# Multi-Round Cooperative Search Games with Multiple Players

## Amos Korman
Université de Paris, IRIF, CNRS, F-75013 Paris, France
amos.korman@irif.fr

## Yoav Rodeh
Ort Braude College, Karmiel, Israel
yoav.rodeh@braude.ac.il

—— **Abstract** ——

Assume that a treasure is placed in one of $M$ boxes according to a known distribution and that $k$ searchers are searching for it in parallel during $T$ rounds. We study the question of how to incentivize selfish players so that group performance would be maximized. Here, this is measured by the *success probability*, namely, the probability that at least one player finds the treasure. We focus on *congestion policies* $C(\ell)$ that specify the reward that a player receives if it is one of $\ell$ players that (simultaneously) find the treasure for the first time. Our main technical contribution is proving that the *exclusive policy*, in which $C(1) = 1$ and $C(\ell) = 0$ for $\ell > 1$, yields a *price of anarchy* of $(1 - (1 - 1/k)^k)^{-1}$, and that this is the best possible price among all symmetric reward mechanisms. For this policy we also have an explicit description of a symmetric equilibrium, which is in some sense unique, and moreover enjoys the best success probability among all symmetric profiles. For general congestion policies, we show how to polynomially find, for any $\theta > 0$, a symmetric multiplicative $(1 + \theta)(1 + C(k))$-equilibrium.

Together with an appropriate reward policy, a central entity can suggest players to play a particular profile at equilibrium. As our main conceptual contribution, we advocate the use of symmetric equilibria for such purposes. Besides being fair, we argue that symmetric equilibria can also become highly robust to crashes of players. Indeed, in many cases, despite the fact that some small fraction of players crash (or refuse to participate), symmetric equilibria remain efficient in terms of their group performances and, at the same time, serve as approximate equilibria. We show that this principle holds for a class of games, which we call *monotonously scalable* games. This applies in particular to our search game, assuming the natural *sharing policy*, in which $C(\ell) = 1/\ell$. For the exclusive policy, this general result does not hold, but we show that the symmetric equilibrium is nevertheless robust under mild assumptions.

## 1    Introduction

Searching in groups is ubiquitous in multiple contexts, including in the biological world, in human populations as well as on the internet [8, 11, 16]. In many cases there is some prior on the distribution of the searched target. Moreover, when the space is large, each searcher typically needs to inspect multiple possibilities, which in some circumstances can only be done sequentially. This paper introduces a game theoretic perspective to such multi-round treasure hunt searches, generalizing a basic collaborative Bayesian framework previously introduced in [8].

Consider the case that a treasure is placed in one of $M$ boxes according to a known distribution $f$ and that $k$ searchers are searching for it in parallel during $T$ rounds, each specifying a box to visit in each round. Assume w.l.o.g. that the boxes are ordered such that lower index boxes have higher probability to host the treasure, i.e., $f(x) \geq f(x + 1)$. We evaluate the group performance by the *success probability*, that is, the probability that the treasure is found by at least one searcher.

If coordination is allowed, letting searcher $i$ visit box $(t - 1)k + i$ at time $t$ will maximize success probability. However, as simple as this algorithm is, it is very sensitive to faults of all sorts. For example, if an adversary that knows where the treasure is can crash a searcher before the search starts (i.e., prevent it from searching), then it can reduce the search probability to zero.

The authors of [8] suggested the use of *identical non-coordinating* algorithms. In such scenarios all processors act independently, using no communication or coordination, executing the same probabilistic algorithm, differing only by the results of their coin flips. As argued in [8], in addition to their economic use of communication, identical non-coordinating algorithms enjoy inherent robustness to different kind of faults. For example, assume that there are $k + k'$ searchers, and that an adversary can fail up to $k'$ searchers. Letting all searchers run the best non-coordinating algorithm for $k$ searchers guarantees that regardless of which $\ell \leq k'$ searchers fail, the overall search efficiency is at least as good as the non-coordinating one for $k$ players. Of course, since $k'$ players might fail, any solution can only hope to achieve the best performance of $k$ players. As it applies to the group performance we term this property as *group robustness*. Among the main results in [8] is identifying a non-coordinating algorithm, denoted $A^\star$, whose expected running time is minimal among non-coordinating algorithms. Moreover, for every given $T$, if this algorithm runs for $T$ rounds, it also maximizes the success probability.

The current paper studies the game theoretic version of this multi-round search problem[1]. The setting of [8] assumes that the searchers adhere fully to the instructions of a central entity. In contrast, in a game theoretical context, searchers are self-interested and one needs to incentivize them to behave as desired, e.g., by awarding those players that find the treasure first. For many real world contexts, the competitive setting is in fact the more realistic one to assume. Applications range from crowd sourcing [26], multi-agent searching on the internet [25], grant proposals [18], to even contexts of animals [17]. See the full version for more details.

---

[1] We concentrate on the normal form version in which players do not receive any feedback during the search (except when the treasure is found in which case the game ends). In particular, we assume that players cannot communicate with each other.

In the competitive setting, choosing a good *rewarding policy* becomes a problem in algorithmic mechanism design [24]. Typically, a reward policy is evaluated by its *price of anarchy (PoA)*, namely, the ratio between the performances of the best collaborative algorithm and the worst equilibrium [19]. Aiming to both accelerate the convergence process to an equilibrium and obtain a preferable one, the announcement of the reward policy can be accompanied by a proposition for players to play particular strategies that form a profile at equilibrium.

This paper highlights the benefits of suggesting (non-coordinating) *symmetric equilibria* in such scenarios, that is, to suggest the same non-coordinating strategy to be used by all players, such that the resulting profile is at equilibrium. This is of course relevant assuming that the *price of symmetric stability (PoSS)*, namely, the ratio between the performances of the best collaborative algorithm and the best symmetric equilibrium, is low. Besides the obvious reasons of fairness and simplicity, from the perspective of a central entity who is interested in the overall success probability, we obtain the group robustness property mentioned above, by suggesting that the $k + k'$ players play according to the strategy that is a symmetric equilibrium for $k$ players. Obviously, this group robustness is valid only provided that the players indeed play according to the suggested strategy. However, the suggested strategy is guaranteed to be an equilibrium only for $k$ players, while in fact, the adversary may keep some of the extra $k'$ players alive. Interestingly, however, in many cases, a symmetric equilibrium for $k$ players also serves as an approximate equilibrium for $k + k'$ players, as long as $k' \ll k$. As we show, this *equilibrium robustness* property is rather general, holding for a class of games, that we call *monotonously scalable* games.

