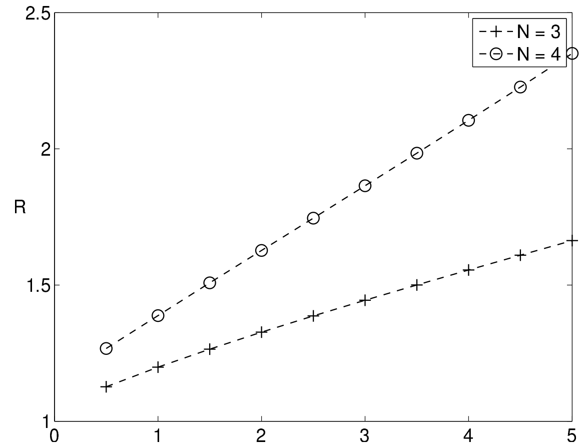

Figure 5: Market convention: Price of anonymity for $C = 1$, $K = N$, $\delta = 0.9$ and varying γ.

supply. This is often achieved through markets. Shneidman et al. (2005) present some of the reasons why markets might be appropriate for resource allocation.

We assume the following:

- Agents can observe $K \geq 1$ coordination signals.
- Agents have a decreasing marginal utility when they access a resource more often.
- They pay a fixed price per each successful access, to the point that each agent prefers to access a resource only for one signal out of K. In practice, this could be implemented by a central authority that observes the convergence rate of the agents, and dynamically increases or decreases the price to achieve convergence.

Such assumptions define what we call "market" convention, where the winners only access their claimed resource for the signals they observed when they first claimed it. The price the agents have to pay serves to decrease the demand. The coordination signal effectively increases the supply of resources K-times, because the resource allocation may be different for each of the signal values.

We know that we can implement this convention for $C \geq 1$ resources using symmetric play (see Section 3). For small K, we can also use Algorithm 1 to calculate the access probabilities. For the ease of exposition, we will first describe the market convention for $C = 1$ resource. Then we will generalize the description to $C > 1$ resources.

4.2.1 One Resource

When each agent only accesses the resource for one signal, we need $K = N$ signals to make sure everyone gets to access once.

In the N-agent, 1-resource case, imagine there are still n agents playing and $(N - n)$ agents who have already claimed the resource for some signal. Imagine that the n agents observe one of the n signals for which no resource has been claimed.

Assume that all agents access the resource with probability p_n. The expected payoff of accessing a resource for agent α is

$$
\begin{aligned}
E_A(p_n, n) := (1 - p_n)^{n-1} \cdot \left(1 + \frac{\delta}{N} \cdot \frac{1}{1-\delta}\right) \\
+ \left[1 - (1 - p_n)^{n-1}\right] \cdot \left[-\gamma + \frac{\delta n}{N - \delta(N-n)} E_A(p_n, n)\right]
\end{aligned}
\tag{14}
$$

The expected payoff of yielding for agent α is

$$
\begin{aligned}
E_Y(p_n, n) := (n-1)p_n(1-p_n)^{n-2} E(n-1) \\
+ \left[1 - (n-1)p_n(1-p_n)^{n-2}\right] \frac{\delta n}{N - \delta(N-n)} E_Y(p_n, n)
\end{aligned}
\tag{15}
$$

When $p_n = 1$, accessing a resource will always lead to a collision, so the payoff of accessing will be negative. When $p_n = 0$, accessing a resource will always claim it, so the payoff of accessing will be positive. In the equilibrium, the agents should be indifferent between accessing and yielding. Therefore, we want to find p_n^* such that $E_A(p_n^*, n) = E_Y(p_n^*, n) = E(n)$.

Finding a closed form expression for p_n^* is difficult, but we can use Algorithm 1 to calculate this probability, as well as the expected payoff $E(n)$, numerically (albeit in practice only for small K).

Figures 4 and 5 show the price of anonymity of the market convention (as defined in Definition 16) for varying discount factor δ, and varying cost of collision γ, respectively. From Section 4.1, the price of anonymity for $C = 1$ is ∞. In contrast, for the market convention this price is in both cases finite and relatively small.

4.2.2 Multiple Resources

Assume now that $C \geq 1$. In any given round, we will denote $\mathbf{c} := (c_1, c_2, \ldots, c_K)$ the vector of resources which have not been claimed yet for each value of the coordination signal $k \in \{1, 2, \ldots, K\}$. We will denote n the number of players who have not claimed any resource yet for any signal value. Finally, let $\mathbf{p} := (p_{n,c_1}, p_{n,c_2}, \ldots, p_{n,c_K})$ be the vector of probabilities, where p_{n,c_k} denotes the probability that a loser will access some resource for signal $k \in \mathcal{K}$, given that c_k resources are available.

From Corollary 1, we know that for the market convention, there exists an equilibrium implementation. But what does it look like exactly? In order to be able to express it (albeit only numerically), we need to define the expected payoff functions the players receive for each action.

For number of losers n, observed coordination signal k and vectors $\mathbf{p}$ and $\mathbf{c}$, we can define expected payoff functions when a player α takes action A and Y, respectively:

$$\begin{aligned}
&E_A(\mathbf{p}, \mathbf{c}, n, k) := \Pr(\alpha \text{ wins}|A) \cdot w + (1 - \Pr(\alpha \text{ wins}|A)) \cdot (-\gamma) \\
&+ \sum_{n_w=1}^{\min(c_k, n)} \Pr(\alpha \text{ loses, } n_w \text{ winners}|A) \cdot E(\mathbf{p}, (c_k - n_w, c_{-k}), n - n_w) \\
&+ \Pr(n_w = 0|A) \cdot \frac{\delta}{K} \cdot E_A(\mathbf{p}, \mathbf{c}, n, k) \\
&+ \Pr(n_w = 0|A) \cdot \left[\frac{\delta}{K} \cdot \sum_{\substack{l=1 \\ l \neq k}}^{K} (p_{n,c_l} E_A(\mathbf{p}, \mathbf{c}, n, l) + (1 - p_{n,c_l}) E_Y(\mathbf{p}, \mathbf{c}, n, l)) \right]
\end{aligned} \tag{16}$$

$$\begin{aligned}
&E_Y(\mathbf{p}, \mathbf{c}, n, k) := \sum_{n_w=1}^{\min(n,c)} \Pr(n_w \text{ winners}|Y) \cdot E(\mathbf{p}, (c_k - n_w, c_{-k}), n - n_w) \\
&+ \Pr(n_w = 0|Y) \cdot \frac{\delta}{K} \left(E_Y(\mathbf{p}, \mathbf{c}, n, k) + \sum_{\substack{l=1 \\ l \neq k}}^{K} (p_{n,c_l} E_A(\mathbf{p}, \mathbf{c}, n, l) + (1 - p_{n,c_l}) E_Y(\mathbf{p}, \mathbf{c}, n, l)) \right)
\end{aligned} \tag{17}$$

