



Intense Competition can Drive Selfish Explorers to Optimize Coverage

Simon Collet, Amos Korman

► To cite this version:

Simon Collet, Amos Korman. Intense Competition can Drive Selfish Explorers to Optimize Coverage. SPAA, Jul 2018, Vienna, Austria. <hal-01783769>

HAL Id: hal-01783769

<https://hal.science/hal-01783769v1>

Submitted on 2 May 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Intense Competition can Drive Selfish Explorers to Optimize Coverage*

Simon Collet[†]

Amos Korman[‡]

Abstract

We consider a game-theoretic setting in which selfish individuals compete over resources of varying quality. The motivating example is a group of animals that disperse over patches of food of different abundances. In such scenarios, individuals are biased towards selecting the higher quality patches, while, at the same time, aiming to avoid costly collisions or overlaps. Our goal is to investigate the impact of collision costs on the parallel coverage of resources by the whole group.

Consider M sites, where a site x has value $f(x)$. We think of $f(x)$ as the reward associated with site x , and assume that if a single individual visits x exclusively, it receives this exact reward. Typically, we assume that if $\ell > 1$ individuals visit x then each receives at most $f(x)/\ell$. In particular, when competition costs are high, each individual might receive an amount strictly less than $f(x)/\ell$, which could even be negative. Conversely, modeling cooperation at a site, we also consider cases where each one gets more than $f(x)/\ell$. There are k identical players that compete over the rewards. They independently act in parallel, in a one-shot scenario, each specifying a single site to visit, without knowing which sites are explored by others. The group performance is evaluated by the expected coverage, defined as the sum of $f(x)$ over all sites that are explored by at least one player. Since we assume that players cannot coordinate before choosing their site we focus on symmetric strategies.

The main takeaway message of this paper is that the optimal symmetric coverage is expected to emerge when collision costs are relatively high, so that the following “Judgment of Solomon” type of rule holds: If a single player explores a site x then it gains its full reward $f(x)$, but if several players explore it, then neither one receives any reward. Under this policy, it turns out that there exists a unique symmetric Nash Equilibrium strategy, which is, in fact, evolutionary stable. Moreover, this strategy yields the best possible coverage among all symmetric strategies. Viewing the coverage measure as the social welfare, this policy thus enjoys a (Symmetric) Price of Anarchy of precisely 1, whereas, in fact, any other congestion policy has a price strictly greater than 1.

Our model falls within the scope of mechanism design, and more precisely in the area of incentivizing exploration. It finds relevance in evolutionary ecology, and further connects to studies on Bayesian parallel search algorithms.

1 Introduction

Studying the way humans and other animals distribute themselves in their environment is a cornerstone of ecology and the social sciences [4, 17, 19, 21]. In many of these contexts, dispersal is governed by two contradicting forces. On the one hand, individuals are biased towards selecting the higher quality patches, and, on the other hand, they aim to avoid costly collisions or overlaps, which can significantly deteriorate the value of a patch [16, 22, 44]. In nature, collision costs can be caused by various factors, including aggressive behavior, or merely due to sharing the patch equally between colliding individuals

*This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 648032).

[†]IRIF, CNRS and University Paris Diderot, Paris, France. E-mail: Simon.Collet@irif.fr.

[‡]IRIF, CNRS and University Paris Diderot, Paris, France. E-mail: Amos.Korman@irif.fr.

(a.k.a., scramble competition [32]). The purpose of this paper is to investigate the impact of collision costs on the overall coverage of resources by the whole group.

Understanding the impact of collisions on the overall parallel coverage can also have implications to the way centralized entities incentive individuals in human organizations. For example, research foundations are often interested in promoting specific research topics, and offer grants to attract researchers to such topics. The question of how to properly define a grant policy that is expected to yield a desirable distribution of researchers can have significant implication on the scientific community [23].

This paper proposes a framework to study the impact of collision on the group performances through an *algorithmic mechanism design* approach. We focus on a relatively simple scenario, which already provides interesting, and perhaps surprising, insights. See Section 5.1 for suggested generalizations left for future work.

1.1 The Dispersal Game

Think about the following imaginary scenario. A group of individuals are searching for resources in some space. Eventually, they gather all resources found for the benefit of the public, and hence the group performances is defined as the sum of the resources values. In order to avoid the free-rider problem, the group incentivizes individuals by granting those who find more resource with higher social status. When an individual exclusively finds a resource of value f , a simple mechanism is to let its social status be proportional to f . However, how should the group define the social status of each of ℓ individuals in case they all found a resource simultaneously?

Formally, we have M sites indexed by $x \in [1, M]$. Each site x has some *importance value* $f(x)$, and we assume without loss of generality that sites are ordered such that lower index sites have higher values, i.e., $f(x) \geq f(x+1)$ for each $x < M$.

We have k players that act a one-shot scenario, each specifying a site x to visit. Crucially, they cannot coordinate, and each player must commit to a site x , without knowing which sites are selected by other players [14]. Moreover, we assume that players cannot change their mind after committing. Formally, each player specifies an index x according to some law, called *strategy*. This is simply a probability distribution, indicating, for each x , the probability $p(x)$ to explore x . A *strategy profile* is a vector of k strategies, one for each player. It is called *symmetric* when all players play the same strategy. For short, we shall use the term *symmetric strategy* to refer to the strategy being played in the case of a symmetric strategy profile.

The group performance is evaluated by the *weighted coverage*, defined as the sum of the values $f(x)$ over visited sites x . With the lack of coordination between users, the most the group can hope for is the best achievable coverage among all symmetric strategies, that is, when assuming that all players play the same strategy. Formally, the *coverage* of a strategy p is defined as:

$$\text{Cover}(p) = \sum_{x=1}^M f(x)(1 - (1 - p(x))^k) \quad (1)$$

Let p^* be a strategy that maximizes $\text{Cover}(p)$. Using a compactness argument, it is not difficult to show that such a strategy exists. The following observation implies that the optimal symmetric strategy p^* has coverage which is, up to a small constant factor, as good as the best achievable coverage in the full coordination case, when the k users can be assigned to precisely cover the k highest value sites.

Observation 1. $\text{Cover}(p^*) > (1 - \frac{1}{e}) \sum_{x=1}^k f(x)$.

Proof. Let \hat{p} be the uniform distribution in $[k]$, that is, we have $\hat{p}(x) = 1/k$ for each $x \leq k$. Then, $\text{Cover}(p^*) \geq \text{Cover}(\hat{p}) = \sum_{x=1}^k f(x) \left(1 - (1 - \frac{1}{k})^k\right) > (1 - \frac{1}{e}) \sum_{x=1}^k f(x)$. \square

A *reward policy* is a function $I(x, \ell)$ specifying the *payoff* (or reward) that a player exploring x receives, given that $\ell - 1$ other players are also exploring x . We shall be particularly interested in *congestion* reward policies, which can be written as:

$$I(x, \ell) = f(x) \cdot C(\ell)$$

where $C(\cdot)$ is some *congestion* function, for which $C(1) = 1$, and C is non-increasing. Note that the specification of the policy does not depend on the total number of players k , and is influenced only by the number of players that visit x .

