Social Welfare in Search Games with Asymmetric Information

Bavly, Gilad and Heller, Yuval and Schreiber, Amnon

Bar-Ilan University

27 February 2020

Social Welfare in Search Games with Asymmetric Information

Gilad Bavly, Yuval Heller, Amnon Schreiber*

September 13, 2020

Abstract

We consider games in which players search for a hidden prize, and they have asymmetric information about the prize's location. We study the social payoff in equilibria of these games. We present sufficient conditions for the existence of an equilibrium that yields the first-best payoff (i.e., the highest social payoff under any strategy profile), and we characterize the first-best payoff. The results have interesting implications for innovation contests and R&D races.

Keywords: incomplete information, search duplication, decentralized research, social welfare. **JEL Codes**: C72, D82, D83.

1 Introduction

Various real-life situations involve agents exploring different routes to making a discovery. These situations often have the following three key properties: (1) heterogeneity: agents may differ in terms of their information, search methods, search costs, etc., (2) competitive environment: the agents work separately, and compete to be the successful discoverer, and (3) externality: both the discoverer and society gain from the discovery, but these gains may differ. For concreteness, consider the following motivating example.

*Department of Economics, Bar-Ilan University, Israel. Email addresses: gilad.bavly@gmail.com, yuval.heller@biu.ac.il, amnon.schreiber@biu.ac.il. This manuscript replaces an obsolete working paper titled "The social payoff in differentiation games." We thank Sergiu Hart, Igal Milchtaich, Abraham Neyman, Ron Peretz, Alon Raviv, Dov Samet, and Eilon Solan for various helpful comments. Bavly and Heller are grateful to the European Research Council for its financial support (ERC starting grant #677057). Bavly acknowledges support from the Department of Economics and the Sir Isaac Wolfson Chair at Bar-Ilan University, and ISF grant 1626/18.

Example 1. Society faces a problem of quickly developing a vaccine for a new infectious disease, such as COVID-19. There are various possible research directions that may lead to success. Different research labs (or pharmaceutical R&D divisions) have heterogeneous private information about the most promising route to quickly develop a vaccine. Society gains if the vaccine is found by at least one lab. A lab that discovers the vaccine gains from the discovery (credit or reward for the scientists, or profits for the pharmaceutical firm), and this gain is reduced if multiple labs jointly make the discovery.

These situations (henceforth, *search games*) are common in various important areas such as R&D races in oligopolistic markets (e.g., Loury, 1979; Chatterjee & Evans, 2004; Akcigit & Liu, 2015; Letina, 2016), design of innovation contests (e.g., Erat & Krishnan, 2012; Bryan & Lemus, 2017; Letina & Schmutzler, 2019), and scientific research (e.g., Kleinberg & Oren, 2011). While the effects of most of the above-mentioned properties have been studied extensively in the literature, the idea that agents might have private information has not gotten much attention. Thus, the main methodological innovation of the present model is the introduction of asymmetric information into search games (as discussed in Section 6).[1]

The expected social gain (from a successful discovery) is clearly constrained by the information structure, as we assume that players are competitive and do not share their private information. The social gain may also be constrained by the fact that players' individual preferences can differ from society's, and players have strategic considerations as well. Thus, the main question we study is: what is the highest social payoff in equilibrium?

Highlights of the Model There are n players who search for a prize hidden in one of a finite set of locations.[2] Player i is able to search in at most K_i locations (all at once). Searching incurs a private cost, which is a convex function of the number of locations in which the player searches. Each player receives some private coarse signal about the actual location of the prize, and chooses which locations to search. Specifically, for each player there is a collection of disjoint subsets of locations (namely, a partition), such that her private signal informs the player in which of these subsets the prize resides (a more general information structure is analyzed in Section 5).

We study a one-shot game (i.e., if the prize is not found, players do not get to search again) with simultaneous actions. This assumption, which differs from the dynamic models

[1]We are aware of one related existing model of a search game with asymmetric information, that of Chen *et al.* (2015). The key difference between our model and theirs is that Chen *et al.* rely on enforceable mechanisms, which allow players to safely share their asymmetric information, as all players must follow a contract once it has been signed. By contrast, we consider a setup in which players cannot rely on enforceable mechanisms, and, thus, they are limited to playing Nash equilibria.

[2]The assumption of having a single prize is common in the literature; see, e.g., Fershtman & Rubinstein (1997); Konrad (2014); Liu & Wong (2019).

studied in many of the papers cited above, may be reasonable in situations in which there is severe urgency to make the discovery, such as in the motivating example (see Section 6 for further discussion, and Section 3.3 for examples of what happens when this assumption is relaxed).

We allow the prize's value to depend on the location. Also, the value for society and the individual values for players may all be different. When multiple players search in the same location ("search duplication"), it reduces the reward that each player will receive in case the prize is indeed there. By contrast, the social value of the prize is unaffected by the number of finders (and it is not affected by the players' search costs).

First Main Result Our answer to the question of what society can achieve in equilibrium consists of two main results. The first states that there exists a (pure) equilibrium that yields the first-best social payoff (namely, the highest social payoff that any strategy profile can yield) if the following two conditions hold for any two locations ω and ω' that a player considers possible (after observing her own private signal): (1) *ordinal consistency*: the player and society have the same ordinal ranking between searching (by herself) in ω and in ω', and (2) *solitary-search dominance*: the player always prefers searching ω by herself to searching ω' with other players, or to not searching at all.

It is relatively easy to see that neither condition can be dropped (see the examples presented in Section 3.3), and that the conditions are sufficient in a simple setup without asymmetric information. Our result shows that, perhaps surprisingly, these two conditions are sufficient in the richer setup with asymmetric information as well. The intuition is that no player has an incentive to "spoil" society's payoff by moving from a socially better location to a worse one, nor by moving from a location that she searches alone to a location that others search. We discuss the implications of this result on the design of innovation contests in Section 3.4.