In the equations above, $E(\mathbf{p}, \mathbf{c}, n)$ is the expected payoff *before* the players observe the coordination signal. It is defined as

$$E(\mathbf{p}, \mathbf{c}, n) := \frac{1}{K} \sum_{k=1}^{K} E(\mathbf{p}, \mathbf{c}, n, k).$$

The winner payoff w is defined as $w := 1 + \frac{\delta}{K \cdot (1-\delta)}$. This is because the winner will access for only one signal: once in the current round, and than in any future round with probability $\frac{1}{K}$.

What are the probabilities that there will be n_w winners in each of the cases? We will start with the simplest case, $\Pr(n_w \text{ winners}|Y)$, given that there are n agents (including agent α), c_k resources and all agents except α play action A with probability p_k.

The problem of calculating this probability is very similar to the well-known *balls-and-bins* problem (Raab & Steger, 1998). In the balls-and-bins problem we assume that we have n balls who are each randomly assigned into one of the c bins. The goal is to find a probability that i bins will have exactly one ball in them. We will express this probability as $\phi(n, c, i)$.

There are N_i ways to pick some i balls, place them into some i bins so that each bin has one ball, and place the remaining $n - i$ balls into the remaining $c - i$ bins randomly,

$$N_i := \binom{c}{i} \binom{n}{i} \cdot i! \cdot (c - i)^{n-i} \tag{18}$$

There are a total of c^n ways to arrange n balls into c bins. Therefore, the probability $\phi(n, c, i)$ is The total number of ways to place n balls in c bins so that exactly i have one ball can be then obtained from the *generalized inclusion-exclusion* principle:

$$\begin{aligned}
\phi(n, c, i) &:= \frac{1}{c^n} \sum_{j=i}^{\min(c,n)} (-1)^{j-i} \binom{j}{i} N_j \\
&= \frac{1}{c^n} \sum_{j=i}^{\min(c,n)} (-1)^{j-i} \binom{j}{i} \binom{c}{j} \binom{n}{j} \cdot j! \cdot (c-j)^{n-j} \qquad (19) \\
&= \frac{n!}{c^n} \binom{c}{i} \sum_{j=i}^{\min(c,n)} (-1)^{j-i} \binom{c-i}{j-i} \frac{(c-j)^{n-j}}{(n-j)!}.
\end{aligned}$$

In the simplification above, we use the *absorption identity* $\binom{j}{i}\binom{c}{j} = \binom{c}{i}\binom{c-i}{j-i}$.

How can we use the function ϕ to calculate $\Pr(n_w \text{ winners}|Y)$? The $n-1$ agents (other than α) decide to play action A with probability p_k, and then choose the resource to access randomly. The agents who choose to access a resource correspond to the balls-and-bins problem. Therefore,

$$\Pr(n_w \text{ winners}|Y) := \sum_{i=0}^{n-1} \binom{n-1}{i} p_k^i \cdot (1-p_k)^{n-1-i} \cdot \phi(i, c, n_w). \qquad (20)$$

To calculate the probability $\Pr(\alpha \text{ wins}|A)$, we can proceed as follows. We assume w.l.o.g that α accesses resource 1. There will be some i agents (out of $n-1$) who will choose action A. We then need all of them to choose other resource than 1. Therefore,

$$\Pr(\alpha \text{ wins}|A) := \sum_{i=0}^{n-1} \binom{n-1}{i} \cdot p_k^i \cdot (1-p_k)^{n-1-i} \left(1 - \frac{1}{c}\right)^i \qquad (21)$$

Finally, to calculate the probability $\Pr(\alpha \text{ loses}, n_w \text{ winners}|A)$, we can use again the balls-and-bins problem. Given that there are $0 \le i \le n-1$ agents who choose action A, there will be $0 \le j \le i$ agents who choose the same resource as agent α. The remaining $(i-j)$ agents face the same balls-and-bins problem for $c-1$ bins (1 bin is already occupied by agent α). Therefore,

$$\Pr(\alpha \text{ loses}, n_w \text{ winners}|A) :=$$
$$\sum_{i=1}^{n-1} \binom{n-1}{i} p_k^i (1-p_k)^{n-1-i} \sum_{j=1}^{i} \binom{i}{j} \left(\frac{1}{c}\right)^j \left(1 - \frac{1}{c}\right)^{i-j} \cdot \phi(i-j, c-1, n_w) \qquad (22)$$

Now that we have expressed the expected payoff functions E_A and E_Y explicitly, we can use Algorithm 1 to calculate the equilibrium access probabilities and expected payoffs.

Figures 6 and 7 show the price of anonymity of the market convention for $C = 3$, $K = 2$ and $N = C \cdot K = 6$. When the discount factor δ grows, the price of anonymity decreases

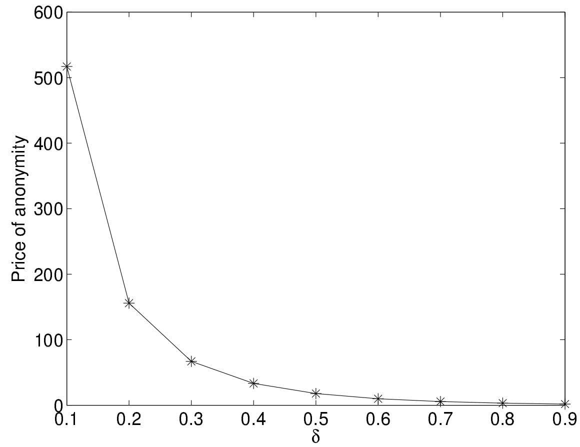

Figure 6: Market convention: Price of anonymity for $N = 6$, $C = 3$, $K = 2$, $\gamma = 0.5$ and varying δ.