## 1.1 The Collaborative Search Game

A treasure is placed in one of $M$ boxes according to a known distribution $f$ and $k$ players are searching for it in parallel during $T$ rounds. Assume w.l.o.g. that $f(x) > 0$ for every $x$ and that $f(x) \geq f(x+1)$.

**Strategies.** An *execution* of $T$ rounds is a sequence of box visitations $\sigma = x(1), x(2), \ldots x(T)$, one for each round $i \leq T$. We assume that a player visiting a box has no information on whether other players have already visited that box or are currently visiting it. Hence, a *strategy* of a player is a probability distribution over the space of executions of $T$ rounds. Note that the probability of visiting a box $x$ in a certain round may depend on the boxes visited by the player until this round, but not on the actions of other players. A strategy is *non-redundant* if at any given round it always checks a box it didn't check before (as long as there are such boxes).

A *profile* is a collection of $k$ strategies, one for each player. Special attention will be devoted to *symmetric profiles*. In such profiles all players play the same strategy (note that their actual executions may be different, due to different probabilistic choices).

**Probability Matrix.** While slightly abusing notation, we shall associate each strategy $A$ with its *probability matrix*, $A : \{1, \ldots, M\} \times \{1, \ldots, T\} \to [0, 1]$, where $A(x, t)$ is the probability that strategy $A$ visits $x$ for the first time at round $t$. We also denote $\widetilde{A}(x, t)$ as the probability that $A$ does not visit $x$ by, and including, time $t$. That is, $\widetilde{A}(x, t) = 1 - \sum_{s \leq t} A(x, t)$ and $\widetilde{A}(x, 0) = 1$. For convenience we denote by $\delta_{x,t}$ the matrix of all zeros except 1 at $x, t$. Its dimensions will be clear from context.

**Group Performance.** A profile is evaluated by its *success probability*, i.e., the probability that at least one player finds the treasure by time $T$. Formally, let $\mathbb{P}$ be a profile. Then,

$$\texttt{success}(\mathbb{P}) = \sum_x f(x) \left( 1 - \prod_{A \in \mathbb{P}} \widetilde{A}(x, T) \right).$$

The expected running time in the symmetric case, which is $\sum_x f(x) \sum_t \widetilde{A}(x,t)^k$, was studied in [8]. That paper identified a strategy, denoted $A^\star$, that minimizes this quantity. In fact, it does so by minimizing the term $\sum_x f(x) \widetilde{A}(x,t)^k$ for each $t$ separately. Note that minimizing the case $t = T$ is exactly the same as maximizing the success probability. Thus, restricted to the case where all searchers use the same strategy, $A^\star$ simultaneously optimizes the success probability as well as optimizes the expected running time. For completeness, a description of $A^\star$ is provided below.

**Algorithm $A^\star$.** We note that in [8] the matrix of $A^\star$ is given, and then an algorithm is explicitly described that has its matrix (Section 4.3 in [8]). We describe the matrix only, as its details are necessary for this paper.

Denote $q(x) = f(x)^{-1/(k-1)}$. For each $t$, $\widetilde{A^\star}(x,t) = \min(1, \alpha(t)q(x))$, where $\alpha(t) \geq 0$ is such that $\sum_x \widetilde{A^\star}(x,t) = M - t$. Of course, once $\widetilde{A^\star}$ is known, then so is $A^\star$: $A^\star(x,t) = \widetilde{A^\star}(x,t-1) - \widetilde{A^\star}(x,t)$.

**Congestion Policies.** A natural way to incentivize players is by rewarding those players that find the treasure before others. A *congestion policy* $C(\ell)$ is a function specifying the reward that a player receives if it is one of $\ell$ players that (simultaneously) find the treasure for the first time. We assume that $C(1) = 1$, and that $C$ is non-negative and non-increasing. Due to the fact that the policy $C \equiv 1$ is rather degenerate, we henceforth assume that $C \not\equiv 1$. We shall give special attention to the following policies.

- The *sharing policy* is defined by $C_{\texttt{share}}(\ell) = 1/\ell$, namely, the treasure is shared equally among all those who find it first.
- The *exclusive policy* is defined by $C_{\texttt{ex}}(1) = 1$, and $C_{\texttt{ex}}(\ell) = 0$ for $\ell > 1$, namely, the treasure is given to the first one that finds it exclusively; if more than one discover it, they get nothing[2].

A *configuration* is a triplet $(C, f, T)$, where $C$ is a congestion policy, $T$ is a positive integer, and $f$ is a positive non-increasing probability distribution on $M$ boxes.

**Values, Utilities and Equilibria.** Let $(C, f, T)$ be a configuration. The *value* of box $x$ at round $t$ when playing against a profile $\mathbb{P}$ is the expected gain from visiting $x$ at round $t$. Formally,

$$v_{\mathbb{P}}(x, t) = f(x) \sum_{\ell=0}^{k-1} C(\ell+1) \Pr \left( \begin{array}{c} x \text{ was not visited before time } t, \text{ and} \\ \text{at time } t \text{ is visited by } \ell \text{ players of } \mathbb{P} \end{array} \right)$$

$$= f(x) \sum_{\ell=0}^{k-1} C(\ell+1) \sum_{\substack{I \subseteq \mathbb{P} \\ |I| = \ell}} \prod_{A \in I} A(x, t) \prod_{A \notin I} \widetilde{A}(x, t).$$

---

[2] In the one round game, the exclusive policy yields a utility for a player that equals its marginal contribution to the social welfare, i.e., the success probability [27]. However, this is not the case in the multi-round game.