The assumption that $C(1) = 1$ implies that a player visiting x exclusively will gain exactly¹ $f(x)$. A natural example of a congestion policy is the *sharing policy* in which $C_{share}(\ell) = 1/\ell$, and hence if ℓ players select site x in parallel then each receives a payoff of $f(x)/\ell$. This policy has received a lot of attention in both the ecology literature and algorithmic game theory community due to its simplicity [22, 23, 44]. Cases where $C(\ell) > 1/\ell$ can model cooperation, in the sense that the presence of several players at a site would result in each having more than their relative share. On the other extreme, note that we do not restrict C to be a positive function. Taking $C(\ell)$ to be negative implies that the presence of ℓ players at x damages each of them quite significantly. In the context of animals, this could represent aggressive behavior that increases with the amount of potential benefit $f(x)$ and could result in severe injuries.

Let X_i be the site specified by player i . Given a policy I , the goal of each player i is to maximize the *expected payoff* it receives, defined as

$$\mathbb{E}\left[I(X_i, \sum_j \mathbb{1}_{X_j=X_i})\right]$$

1.2 Evaluating Policies

The game we define is symmetric and possesses both mixed symmetric equilibria and pure non-symmetric equilibria. Most previous works dealing with such situations prefer to focus on pure equilibria (rather than on the mixed equilibria), because of their conceptual simplicity. The classical example of this situation is the well-know game “battle of the sexes” [34]. However we argue otherwise: the number of pure equilibria grows exponentially with the number of players, and choosing an equilibrium among those requires coordination between the players. In some contexts, because coordination is not possible or communication is too costly, it does not seem reasonable to expect that a pure equilibrium arises in practice, and instead, it is more natural to expect the emergence of a symmetric equilibrium. Since we focus on large distributed systems, where players do not always possess individual identities and often have identical behaviors, we will restrict our attention to symmetric equilibria, and disregard pure equilibria².

Viewing p^* as the symmetric strategy that maximizes the social welfare, we adopt an interpretation of the concept of Price of Anarchy to evaluate a policy $I(x, \ell)$, by measuring the coverage of the worst symmetric Nash Equilibrium under $I(x, \ell)$ in comparison with the coverage of p^* . Specifically, let $C(\ell)$ be

¹Alternatively, we could have defined $C(1)$ to be any constant that doesn’t depend on x . This would mean that the reward a single player gets when visiting x exclusively is proportional to the value $f(x)$. (This assumption contrasts with the setting of [23], see more details in Section 1.6.) Note that in this sense, $f(x)$ plays a double role - it has a meaning from the perspective of the group and a meaning from the perspective of an individual.

²At first glance, our focus on symmetric strategies may seem non-compatible with the fact that conspecific animals often appear not to have identical strategies. However, it is important to note that a symmetric strategy profile in the context of ESS (see Section 1.4) does not imply that every player follows the same strategy! Indeed, a symmetric strategy profile can arise also if the population contains pure strategies in proportions as specified by the mixed strategy. In this case, since the framework of ESS assumes that at each time k players are selected at random from the population, each of the selected random players ends up having the same mixed strategy. For this reason, restricting attention to symmetric equilibria is in fact very common in the ecology literature [16, 22, 44].

a congestion function, and $f(x)$ be an importance value function. Define

$$\text{SPoA}(C, f) := \sup_{p \in \mathcal{P}(C, f)} \left\{ \frac{\text{Cover}(p^*)}{\text{Cover}(p)} \right\},$$

where $\mathcal{P}(C, f)$ is the set of symmetric Nash Equilibrium under $I(x, \ell) = f(x) \cdot C(\ell)$. The *Symmetric Price of Anarchy* (SPoA) of the congestion function $C(\ell)$ is defined as the sup of $\text{SPoA}(C, f)$ over all functions $f : [M] \rightarrow \mathbb{R}^+$, where we take the freedom to choose their domain M as well, i.e.,

$$\text{SPoA}(C) := \sup_{\substack{M \in \mathbb{N} \\ f: [M] \rightarrow \mathbb{R}^+}} \{ \text{SPoA}(C, f) \}$$

By definition, $\text{SPoA}(C) \geq 1$ for any congestion function C . On the other hand, there exist congestion functions with much higher SPoA . For example, taking $C(\ell) \equiv 1$ yields SPoA of roughly k . Indeed, under this function, and assuming a strictly decreasing f , a strategy at equilibrium would explore the first site with probability 1, and for distributions f that decrease very slowly (and thus are close to uniform), this would yield a gap $\Omega(k)$, with respect to, e.g., the uniform strategy. Note, however, that this policy is unlikely to occur in nature, as it implies that a value at a site is fully obtained by all visitors.

The PoA of the sharing policy C_{share} was studied by Kleinberg and Oren in a somewhat similar model [23] (see also Section 1.6). Thanks to a more general result from Vetta [45], they showed that the SPoA of the sharing policy, defined as the ratio between the best achievable coverage and the worst coverage among Nash equilibria, is at most 2. Since the worst symmetric equilibrium coverage is at least as good as the worst equilibrium coverage, and the best achievable coverage is at least as good as the best achievable coverage with a symmetric strategy profile, this directly implies that our notion of SPoA also satisfies $\text{SPoA}(C_{\text{share}}) \leq 2$.

1.3 Ideal Free Distribution (IFD)

The setting of animals competing over patches of resources, often referred to as *playing the field*, has been extensively studied in the ecology discipline through the theory of *Ideal Free Distribution (IFD)*, see reviews in [22, 44]. Introduced by Fretwell and Lucas [16], the standard model assumes that each individual follows the same distribution $p(\cdot)$ (i.e., a symmetric strategy) and hence the fraction of the population that is expected to occupy a site x is $p(x)$.

An IFD is a probability distribution $p(\cdot)$ in which every site which is explored with positive probability yields the same payoff, and every other site yields a lower payoff if explored. More precisely, as before, let $I(x, \ell)$ denote the payoff received by an individual selecting a site x together with $\ell - 1$ other individuals. Fix a player and let $P(x, \ell) = \binom{k-1}{\ell} p(x)^\ell (1 - p(x))^{k-\ell-1}$ be the probability that among $k - 1$ players, ℓ of them selected site x . The *value_p* of a site x corresponds to the expected gain for exploring site x and is defined as:

$$\nu_p(x) := \sum_{\ell=1}^k I(x, \ell) \cdot P(x, \ell - 1) \quad (2)$$

By definition, the expected payoff that a player receives is $\sum_x p(x) \cdot \nu_p(x)$. The IFD is a distribution p , that for some integer W , satisfies the following *IFD conditions*.

- $p(x) > 0 \iff x \leq W$
- $\nu_p(1) = \nu_p(2) = \dots = \nu_p(W) = \nu_p$, and for all $W + 1 \leq x$, we have $\nu_p > \nu_p(x)$

The following observation is well-known, see e.g., [8, 16]. For completeness, a sketch of the proof is provided in Appendix A.

Observation 2. Assume that $I(x, \cdot)$ is non-increasing for any x (note that this is indeed the case for congestion policies). There exists a unique strategy satisfying the IFD conditions. Moreover, this IFD is the only symmetric Nash Equilibrium of the game.

1.4 Evolutionary Stable Strategy (ESS)

In addition to the notion of Nash Equilibrium, we shall adopt the stronger concept of an *Evolutionary Stable Strategy (ESS)* [3–5, 18, 28, 43]. This concept has become a cornerstone of evolutionary game theory, and has been extensively used to study the evolution of animal and human behavior [4, 28]. Informally, an ESS is a strategy σ such that, in a population composed of a majority of σ -strategists and a minority of “mutants” playing strategy π , the σ -strategists have a higher payoff, provided that the proportion of mutants is not too high.