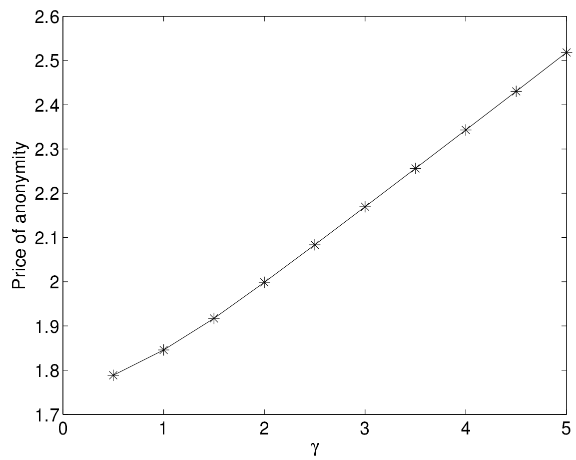

Figure 7: Market convention: Price of anonymity for $N = 6$, $C = 3$, $K = 2$, $\delta = 0.9$ and varying γ.

(note that in Figure 6 the y-axis is logarithmic). This is because for small δ, the benefit of winning the resource right away is much higher than the payoff of winning later. On the other hand, as δ gets closer to 1, the agents don't care whether they win now or later. Since the market convention guarantees that everyone will be able to access some resource for some signal value, when $\delta \to 1$, the expected payoff of winner and losers will be the same. Also, as $\delta \to 1$, the cost the agents have to pay for learning the convention decreases compared to the payoff they obtain after they have learnt it.

When γ increases, the price of anonymity increases. The cost of collision has a direct effect on the expected payoff functions E_A and E_Y. Therefore, the expected equilibrium payoff will be higher if the cost is lower. Changing the γ has no effect on the optimal asymmetric outcome though, since the agents don't have to pay any cost because there are no collisions.

4.2.3 ε-equilibrium

Calculating the equilibrium access probabilities for the market convention is difficult – we need to use a numerical algorithm, and as the number of signals K grows, the number of equations grows exponentially. Therefore, we would like to find access probabilities which are easy to compute and for which the agents' incentive to deviate is *too small*. Indeed, game theory is often interested in ε-*equilibria*, in which no agent can improve her payoff by more than $\varepsilon > 0$.

The market pricing ensures that each agent only wants to access a resource for one signal value. It also doesn't depend on the access probabilities of the agents, only on their utility functions. Once the agents converge to the asynchrony round (i.e. a pure-strategy NE of the resource allocation for every signal value), their future expected payoff will be

$$\frac{K^{-1}}{1-\delta}, \tag{23}$$

and no agent can improve her payoff by deviating since the players are playing a PSNE of the stage game in each round.

If the agents who haven't claimed their resource yet play action A with a constant probability $0 < p_{const} < 1$, the expected time before the reach the asynchrony is finite. (from the properties of the *balls-and-bins* problem, see Section 4.3, or Raab & Steger, 1998). We can prove the following theorem:

Theorem 12. *Suppose that in the N-agent, C-resource allocation game, the agents adopt the market convention with the following implementation: The agents who haven't claimed any resource yet play action A with a constant probability p_{const} (we call this the* constant-probability *implementation). Let $E(\delta)$ be the expected payoff for each agent in this case for a given discount factor δ. Let $E'(\delta)$ be the expected payoff of the best-response strategy to this convention and implementation.*

Then for any $\varepsilon > 0$, there exists $0 < \delta_0 < 1$ such that for all δ, $\delta_0 \leq \delta < 1$,

$$\frac{E(\delta)}{E'(\delta)} > 1 - \varepsilon. \tag{24}$$

Proof. Because of the market pricing, each agent only wants to access one resource for one value of the coordination signal. So the best-response payoff E' is

$$E'(\delta) \leq \frac{K^{-1}}{1-\delta}, \tag{25}$$

no matter what strategy do the other agents play.

When the agents adopt the market convention with the constant-probability implementation, then in every round until they converge to a PSNE, they receive a payoff between $\gamma < 0$ (the collision cost) and 1. After they reach the PSNE, their expected payoff is

$$\frac{K^{-1}}{1-\delta} \tag{26}$$

as stated above. We can therefore say that

$$E(\delta) \geq \sum_{i=0}^{\infty} \Pr(\text{agents reach PSNE in } i \text{ steps}) \cdot \left[\gamma \cdot \frac{1-\delta^i}{1-\delta} + K^{-1} \cdot \frac{\delta^i}{1-\delta}\right] \tag{27}$$

We can define a random variable X such that $X = i$ if the agents reach a PSNE after exactly i steps. From the properties of the expected value, we can ee that

$$E(\delta) \geq \frac{\gamma \cdot (1 - E\left[\delta^X\right]) + K^{-1} \cdot E\left[\delta^X\right]}{1-\delta}. \tag{28}$$

The function $\phi(x) := \delta^x$ is a convex function. From Jensen's inequality (1906), we know that

$$E\left[\delta^X\right] \geq \delta^{E[X]}. \tag{29}$$

Therefore,

$$\frac{E(\delta)}{E'(\delta)} \geq \frac{\gamma \cdot \left(1 - \delta^{E[X]}\right) + K^{-1} \cdot \delta^{E[X]}}{K^{-1}}. \tag{30}$$

The expected time $E[X]$ to reach the PSNE is finite and doesn't depend on δ, so we can treat is as a constant. Because $\delta^{E[X]}$ is continuous in δ, monotonous and $\lim_{\delta \to 1_-} \delta^{E[X]} = 1$, we can see that for a given $\varepsilon > 0$, there exists $0 < \delta_0 < 1$ such that for all δ, $\delta_0 \leq \delta < 1$,

$$\frac{E(\delta)}{E'(\delta)} > 1 - \varepsilon. \tag{31}$$

□

By ensuring that each agent only wants to access some resource for one signal value, the market convention makes the cooperative strategy from our previous work (Cigler & Faltings, 2011) *almost* rational.