The *utility of A in round t* and the *utility of A* are defined as:

$$U_{\mathbb{P}}(A, t) := \sum_x A(x, t) \cdot v_{\mathbb{P}-A}(x, t), \qquad U_{\mathbb{P}}(A) := \sum_t U_{\mathbb{P}}(A, t), \qquad (1)$$

where $\mathbb{P}^{-A}$ is the set of players of $\mathbb{P}$ excluding $A$. Here are some specific cases we are interested in:

- For symmetric profiles, $v_A(x, t)$ denotes the value when playing against $k - 1$ players playing $A$. Then $v_A(x, t) = f(x) \sum_{\ell=0}^{k-1} C(\ell + 1) \binom{k-1}{\ell} A(x, t)^\ell \widetilde{A}(x, t)^{k-\ell-1}$.
- For the exclusive policy, $v_{\mathbb{P}}(x, t) = f(x) \prod_{A \in \mathbb{P}} \widetilde{A}(x, t)$.
- For the exclusive policy in symmetric profiles, $v_A(x, t) = f(x) \widetilde{A}(x, t)^{k-1}$.

A profile $\mathbb{P}$ is a *Nash equilibrium* under a configuration, if for any $A \in \mathbb{P}$ and any other strategy $B$, $U_{\mathbb{P}-A}(A) \geq U_{\mathbb{P}-A}(B)$. Similarly, a strategy $A$ is called a *symmetric equilibrium* if the profile $A^k$ consisting of all $k$ players playing according to $A$ is an equilibrium. We also use the notion of approximate equilibrium. For $\epsilon > 0$, we say a profile $\mathbb{P}$ is a $(1 + \epsilon)$-equilibrium if for every $A \in \mathbb{P}$ and for every strategy $B$, $U_{\mathbb{P}-A}(B) \leq (1 + \epsilon) U_{\mathbb{P}-A}(A)$.

**A Game of Doubly-Substochastic Matrices.** Both the expressions for the success probability and utility solely depend on the values of the probability matrices associated with the strategies in question. Hence we view all strategies sharing the same matrix as *equivalent*. Note that a matrix does not necessarily correspond to a unique strategy, as illustrated by the following equivalent strategies, for which $A(x, t) = B(x, t) = 1/M$ for every $t \leq M$ and $0$ thereafter:

- Strategy $A$ chooses uniformly at every round one of the boxes it didn't choose yet.
- Strategy $B$ chooses once $x \in \{0, \ldots, M - 1\}$. Then, at round $t$ it visits box $(x + t \mod M) + 1$.

Matrices are much simpler to handle than strategies, and so we would rather think of our game as a game of probability matrices than a game of strategies. For this we need to characterize which matrices are indeed probability matrices of strategies. Clearly, a probability matrix is non-negative. Also, by their definition, each row and each column sums to at most 1. Such a matrix is called *doubly-substochastic*. In the extended version we prove the converse, i.e., that every doubly-substochastic matrix is a probability matrix of some strategy. Furthermore, this strategy is implementable as a polynomial algorithm. We will therefore view our game as a game of doubly-substochastic matrices.

**Greediness.** Informally, a strategy is greedy at a round if its utility in this round is the maximum possible in this round. Formally, given a profile $\mathbb{P}$ and some strategy $A$, we say that $A$ is *greedy* w.r.t. $\mathbb{P}$ at time $t$ if for any strategy $B$ such that for every $x$ and $s < t$, $B(x, s) = A(x, s)$, we have $U_{\mathbb{P}}(A, t) \geq U_{\mathbb{P}}(B, t)$. We say $A$ is greedy w.r.t. $\mathbb{P}$ if it is greedy w.r.t. $\mathbb{P}$ for each $t \leq T$. A strategy $A$ is called *self-greedy* (or *sgreedy* for short) if it is greedy w.r.t. the profile with $k - 1$ players playing $A$.

**Evaluating Policies.** Let $(C, f, T)$ be a configuration. Denote by $\text{Nash}(C, f, T)$ the set of equilibria for this configuration, and by $\text{S-Nash}(C, f, T)$ the subset of symmetric ones. Let $\mathcal{P}(T)$ be the set of all profiles of $T$-round strategies. We are interested in the following measures.

- The *Price of Anarchy (PoA)* is $\mathtt{PoA}(C, f, T) := \frac{\max_{\mathbb{P} \in \mathcal{P}(T)} \mathtt{success}(\mathbb{P})}{\min_{\mathbb{P} \in \mathrm{Nash}(C, f, T)} \mathtt{success}(\mathbb{P})}$.

- The *Price of Symmetric Stability (PoSS)* is $\mathtt{PoSS}(C, f, T) := \frac{\max_{\mathbb{P} \in \mathcal{P}(T)} \mathtt{success}(\mathbb{P})}{\max_{A \in \mathrm{S\text{-}Nash}(C, f, T)} \mathtt{success}(A^k)}$.

- The *Price of Symmetric Anarchy (PoSA)* is $\mathtt{PoSA}(C, f, T) = \frac{\max_{\mathbb{P} \in \mathcal{P}(T)} \mathtt{success}(\mathbb{P})}{\min_{A \in \mathrm{S\text{-}Nash}(C, f, T)} \mathtt{success}(A^k)}$.

**On the Difficulty of the Multi-Round Game.** The setting of multi-rounds poses several challenges that do not exist in the single round game. An important one is the fact that, in contrast to the single round game, the multi-round game is not a potential game. Indeed, being a potential game has several implications, and a significant one is that such a game has always a pure equilibrium. However, we show that multi-round games do not always have pure equilibria and hence they are not potential games. Another important difference is that for policies that incur high levels of competition (such as the exclusive policy), profiles that maximize the success probability are at equilibrium in the single round case, whereas they are not in the multi-round game. See the full version of the paper for more details.

## 1.2 Our Results

**Equilibrium Robustness.** We first provide a simple, yet general, robustness result, that holds for symmetric (approximate) equilibria in a family of games, termed *monotonously scalable*. Informally, these are games in which the sum of utilities of players can only increase when more players are added, yet for each player, its individual utility can only decrease. Our search game with the sharing policy is one such example.