The standard view of ESS is to consider an infinite population whose individuals play repeatedly against each other in pairs, chosen uniformly at random. In our setting, individuals play in groups of k , and hence a generalization is required. One possible generalization is to consider a fixed population of k individuals that play the game between themselves. Unfortunately, generalizing ESS to finite populations is quite challenging, and the known generalizations are often too restrictive to be applicable. Instead, we chose the generalization in which the population remains infinite and players are randomly matched in k -tuples, for any fixed $k \geq 2$, see e.g., the book [4]. Modeling biological settings, this generalization may correspond, for example, to scenarios in which a large colony of bats breaks on a daily basis into smaller groups, each of which forages in a different field of patched resources, see [10, 29].

Formally, let $\mathcal{E}(\rho; \sigma^\ell, \pi^{k-\ell-1})$ denote the payoff received by an individual playing strategy ρ against ℓ players playing strategy σ and $(k - \ell - 1)$ players playing strategy π . When $\sigma = \pi$, we use the abridged notation $\mathcal{E}(\rho; \sigma^{k-1})$. Consider a population in which a proportion $(1 - \epsilon)$ of it are playing strategy σ and a proportion of ϵ are playing π . The average payoff of a player playing ρ matched against $k - 1$ opponents drawn uniformly at random from this population is denoted by $\mathcal{U}[\rho; (1 - \epsilon)\sigma + \epsilon\pi]$, and is equal to

$$\mathcal{U}[\rho; (1 - \epsilon)\sigma + \epsilon\pi] = \sum_{\ell=0}^{k-1} \binom{k-1}{\ell} (1 - \epsilon)^\ell \epsilon^{k-\ell-1} \mathcal{E}(\rho; \sigma^\ell, \pi^{k-\ell-1}) \quad (3)$$

A strategy σ is an ESS if for every other strategy π , there exist $\epsilon_\pi > 0$ such that for every $\epsilon \leq \epsilon_\pi$, $\mathcal{U}[\sigma, (1 - \epsilon)\sigma + \epsilon\pi] > \mathcal{U}[\pi, (1 - \epsilon)\sigma + \epsilon\pi]$, i.e., σ cannot be “invaded” by π , as long as the proportion of π -strategists remains below ϵ_π .

We also recall the following *characterization of ESS* [4]: a strategy σ is an ESS if for every other strategy π , there is an integer $0 \leq m_\pi \leq k - 1$, such that both conditions below hold:

- $\mathcal{E}(\sigma; \sigma^{k-m_\pi-1}, \pi^{m_\pi}) > \mathcal{E}(\pi; \sigma^{k-m_\pi-1}, \pi^{m_\pi})$
- $\forall \ell, 0 \leq \ell < m_\pi$, we have $\mathcal{E}(\sigma; \sigma^{k-\ell-1}, \pi^\ell) = \mathcal{E}(\pi; \sigma^{k-\ell-1}, \pi^\ell)$

It should be clear from this definition that every ESS is a symmetric Nash equilibrium of the k -player game. Indeed, these conditions ensure that $\mathcal{E}(\sigma; \sigma^{k-1}) \geq \mathcal{E}(\pi; \sigma^{k-1})$ for every $\pi \neq \sigma$. Further details about the ESS notion appear in Section 5.1.

1.5 Our Results

We study the interplay between the congestions and the expected coverage of strategies at equilibrium. Our main takeaway message is that optimal coverage is expected to emerge by a congestion policy which is defined by a “Judgment of Solomon” type of rule: If a single player attempts to exploit a site then it gains the full reward $f(x)$, but when several players attempt to exploit it, they all receive nothing. Formally, define

the *exclusive congestion function* as follows:

$$C_{exc}(\ell) = \begin{cases} 1 & \text{if } \ell = 1 \\ 0 & \text{otherwise} \end{cases}$$

and define the *exclusive reward policy* as

$$I_{exc}(\ell, x) = f(x) \cdot C_{exc}(\ell)$$

We identify I_{exc} as a critical reward policy, in the sense that its IFD yields optimal coverage, while any other congestion policy yields strictly worse coverage.

More formally, let σ^* be the IFD for the exclusive reward policy. We first claim that under this policy, σ^* is not only a Nash Equilibrium but also an ESS.

Theorem 3. σ^* is an ESS under I_{exc} .

Next, we claim that σ^* yields the best possible coverage, among all symmetric strategies.

Theorem 4. For any strategy σ , $\text{Cover}(\sigma^*) \geq \text{Cover}(\sigma)$, with equality if and only if $\sigma = \sigma^*$.

Since the IFD is the only symmetric Nash Equilibrium under I_{exc} , the price of anarchy equals 1.

Corollary 5. $\text{SPoA}(C_{exc}) = 1$.

The criticality of I_{exc} follows by the fact that among all congestion policies, it is the only one whose IFD yields optimal coverage for all value functions f .

Theorem 6. For any congestion function $C \neq C_{exc}$, we have $\text{SPoA}(C) > 1$.

We find the fact that $\text{SPoA}(C_{exc}) = 1$ rather surprising. Indeed, although it appears intuitive that increasing the competition will result in better coverage, the exclusive policy C_{exc} is, in fact, not the one with highest levels of competition. Indeed, one could define a congestion policy where in case of conflict, each of the colliding players receives a negative payoff (i.e., $C(\ell)$ being negative, see Section 1.1). This means that the competition level could significantly exceed the level of C_{exc} . The fact that the best coverage occurs exactly at that level is thus quite surprising. A second important factor to note is that even if one identifies the best policy (i.e., the one for which the SPoA is smallest), it is a priori unclear that the resulted coverage would actually be optimal among all symmetric strategies (including the non-competing ones), i.e., that $\text{SPoA} = 1$.

An informal discussion regarding the implications of our results in the context of animal behavior is presented in Section 5.2.

See Figure 1 for an illustration of the coverage as a function of the competition extent, in specific instances of 2 players and 2 sites.

1.6 Related Work

A research theme that is currently quite popular in the Operations Research, and Economic Theory research communities goes by the name “incentivizing exploration”. Similarly to us, such papers tend to focus on the problem of designing mechanisms to coordinate the activities of independent selfish agents exploring a space of alternatives [15, 25, 26, 35]. These papers however typically focus on sequential rather than simultaneous exploration, leading them to model the problem as a multi-armed bandit problem rather than a congestion game. For this reason, there is no real technical connection between the two lines of work, but there is nevertheless a thematic connection.

Closest to our work is the paper by Kleinberg and Oren [23] which considers a general model that is motivated by research foundations offering grants to incentivize researchers to work on particular topics.