4.3 Expected Convergence Time

In this section, we will analyze what is the expected number of rounds the agents need to converge to a perfect allocation of resources (one where all resources are used by exactly one, and there are no collisions). We will first prove an upper bound on the expected number of steps to the convergence for the bourgeois convention, and then present experiments for the market convention.

4.3.1 Bourgeois Convention

In order to prove the convergence of the bourgeois convention, we will describe its execution as a Markov chain. A Markov chain describing the execution of the bourgeois convention in N-agent, C-resource allocation game is a chain whose state at round t is $X_t \in \{0, 1, \ldots, C\}$, where $X_t = c$ means that there are c unclaimed resources at round t.

We are interested in the expected number of rounds it will take to this Markov chain to reach state 0 if it started in state C. This is called the expected hitting time:

Definition 29. (Norris, 1998) **(Hitting time)** Let $(X_t)_{t\geq 0}$ be a Markov chain with state space I. Given a probability space $(\Omega, \Sigma, \Pr)$, the *hitting time* of a subset $A \subset I$ is a random variable $H^A : \Omega \to \{0, 1, \dots\} \cup \{\infty\}$ given by

$$H^A(\omega) = \inf\{t \geq 0 : X_t(\omega) \in A\}$$

In general, the expected hitting time of a set of states A can be found by solving a system of linear equations:

Theorem 13. *The vector of mean hitting times $k^A = E(H^A) = (k_i^A : i \in I)$ is the minimal non-negative solution to the system of linear equations*

$$\begin{cases} k_i^A = 0 & \text{for } i \in A \\ k_i^A = 1 + \sum_{j \notin A} p_{ij} k_j^A & \text{for } i \notin A \end{cases} \tag{32}$$

Solving them analytically for our Markov chain is however difficult. Fortunately, when the Markov chain has only one absorbing state $i = 0$, and it can only move from state i to j if $i \geq j$, we can use the following theorem to derive an upper bound on the hitting time (proved by Rego, 1992):

Theorem 14. *Let $A = \{0\}$. If*

$$\forall i \geq 1 : E(X_{t+1} | X_t = i) < \frac{i}{\beta}$$

for some $\beta > 1$, then

$$k_i^A < \lceil \log_\beta i \rceil + \frac{\beta}{\beta - 1}$$

In order to use the Theorem 14, we need to calculate the expected state $E(X_{t+1}|X_t = c)$.

Lemma 15. *Let $X_t = c$, and let there be $n := N - C + c$ agents who have not claimed a resource yet. Let us denote $q(n, c) = \frac{p}{c} \cdot n \cdot \left(1 - \frac{p}{c}\right)^{n-1}$ that a resource i will be claimed in round t if the agents play the subgame-perfect equilibrium strategy vector described above.*

Then the next expected state is

$$E(X_{t+1}|X_t = c) := (1 - q(n, c)) \cdot c$$

Proof. For a resource i, we can denote W_i the random variable, where $W_i = 1$ if the resource i has been claimed in round t, and $W_i = 0$ otherwise. The random variable W_i is Bernoulli-distributed with probability $q(n, c)$.

The next expected state is then

$$E(X_{t+1}|X_t = c) = c - E\left[\sum_{i=1}^{c} W_i\right] = c - \sum_{i=1}^{c} E[W_i] = (1 - q(n, c)) \cdot c, \tag{33}$$

because $E[W_i] = q(n, c)$. □

In the following lemmas, we will denote

$$\lambda := \frac{|\gamma|}{|\gamma| + \frac{1}{1-\delta}} \tag{34}$$

Lemma 16. *For a given collision cost γ and discount factor δ, there exists a constant $0 < \mu < 1$ such that for $c \leq \mu \cdot n$, $p^* < 1$.*

Proof. According to the definition of the subgame-perfect equilibrium strategy, $p^* := c \cdot \left(1 - \sqrt[n-1]{\lambda}\right)$.

We want $p^* < 1$, which is equivalent to

$$c \cdot \left(1 - \sqrt[n-1]{\lambda}\right) < 1 \tag{35}$$

$$\left(1 - \frac{1}{c}\right)^{n-1} < \lambda \tag{36}$$

$$\tag{37}$$

We know that $c \leq \mu \cdot n$, so

$$\left(1 - \frac{1}{c}\right)^{n-1} \leq \left(1 - \frac{1}{\mu \cdot n}\right)^{n-1} \leq e^{-\mu}. \tag{38}$$

If we therefore set μ such that $e^{-\mu} < \lambda$, the access probability $p^* < 1$. □

Lemma 17. *For given γ and δ, there exists $0 < \eta < 1$ such that for any c,*

$$E(X_{t+1} | X_t = c) \leq (1 - \eta) \cdot c$$

Proof. We will prove the lemma for two cases: when $p^* < 1$ and when $p^* = 1$.

First, let us prove the case $p^* < 1$, that is $p^* = c \cdot \left(1 - \sqrt[n-1]{\lambda}\right)$. Therefore, $q(n,c) = \left(1 - \sqrt[n-1]{\lambda}\right) \cdot n \cdot \lambda$. It can be shown that for any n,

$$q(n,c) = \left(1 - \sqrt[n-1]{\lambda}\right) \cdot n \cdot \lambda \geq -\lambda \log \lambda. \tag{39}$$

Now let $p^* = 1$. From Lemma 16 it must be that $c > \mu \cdot n$. Then

$$q(n,c) := \frac{c}{n} \cdot \left(1 - \frac{1}{c}\right)^{n-1} \geq \mu \cdot \left(1 - \frac{1}{\mu \cdot n}\right)^{n-1}, \tag{40}$$

because $q(n,c)$ is increasing with c.