▶ **Theorem 1.** *Consider a symmetric monotonously scalable game. If A is a symmetric $(1 + \epsilon)$-equilibrium for k players, then it is an $(1 + \epsilon)(1 + \ell/k)$-equilibrium when played by $k + \ell$ players.*

Theorem 1 is applicable in fault tolerant contexts. Consider a monotonously scalable game with $k + k'$ players out of which at most $k'$ may fail. Let $A_k$ be a symmetric (approximate) equilibrium designed for $k$ players and assume that its social utility is high compared to the optimal profile with $k$ players. The theorem implies that if players play $A_k$, then regardless of which $\ell \leq k'$ players fail (or decline to participate), the incentive to switch strategy would be very small, as long as $k' \ll k$. Moreover, due to symmetry, if the social utility of the game is monotone, then the social utility of $A_k$ when played with $k$ players is guaranteed when playing with more. Thus, in such cases we obtain both group robustness and equilibrium robustness.

**General Congestion Policies.** Coming back to our search game, we consider general policies, focus on symmetric profiles, and specifically, on the properties of sgreedy strategies.

▶ **Theorem 2.** *For every policy C there exists a non-redundant sgreedy strategy. Moreover, all such strategies are equivalent and are symmetric $(1 + C(k))$-equilibria.*

When $C(k) = 0$ this shows that a non-redundant sgreedy strategy is actually a symmetric equilibrium. We next claim that this is the only symmetric equilibrium (up to equivalence).

▷ Claim 3. For any policy such that $C(k) = 0$, all symmetric equilibria are equivalent.

Theorem 2 is non-constructive because it requires calculating the inverse of non-trivial functions. Therefore, we resort to an approximate solution.

▶ **Theorem 4.** *Given $\theta > 0$, there exists an algorithm that takes as input a configuration, and produces a symmetric $(1 + C(k))(1 + \theta)$-equilibrium. The algorithm runs in polynomial time in $T$, $k$, $M$, $\log(1/\theta)$, $\log(1/(1 - C(k)))$, and $\log(1/f(M))$.*

**The Exclusive Policy.** Recall that the exclusive policy is defined by $C_{\text{ex}}(1) = 1$ and $C_{\text{ex}}(\ell) = 0$ for every $\ell > 1$. We show that $A^\star$ is a non-redundant and sgreedy strategy in the exclusive policy. Hence, Theorem 2 implies the following.

▶ **Theorem 5.** *Under the exclusive policy, Strategy $A^\star$ of [8] is a symmetric equilibrium.*

Claim 3 together with the fact (established in [8]) that $A^\star$ has the highest success probability among symmetric profiles, implies that both the PoSS and the PoSA of $C_{\text{ex}}$ are optimal (and equal) on any $f$ and $T$ when compared to any other policy. The next theorem considers general equilibria.

▶ **Theorem 6.** *Consider the exclusive policy. For any profile $\mathbb{P}_{\text{nash}}$ at equilibrium and any symmetric profile $A$, $\text{success}(\mathbb{P}_{\text{nash}}) \geq \text{success}(A)$.*

Observe that, as $A^\star$ is a symmetric equilibrium, Theorem 6 provides an alternative proof for the optimality of $A^\star$ (established in [8]). Interestingly, this alternative proof is based on game theoretic considerations, which is a very rare approach in optimality proofs.

Combining Theorems 5 and 6, we obtain:

▶ **Corollary 7.** *For any $f$ and $T$, $\text{PoA}(C_{\text{ex}}, f, T) = \text{PoSA}(C_{\text{ex}}, f, T)$. Moreover, for any policy $C$, $\text{PoA}(C_{\text{ex}}, f, T) \leq \text{PoA}(C, f, T)$.*

At first glance the effectiveness of $C_{ex}$ might not seem so surprising. Indeed, it seems natural that high levels of competition would incentivize players to disperse. However, it is important to note that $C_{\text{ex}}$ is not extreme in this sense, as one may allow congestion policies to also have negative values upon collisions. Moreover, one could potentially define more complex kinds of policies, e.g., policies that depend on time, and reward early finds more. However, the fact that $A^\star$ is optimal among all symmetric profiles combined with the fact that any symmetric policy has a symmetric equilibrium [23] implies that no symmetric reward mechanism can improve either the PoSS, the PoSA, or the PoA of the exclusive policy.

We proceed to show a tight upper bound on the PoA of $C_{\text{ex}}$. Note that as $k$ goes to infinity the bound converges to $e/(e-1) \approx 1.582$.

▶ **Theorem 8.** *For every $T$, $\sup_f \text{PoA}(C_{\text{ex}}, f, T) = (1 - (1 - 1/k)^k)^{-1}$.*

Concluding the results on the exclusive policy, we study the robustness of $A^\star$ in the full version of the paper. Let $A_k^\star$ denote algorithm $A^\star$ when set to work for $k$ players. Unfortunately, for any $\epsilon$, there are cases where $A_k^\star$ is not a $(1+\epsilon)$-equilibrium even when played by $k+1$ players. However, as indicated below, $A^\star$ is robust to failures under reasonable assumptions regarding the distribution $f$.

▶ **Theorem 9.** *If $\frac{f(1)}{f(M)} \leq (1+\epsilon)^{\frac{k-1}{k'}}$, then $A_k^\star$ is a $(1+\epsilon)$-equilibrium when played by $k + k'$ players.*

**The Sharing Policy.** Another important policy to consider is the sharing policy. This policy naturally arises in some circumstances, and may be considered as a less harsh alternative to the exclusive one. Although not optimal, it follows from Vetta [31] that its PoA is at most 2 (see the full version). Furthermore, as this policy yields a monotonously scalable game, a symmetric equilibrium under it is also robust. Therefore, the existence of a symmetric profile which is both robust and has a reasonable success probability is guaranteed.

Unfortunately, we did not manage to find a polynomial algorithm that generates a symmetric equilibrium for this policy. However, Theorem 4 gives a symmetric $(1+\theta)(1+1/k)$-equilibrium in polynomial time for any $\theta > 0$. This strategy is also robust thanks to Theorem 1.

Moreover, the proof in [31] regarding the PoA can be extended to hold for approximate equilibria. In particular, if $\mathbb{P}$ is some $(1 + \epsilon)$-equilibrium in the sharing policy, then for every $f$ and $T$, $\texttt{success}(\mathbb{P}) \geq \frac{1}{2+\epsilon} \max_{\mathbb{P}' \in \mathcal{P}(T)} \texttt{success}(\mathbb{P}')$. See the full version of the paper for the proof.