Second Main Result Our second main result characterizes the first-best social payoff. We show that the first-best payoff is constrained only by compatibility with the information structure, where the compatibility condition is in the spirit of Hall's marriage theorem (Hall, 1935). Our proof relies on representing a search game as a bipartite graph and adapting and extending classic results from graph theory, the max-flow min-cut theorem (Ford & Fulkerson, 1956) and the Birkhoff–von Neumann theorem (Birkhoff, 1946; Von Neumann, 1953), to our setup.[3]

[3]Recent economic applications (and extensions) of these graph-theory results have appeared in matching mechanisms (e.g., Budish *et al.*, 2013; Bronfman *et al.*, 2018), large anonymous games (e.g., Blonski, 2005), public good games with multiple resources (e.g., Tierney, 2019), and auctions of multiple discrete items (e.g.,

One interesting implication of this result (presented in Section 4.1) is that the first-best payoff would not increase if we modified our setup and allowed players to coordinate partial search efforts within locations, so that their efforts do not overlap (i.e., when two players each assign an effort of 50% to location ω, the prize is always found, if it is in ω).

Structure Section 2 presents our model. We study the existence of an equilibrium with a first-best social payoff in Section 3. Section 4 characterizes the first-best payoff. In Section 5 we consider more general information structures. We conclude and discuss the relations with the literature in Section 6. Appendix A applies our results to a special class of search games. Appendix B presents the formal proofs.

2 Model

Setup Let $N = \{1, 2, ..., n\}$ be a finite set of players. A typical player is denoted by i. We use $-i$ to denote the set of all players except player i. We describe the private information of the players in terms of knowledge partitions (Aumann, 1976). Let Ω be the set of the states of the world (henceforth, *states*). Nature chooses one state $\omega \in \Omega$ that is the true state of the world. Each player i is endowed with Π_i, which is a partition of Ω, namely, a list of disjoint subsets of Ω whose union is the whole Ω. We refer to the elements of player i's partition (i.e., the subsets) as player i's *cells*. For each state ω, let $\pi_i(\omega)$ denote the cell of player i that contains the state ω. If the true state is ω, then player i knows that the true state is one of the states in $\pi_i(\omega)$.

Note that the knowledge partitions framework is equivalent to a model in which each player observes a private random signal. Each cell of player i's partition corresponds to a different realization of her private signal. W.l.o.g., one may view the partition Π_i as the set of possible realizations of player i's private signal itself; i.e., each cell in Π_i is a possible signal, and if the state of the world is ω then player i observes the signal $\pi_i(\omega)$.

The players search for a prize hidden in one of a finite set of possible locations. Importantly, in the baseline model we assume that the location of the prize determines the private signal of each player (in other words, the signals that players observe are a deterministic function of the prize's location). This implies that w.l.o.g. each state of the world in our model corresponds to a different location of the prize.[4] Hence, we identify the finite set of *locations* with the set of states Ω. When a player searches in location (i.e., state) $\omega \in \Omega$, she finds the prize if the location of the prize is ω (i.e., if the true state of the world is ω).

Ben-Zwi, 2017).

[4]In Section 5 we discuss a more general model that dispenses with the above assumption.

Figure 1 demonstrates an information structure in a two-player search game.

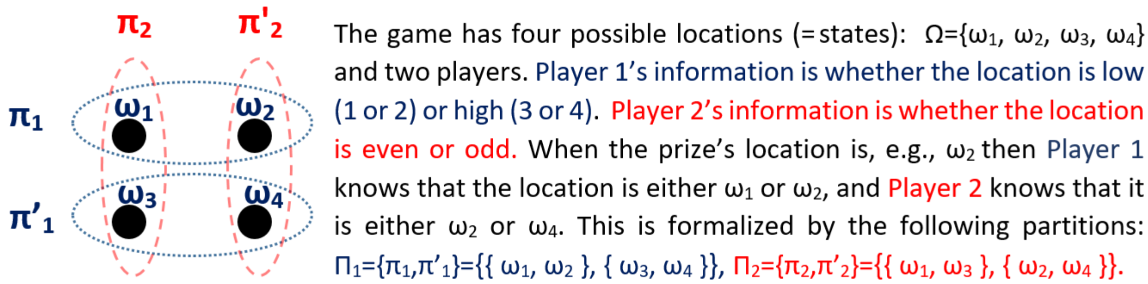

Figure 1: Illustration of information structure of a two-player search game

We say that player i receives no information at all if her information partition Π_i is trivial, i.e., $\Pi_i = \{\Omega\}$ contains a single element, which is the whole Ω. A setting that does not allow for asymmetric information corresponds to the degenerate case in our model where all players have trivial partitions.

Let $\mu \in \Delta(\Omega)$ denote the (common) prior belief about the prize's location, where $\Delta(\Omega)$ denotes the set of distributions over Ω. For a subset of locations $E \subseteq \Omega$, let $\mu(E) = \sum_{\omega \in E} \mu(\omega)$ denote the prior probability of E. For non-triviality, we assume that every cell has a positive prior probability, i.e., $\mu(\pi_i) > 0$ for every $\pi_i \in \Pi_i$ and $i \in N$. When the (unknown) location of the prize is ω, each player i assigns a posterior belief of $\mu(\omega'|\pi_i(\omega))$ to the location being ω', where

$$\mu\left(\omega'|\pi_i(\omega)\right) = \begin{cases} \mu(\omega')/\mu(\pi_i(\omega)) & \omega' \in \pi_i(\omega) \\ 0 & \omega' \notin \pi_i(\omega). \end{cases}$$

We allow heterogeneity in the maximal number of locations that each player can search. Specifically, each player i chooses up to $K_i \in \mathbb{N}$ locations in which she searches, where K_i is the player's search *capacity*. A (pure) *strategy* of player i is a function s_i that assigns to each cell $\pi_i \in \Pi_i$ a subset of π_i with at most K_i elements. We interpret $s_i(\pi_i)$ as the set of up to K_i locations in which player i searches when she observes the signal π_i. If no ambiguity can arise, we may also say that player i (ex-ante) searches in location ω, if $\omega \in s_i(\pi_i(\omega))$, i.e., if player i searches in ω when the prize is located in[5] ω.