Now

$$\mu \cdot \left(1 - \frac{1}{\mu \cdot n}\right)^{n-1} \geq \mu \cdot e^{-\mu}. \tag{41}$$

For fixed γ, δ, the μ and λ are constants, so we can set η as

$$\eta := \min(\mu \cdot e^{-\mu}, -\lambda \log \lambda). \tag{42}$$

From above, this proves the lemma. □

Theorem 18. *The expected time for the agents to converge to a resource allocation where all the resources are claimed is $O(\log C)$.*

Proof. We have shown how we can express the expected convergence time as expected hitting time of a certain Markov chain.

From Lemma 17 we saw that there exists η such that for any c,

$$E(X_{t+1}|X_t = c) \leq (1 - \eta) \cdot c.$$

We can now combine this result with Theorem 14 to show that the expected hitting time from the state C to state 0 is

$$k_N^0 < \lceil \log_{\frac{1}{1-\eta}} C \rceil + \frac{1}{\eta} \approx O\left(\frac{1}{\eta} \cdot \log C\right) = O(\log C), \tag{43}$$

because η is a constant.

□

4.3.2 Market Convention

For the market convention, it is unfortunately very difficult to express the expected number of convergence steps in a closed-form expression. However, we can use linear programming to calculate the expected number of convergence steps for a given parameters N, C, K, γ and δ.

The Markov chain for the market convention for K signals and C resources looks as follows: Its state at time t is $V_t \in \{0, 1, \ldots, C\}^K$, where V_{t_k} denotes how many resources have not been claimed for signal k. The initial state V_0 is such that $V_{0_k} = C$ for all $k \in \{1, \ldots, K\}$. If $N \geq C \cdot K$, the final state is when $V_{t_k} = 0$ for all k. When $N < C \cdot K$, the final states are such that $\sum_{k \in \{1,\ldots,K\}} V_{t_k} = C \cdot K - N$.

The transition probabilities between two states V_i and V_j, $V_i \neq V_j$, are the following: Suppose $\exists k : V_{j_k} < V_{i_k}$ and $\forall l \neq k : V_{j_l} \neq V_{i_l}$. Let us denote $c := V_{i_k}$, i.e. the number of unclaimed resources in state V_i for signal k, and $n := N - (C - V_{i_k})$ the number of agents who have not claimed any resource for signal k in state V_i.

$$\Pr(V_{t+1} = V_j | V_t = V_i) := \frac{1}{K} \sum_{m=0}^{n} \binom{n}{m} p_k^m (1 - p_k)^{n-m} \cdot \phi(m, c, V_{i_k} - V_{j_k}) \tag{44}$$

Otherwise if $V_j \neq V_i$, $\Pr(V_{t+1} = V_j | V_t = V_i) := 0$.

Figure 8 shows the expected number of rounds to converge for varying discount factor δ. Generally, we would expect the access probability to increase with increasing δ, since the profit from winning the resource increases. This should then increase the convergence time. However, in our experiments the influence of δ on the convergence time is negligible, although we can observe a slight increase as δ increases. Figure 9 shows the convergence for varying collision cost γ. For γ close to 0, the convergence time remains stable. However, for very high cost γ, the convergence time increases linearly with γ. In this case, the high cost of collision drives the resource access probability low, because agents try to avoid collisions "at all costs".

Figure 10 shows the expected convergence when we increase number of resources C and number of agents N proportionally. The increase in convergence time is still sub-linear to the increase in C.

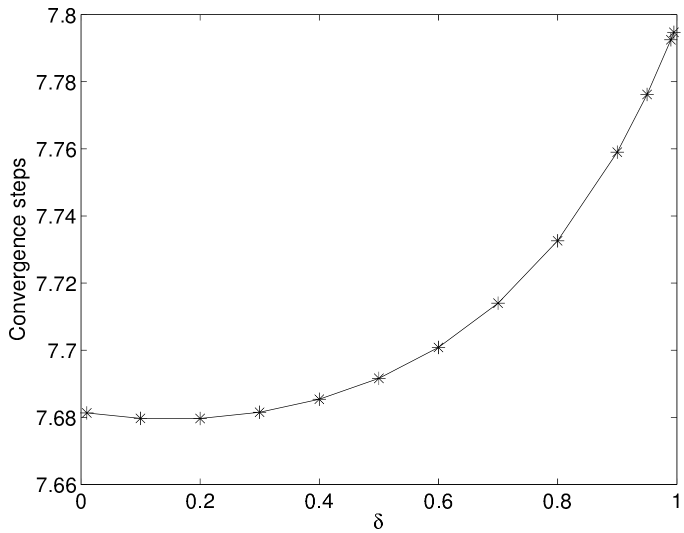

Figure 8: Market convention: Expected number of convergence steps given $N = 6$, $C = 3$, $K = 2$, $\gamma = 1.0$ and varying δ.

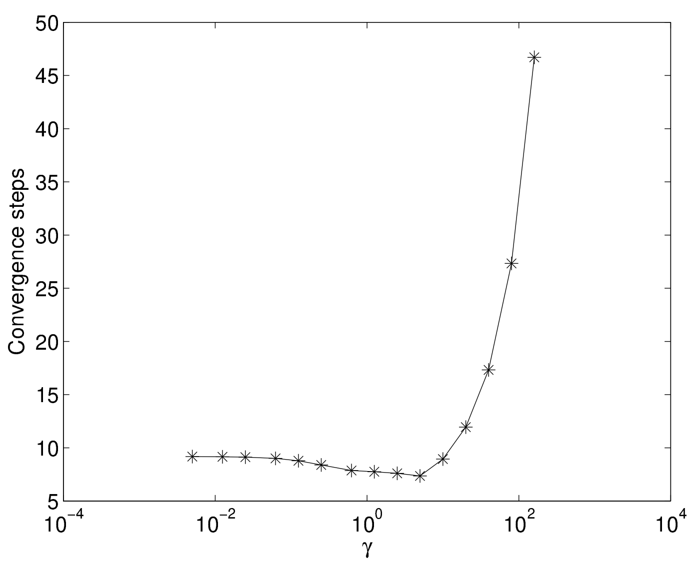

Figure 9: Market convention: Expected number of convergence steps given $N = 6$, $C = 3$, $K = 2$, $\delta = 0.9$ and varying γ.