## 1.3   Related Works

Fault tolerance has been a major topic in distributed computing for several decades, and in recent years more attention has been given to these concepts in game theory [13, 14]. For example, Gradwohl and Reingold studied conditions under which games are robust to faults, showing that equilibria in anonymous games are fault tolerant if they are "mixed enough" [12].

Restricted to a single round the search problem becomes a coverage problem, which has been investigated in several papers. For example, Collet and Korman studied in [30] (one-round) coverage while restricting attention to symmetric profiles only. The main result therein is that the exclusive policy yields the best coverage among symmetric profiles. Gairing [10] also considered the single round setting, but studied the optimal PoA of a more general family of games called *covering games* (see also [27, 28]). Motivated by policies for research grants, Kleinberg and Oren [18] considered a one-round model similar to that in [30]. Their focus however was on pure strategies only. The aforementioned papers give a good understanding of coverage games in the single round setting. As mentioned, however, the multi-round setting studied here is substantially more complex than the single-round setting.

The area of "incentivizing exploration" also studies the tradeoff between exploration, exploitation and incentives [21, 9, 26, 20]. This area often focuses on different variants of the Multi-Armed Bandit problem. The settings of selfish routing, job scheduling, and congestion games [22, 29] all bear similarities to the search game studied here, however, the social welfare measurements of success probability or running time are very different from the measures studied in these frameworks, such as makespan or latency [1, 3, 24, 2].

## 2   Robustness in Symmetric Monotonously Scalable Games

Consider a symmetric game where the number of players is not fixed. Let $U_{\mathbb{P}}(A)$ denote the utility that a player playing $A$ gets if the other players play according to $\mathbb{P}$ and let $\sigma(\mathbb{P}) = \sum_{A \in \mathbb{P}} U_{\mathbb{P}-A}(A)$. We say that such a game is *monotonously scalable* if:

1. Adding more players can only increase the sum of utilities, i.e., if $\mathbb{P} \subseteq \mathbb{P}'$ then $\sigma(\mathbb{P}) \leq \sigma(\mathbb{P}')$.
2. Adding more players can only decrease the individual utilities, i.e., if $\mathbb{P} \subseteq \mathbb{P}'$ then for all $A \in \mathbb{P}$, $U_{\mathbb{P}-A}(A) \geq U_{\mathbb{P}'-A}(A)$.

▶ **Theorem 1.** *Consider a symmetric monotonously scalable game. If $A$ is a symmetric $(1 + \epsilon)$-equilibrium for $k$ players, then it is an $(1 + \epsilon)(1 + \ell/k)$-equilibrium when played by $k + \ell$ players.*

**Proof.** On the one hand by symmetry,

$$U_{A^{k+\ell-1}}(A) = \frac{\sigma(A^{k+\ell})}{k+\ell} \geq \frac{\sigma(A^k)}{k+\ell},$$

where the last step is because $\sigma$ is non-decreasing. On the other hand, if $B$ is some other strategy,

$$U_{A^{k+\ell-1}}(B) \leq U_{A^{k-1}}(B) \leq (1+\epsilon)U_{A^{k-1}}(A) = (1+\epsilon)\frac{\sigma(A^k)}{k}.$$

The first inequality is because $U$ is non-increasing, and the second is because $A$ is a $(1+\epsilon)$-equilibrium for $k$ players. Therefore, what a player can gain by switching from $A$ to $B$ is at most a multiplicative factor of $(1+\epsilon)(k+\ell)/k = (1+\epsilon)(1+\ell/k)$. ◀

An example of such a game is our setting with the sharing policy. Note however, that our game with the exclusive policy does not satisfy the first property, as adding more players can actually deteriorate the sum of utilities. Another example is a generalization known as *covering games* [10]. This sort of game is the same as our game in a single-round version, except that each player chooses not necessarily one element, but a set of elements, from a prescribed set of sets. Again, to be a monotonously scalable game, the congestion policy should be the sharing policy. Note that one may consider a multi-round version of these games, which will be monotonously scalable as well.

## 3 General Policies

The proofs of this section appear in the full version of the paper.

### 3.1 Non-Redundancy and Monotonicity

A doubly-substochastic matrix $A$ is called *non-redundant at time $t$* if $\sum_x A(x,t) = 1$. It is *non-redundant* if it is non-redundant for every $t \leq M$. In the algorithmic view, as $\sum_x A(x,t)$ is the probability that a new box is opened at time $t$, then a strategy's matrix is non-redundant iff it never checks a box twice, unless it already checked all boxes.

▶ **Lemma 10.** *If a profile $\mathbb{P}$ is at equilibrium and $\texttt{success}(\mathbb{P}) < 1$ then every player is non-redundant.*

We will later see that in the symmetric case the condition in the lemma is not needed. However, the following example shows it is necessary in general. Let $M = k$, $T > 1$, and assume that for every $1 \leq i \leq k$, player $i$ goes to box $i$ in every round. Under the exclusive policy, this strategy is an equilibrium, whereas each player is clearly redundant. The following monotonicity lemmas hold under any congestion policy $C$.

▶ **Lemma 11.** *Consider two doubly-substochastic matrices $A$ and $B$. If $A(x,t) > B(x,t)$, and for all $s < t$, $A(x,s) = B(x,s)$ then $v_A(x,t) < v_B(x,t)$.*

▶ **Lemma 12.** *Let $A$ be doubly-substochastic. For every $x$ and $t$, $v_A(x,t+1) \leq v_A(x,t)$. Moreover, if $A(x,t+1) > 0$ then the inequality is strict.*

Using the above, we prove a stronger result than Lemma 10 for the symmetric case:

▶ **Lemma 13.** *If $A$ is a symmetric equilibrium then it is non-redundant.*

**Proof.** Because of Lemma 10 it is sufficient to consider only the case where $\texttt{success}(A) = 1$. Let $T' = \min\{M, T\}$, and assume by contradiction that $A$ is redundant. Thus there is some $t \leq T'$ where $\sum_x A(x,t) < 1$. Hence, $\sum_{s \leq T'} \sum_x A(x,s) < T'$. Therefore, there is some $x$ such that $\sum_{s \leq T'} A(x,s) < 1$ and so $\sum_{s \leq t} A(x,s) < 1$. As $\texttt{success}(A) = 1$, there is some $t' > t$ such that $A(x,t') > 0$. Define $A' = A + \epsilon(\delta_{x,t} - \delta_{x,t'})$. Taking $\epsilon > 0$ small enough, $A'$ is doubly-substochastic. Also, $U_A(A') - U_A(A) = \epsilon(v_A(x,t) - v_A(x,t'))$, which is strictly positive by Lemma 12. Contradicting the fact that $A$ is an equilibrium. ◀