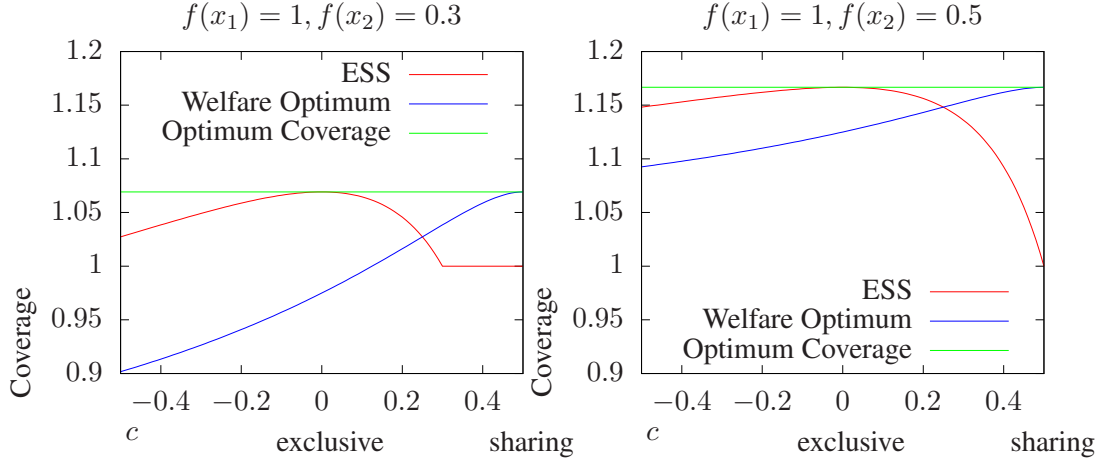


Figure 1: Coverage as a function of the extent of competition (parametrized by c) for the case where two players complete over two sites x_1 and x_2 . On the left (respectively, right) we consider the case where $f(x_1) = 1$ and $f(x_2) = 0.3$ (respectively, $f(x_2) = 0.5$). The X-axis represents a range of congestion functions C_c , defined as $C_c(1) = 1$ and $C_c(2) = c$, where $c \in [-0.5, 0.5]$. The case $c = 0$ corresponds to the exclusive function and the case $c = 0.5$ corresponds to the sharing function. The Y-axis represents the coverage. The red line corresponds to the coverage of the ESS, the green corresponds to the optimum coverage, taken over all symmetric strategies, and the blue corresponds to the coverage of the symmetric strategy that maximizes the individual payoffs.

That paper proposes two mechanisms for selecting grant policies that are expected to incentivize researchers to adopt an optimal distribution, while assuming that researchers are restricted to the sharing policy. The first mechanism relies on the ability to freely choose the rewards (grants) associated with sites (topics), despite the fact that the significance of a site from the perspective of the group is fixed. In contexts of animals, however, the overall rewards (termed here as $f(x)$) correspond to the amount of food in patch x and therefore cannot be modified. Moreover, implementing their mechanism of choosing rewards requires the knowledge of the number of players k , whereas the specification of the congestion policies studied here does not require this knowledge (see Section 1.1). The second mechanism proposed in [23] requires that the players receive different rewards, set in a particular way by the central entity, despite doing the same actions. This might be seen as unfair and unsatisfying. Moreover, implementing this would require high levels of coordination which again seem unlikely in the context of animals.

The settings of selfish routing, job scheduling, and congested games [30, 42] all bear similarities to the dispersal game, however, the measurement of the global performances considered here, namely, the coverage, is very different from the measures studied in the former frameworks, such as makespan or latency [1, 2, 9, 33].

Finally, many articles have informally mentioned the IFD as an example of an ESS [22, 31, 36, 37, 39, 40]. However, on a formal level, it is only recently that a rigorous proof was given [7], and even this proof was with respect a limited definition of ESS (essentially, in our notation, they proved the first item of the characterization with $m_\pi = k - 1$, without considering the second item). Since that paper, several other works have also addressed this question rigorously, but, to the best of our knowledge, they all considered either weaker versions than our ESS definition, or other contexts, such as continuous populations dynamics [6, 41].

2 The Exclusive Reward Policy

This section considers the exclusive reward policy I_{exc} . In Section 2.1 we explicitly calculate the IFD for this policy, called σ^* . Then, in Section 2.2, we prove that σ^* yields optimal coverage.

2.1 Algorithm σ^*

Consider a strategy p that satisfies the IFD conditions for the exclusive reward policy. We know that it exists by Observation 2 and we calculate it explicitly. In the context of C_{exc} , the $value_p$ of a site x , as given by Eq. (2), is:

$$\nu_p(x) = f(x)(1 - p(x))^{k-1}$$

The second IFD condition implies that for any two $x, y \in [1, W]$, we have $f(x)(1 - p(x))^{k-1} = f(y)(1 - p(y))^{k-1}$, or in other words,

$$\frac{1 - p(x)}{1 - p(y)} = \left(\frac{f(y)}{f(x)} \right)^{1/(k-1)} \quad (4)$$

A natural guess for a distribution that satisfies this is the following Pareto distribution:

$$p(x) = \begin{cases} 1 - \frac{\alpha}{f(x)^{\frac{1}{k-1}}} & \text{if } x \leq W \\ 0 & \text{otherwise} \end{cases}$$

Since p is a distribution, we must have $\sum_{x=1}^W p(x) = 1$, from which we can extract

$$\alpha = \frac{W - 1}{\sum_{x \leq W} f(x)^{-\frac{1}{k-1}}}.$$

To complete the description of our candidate IFD strategy, it remains to calculate W . By the second IFD condition, we get that if $W < M$, then for every $x \leq W$:

$$f(x)(1 - p(x))^{k-1} = \nu_p(x) > \nu_p(W + 1) = f(W + 1) \quad (5)$$

Rearranging, we get

$$p(x) < 1 - \left(\frac{f(W + 1)}{f(x)} \right)^{\frac{1}{k-1}}.$$

Since $p(x)$ is a distribution whose support is $[W] := \{1, \dots, W\}$ we get

$$1 = \sum_{x \leq W} p(x) < \sum_{x \leq W} \left(1 - \left(\frac{f(W + 1)}{f(x)} \right)^{\frac{1}{k-1}} \right).$$

This means that we can define W as the first index that satisfies the equation above, if one exists, or M otherwise. Equivalently, we can define W as the largest index that satisfies

$$\sum_{x \leq W} \left(1 - \left(\frac{f(W)}{f(x)} \right)^{\frac{1}{k-1}} \right) \leq 1$$

Altogether, this leads to strategy σ^* , as defined below.

$$\sigma^*(x) = \begin{cases} 1 - \frac{\alpha}{f(x)^{\frac{1}{k-1}}} & \text{if } x \leq W \\ 0 & \text{otherwise} \end{cases}$$

where W and the normalization factor α are calculated as follows.

$$W = \arg \max_{y \leq M} \left\{ \sum_{x \leq y} \left(1 - \left(\frac{f(y)}{f(x)} \right)^{\frac{1}{k-1}} \right) \leq 1 \right\}$$

$$\alpha = \frac{W - 1}{\sum_{x \leq W} f(x)^{-\frac{1}{k-1}}}$$

Pseudocode of Algorithm $\sigma^*(x)$

Interestingly, it turns out that algorithm σ^* is actually identical to the first round in the algorithm A^* used in [24] in the context of Bayesian search.

The following claim verifies that the constructed strategy is indeed the IFD.

Claim 7. σ^* satisfies the IFD conditions under the exclusive reward policy I_{exc} .

Proof. The first IFD condition is satisfied by σ^* by construction. It remains to show that $\nu_p(1) = \nu_p(2) = \dots = \nu_p(W)$, and for all $x > W$, we have $\nu_p(W) > \nu_p(x)$. Under the exclusive reward policy I_{exc} , the $value_p$ of a site $x \leq W$, as given by Eq. (2), is:

$$\nu_p(x) = f(x)(1 - \sigma^*(x))^{k-1} = \alpha^{k-1}$$

which proves the first part. For the second part, we first derive $f(W+1) < \alpha^{k-1}$ as follows:

$$\begin{aligned} \sum_{x \leq W+1} \left(1 - \left(\frac{f(W+1)}{f(x)} \right)^{\frac{1}{k-1}} \right) &> 1 \\ \sum_{x \leq W} \left(1 - \left(\frac{f(W+1)}{f(x)} \right)^{\frac{1}{k-1}} \right) &> 1 \\ \sum_{x \leq W} \left(\frac{f(W+1)}{f(x)} \right)^{\frac{1}{k-1}} &< W - 1 \\ f(W+1)^{\frac{1}{k-1}} &< \frac{W - 1}{\sum_{x \leq W} f(x)^{-\frac{1}{k-1}}} \\ f(W+1) &< \alpha^{k-1} \end{aligned}$$

Now it suffices to observe that the $value_p$ of a site $x > W$ is $\nu_p(x) = f(x) \leq f(W+1) < \alpha^{k-1} = \nu_p(W)$. \square

As claimed in Theorem 3, it turns out that under the exclusive policy I_{exc} , strategy σ^* is not only a Nash Equilibrium, but in fact an ESS.