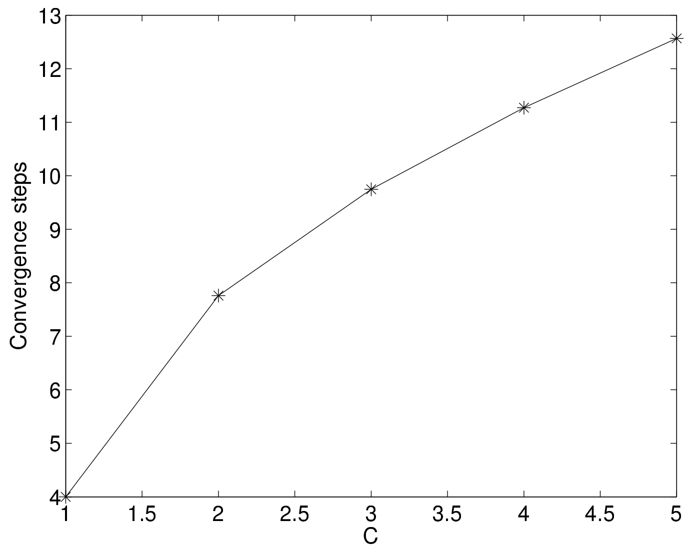

Figure 10: Market convention: Expected number of convergence steps given $K = 2$, $\delta = 0.9$, $\gamma = 1.0$ and varying number of resources C and agents $N = 2 \cdot C$.

	ex-post fair	efficient	equilibrium
C&F'11	(✓)[1]	✓	*no*
Bourgeois	*no*	*no*	✓
Egalitarian[2]	✓	✓	✓
Market	✓	?	✓

Table 1: Properties of conventions

4.4 Convention Properties

We compare the properties of the following conventions: *C&F'11*, a channel allocation algorithm presented in our previous work (Cigler & Faltings, 2011); *bourgeois* and *egalitarian* conventions, presented by Bhaskar (2000); and *market* convention, presented in above.

We compare the conventions according to the following properties:

Ex-post fairness Is the expected payoff to all agents the same *even after asynchrony*?

Efficiency Does the convention maximize social welfare among all possible conventions?

Equilibrium Does the convention have an equilibrium implementation?

Table 1 summarizes the properties of the conventions. The *C&F'11* convention is only approximately ex-post fair. The fairness is improving as the number of coordination signals increases, but some agents might have a worse payoff than others. On the other hand, it is efficient, at least with no discounting ($\delta = 1$). However, it is not an equilibrium

1. Fair asymptotically, as $N \to \infty$.
2. Only for 2-agents, 1-resource games.

of the repeated game. The bourgeois convention is neither fair nor efficient, in fact the expected payoff to the agents is 0 (for a small number of resources). It is has an equilibrium implementation though, since the agents are indifferent between being a winner and a loser. The egalitarian convention is fair, efficient and has an equilibrium implementation. However, it is only defined for games of 2 agents and 1 resource. Finally, the market convention is fair and also has an equilibrium implementation. It is clearly more efficient than the bourgeois convention. Nevertheless, finding the most efficient convention remains an open problem.

5. Folk Theorems and Symmetric Equilibria

In the previous sections, we have analyzed a special kind of symmetric equilibria of the resource allocation game. The agents first followed a Markovian implementation, and as soon as they play a pure-strategy NE, they adopted a convention. For an infinitely repeated game with discounting, the set of Nash equilibria can be characterized using the so-called *folk theorem*. While their name indicates that they have been known and used well before they were first published, we will follow the version described by Fudenberg and Maskin (1986). Informally, the folk theorem states that in the infinitely repeated game, for every feasible and individually rational payoff vector of the stage game, there exists a Nash equilibrium of the repeated game where the average payoffs per round correspond to the stage game payoff vector.

A payoff vector is individually rational if it Pareto-dominates the minimax payoff of the stage game. For player i, the minimax payoff is

$$v_i^* := \min_{\chi_{-i}} \max_{\chi_i} u_i(\chi_i, \chi_{-i}). \tag{45}$$

To simplify the notation, Fudenberg and Maskin (1986) normalize the payoffs so that for the minimax payoffs, $(v_1^*, v_2^*, \ldots, v_N^*) = (0, 0, \ldots, 0)$. Let

$$U := \{(v_1, \ldots, v_N) | \exists (a_1, \ldots, a_N) \in \mathcal{A}_\infty \times \ldots \times \mathcal{A}_\mathcal{N} \text{ s.t. } u(a_1, \ldots, a_N) = (v_1, \ldots, v_N)\},$$

$$V := \text{convex hull of } U,$$

$$V^* := \{(v_1, \ldots, v_N) \in V | v_i > 0 \text{ for all } i\}.$$

The set V is the set of feasible payoffs in the stage games (that is, payoffs which can be achieved by playing some mixed or correlated strategy). The set V^* is the subset of feasible payoffs which are also individually rational.

Theorem 19. *(Fudenberg & Maskin, 1986) For any $(v_1, \ldots, v_N) \in V^*$, if the discount factor δ is close enough to 1, there exists a Nash equilibrium of the infinitely repeated game with discounting where, for all i, the average payoff to player i is v_i.*

The idea of the proof is as follows: The agents cycle through a prescribed sequence of game outcomes so that they achieve the desired payoffs. If one player deviates, the others punish him by playing the minimax strategy forever after.