We focus in the present paper on pure strategies. Let $S_i \equiv S_i(G)$ denote the set of all (pure) strategies of player i, and let $S \equiv S(G) = \prod_{i \in N} S_i$ be the set of strategy profiles in the game G. For example, in Figure 1 Player 1, with a capacity of one, has four pure strategies. One such strategy, denoted by s_1, is given by $(s_1(\pi_1) = \omega_2; s_1(\pi_1') = \omega_3)$; i.e., a player following s_1 searches in location ω_2 upon observing signal π_1 and searches in ω_3 upon

[5]Equivalently, the locations in which player i (ex-ante) searches are $\cup_{\pi_i \in \Pi_i} s_i(\pi_i)$, namely, the union (across all her cells) of the locations she searches within each cell when that cell happens to be her signal.

observing π'_1. Suppose that Player 1 follows s_1 and the location of the prize is ω_4. Then she will observe the signal π'_1 and search in ω_3 (and hence she will not find the prize).

Remark 1. All of our results hold in a more general setup with either of the following extensions (with minor modifications to the proofs):

1. *Heterogeneous priors*: each player i has a different prior μ_i.

2. *Heterogeneous restricted locations: each player i is allowed to search only in a subset $\Omega_i \subseteq \Omega$ of the locations.*

Costs, Rewards, and Duplication Searching incurs a private cost, which is a convex function of the number of locations in which a player searches.[6] Specifically, each player i bears a cost $c_i(k) \geq 0$ when searching within k locations, where $c_i(0) = 0$ and $c_i(k+1) - c_i(k) \geq c_i(k) - c_i(k-1)$ for any $k \in \{1, .., K_i - 1\}$. We say that a game has a *costless search* (up to the capacity constraints) if $c_i \equiv 0$ (i.e., if $c_i(k) = 0$ for every $k \in \{1, .., K_i\}$ and every player i).

For any location ω, let $v_i^m(\omega) \in \mathbb{R}^+$ denote the *reward* for player i when m players, including player i, find the prize in ω. The reward for finding the prize alone, $v_i^1(\omega)$, is also called the *private value* of player i (at location ω). We assume that the finder's reward is weakly decreasing in the number of joint finders (i.e., $v_i^{m+1}(\omega) \leq v_i^m(\omega)$ for any m and ω), which reflects the negative impact of search duplication. Two examples of such decreasing rewards that are commonly used in the literature are (1) $v_i^m(\omega) = \frac{1}{m} \cdot v_i^1(\omega)$, which may correspond to a setup in which one of the players who search in the prize's location is randomly chosen to be its undisputed owner, and she gains the prize's full value (see, e.g., Fershtman & Rubinstein, 1997), and (2) $v_i^m(\omega) = 0$ for any $m \geq 2$ and any ω, which corresponds to a setup in which a (Bertrand) price competition between the pharmaceutical firms or a "credit war" between the research labs destroys the finder's reward in case of a joint discovery (e.g., Chatterjee & Evans, 2004).

In addition to the players, we introduce an external entity, *society*, who is not one of the players and is indifferent to the identity of the prize finder, as long as the prize is found. In our normative analysis we set the objective of maximizing society's payoff. One can think of society as representing a government who cares for the welfare of those in society (e.g., consumers or patients) who will be affected by the discovery. For any location ω, let $v_{\mathfrak{s}}(\omega) \in \mathbb{R}^+$ denote the prize's social value for society when the prize is found in ω. Note that the social value does not depend on the identity or the number of the prize's finders. In

[6]Extending the costs to depend also on which locations, not just how many, are being searched may be an interesting direction for future research.

particular, the social value of the prize is not reduced when there are multiple finders, which seems plausible in various setups. For example, it seems plausible that price competition between competing pharmaceutical firms will not harm society (it might even benefit the consumers), and that the social gain from a new discovery is not likely to be reduced when two scientists fight over the credit.

Further note that in our model society disregards the players' search costs. This modeling choice seems reasonable in setups where the potential social impact of a discovery overshadows (in society's eyes) the player's individual gains and costs, as in the motivating example of finding a vaccine. In other setups this assumption might be less appropriate, and we leave for future research the interesting question of how to extend our model and our results to a social value that accounts for players' costs (bearing in mind that players may differ both in their costs and in their reward function). Chatterjee & Evans (2004) study the efficiency of the equilibrium outcomes with a social value that accounts for players' costs, but their scope is limited to two players and two locations.

We say that the game has *common values* if $v_i^1(\omega) = v_j^1(\omega) = v_{\mathfrak{s}}(\omega)$ for every two players $i, j \in N$ and every location $\omega \in \Omega$.

Summarizing all the above components allows us to define a *search game* as a tuple $G = (N, \Omega, \Pi, \mu, K, c, v)$, with the various components as defined above.

Private Payoffs and Equilibrium Fix a strategy profile $s \in S$. Let $m_s(\omega)$ denote the number of players who search in ω when the prize's location is ω, i.e.,

$$m_s(\omega) = \sum_{i \in N} \mathbf{1}_{\omega \in s_i(\pi_i(\omega))}.$$

The reward (resp., cost) of player i conditional on the prize's location being ω is equal to $1_{\omega \in s_i(\pi_i(\omega))} v_i^{m_s(\omega)}(\omega)$ (resp., $c_i(|s_i(\pi_i(\omega))|)$). Thus, the (net) payoff of player i conditional on the prize's location being ω, denoted by $u_i(s|\omega)$, is

$$u_i(s|\omega) = 1_{\omega \in s_i(\pi_i(\omega))} v_i^{m_s(\omega)}(\omega) - c_i(|s_i(\pi_i(\omega))|).$$

The players and society are both risk neutral with respect to their payoffs. The (ex-ante) expected (net) payoff of player i is given by $u_i(s) = \sum_{\omega \in \Omega} \mu(\omega) \cdot u_i(s|\omega)$.