2.2 σ^* has Optimal Coverage

In this subsection we prove Theorem 4, that is, we prove that for any strategy σ , $\text{Cover}(\sigma^*) \geq \text{Cover}(\sigma)$, with equality if and only if $\sigma = \sigma^*$. Before dwelling into the proof we note that algorithm A^* was shown in [24] to be optimal for the setting in which k independent searchers are searching for a treasure placed in one of M boxes, according to a Bayesian distribution. The proof therein is long and relies on deep techniques that were developed in [14]. Since σ^* agrees with A^* on the first round of A^* and since there are some similarities between the search and the coverage objectives, it might be possible to translate the proof in [14, 24] to show the optimality of σ^* in terms of coverage. However, we managed to find a simple and concise proof of this fact, that can also be used to show the uniqueness of σ^* . We provide it below.

Note that maximizing $\text{Cover}(p)$ is equivalent to minimizing $\mathcal{T}(p) = \sum_x f(x)(1 - p(x))^k$. We next show that σ^* minimizes the latter expression.

First observe that since the number of resources is finite, it follows by a compactness argument that the infimum of $\{\mathcal{T}(\sigma) \mid \sigma \text{ is a distribution}\}$ is attained by some distribution σ_{\min} . Assume by contradiction that $\sigma_{\min} \neq \sigma^*$. Since both σ_{\min} and σ^* are distributions, we must have some $x_0 \leq W$ such that $\sigma^*(x_0) > \sigma_{\min}(x_0)$ and some x_1 such that $\sigma^*(x_1) < \sigma_{\min}(x_1)$. For sufficiently small $\epsilon > 0$, let us define $\sigma'(x)$ which equals σ_{\min} everywhere except that $\sigma'(x_1) = \sigma_{\min}(x_1) - \epsilon$ and $\sigma'(x_0) = \sigma_{\min}(x_0) + \epsilon$. In other words, we create σ' from σ_{\min} by shifting a small mass of probability from x_1 to x_0 .

Aiming to contradict the minimality of σ_{\min} , our goal is to show that $\mathcal{T}(\sigma') < \mathcal{T}(\sigma_{\min})$. In other words, we aim to show that the following expression, which equals to $\mathcal{T}(\sigma_{\min}) - \mathcal{T}(\sigma')$, is positive.

$$f(x_0) \left((1 - \sigma_{\min}(x_0))^k - (1 - \sigma_{\min}(x_0) - \epsilon)^k \right) + f(x_1) \left((1 - \sigma_{\min}(x_1))^k - (1 - \sigma_{\min}(x_1) + \epsilon)^k \right) \quad (6)$$

By a Taylor expansion, for a small $\epsilon > 0$, we obtain:

$$\mathcal{T}(\sigma_{\min}) - \mathcal{T}(\sigma') \approx \epsilon k \left(f(x_0)(1 - \sigma_{\min}(x_0))^{k-1} - f(x_1)(1 - \sigma_{\min}(x_1))^{k-1} \right) \quad (7)$$

By Eq. (5) and the definitions of x_0 and x_1 , we know that

$$f(x_0)(1 - \sigma_{\min}(x_0))^{k-1} > f(x_0)(1 - \sigma^*(x_0))^{k-1} \geq f(x_1)(1 - \sigma^*(x_1))^{k-1} > f(x_1)(1 - \sigma_{\min}(x_1))^{k-1} \quad (8)$$

Hence, for sufficiently small $\epsilon > 0$, we get that $\mathcal{T}(\sigma_{\min}) > \mathcal{T}(\sigma')$, contradicting the minimality of σ_{\min} . This establishes the fact that $\sigma^* = \sigma_{\min}$, implying that σ^* maximizes the coverage $\text{Cover}(p)$.

Finally, the fact that $\sigma^* = \sigma_{\min}$ also implies that the distribution that minimizes $\mathcal{T}(p)$ (and hence maximizes $\text{Cover}(p)$) is unique. Hence, $\text{Cover}(\sigma^*) = \text{Cover}(\sigma)$ can occur only if $\sigma = \sigma^*$. This concludes the proof of Theorem 4. \square

3 The Ideal Free Distribution is an ESS

The goal of this section is to prove Theorem 3, i.e, we aim to show that under I_{exc} , strategy σ^* is an ESS.

Let σ be a strategy, distinct from σ^* . First note that, when playing against strategy σ^* , because it is the IFD, the expression of $\mathcal{E}_k(\sigma; \sigma^{*(k-1)})$ becomes :

$$\mathcal{E}_k(\sigma; \sigma^{*(k-1)}) = \sum_{x=1}^M \sigma(x) \nu_{\sigma^*}(x) = \sum_{x=1}^W \sigma(x) \nu_{\sigma^*}(W) + \sum_{x=W+1}^M \sigma(x) f(x) \quad (9)$$

Since for every site $x > W$, we have $f(x) < \nu_{\sigma^*}(W)$, it is clear that every strategy σ whose support is not a subset of $[1, W]$ has

$$\mathcal{E}_k(\sigma; \sigma^{*(k-1)}) < \mathcal{E}_k(\sigma^*; \sigma^{*(k-1)})$$

and therefore satisfies the ESS characterization conditions with $m_\sigma = 0$.

Next, let us assume that the support of σ is included in $[1, W]$, and show that the ESS characterization conditions are satisfied for $m_\sigma = 1$, i.e., we aim to show that

$$\forall \ell, 1 \leq \ell \leq k-2, \text{ we have} \\ \mathcal{E}_k(\sigma^*; \sigma^\ell, \sigma^{*(k-\ell-1)}) > \mathcal{E}_k(\sigma; \sigma^\ell, \sigma^{*(k-\ell-1)}), \quad (10)$$

$$\text{and } \mathcal{E}_k(\sigma^*; \sigma^{*(k-1)}) = \mathcal{E}_k(\sigma; \sigma^{*(k-1)}). \quad (11)$$

In fact, to show that the ESS characterization conditions are satisfied for $m_\sigma = 1$, we need only the case $\ell = 1$ in Inequality (10), but we prove here a stronger stability criteria, showing that this inequality holds for every $1 \leq \ell \leq k-2$.

Eq. (11) follows quite simply from the observation above, coupled with the assumption that the support of σ is in $[1, W]$, which gives both $\mathcal{E}_k(\sigma^*; \sigma^{*(k-1)}) = \nu_{\sigma^*}$ and $\mathcal{E}_k(\sigma; \sigma^{*(k-1)}) = \nu_{\sigma^*}$.

Eq. (10) requires more elaboration. We first show that for every $1 \leq \ell \leq k-2$:

$$\mathcal{E}_k(\sigma^*; \sigma^\ell, \sigma^{*(k-\ell-1)}) \geq \mathcal{E}_k(\sigma; \sigma^\ell, \sigma^{*(k-\ell-1)})$$

and then show that equality holds iff $\sigma = \sigma^*$. We start with the following claim.