Our focus so far has been on finding symmetric equilibrium strategies. The folk theorem doesn't say anything about whether the equilibrium strategy will be symmetric, even if the payoff vector is symmetric. Nevertheless, we can define another class of symmetric strategies of the infinitely repeated game, than the one based on conventions and their implementations. The strategies have the following form: The players follow a symmetric (mixed) strategy of the stage game. If one player deviates from this strategy, other players punish her by following the minimax strategy. For the resource allocation game, the minimax payoff is $(0, 0, \ldots, 0)$ and is achieved in the mixed strategy Nash equilibrium. From the Folk theorem, such strategy can be sustained as the Nash equilibrium of the repeated game (though not necessarily a subgame-perfect equilibrium).

A symmetric strategy of the stage resource allocation game is a vector of access probabilities $\mathbf{q} = (q_1, q_2, \ldots, q_C)$ where q_c is the probability that each agent will access resource c. We are interested in finding access probability vector $\mathbf{q}^*$ which achieves the highest expected payoff.

For a given access probability vector $\mathbf{q}$, the expected payoff that an agent receives is as follows:

$$E(\mathbf{q}) := \sum_{j=1}^{C} q_j \cdot \left[(1 - q_j)^{N-1} \cdot 1 - \left(1 - (1 - q_j)^{N-1}\right) \cdot |\gamma|\right] \tag{46}$$

Theorem 20. *For a resource allocation game with $N = 2$ agents and $C = 1$ resource, the resource access probability which maximizes the expected payoff of the stage game is*

$$q^* = \frac{1}{2} \cdot \frac{1}{1 + |\gamma|}. \tag{47}$$

Proof. We calculate the derivative of expected payoff function from Equation 46 for $N = 2$ and $C = 1$:

$$\frac{\partial E(q)}{\partial q} = 1 - 2q \cdot (1 + |\gamma|)$$

Setting the first derivative equal to 0, we get

$$q^* = \frac{1}{2} \cdot \frac{1}{1 + |\gamma|}.$$

Since the second derivative is

$$\frac{\partial^2 E(q)}{\partial^2 q} = -1 - |\gamma| < 0,$$

the probability q^* is a local maximum of the expected payoff function $E(q)$. □

Corollary 2. *For a resource allocation game with $N = 2$ agents and $C = 1$ resource, the highest expected payoff of a symmetric strategy in the stage game is*

$$E^* = \frac{1}{4} \cdot \frac{1}{1 + |\gamma|}. \tag{48}$$

Corollary 3. *For a resource allocation game with $N = 2$ agents and $C = 1$ resource, the price of anonymity of the strategy which accesses with the optimal access probability q^* is $4 \cdot (1 + |\gamma|)$.*

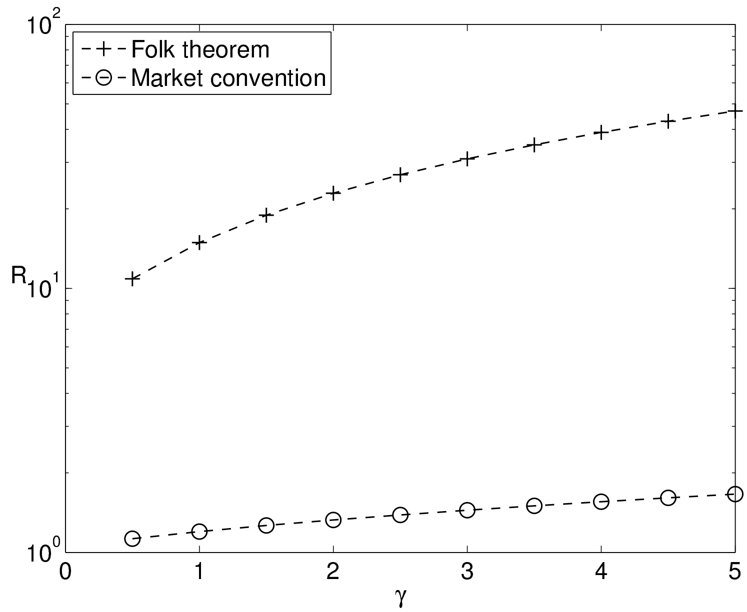

Figure 11: Price of anonymity of the symmetric strategy following from the folk theorem, compared to the price of anonymity of the market convention for $N = 3$, $C = 1$ and varying cost of collision γ.

For the general case of resource allocation game with N agents and C resources, we can find the probability vector which maximizes the Equation 46 (given the constraint $\sum_{j=1}^{C} q_j \leq 1$) using the method of Lagrangian multipliers[2].

2. Our goal is to maximize $E(q_1, q_2, \ldots, q_C)$ such that $q_0 + \sum_{i=1}^{C} q_i = 1$ and $q_i \geq 0$, where q_0 is the probability that an agent yields. The Lagrangian function is then

$$\Lambda(q_0, q_1, \ldots, q_C, \lambda) := E(q_1, q_2, \ldots, q_C) + \lambda \cdot \left(q_0 + \sum_{i=1}^{C} q_i - 1 \right).$$

The first partial derivatives are

$$\frac{\partial \Lambda}{\partial q_0} := \lambda \tag{49}$$

$$\frac{\partial \Lambda}{\partial q_i} := \frac{\partial E}{\partial q_i} + \lambda \tag{50}$$

$$= (1 - |\gamma|) \cdot (1 - q_i)^{N-1} - q_i \cdot (1 + |\gamma|) \cdot (N - 1) \cdot (1 - q_i)^{N-2} - |\gamma| + \lambda \text{ for } 1 \leq i \leq C \tag{51}$$

$$\frac{\partial \Lambda}{\partial \lambda} := q_0 + \sum_{i=1}^{C} -1. \tag{52}$$

A necessary condition for a solution to be a maximum is that the partial derivatives of the Lagrangian function are 0. Therefore, $\lambda := 0$ and $q_0 := 1 - \sum_{i=1}^{C} q_i$. For q_i, $1 \leq i \leq C$ we can find the solution to $\frac{\partial \Lambda}{\partial q_i} = 0$ using a numerical root-finding algorithm. If there is no point where the partial derivatives are zero, we have to compare $E(\mathbf{q})$ for a (finite) number of extreme points.