A strategy profile $s = (s_1, ..., s_n)$ is *a (Bayesian) Nash equilibrium* of search game G if no player can gain by unilaterally deviating from the equilibrium; i.e., if for every player i and every strategy s_i' the following inequality holds: $u_i(s) \geq u_i(s_i', s_{-i})$, where s_{-i} describes the strategy profile played by all players except player i.

Social Payoff Fix a strategy profile $s \in S$. Let $U(s|\omega) = v_{\mathfrak{s}}(\omega) \cdot \mathbf{1}_{m_s(\omega) \geq 1}$ denote the social payoff, conditional on the prize's location being ω. The expected social payoff is equal to $U(s) = \sum_{\omega \in \Omega} \mu(\omega) \cdot U(s|\omega)$. Let U_{opt} denote the *socially optimal payoff* (or the *first-best* payoff): $U_{\text{opt}} = \max_{s \in S} U(s)$. A strategy profile s is *socially optimal* if it achieves the socially optimal payoff, i.e., if $U(s) = U_{\text{opt}}$.

A strategy profile is *location-maximizing* if it maximizes the number of locations in which the prize is found; i.e., if for any strategy profile $s' \in S$,

$$\sum_{\omega \in \Omega} \mathbf{1}_{\{m_s(\omega) \geq 1\}} \geq \sum_{\omega \in \Omega} \mathbf{1}_{\left\{m_{s'}(\omega) \geq 1\right\}}.$$

The set of socially optimal strategy profiles is typically different from the set of location-maximizing strategy profiles. The two notions coincide if society assigns the same value to every location, i.e., if $v_{\mathfrak{s}}(\omega) = v_{\mathfrak{s}}(\omega')$ for any two locations $\omega, \omega' \in \Omega$. A strategy profile is *exhaustive* if the prize is always found, i.e., if $m_s(\omega) \geq 1$ for every $\omega \in \Omega$. It is immediate that an exhaustive strategy profile is both socially optimal and location-maximizing.

3 Socially Optimal Equilibrium

In this section we present conditions under which the strategic constraints (namely, each player maximizing her private payoff) do not limit the social payoff; that is, we give sufficient conditions for the existence of socially optimal equilibria.

3.1 Search Games are Weakly Acyclic

A sequence of strategy profiles is an improvement path (Monderer & Shapley, 1996) if each strategy profile differs from its preceding profile by the strategy of a single player, who obtained a lower payoff in the preceding profile.

Definition 1. A sequence of strategy profiles $(s^1, ..., s^T)$ is an *improvement path* if for every $t \in \{1, ..., T-1\}$ there exists a player $i_t \in N$ such that: (1) $s_j^t = s_j^{t+1}$ for every player $j \neq i_t$, and (2) $u_{i_t}\left(s^{t+1}\right) > u_{i_t}\left(s^t\right)$.

We begin by presenting an auxiliary result, which states that any search game is weakly acyclic: starting from any strategy profile, there exists an improvement path that ends in a Nash equilibrium.[7]

[7]The proof introduces an agent-normal form representation of our game (in the spirit of Selten, 1975), which is similar to matroid congestion games with player-specific payoffs. Ackermann *et al.* (2009, Theorem 8) show that these games are weakly acyclic. Their result cannot be directly applied to our setup, as there

Definition 2 (Milchtaich, 1996). A game is *weakly acyclic* if for any $s^1 \in S$, there exists an improvement path $(s^1, ..., s^T)$, such that s^T is a (pure) Nash equilibrium.

Proposition 1. *Any search game is weakly acyclic.*

Sketch of proof; formal proof is in Appendix B.1. Define the payoff of a cell $\pi_i \in \Pi_i$ as the expected payoff of player i given that her signal is π_i. Note that player i is best-responding iff every cell of i is best-responding. Player i has K_i units of capacity, which we index by $j = 1, \ldots, K_i$. A *cell-unit* of player i is a pair (π_i, j), where $\pi_i \in \Pi_i$ is a cell, and j a unit index. W.l.o.g. we assume that a strategy chooses a specific location for every cell-unit α, or chooses that α be inactive. We define the payoff of a cell-unit (π_i, j) as the payoff of the cell π_i. Note that this payoff equals the sum of the (interim) expected rewards in the locations of π_i's active cell-units, minus the cost of activating that many cell-units.

Given a strategy profile, suppose that there is no single inactive cell-unit whose activation improves its own (i.e., the cell's) payoff. Then activating multiple cell-units does not improve the cell's payoff either, because of the convexity of the cost function. The case of deactivation is similar. Therefore, we can show that a cell π_i is best-responding iff every cell-unit of π_i is best-responding.

The key part is Lemma 1 that says that if the members of a set B of cell-units (of various players) are best-responding, and $\alpha \notin B$ is another cell-unit, then there is a sequence of cell-unit improvements that ends with all the members of $B \cup \{\alpha\}$ best-responding. To prove weak acyclicity, start from any profile s^1, and using this lemma inductively add one cell-unit at a time, until eventually everyone is best-responding.

To prove the lemma, we construct a sequence of improvements by the members of $B \cup \{\alpha\}$. First, let α switch from its current choice to its best-response. If α was active before the switch, we add a dummy player in the location ω^1 that α left. Now begins a sequence we call Phase I. Suppose that α switched to some location ω^2. While cell-units (of $B \cup \{\alpha\}$) not located in ω^2 are still best-responding, those in ω^2 may now prefer to switch because of the extra cell-unit in ω^2 (call ω^2 the current "plus location"). Let one of them switch to its best-response ω^3, and then another cell-unit may switch from ω^3, etc. Phase I goes on until everyone is best-responding, unless someone switches to ω^1, in which case Phase I is immediately terminated.