Claim 8. *Expanding each side in Eq. (10), we obtain:*

$$\mathcal{E}_k(\sigma^*; \sigma^\ell, \sigma^{*(k-\ell-1)}) = \alpha^{k-\ell-1} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^\ell - \alpha \sum_{x \leq W} f(x)^{\frac{\ell-1}{k-1}} (1 - \sigma(x))^\ell \right), \quad (12)$$

$$\mathcal{E}_k(\sigma; \sigma^\ell, \sigma^{*(k-\ell-1)}) = \alpha^{k-\ell-1} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^\ell - \sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1} \right) \quad (13)$$

Proof of Claim 8. We start by expanding the left hand side of Eq. (10).

$$\begin{aligned} \mathcal{E}_k(\sigma^*; \sigma^\ell, \sigma^{*(k-\ell-1)}) &= \sum_x f(x) \sigma^*(x) (1 - \sigma^*(x))^{k-\ell-1} (1 - \sigma(x))^\ell \\ &= \sum_{x \leq W} (f(x) - \alpha f(x)^{\frac{k-2}{k-1}}) \alpha^{k-\ell-1} f(x)^{\frac{\ell+1-k}{k-1}} (1 - \sigma(x))^\ell \\ &= \alpha^{k-\ell-1} \sum_{x \leq W} (f(x)^{\frac{\ell}{k-1}} - \alpha f(x)^{\frac{\ell-1}{k-1}}) (1 - \sigma(x))^\ell \\ &= \alpha^{k-\ell-1} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^\ell - \alpha \sum_{x \leq W} f(x)^{\frac{\ell-1}{k-1}} (1 - \sigma(x))^\ell \right) \end{aligned}$$

This establishes the first part of the claim.

Next, we expand the right hand side of Eq. (10).

$$\begin{aligned}
\mathcal{E}_k(\sigma; \sigma^\ell, \sigma^{*(k-\ell-1)}) &= \sum_x f(x) \sigma(x) (1 - \sigma^*(x))^{k-\ell-1} (1 - \sigma(x))^\ell \\
&= \sum_{x \leq W} f(x) \sigma(x) (1 - \sigma^*(x))^{k-\ell-1} (1 - \sigma(x))^\ell \\
&= \sum_{x \leq W} f(x) \sigma(x) \alpha^{k-\ell-1} f(x)^{\frac{\ell+1-k}{k-1}} (1 - \sigma(x))^\ell \\
&= \alpha^{k-\ell-1} \sum_{x \leq W} \sigma(x) f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^\ell \\
&= \alpha^{k-\ell-1} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^\ell - \sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1} \right)
\end{aligned}$$

Concluding the proof of Claim 8.

By Claim 8, it follows that the desired inequality is equivalent to:

$$\alpha \sum_{x \leq W} f(x)^{\frac{\ell-1}{k-1}} (1 - \sigma(x))^\ell \leq \sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1}$$

By definition of α , we rearrange this inequality into

$$(W-1) \sum_{x \leq W} f(x)^{\frac{\ell-1}{k-1}} (1 - \sigma(x))^\ell \leq \sum_{x \leq W} \frac{1}{f(x)^{\frac{1}{k-1}}} \sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1}$$

Decomposing $f(x)^{\frac{\ell-1}{k-1}} = \frac{f(x)^{\frac{\ell}{\ell+1} \cdot \frac{\ell}{k-1}}}{f(x)^{\frac{1}{(\ell+1)(k-1)}}$, and applying Hölder's Inequality, we obtain:

$$\sum_{x \leq W} f(x)^{\frac{\ell-1}{k-1}} (1 - \sigma(x))^\ell \leq \left(\sum_{x \leq W} \frac{1}{f(x)^{\frac{1}{k-1}}} \right)^{\frac{1}{\ell+1}} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1} \right)^{\frac{\ell}{\ell+1}}$$

This implies that to prove that Eq. (10) holds, it is sufficient to establish the following inequality:

$$(W-1) \left(\sum_{x \leq W} \frac{1}{f(x)^{\frac{1}{k-1}}} \right)^{\frac{1}{\ell+1}} \leq \sum_{x \leq W} \frac{1}{f(x)^{\frac{1}{k-1}}} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1} \right)^{\frac{1}{\ell+1}}$$

Using the fact that σ is a distribution, and hence $\sum_{x \leq W} (1 - \sigma(x)) = W - 1$, it suffices to prove that

$$\sum_{x \leq W} (1 - \sigma(x)) \left(\sum_{x \leq W} \frac{1}{f(x)^{\frac{1}{k-1}}} \right)^{\frac{1}{\ell+1}} \leq \sum_{x \leq W} \frac{1}{f(x)^{\frac{1}{k-1}}} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1} \right)^{\frac{1}{\ell+1}}$$

Rearranging, this is equivalent to proving that

$$\sum_{x \leq W} (1 - \sigma(x)) \leq \left(\sum_{x \leq W} \frac{1}{f(x)^{\frac{1}{k-1}}} \right)^{\frac{\ell}{\ell+1}} \left(\sum_{x \leq W} f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1} \right)^{\frac{1}{\ell+1}} \quad (14)$$

Which finally follows by applying Hölder's Inequality.

The fact that the inequality is strict for $\ell \geq 1$ follows from the following.

Claim 9. For $\ell \geq 1$, $\mathcal{E}_k(\sigma^*; \sigma^\ell, \sigma^{*(k-\ell-1)}) = \mathcal{E}_k(\sigma; \sigma^\ell, \sigma^{*(k-\ell-1)}) \implies \sigma = \sigma^*$.

Proof of Claim 9. The proof of the claim is based on the observation that the former equality implies, in particular, an equality in Hölder's inequality in Eq. (14). For $\ell \geq 1$, this can be the case only if the two sides are linearly dependent, i.e., for each $x \leq W$, we have

$$f(x)^{\frac{\ell}{k-1}} (1 - \sigma(x))^{\ell+1} = \frac{\beta^{\ell+1}}{f(x)^{\frac{1}{k-1}}}$$

where β is some fixed constant. Together with the fact that the support of σ is included in $[W]$, the strategy σ must therefore be of the form

$$\sigma(x) = \begin{cases} 1 - \frac{\beta}{f(x)^{\frac{1}{k-1}}} & \text{if } x \leq W \\ 0 & \text{otherwise} \end{cases}$$

Finally, we know that

$$\sum_{x \leq W} \sigma(x) = \sum_{x \leq W} \left(1 - \frac{\beta}{f(x)^{\frac{1}{k-1}}} \right) = 1$$

Solving this for β gives

$$\beta = \frac{W - 1}{\sum_{x \leq W} f(x)^{-\frac{1}{k-1}}} = \alpha$$

and therefore $\sigma = \sigma^*$. This completes the proof of Claim 9, and thus established the proof for Theorem 3. \square

4 The Criticality of the Exclusive Collision Cost Function

The goal in this section is to prove Theorem 6. That is, fixing a congestion function $C(\cdot) \neq C_{exc}(\cdot)$, we aim to show that $\text{SPoA}(C) > 1$.

First recall that by definition of congestion functions, $C(\cdot)$ is non-increasing. Hence Observation 2 applies, and the existence of the IFD is guaranteed for every value function f . Our plan is to show that there exists a value function f , for which this IFD is different than σ^* . Once this is established, the fact that $\text{SPoA}(C) > 1$ follows from Theorem 4.

Assume, towards contradiction, that for every f , the IFD under $C(\cdot)$ is precisely σ^* . Let $M \gg k$, and let f be a strictly decreasing function, that, nevertheless, decreases very slowly, such that, e.g., for every $x \leq y$ in $[1, M]$, we have $\frac{f(y)}{f(x)} \geq \frac{f(M)}{f(1)} > (1 - \frac{1}{2k})^{k-1}$. Together with the definition of W (see the pseudocode of Algorithm $\sigma^*(x)$), we obtain that $W \geq 2k$.