Figure 11 compares the price of anonymity of the folk-theorem-based symmetric strategy with the price of anonymity of the market convention, for $N = 3$ agents and $C = 1$ resource. Since the price of anonymity of the folk theorem strategy doesn't depend on the discount factor δ (it only needs to be high enough for the strategy to be an equilibrium), we only show the graph for varying collision cost γ. The price of anonymity of the folk-theorem strategy is an order of magnitude higher than the price of anonymity of the market convention.

6. Conclusions

In this paper, we considered symmetric protocols to rationally coordinate on an asymmetric, efficient allocation infinitely repeated resource allocation game with discounting of N agents and C resources. We assumed that the agents are identical, and that the resources are homogeneous. We based our work on the idea of Bhaskar (2000): we let the agents choose their actions randomly, after which they adopt a certain convention. We generalize the work of Bhaskar to an arbitrary resource allocation problem with N agents and C resources. We show that for any convention, there exists a symmetric subgame-perfect equilibrium that implements it. We presented two such conventions for the repeated resource allocation game: bourgeois and market convention. We defined the price of anonymity as the ratio between the expected social payoff of the best asymmetric strategy profile and the expected social payoff of a given symmetric Nash equilibrium. We showed that while the price of anonymity for the bourgeois convention is infinite (at least for small number of resources), the price of anonymity of the market convention is finite and relatively small. This is because the market convention reduces the demand by imposing a price on successful access, while at the same time increasing the supply by having the agents condition their strategy on a global coordination signal $k \in \{1, \ldots, K\}$. This way, the conflict between the agents is reduced. We also showed analytically that when the agents adopt the bourgeois convention, they will converge to a perfect resource allocation in polynomial time.

For the market convention, calculating the equilibrium access probabilities is difficult. We need to use a numerical algorithm, whose complexity is exponential in the number of coordination signals K. However, the market mechanism already makes sure that no agent wants to access some resource for more than one coordination signal. Therefore, we showed that the cooperative solution where agents access the resources with constant probability is an ε-equilibrium, given that the discount factor δ is high enough.

In the future work, we would like to investigate whether there exist more efficient conventions than the market convention (i.e. conventions with smaller price of anonymity). In general, finding an optimal convention is an NP-hard problem (Balan, Richards, & Luke, 2011), but for a more restricted set of infinitely repeated resource allocation games, we might be able to find the optimal convention, similar to the Thue-Morse sequence (Richman, 2001) used by Kuzmics et al. (2010) in the Nash demand game.

Acknowledgements

We are particularly thankful to David Parkes for giving the first author the unique opportunity to spend a few weeks in his lab at Harvard. David's openness and unparalleled knowledge is what really helped originate this work. We would also like to thank Kate

Larson for reading the draft of this paper and helping make the theoretical analysis much more readable.

References

Aumann, R. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, *1*(1), 67–96.

Balan, G., Richards, D., & Luke, S. (2011). Long-term fairness with bounded worst-case losses. *Autonomous Agents and Multi-Agent Systems*, *22*(1), 43–63.

Bhaskar, V. (2000). Egalitarianism and efficiency in repeated symmetric games. *Games and Economic Behavior*, *32*(2), 247–262.

Bonnet, F., & Raynal, M. (2011). The price of anonymity: Optimal consensus despite asynchrony, crash, and anonymity. *ACM Trans. Auton. Adapt. Syst.*, *6*(4), 23:1–23:28.

Chothia, T., & Chatzikokolakis, K. (2005). A survey of anonymous peer-to-peer file-sharing. In *Embedded and Ubiquitous Computing–EUC 2005 Workshops*, pp. 744–755. Springer.

Cigler, L., & Faltings, B. (2011). Reaching correlated equilibria through multi-agent learning. In *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pp. 509–516. International Foundation for Autonomous Agents and Multiagent Systems.

Durresi, A., Paruchuri, V., Durresi, M., & Barolli, L. (2005). A hierarchical anonymous communication protocol for sensor networks. In *Embedded and Ubiquitous Computing–EUC 2005*, pp. 1123–1132. Springer.

Fudenberg, D., & Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, *54*(3), 533–554.

Jensen, J. L. (1906). Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, *30*(1), 175–193.

Koutsoupias, E., & Papadimitriou, C. (1999). Worst-case equilibria. In *Proceedings of the 16th annual conference on Theoretical aspects of computer science*, STACS'99, pp. 404–413, Berlin, Heidelberg. Springer-Verlag.

Kuzmics, C., Palfrey, T., & Rogers, B. (2010). Symmetric players in repeated games: Theory and evidence..

Leyton-Brown, K., & Shoham, Y. (2008). *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan & Claypool, San Rafael, CA.

Monderer, D., & Shapley, L. S. (1996). Potential games. *Games and economic behavior*, *14*(1), 124–143.

Norris, J. R. (1998). *Markov Chains (Cambridge Series in Statistical and Probabilistic Mathematics)*. Cambridge University Press.

Raab, M., & Steger, A. (1998). *Balls into Bins A Simple and Tight Analysis*, Vol. 1518 of *Lecture Notes in Computer Science*, chap. 13, pp. 159–170. Springer Berlin Heidelberg, Berlin, Heidelberg.

Rego, V. (1992). Naive asymptotics for hitting time bounds in markov chains. *Acta Informatica*, *29*(6), 579–594.

Richman, R. M. (2001). Recursive binary sequences of differences. *Complex Systems*, *13*(4), 381–392.

Rosenthal, R. W. (1973). A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, *2*(1), 65–67.

Shneidman, J., Ng, C., Parkes, D. C., Auyoung, A., Snoeren, A. C., Vahdat, A., & Chun, B. (2005). Why markets could (but don't currently) solve resource allocation problems in systems. In *In USENIX 05: Proceedings of the 10th USENIX Workshop on Hot Topics in Operating Systems*, p. 7.

Taussky, O. (1949). A recurring theorem on determinants. *The American Mathematical Monthly*, *56*(10), 672–676.

Wang, L., Wu, K., Hamdi, M., & Ni, L. M. (2011). Attachment learning for multi-channel allocation in distributed OFDMA networks. *Parallel and Distributed Systems, International Conference on*, *0*, 520–527.