If a cell-unit is deactivated on stage t, it will not incentivize another cell-unit to deactivate on stage $t+1$, because of the convexity of costs. Moreover, Phase I will end after stage t, since there would not be any plus location.

are some technical differences; most notably, our cost function being non-linear (while in Ackermann *et al.*'s setup the cost of searching in two locations must be the sum of the costs in each location). Nevertheless, the proofs turn out to be similar.

To see that Phase I cannot go on forever, consider a cell-unit β that switches from location ω to location ω', making ω' the new plus location. The switch must strictly increase β's expected reward, and later the expected reward in ω' cannot drop below its current level; it may only be higher (if the plus is somewhere else). Thus, β's expected reward will never drop back to the level it was at before the switch, even if β does not improve again. Therefore, Phase I cannot enter a cycle; hence, it must end.

Let σ^* denote the strategy profile when Phase I ends. At this point we remove the dummy from ω^1, and denote the resulting profile by s^*. If Phase I ended because someone switched to ω^1, then everyone is best-responding under s^*, and we are done. Otherwise, Phase I ended because everyone was best-responding under σ^*, and now follows Phase II.

While Phase I can be described as restabilizing after one cell-unit is added, the analogous Phase II restabilizes after one cell-unit is removed. First, one cell-unit switches from some location ω' to the current "minus location" ω^1, then another switches to ω', etc. On each stage we choose a cell-unit switch that is best for its cell, i.e., there exists no cell-unit switch that yields a higher increase in that cell's payoff.

Phase II must eventually end, by the argument analogous to that of Phase I. Then everyone is best-responding, and the lemma is proven. □

In particular, Proposition 1 implies that:

Corollary 1. *Any search game admits a pure Nash equilibrium.*

3.2 Existence of a Socially Optimal Equilibrium

We begin by defining two properties required for our first main result (Theorem 1).

Ordinal Consistency Our first property requires that the ordinal ranking of any player over her expected private values within a cell is (weakly) compatible with society's ranking. That is, we say that a search game has ordinally consistent payoffs if for any two locations ω and ω' in the same cell of player i, if the expected private value of player i is strictly lower in ω than in ω', then the expected social value is weakly lower in ω.

Definition 3. Search game G has *ordinally consistent payoffs* if for any player i, any cell $\pi_i \in \Pi_i$, and any two locations $\omega, \omega' \in \pi_i$, the following implication holds:

$$\mu(\omega) \cdot v_i^1(\omega) < \mu\left(\omega'\right) \cdot v_i^1\left(\omega'\right) \Rightarrow \mu(\omega) \cdot v_{\mathfrak{s}}(\omega) \leq \mu\left(\omega'\right) \cdot v_{\mathfrak{s}}\left(\omega'\right).$$

Observe that having common values implies that the search game has ordinally consistent payoffs. Further observe that if society has uniform expected values (i.e., if $\mu(\omega) \cdot v_{\mathfrak{s}}(\omega) =$

$\mu(\omega') \cdot v_{\mathfrak{s}}(\omega')$ for any two locations $\omega, \omega' \in \Omega$), then the search game has ordinally consistent payoffs regardless of what the players' private payoffs are.

Solitary-Search Dominance Solitary-search dominance requires that any player always prefer searching alone in any location to (1) searching jointly with other players in another location within the same cell, or (2) leaving some of her search capacity unused. Formally:

Definition 4. Search game G has *solitary-search dominant payoffs* if

$$\mu(\omega|\pi_i) \cdot v_i^1(\omega) \geq \mu\left(\omega'|\pi_i\right) \cdot v_i^2\left(\omega'\right) \tag{1}$$

and

$$\mu(\omega|\pi_i) \cdot v_i^1(\omega) \geq c_i(K_i) - c_i(K_i - 1), \tag{2}$$

for any player i, any cell $\pi_i \in \Pi_i$, and any pair[8] $\omega, \omega' \in \pi_i$.

Suppose first that there is no asymmetric information (which corresponds to the case where all players have trivial information in our model). If either ordinal consistency or solitary-search dominance are not assumed, it is relatively easy to construct a game that does not admit a socially optimal equilibrium (we construct such examples on Section 3.3). It is also not hard to show, on the other hand, that ordinal consistency and solitary-search dominance imply the existence of a socially optimal equilibrium. By virtue of Proposition 1, we can now show that this remains true with asymmetric information, as these two conditions are sufficient for any search game.

Theorem 1. *Let G be a search game with ordinally consistent and solitary-search dominant payoffs. Then there exists a socially optimal (pure) equilibrium.*

Proof. Consider a pure strategy profile that maximizes the social payoff. Proposition 1 implies that there is a finite sequence of unilateral improvements that ends in a Nash equilibrium. In what follows we show that the properties of ordinal consistency and solitary-search dominance jointly imply that the social payoff cannot decrease along that sequence. Without loss of generality we can assume that each unilateral improvement consists of changing merely a single choice within a single cell, since this is in fact what the proof of Proposition 1 shows.

[8]Due to the assumption of the cost function being convex, (2) implies that $\mu(\omega|\pi_i) \cdot v_i^1(\omega) \geq c_i(k) - c_i(k-1)$ for any $1 \leq k \leq K_i$. A similar assumption of the search cost being sufficiently small so that players always prefer searching alone to not using their search capacity appears in Chatterjee & Evans (2004).

First we note that in each improvement, if the improving player leaves a location in which there were multiple searchers, then the social payoff cannot decrease. Next, solitary-search dominance implies that if she leaves a location in which she is the sole searcher, then she moves to an unoccupied location, as moving to an occupied location would contradict (1), and "quitting" (namely, deactivating that unit of capacity) would contradict (2). Finally, ordinal consistency implies that if she moves to being the sole searcher in another location, then the social payoff must weakly increase. □

In the socially optimal equilibrium, search costs may sometimes deter a player from searching in some location ω if other players might search there as well. Inequality (2) merely states that she will never be deterred by costs if she can search in ω alone.