Because σ^* is the IFD, we have that $\forall x \leq W$, $\nu_p(x) = \nu_p(W)$. Let us develop this expression:

$$\begin{aligned} \nu_p(W) &= f(x) \sum_{\ell=1}^k C(\ell) P(x, \ell-1) \\ &= f(x) \sum_{\ell=1}^k C(\ell) \binom{k-1}{\ell-1} \sigma^*(x)^{\ell-1} (1 - \sigma^*(x))^{k-\ell} \\ &= f(x) \sum_{\ell=1}^k C(\ell) \binom{k-1}{\ell-1} \left(1 - \alpha f(x)^{-\frac{1}{k-1}} \right)^{\ell-1} \alpha^{k-\ell} f(x)^{-\frac{k-\ell}{k-1}} \\ &= \sum_{\ell=1}^k C(\ell) \binom{k-1}{\ell-1} \left(1 - \alpha f(x)^{-\frac{1}{k-1}} \right)^{\ell-1} \alpha^{k-\ell} f(x)^{\frac{\ell-1}{k-1}} \end{aligned}$$

The last line above is a polynomial equation in $f(x)^{\frac{1}{k-1}}$ of degree at most $k - 1$. Note that this polynomial is not a constant. Indeed, since $C \neq C_{exc}$ and $C(1) = C_{exc}(1) = 1$, there must exist $\ell \geq 2$, such that $C(\ell) \neq 0$. Hence, the $\ell - 1$ 'st coefficient of the polynomial is non-zero, implying that this polynomial is not a constant. It follows that there are at most $k - 1$ of values for $f(x)$ that would satisfy this equation (recall that f is strictly decreasing). This yields a contradiction since the equation is supposed to be true for every site $x \leq W$, and $W \geq 2k$. This concludes the proof of Theorem 6. \square

5 Discussion

5.1 Conclusions and Future Directions

This paper focuses on the mechanism design challenges in the basic dispersal game inspired by animal behavior. Our focus is on congestion policies, where there is no control over the utility associated with a site, and only the price associated with congestion can vary. Indeed, for natural scenarios in particular, it makes sense that evolution would impact the competition levels but the values of sites would be determined by the environment. The definition of congestion policies neglects, however, a factor that might play a significant role in several scenarios, namely, the cost incurred when in visiting a site x (e.g., the energetic cost consumed while traveling to x). Studying the more general model that takes into account this extra cost is left for future work.

The coverage measure can find relevance in contexts of animal foraging. This is because effective group coverage can indirectly contribute to the fitness of individuals, especially when the group is in competition with other groups over the same set of resources [17]. Indeed, in this case, the consumption of many resources by conspecifics reduces the fitness of individuals in the competing group and vice versa. It is therefore of interest to develop a more comprehensive game-theoretic perspective that would integrate into the individual fitness both the competition inside the group, and the competition between groups [11]. A plausible insight that can be learned from our results is that aggressive behavior between conspecifics, which appears to be wasteful or even harmful from the group's perspective, can in fact be effective on the group level, as it allows for better coverage. For further discussion on the interpretation of our results in the context of animals see Section 5.2.

Our definition of coverage assumes that a single player in a site x suffices to consume the full utility $f(x)$. In natural settings, e.g., in animal foraging scenarios, this assumption might be too strong, and it is of interest to study relaxations of it. One possibility is to assume that there is a maximum capacity of consumption per individual. In some other cases, a more complex definition of coverage might be appropriate. We leave the study of such generalizations for future work.

Although we formally discuss only the one-shot scenario, we wish to stress that the ESS, by definition, is stable when played repeatedly in an infinite population. We did not, however, investigate other forms of repetition, which are left for future work.

Finally, game theory in computer science typically studies the notion of Nash Equilibrium as the main concept to capture stability. This notion was originally designed for games with small number of players. When considering games with thousands or even millions of players, it becomes less plausible to assume that all players are rational, and one may seek for stronger notions of stability. ESS, which was developed in the evolutionary game theoretical community, appears to be a good candidate to capture stability in large populations. We believe that this notion should be applied more in game theory works that correspond to large distributed settings

5.2 Informal Discussion on Animal Dispersal

From the perspective of the group, consuming a large amount of food by all members together can indirectly increase the fitness of the group members and hence become significant for their survival. For example, when multiple species compete over the same patched resource, a thorough consumption of patches by one species implies less food remaining for the other. In this respect, it would be interesting to experiment on the

interplay between two similar species that compete over the same set of resources, but differ in their level of aggressiveness toward conspecifics [38]. In cases where there is no direct contact between the species (e.g., when they feed at different times of the day), it may appear that the more aggressive species would be inferior as it induces unnecessary waste of energy and risks of injury. However, our results suggest that perhaps it is the aggressive species that would turn out to be the superior one. Indeed, its aggressive behavior incurs higher collisions costs, which may drive individuals to better cover the resources, on the expense of the more peaceful species.

During foraging, both selfish individuals and collaborating groups (e.g., ants or bees) share the basic challenge of balancing the need to exploit the more promising resources, while trying to avoid overlaps (or collisions) with conspecifics [12, 17]. This tradeoff, however, is manifested in the two cases in a manner that appears to be very different: While collaborative groups aim to optimize the tradeoff by setting the parameters in a somewhat “centralized” manner, the competing dynamics is governed by a selection process yielding an evolutionary stable strategy which need not be optimal for individuals. As suggested here, despite these differences, when colliding individuals are punished severely, the two perspectives yield similar strategies.

Finally, the alignment in behavior between collaborating and competing individuals, as reported here, may suggest a modest contribution to the theory of evolution of eusocial species. The currently dominated theory proposes that eusocial species have evolved from self motivated individuals due to kin selection [20]. An issue that is discussed less is that a shift in the “motivation focus” can potentially be accompanied by a shift in function and hence, perhaps, also in structure [27]. Indeed, in principle, while evolving from self-interested to altruistic, an individual which had evolved in a way that was tailored to its own needs, may now need to adapt to the collective needs. As these needs are on different scales, they may not necessarily be akin [13], and hence a long evolutionary process may be required in order to bridge this gap. In this respect, that seemingly cooperative behavior can arise in competing systems, as demonstrated here, suggests that the behavioral transition, at least in the context of dispersal or forage, may have been smoother than what one might expect.

Acknowledgements. The authors are thankful for Yoav Rodeh, Yossi Yovel, Yuval Emek, Lucas Boczkowski, Emanuele Natale, and Ofer Feinerman for useful discussions concerning both the model and its applications to biological contexts.

References

- [1] Susanne Albers and Matthias Hellwig. Online makespan minimization with parallel schedules. *CoRR*, abs/1304.5625, 2013.
- [2] Sayan Bhattacharya, Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Coordination mechanisms from (almost) all scheduling policies. In *Innovations in Theoretical Computer Science, ITCS’14, Princeton, NJ, USA, January 12-14, 2014*, pages 121–134, 2014.
- [3] C. Cannings Broom, M. and G. Vickers. Multi-player matrix games. *Bulletin of Mathematical Biology*, 59(3):931–952, 1997.
- [4] Mark Broom and Jan Rychtár. *Game-theoretical models in biology*. CRC Press, 2013.
- [5] Maciej Bukowski and Jacek Miekisz. Evolutionary and asymptotic stability in symmetric multi-player games. *International Journal of Game Theory*, 33(1):41–54, 2004.
- [6] Chris Cosner, Juan Dávila, and Salomé Martínez. Evolutionary stability of ideal free nonlocal dispersal. *Journal of Biological Dynamics*, 6(2):395–405, 2012.