## 3.2  Greedy Strategies

▶ **Lemma 14.** *A non-redundant strategy $A$ is greedy w.r.t. $\mathbb{P}$ at time $t$ iff for every $x$ and $y$, if $A(x,t) > 0$ and $v_{\mathbb{P}}(x,t) < v_{\mathbb{P}}(y,t)$ then $\widetilde{A}(y,t) = 0$.*

The lemma above gives a useful equivalent definition for greediness. We can then prove:

▶ **Theorem 2.** *For every policy $C$ there exists a non-redundant sgreedy strategy. Moreover, all such strategies are equivalent and are symmetric $(1 + C(k))$-equilibria.*

**Proof.** See the full version for a proof for the existence of a strategy $A$ that is non-redundant and sgreedy. We prove here that such a strategy is a $(1 + C(k))$-equilibrium. Consider a strategy $B$. We compare the utility of $B$ to that of $A$ when both play against $k - 1$ players playing $A$. By non-redundancy, all of $v_A(x,t)$ are 0 when $t > M$, and so we can assume $T \leq M$.

Denote $\mathtt{maxv}(t) = \max_x v_A(x,t)$. Since the utility of $B$ in any round $t$ is a convex combination of $v_A(x,t)$, we have $U_A(B,t) \leq \mathtt{maxv}(t)$. We say that $A$ *fills* box $x$ at round $t$ if $A(x,t) > 0$ and $\widetilde{A}(x,t) = 0$. The following four claims hold for any round $t$:

1. If $A$ does not fill any box at round $t$ then $U_A(A,t) = \mathtt{maxv}(t)$. This is because $U_A(A,t)$ is a convex combination of $v_A(x,t)$ for the boxes where $A(x,t) > 0$, which by the characterization of greediness in Lemma 14, all have the same value at time $t$.
2. $U_A(A,1) = \mathtt{maxv}(1)$. Why? if no box is filled in round 1, then Item 1 applies. Otherwise, for some box $x$, $A(x,1) = 1$, and all other boxes have $A(\cdot,1) = 0$. The result follows again by Lemma 14.
3. For any $s < t$, $U_A(A,s) \geq \mathtt{maxv}(t)$. We prove this by showing that for every $x$, $U_A(A,s) \geq v_A(x,t)$. If $v_A(x,t) = 0$, then the claim is clear. Otherwise, $A(x,t) > 0$ or $\widetilde{A}(x,t) > 0$ or both. Either way, $\widetilde{A}(x,s) > 0$. Therefore, as $A$ is sgreedy, for every $y$ such that $A(y,s) > 0$, $v_A(y,s) \geq v_A(x,s) \geq v_A(x,t)$. The last inequality follows from monotonicity, i.e., Lemma 12. As $v_A(A,s)$ is a convex combination of such $y$'s we conclude.
4. If $A$ fills box $x$ at time $t > 1$ then for any $s < t$, $v_A(x,t) \leq C(k)v_A(x,s)$. To see why, first note that $v_A(x,s) \geq f(x)C(1)\widetilde{A}(x,s)^{k-1} = f(x)\widetilde{A}(x,s)^{k-1}$. On the other hand, since $\widetilde{A}(x,t) = 0$, $v_A(x,t) = f(x)C(k)A(x,t)^{k-1} \leq f(x)C(k)\widetilde{A}(x,s)^{k-1}$, because $A(x,t) \leq \widetilde{A}(x,t-1) \leq \widetilde{A}(x,s)$. Combining the above two inequalities gives the result.

Denote by $X_1$ the set of rounds for which there is no box $x$ that is filled by $A$. Let $X_2$ be the rest of the rounds, except for $t = 1$ which is in neither. Also denote $t_0 = \min X_2$, and $t_1 = \max X_2$. Since $U_A(B) \leq \sum_t \mathtt{maxv}(t)$, by Items 1,2 and 3 above,

$$U_A(B) \leq \sum_{t \in X_1 \cup \{1\}} U_A(A,t) + \mathtt{maxv}(t_0) + \sum_{t \in X_2 \setminus \{t_1\}} U_A(A,t) \leq U_A(A) + \mathtt{maxv}(t_0).$$

We conclude by using Items 4 and 2 and showing:

$$\mathtt{maxv}(t_0) = \max_x v_A(x,t_0) \leq \max_x C(k)v_A(x,1) = C(k)U_A(A,1) \leq C(k)U_A(A). \qquad \blacktriangleleft$$

In the full version we provide an example showing that in the sharing policy, a non-redundant sgreedy strategy is not necessarily at equilibrium. On the other hand, it is worth noting that for any policy, the existence of a symmetric equilibrium follows from [23], and for $C(k) = 0$ we can get a full characterization of such equilibria:

▷ **Claim 15.** For any policy such that $C(k) = 0$, all symmetric equilibria are equivalent.

Interestingly, this result does not extend to non-symmetric profiles even for the exclusive policy, as is demonstrated by the following example of a non-greedy non-redundant equilibrium. Consider three players and two rounds. $f(1) = f(2) = f(3) = (1 - \epsilon)/3, f(4) = \epsilon$, for some small positive $\epsilon$. Player 1 plays first 4 and then 1. Player 2 plays 2 and then 3, and player 3 plays 3 and then 2. This can be seen to be an equilibrium, yet player 1 is not greedy.

Finally, the proof of Theorem 4, which shows how to construct an approximate equilibrium in polynomial time, is presented in the full version of the paper. This proof involves defining notions of approximate greediness and non-redundancy, proving an equivalent of Theorem 2 for them, and then using bounds on the rate of change that $v_A(x,t)$ goes through as a function of $A(x,t)$. This allows us to polynomially find an approximate sgreedy and non-redundant matrix, thus giving a polynomial strategy with our use of the Birkhoff von-Neumann theorem (see the full version).