Our next result states that even without the ordinal consistency assumption, some efficiency is still guaranteed, in the sense that there exists an equilibrium that maximizes the number of locations in which the players search. Formally:

Corollary 2. *Every search game G with solitary-search dominant payoffs admits a location-maximizing equilibrium.*

Proof. Let $\hat{G} = (N, \Omega, \Pi, \mu, K, c, \hat{v})$ be a search game similar to $G = (N, \Omega, \Pi, \mu, K, c, v)$, except that $\hat{v}_{\mathfrak{s}}(\omega) = {}^{1}/_{\mu(\omega)}$ for any locations $\omega \in \Omega$. Observe that $\hat{G}$ is a search game with ordinally consistent and solitary-search dominant payoffs. This implies that $\hat{G}$ admits a socially optimal equilibrium $\hat{s}$. Observe that the definition of $\hat{v}_{\mathfrak{s}}$ implies that $\hat{s}$ is a location-maximizing strategy profile. Further observe that $\hat{s}$ is also an equilibrium of G (as G and $\hat{G}$ differ only in the social payoff). □

In particular, any game with solitary-search dominant payoffs that admits an exhaustive strategy profile, also admits an exhaustive equilibrium.

Price of Stability/Anarchy Theorem 1 states that there is an equilibrium that maximizes the social payoff (i.e., that the price of stability is 1)[9] in any search game with ordinally consistent and solitary-search dominant payoffs. By contrast, Figure 2 demonstrates that the social payoff might be substantially lower in other Nash equilibria (i.e., that the price of anarchy can be more than 1).

[9]The price of stability (resp., anarchy) is defined as the ratio between the socially optimal payoff U_{opt} and the maximal (resp., minimal) social payoff induced by a Nash equilibrium; i.e., $PoS = \frac{U_{\text{opt}}}{\max_{s \in NE(G)} U(s)}$ and $PoA = \frac{U_{\text{opt}}}{\min_{s \in NE(G)} U(s)}$, where $NE(G)$ is the set of Nash equilibria.

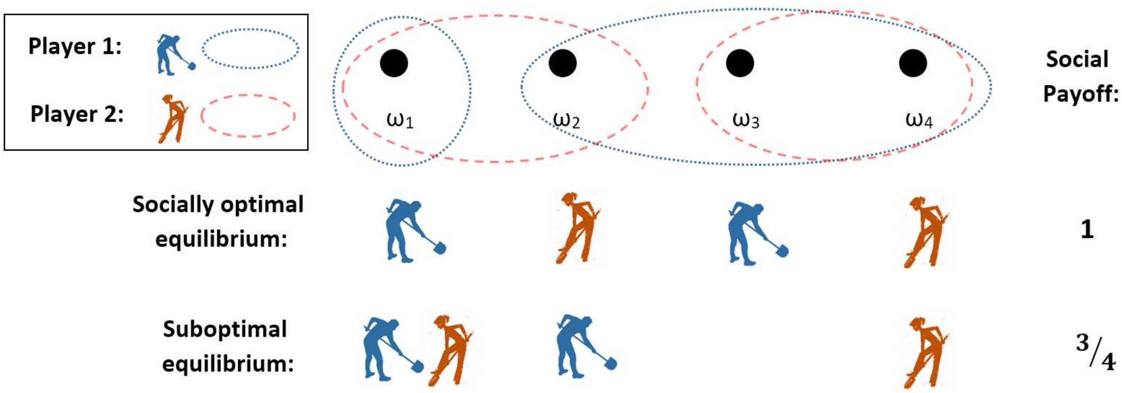

Figure 2: **Example for the price of anarchy**. The figure presents two equilibria in a two-player search game with ordinally consistent and solitary-search dominant payoffs (the ellipses represent the partition elements), uniform prior, costless search ($c \equiv 0$), reward of $v_i^m \equiv \frac{1}{m}$, social value $v_{\mathfrak{s}} \equiv 1$, and a capacity of one for every player. The figure shows the location searched by each player for each possible signal. For example, in the socially optimal equilibrium Player 1 searches in location ω_1 when observing the signal $\{\omega_1\}$ and searches in ω_3 when observing the signal $\{\omega_2,\omega_3,\omega_4\}$. The first (resp., second) equilibrium is (resp., is not) socially optimal with a social payoff of 1 (resp., 0.75).

3.3 Necessity of All Assumptions in Theorem 1

The following three examples demonstrate that all the assumptions of Theorem 1 are necessary to guarantee the existence of a socially optimal equilibrium. We postpone the discussion of the necessity of deterministic signals to Section 5.

Necessity of Solitary-Search Dominance Example 2 demonstrates that solitary-search dominance is necessary for Theorem 1.

Example 2. For any $r \in (0,1)$ let

$$G = \left(N = \{1,2\}, \Omega = \left\{\omega, \omega'\right\}, \Pi \equiv \{\Omega\}, \mu, K \equiv 1, c \equiv 0, \left(v_i^1 \equiv 1, v_i^2 \equiv r, v_{\mathfrak{s}} \equiv 1\right)\right)$$

be a two-player search game with trivial information partitions (namely, each partition Π_i contains a single element, which is the whole Ω), and a common prior μ defined as follows: $\mu(\omega) = \frac{2}{3}$ and $\mu(\omega') = \frac{1}{3}$. Both locations induce a private value of 1 to a sole searcher and a private value of $r \in (0,1)$ in case of simultaneous searches. Note that G has ordinally consistent payoffs, and that it satisfies solitary-search dominance iff $r \leq 0.5$. In what follows we show that for any $r > 0.5$ the unique best-reply against an opponent who searches in location ω is to search in ω as well (which implies that searching in ω is a dominant strategy). This is so because searching in ω yields an expected payoff of $\frac{2}{3} \cdot r$, while searching in ω' yields

$\frac{1}{3} \cdot 1$. This, in turn, implies that the unique equilibrium is both players searching in ω, which is suboptimal.