- [7] Ross Cressman and Vlastimil Křivan. Migration dynamics for the ideal free distribution. *The American Naturalist*, 168(3):384–397, 2006.
- [8] Ross Cressman and Vlastimil Křivan. The ideal free distribution as an evolutionarily stable state in density-dependent population games. *Oikos*, 119(8):1231–1242, 2010.
- [9] Artur Czumaj and Berthold Vöcking. Tight bounds for worst-case equilibria. *ACM Trans. Algorithms*, 3(1):4:1–4:17, February 2007.
- [10] Dina K.N. Dechmann, Silke L. Heucke, Luca Giuggioli, Kamran Safi, Christian C. Voigt, and Martin Wikelski. Experimental evidence for group hunting via eavesdropping in echolocating bats. *Proceedings of the Royal Society of London B: Biological Sciences*, 276(1668):2721–2728, 2009.
- [11] Avigdor Eldar. Social conflict drives the evolutionary divergence of quorum sensing. *Proceedings of the National Academy of Sciences*, 108(33):13635–13640, 2011.
- [12] Ofer Feinerman and Amos Korman. The ANTS problem. *Distributed Computing*, 30(3):149–168, 2017.
- [13] Ofer Feinerman and Amos Korman. Individual versus collective cognition in social insects. *Journal of Experimental Biology*, 220(1):73–82, 2017.
- [14] Pierre Fraigniaud, Amos Korman, and Yoav Rodeh. Parallel exhaustive search without coordination. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 312–323, 2016.
- [15] Peter Frazier, David Kempe, Jon Kleinberg, and Robert Kleinberg. Incentivizing exploration. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 5–22. ACM, 2014.
- [16] Stephen Dewitt Fretwell and Henry L. Lucas Jr. On territorial behavior and other factors influencing habitat distribution in birds. *Acta biotheoretica*, 19(1):16–32, 1969.
- [17] Luc-Alain Giraldeau and Thomas Caraco. Social foraging theory. *Princeton University Press*.
- [18] Chaitanya S Gokhale and Arne Traulsen. Evolutionary games in the multiverse. *Proceedings of the National Academy of Sciences*, 107(12):5500–5504, 2010.
- [19] Thomas Gregor, Koichi Fujimoto, Noritaka Masaki, and Satoshi Sawai. The onset of collective behavior in social amoebae. *Science*, 328(5981):1021–1025, 2010.
- [20] William D. Hamilton. Narrow roads of geneland. freeman. *Freeman*, 1996.
- [21] Thomas T Hills, Peter M Todd, David Lazer, A David Redish, Iain D Couzin, Cognitive Search Research Group, et al. Exploration versus exploitation in space, mind, and society. *Trends in cognitive sciences*, 19(1):46–54, 2015.
- [22] Martyn Kennedy and Russell D. Gray. Can ecological theory predict the distribution of foraging animals? a critical analysis of experiments on the ideal free distribution. *Oikos*, 68(1):158–166, 1993.
- [23] Jon Kleinberg and Sigal Oren. Mechanisms for (mis) allocating scientific credit. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 529–538. ACM, 2011.
- [24] Amos Korman and Yoav Rodeh. Parallel search without coordination. In *Proceedings of the 24th International Colloquium on Structural Information and Communication Complexity, SIROCCO 2017, Porquerolles, France, June 19-22, 2017*, 2017.

- [25] Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the gwisdom of the crowdh. *Journal of Political Economy*, 122(5):988–1012, 2014.
- [26] Yishay Mansour, Aleksandrs Slivkins, and Vasilis Syrgkanis. Bayesian incentive-compatible bandit exploration. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 565–582. ACM, 2015.
- [27] John M. McNamara and Alasdair I. Houston. Integrating function and mechanism. *Trends in Ecology & Evolution*, 24(12):670 – 675, 2009.
- [28] Linda Mealey. The sociobiology of sociopathy: An integrated evolutionary model. *Behavioral and Brain Sciences*, 18(3):523–541, 1995.
- [29] Rodrigo A Medellin, Marina Rivero, Ana Ibarra, J Antonio de la Torre, Tania P Gonzalez-Terrazas, Leonora Torres-Knoop, and Marco Tschapka. Follow me: foraging distances of leptonycteris yerbabue-nae (chiroptera: Phyllostomidae) in sonora determined by fluorescent powder. *Journal of Mammalogy*, 99(2):306–311, 2018.
- [30] Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
- [31] Douglas W Morris. Shadows of predation: habitat-selecting consumers eclipse competition between coexisting prey. *Evolutionary Ecology*, 17(4):393–422, 2003.
- [32] AJ Nicholson. An outline of the dynamics of animal populations. *Australian Journal of Zoology*, pages 9–65, 1954.
- [33] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [34] Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.
- [35] Yiangos Papanastasiou, Kostas Bimpikis, and Nicos Savva. Crowdsourcing exploration. *Management Science*, 2017.
- [36] Geoffrey A Parker. Searching for mates. *Behavioural ecology: an evolutionary approach*, 1:214–244, 1978.
- [37] Geoffrey A Parker. Evolutionary stable strategies. *Behavioural ecology: an evolutionary approach*, pages 30–61, 1984.
- [38] Noa Pinter-Wollman, Tamar Dayan, David Eilam, and Noga Kronfeld-Schor. Can aggression be the force driving temporal separation between competing common and golden spiny mice? *Journal of Mammalogy*, 87(1):48–53, 2006.
- [39] H Ronald Pulliam and Thomas Caraco. Living in groups: is there an optimal group size. *Behavioural ecology: an evolutionary approach*, 2:122–147, 1984.
- [40] H Ronald Pulliam and Brent J Danielson. Sources, sinks, and habitat selection: a landscape perspective on population dynamics. *The American Naturalist*, 137:S50–S66, 1991.
- [41] Nicanor Quijano and Kevin M Passino. The ideal free distribution: Theory and engineering application. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(1):154–165, 2007.

- [42] Robert W Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [43] J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246(5427):15–18, 1973.
- [44] Tom Tregenza. Building on the ideal free distribution. *Adv. Ecol. Res*, pages 253–307, 1995.
- [45] Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 416–425. IEEE, 2002.

APPENDIX

A Proof Sketch for Observation 2

The existence of the IFD follows by the fact that $\nu_p(x)$ is non-increasing with $p(x)$. By definition, the IFD is a symmetric Nash Equilibrium of the game. We now prove its uniqueness. Imagine there are two different symmetric Nash Equilibria, π and π' , respectively associated with their values ν_π and $\nu_{\pi'}$. Note that for any fixed x , the *value* $\nu_p(x)$ is a strictly decreasing function of $p(x)$. Therefore

$$\pi = \pi' \iff \nu_\pi = \nu_{\pi'}$$

Assume by contradiction that $\nu_\pi > \nu_{\pi'}$. This implies that for any $x \in [\pi]$, $\pi(x) < \pi'(x)$, where $[\pi]$ denotes the support of π , i.e the set of sites explored with positive probability. Summing over all sites in $[\pi]$, we get

$$\sum_{x \in [\pi]} \pi(x) < \sum_{x \in [\pi]} \pi'(x)$$

Since the left-hand size is equal to 1, we get a contradiction. This proves the uniqueness of the symmetric Nash Equilibrium, and of the IFD. \square