## 4 The Exclusive Policy

Missing proofs of this section appear in the full version of this paper. There, we first prove that under the exclusive policy, $A^\star$ is sgreedy and non-redundant. Hence, according to Theorem 2,

▶ **Theorem 5.** *Under the exclusive policy, Strategy $A^\star$ of [8] is a symmetric equilibrium.*

According to Claim 3 all symmetric equilibria under the exclusive policy are equivalent, and thus equivalent to $A^\star$. Hence, the optimality of $A^\star$ (w.r.t. symmetric profiles) implies that both the PoSA and PoSS of the exclusive policy are optimal. That is, for every $f, T$, and policy $C$,

$$\mathtt{PoSA}(C_{\mathtt{ex}}, f, T) = \mathtt{PoSS}(C_{\mathtt{ex}}, f, T) \leq \mathtt{PoSS}(C, f, T).$$

Our next goal is to establish the PoA of the exclusive policy. For this purpose, we first prove that the success probability of any equilibrium is at least as large as that of any symmetric profile. Since $A^\star$ is a symmetric equilibrium, its optimality among symmetric profiles follows. Hence, the proof provides as alternative proof to the one in [8].

▶ **Theorem 6.** *Consider the exclusive policy. For any profile $\mathbb{P}_{\mathtt{nash}}$ at equilibrium and any symmetric profile $A$, $\mathtt{success}(\mathbb{P}_{\mathtt{nash}}) \geq \mathtt{success}(A)$.*

**Proof.** Let $A$ be a strategy and $\mathbb{P}_{\mathtt{nash}}$ be a profile at equilibrium with respect to $C_{\mathtt{ex}}$. If $\mathtt{success}(\mathbb{P}_{\mathtt{nash}}) = 1$, then the inequality is trivial. According to Lemma 10, we can therefore assume that all players of $\mathbb{P}_{\mathtt{nash}}$ are non-redundant and that $T \leq M$. Denote the probability of visiting $x$ in profile $\mathbb{P}$ by

$$\mathtt{success}(\mathbb{P}, x) = 1 - \prod_{B \in \mathbb{P}} \widetilde{B}(x, T).$$

We say that box $x$ is *high* with respect to a profile $\mathbb{P}$ if $\mathtt{success}(\mathbb{P}, x) > \mathtt{success}(A, x)$, *low* if $\mathtt{success}(\mathbb{P}, x) < \mathtt{success}(A, x)$, and *saturated* if they are equal. The next lemma uses the fact that $A$ is symmetric.

▶ **Lemma 16.** *If a profile $\mathbb{P}$ is non-redundant and contains no high boxes, then all boxes are saturated.*

We proceed to prove a weak greediness property for equilibria. Denote a box $x$ *full* for player $B$ if $\sum_t B(x, t) = 1$. Also, for readability of what follows, when $\mathbb{P}$ is clear from the context, we shall denote $v_B(x, t) = v_{\mathbb{P} - B}(x, t) = f(x) \cdot \prod_{A \in \mathbb{P} \setminus \{B\}} \widetilde{A}(x, t)$.

▶ **Lemma 17.** *Consider a profile* $\mathbb{P}_{\texttt{nash}}$ *at equilibrium. For every* $B \in \mathbb{P}_{\texttt{nash}}$ *and* $t, x, y$ *such that* $y$ *is not full in* $B$, *if* $B(x,t) > 0$ *then* $v_B(x,t) \geq v_B(y,t)$.

**Proof.** Assume otherwise. Define an alternative matrix $B'$ for player $B$, as $B' = B + \epsilon(\delta_{y,t} - \delta_{x,t})$. For a sufficiently small $\epsilon > 0$, $B'$ is a doubly-substochastic matrix because $y$ is not full in $B$. Then, $U_{\mathbb{P}_{\texttt{nash}}^{-B}}(B') - U_{\mathbb{P}_{\texttt{nash}}^{-B}}(B) = \epsilon(v_B(y,t) - v_B(x,t)) > 0$, in contradiction. ◀

Let us define a process that starts with the profile $\mathbb{P}_{\texttt{nash}}$ and changes it by a sequence of *alterations*, each shifting some amount of probability between two boxes. Importantly, we make sure that each alteration can only decrease the success probability. Hence, the proof is concluded once we show that the final profile has a success probability that is higher than that of $A$.

We first describe the alternations. Each alteration considers the current profile $\mathbb{P}$, and changes it to $\mathbb{P}'$. It takes some high box $x$, some low box $y$ (both w.r.t. $\mathbb{P}$), and the maximal $t$ such that there is a player $B \in \mathbb{P}$ with $B(x,t) > 0$. It defines $B' = B + \epsilon(\delta_{y,t} - \delta_{x,t})$, and lets the player that played $B$ play $B'$ instead. This $\epsilon$ is taken to be the largest so that $x$ does not become low, $y$ does not become high, and such that $\epsilon \leq B(x,t)$, so that the entries remain non-negative. Note that $B'$ is doubly sub-stochastic, because taking care that $y$ remains low, also means that $y$'s row in $B'$ still sums to less than 1.

After this alteration, either $x$ is saturated, $y$ is saturated, or $B'(x,t) = 0$. Clearly, in a finite number of alterations a profile $\mathbb{P}_{\texttt{final}}$ is obtained, for which either no box is high or no box is low.

▶ **Lemma 18.** $\texttt{success}(\mathbb{P}_{\texttt{final}}) \geq \texttt{success}(A)$.

**Proof.** By Lemma 16, $\mathbb{P}_{\texttt{final}}$ can only contain high and saturated boxes, that is, for every $x$, $\texttt{success}(\mathbb{P}_{\texttt{final}}, x) \geq \texttt{success}(A, x)$. However, $\texttt{success}(\mathbb{P}) = \sum_x f(x)\texttt{success}(\mathbb{P}, x)$, and so $\texttt{success}(\mathbb{P}_{\texttt{final}}) \geq \texttt{success}(A)$. ◀

Lastly, the following lemma concludes the proof of Theorem 6.