Necessity of Ordinal Consistency Example 3 demonstrates that the consistency requirement is necessary to guarantee the existence of a socially optimal equilibrium. Specifically, it shows that even for one-player search games, and even when society and the player have the same ordinal ranking over the values of the prize in each location and search is costless, the unique Nash equilibrium is not necessarily socially optimal if the ordinal consistency requirement is not satisfied.

Example 3. Let $G = (N = \{1\}, \Omega = \{\omega, \omega'\}, \Pi_1 = \{\Omega\}, \mu, K_1 = 1, c_1 = 0, v)$ be a one-player search game with a prior $\mu(\omega) = 1/4$, $\mu(\omega') = 3/4$, and with values of $v_{\mathfrak{s}}(\omega) = 2$, $v_{\mathfrak{s}}(\omega') = 1$, $v_1^1(\omega) = 4$, and $v_1^1(\omega') = 1$. Observe that the game's payoffs are trivially solitary-search dominant due to having a single player and a costless search. It is simple to see that the player searches in location ω in the unique equilibrium, although this yields a lower social payoff than searching in ω'.

Necessity of One-Shot, Simultaneous Searches An (implicit) key assumption in our model is that all searches are done simultaneously at a single point in time. The following two examples demonstrate that Theorem 1 is no longer true if players search sequentially (Example 4) or if there are multiple rounds of search (Example 5).

Example 4 (Sequential play, see Figure 3)**.** Let

Figure 3: Illustration of Example 4: Sequential Play

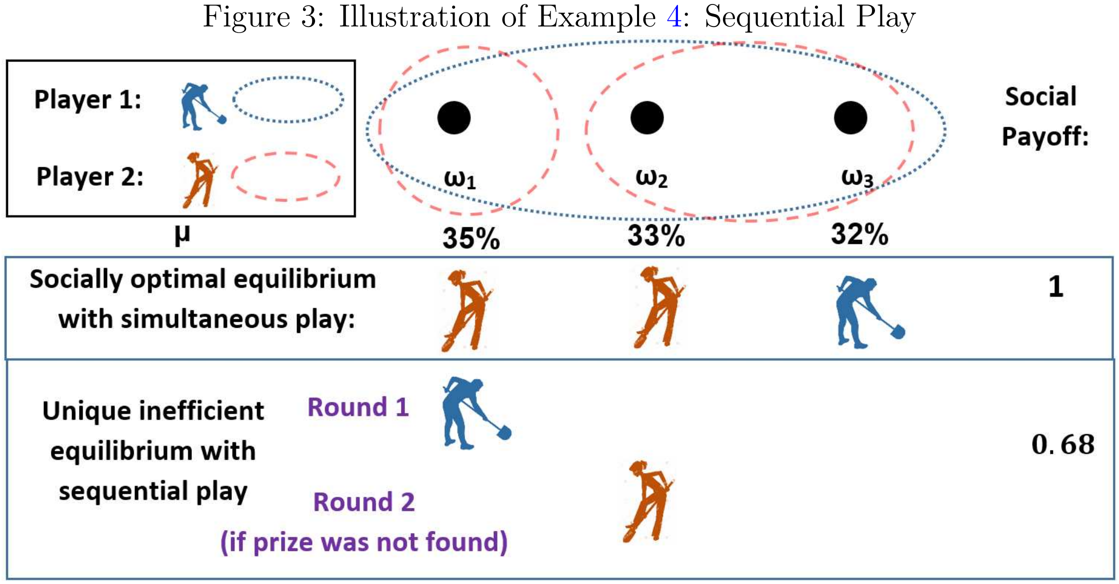

$$G=\left(N=\{1,2\},\Omega=\{\omega_1,\omega_2,\omega_3\},\Pi,\mu,K\equiv 1,c\equiv 0,\left(v_i^1\equiv 1,v_i^2\equiv 0.5,v_{\mathfrak{s}}\equiv 1\right)\right)$$

be a two-player search game with three locations, capacity 1 for each player and costless search. All locations yield a private value of 1, which is equally shared between simultaneous finders. Player 1 has the trivial partition $\Pi_1=\{\Omega\}$, while Player 2 knows if the prize is in location ω_1 or not (i.e., $\Pi_2=\{\{\omega_1\},\{\omega_2,\omega_3\}\}$. The prior assigns slightly higher (resp., lower) probability to location ω_1 (resp., ω_3), i.e., $\mu(\omega_1)=35\%$, $\mu(\omega_2)=33\%$, $\mu(\omega_3)=32\%$. Observe that the game satisfies ordinal consistency and solitary-search dominance. In our model, in which players search simultaneously, the game admits two (pure) Nash equilibria, both of which are socially optimal: Player 1 searches in either location 2 or 3, and Player 2 searches in the remaining two locations.

By contrast, if the game is sequential and Player 1 plays first, then the game admits a unique equilibrium, which is not efficient: Player 1 searches in ω_1, and if the prize has not been found, Player 2 searches in location ω_2 (and no player searches in location ω_3). Note that this profile is the unique equilibrium regardless of whether or not the model lets Player 2 observe the location in which Player 1 searched in the previous round.