▶ **Lemma 19.** *An alteration can only decrease the probability of success.* ◀

Since $A^\star$ is a symmetric equilibrium, we immediately get that for every $f$ and $T$, the PoA is attained by $A^\star$, that is, $\texttt{PoA}(C_{ex}, f, T) = \max_{\mathbb{P} \in \mathcal{P}(T)} \texttt{success}(\mathbb{P})/\texttt{success}(A^\star)$. Since $A^\star$ has the best success probability among symmetric profiles, and that every policy has some symmetric equilibrium, we get Corollary 7. To make this more concrete, we show that in the worst case,

▶ **Theorem 8.** *For every* $T$, $\sup_f \texttt{PoA}(C_{\texttt{ex}}, f, T) = (1 - (1 - 1/k)^k)^{-1}$.

Note that as $k$ goes to infinity the PoA converges to $e/(e-1) \approx 1.582$.

## 5 Future Work and Open Questions

In [8], the main complexity measure was actually the running time and not the success probability. Our results about equilibria are also relevant to this measure, but the social gain is different. For example, it is still true that $A^\star$ is an equilibrium under the exclusive policy, and that all other symmetric equilibria in the exclusive policy are equivalent to it. As $A^\star$ is optimal among symmetric profiles w.r.t. the running time, the PoSA of $C_{\texttt{ex}}$ is equal to the PoSS, and it is also the best among all policies. Furthermore, importing from [8], we know that the PoSA (w.r.t. the running time) is about 4. However, showing the analogue of

Corollary 7, namely, that the PoA of $C_{\mathtt{ex}}$ is that achieved by $A^\star$, seems difficult, especially because general equilibria are not necessarily greedy. Moreover, the results of Vetta [31] do not apply when analyzing the running time, and finding the PoA, PoSA, and PoSS of the sharing policy, for example, remains open.

Another interesting variant would be to consider feedback during the search. For example, assuming that a player visiting a box $x$ knows whether or not other players were there before. Such a feedback can help in the case that the players collaborate [5], but seems to significantly complicate the analysis in the game theoretic variant.

Finally, we would like to encourage game theoretical studies of other frameworks of collaborative search, e.g., [4, 6, 7, 15].

## References

**1** Susanne Albers and Matthias Hellwig. Online Makespan Minimization with Parallel Schedules. *CoRR*, abs/1304.5625, 2013. `arXiv:1304.5625`.

**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. `doi:10.1145/2554797.2554811`.

**3** Artur Czumaj and Berthold Vöcking. Tight Bounds for Worst-case Equilibria. *ACM Trans. Algorithms*, 3(1):4:1–4:17, February 2007. `doi:10.1145/1186810.1186814`.

**4** Dariusz Dereniowski, Yann Disser, Adrian Kosowski, Dominik Pajak, and Przemyslaw Uznanski. Fast collaborative graph exploration. *Inf. Comput.*, 243:37–49, 2015. `doi:10.1016/j.ic.2014.12.005`.

**5** Stefan Dobrev, Rastislav Královic, and Dana Pardubská. Treasure Hunt with Barely Communicating Agents. In *21st International Conference on Principles of Distributed Systems, OPODIS 2017, Lisbon, Portugal, December 18-20, 2017*, pages 14:1–14:16, 2017. `doi:10.4230/LIPIcs.OPODIS.2017.14`.

**6** Yuval Emek, Tobias Langner, Jara Uitto, and Roger Wattenhofer. Solving the ANTS Problem with Asynchronous Finite State Machines. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*, pages 471–482, 2014. `doi:10.1007/978-3-662-43951-7_40`.

**7** Ofer Feinerman and Amos Korman. The ANTS problem. *Distributed Computing*, 30(3):149–168, 2017. `doi:10.1007/s00446-016-0285-8`.

**8** Pierre Fraigniaud, Amos Korman, and Yoav Rodeh. Parallel Bayesian Search with no Coordination. *J. ACM*, 2019.

**9** 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.

**10** Martin Gairing. Covering Games: Approximation through Non-cooperation. In *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*, pages 184–195, 2009. `doi:10.1007/978-3-642-10841-9_18`.

**11** Luc-Alain Giraldeau and Thomas Caraco. Social Foraging Theory. *Princeton University Press*, 2000.

**12** Ronen Gradwohl and Omer Reingold. Fault tolerance in large games. *Games and Economic Behavior*, 86:438–457, 2014. `doi:10.1016/j.geb.2013.06.007`.

**13** Joseph Y. Halpern. A computer scientist looks at game theory. *Games and Economic Behavior*, 45(1):114–131, 2003. `doi:10.1016/S0899-8256(02)00529-8`.

**14** Joseph Y. Halpern. Computer Science and Game Theory: A Brief Survey. *CoRR*, abs/cs/0703148, 2007. `arXiv:cs/0703148`.

**15**   Yuya Higashikawa, Naoki Katoh, Stefan Langerman, and Shin-Ichi Tanigawa. Online Graph Exploration Algorithms for Cycles and Trees by Multiple Searchers. *J. Comb. Optim.*, 28(2):480–495, August 2014. `doi:10.1007/s10878-012-9571-y`.

**16**   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.

**17**   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. URL: `http://www.jstor.org/stable/3545322`.

**18**   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.

**19**   Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009. `doi:10.1016/j.cosrev.2009.04.003`.

**20**   Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the "Wisdom of the Crowd". *Journal of Political Economy*, 122(5):988–1012, 2014.

**21**   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.

**22**   Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.

**23**   John Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951. URL: `http://www.jstor.org/stable/1969529`.

**24**   Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, 2007.

**25**   University of California Berkeley. BOINC. `https://boinc.berkeley.edu/`, 2017.

**26**   Yiangos Papanastasiou, Kostas Bimpikis, and Nicos Savva. Crowdsourcing exploration. *Management Science*, 2017.

**27**   Arthur C Pigou. The Economics of Welfare. *Library of Economics and Liberty*, 1932.

**28**   Vinod Ramaswamy, Dario Paccagnan, and Jason R. Marden. The Impact of Local Information on the Performance of Multiagent Systems. *CoRR*, abs/1710.01409, 2017. `arXiv:1710.01409`.

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

**30**   Christian Scheideler and Jeremy T. Fineman, editors. *Proceedings of the 30th on Symposium on Parallelism in Algorithms and Architectures, SPAA 2018, Vienna, Austria, July 16-18, 2018.* ACM, 2018. `doi:10.1145/3210377`.

**31**   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.