Example 5 (Multiple Rounds, see Figure 4.)**.** Let

Figure 4: Illustration of Example 5: Multiple Rounds

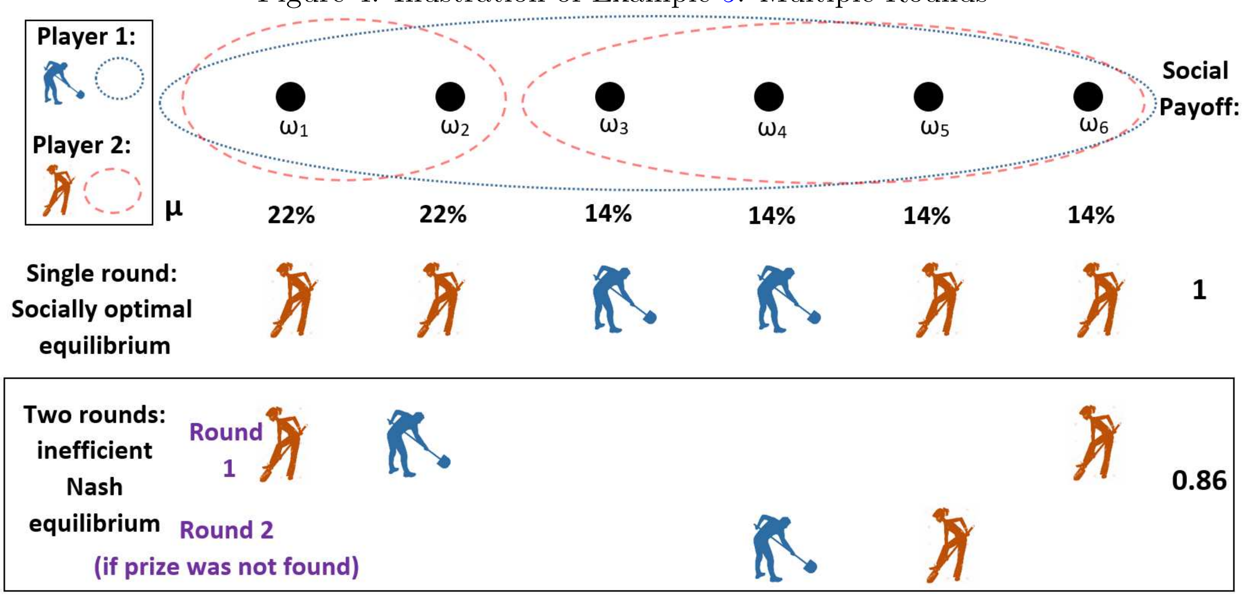

$$G=\left(N=\{1,2\},\Omega=\{\omega_1,\omega_2,\omega_3,\omega_4,\omega_5,\omega_6\},\Pi,\mu,K\equiv 2,c\equiv 0,\left(v_i^1\equiv 1,v_i^2\equiv 0.5,v_{\mathfrak{s}}\equiv 1\right)\right)$$

be a two-player search game with six locations, a capacity of two for each player, and costless search. All locations yield a private value of 1, which is equally shared between simultaneous finders. The prior assigns a higher probability to the first two locations: $\mu(\omega_1) = \mu(\omega_2) = 22\%$, $\mu(\omega_3) = \mu(\omega_4) = \mu(\omega_5) = \mu(\omega_6) = 14\%$. Player 1 has the trivial partition $\Pi_1 = \{\Omega\}$, while Player 2 knows if the prize is in the first two locations or not, i.e., $\Pi_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4, \omega_5, \omega_6\}\}$. Observe that the game satisfies ordinal consistency and solitary-search dominance. In our model, in which there is a single round of play, all Nash equilibria are socially optimal: Player 1 searches in two locations out of $\{\omega_3, \omega_4, \omega_5, \omega_6\}$, while Player 2 searches in the remaining four locations.

By contrast, if there are two rounds of play, where in each round the players search simultaneously and each player uses one unit of her capacity, then in any equilibrium Player 1 searches in either location ω_1 or ω_2 in the first round, because in at least one of these two locations the probability of Player 2 searching there in the first round is at most 50%, and thus searching there yields an expected payoff of at least $\frac{22\%+11\%}{2} = 16.5\% > 14\%$. This, in turn, implies that at least one of the locations $\{\omega_3, \omega_4, \omega_5, \omega_6\}$ is not searched by any player in any of the two rounds, which implies that none of the equilibria is efficient. Here, too, this holds regardless of whether or not each player observes where the opponent searched in the previous round.

3.4 Implications for Innovation Contests

Consider the setup of an innovation contest, in which a contest designer, who wishes to maximize the social payoff, might influence the private payoffs of players by offering a monetary bonus to the prize's finder, which is added to the inherent reward. Kleinberg & Oren (2011) characterize the optimal bonuses in a model of such a contest without private information (additional differences are that in their model each player searches once, search is costless, and each player has a positive probability of failing to find the prize when searching in the prize's location). The following fact plays an important part in their analysis (Kleinberg & Oren, 2011, Claim 2.2): the socially optimal strategy profile is obtained by a simple greedy algorithm, according to which players are assigned to locations one at a time in an arbitrary order, and in each iteration the player is assigned to the location with the greatest expected social value. By contrast, the example illustrated in Figure 2 demonstrates that a greedy algorithm may not yield the socially optimal profile in our setup with asymmetric information, because the above-mentioned greedy algorithm that starts with Player 1 may assign Player 1 to search in location ω_2, and thus it may lead to the suboptimal profile in which no player searches in location ω_3. The failure of the greedy algorithm suggests that the analysis of the current setup may substantially differ from Kleinberg & Oren (2011).

In what follows we sketch a few implications of Theorem 1 in a contest with asymmetric information, while leaving the interesting question of characterizing the optimal bonuses in this setup to future research.

Observe first that if the private payoffs satisfy ordinal consistency and solitary-search dominance, then Theorem 1 implies that the designer can maximize the social payoff without offering any bonus: the designer is only required to be able to give nonenforced recommendations to the players (which allows him to induce the play of the socially optimal Nash equilibrium, rather than other equilibria). In what follows we consider the case in which solitary-search dominance is violated in the search game (without additional monetary bonuses).

Consider first a setup in which the contest designer can only offer a constant bonus, which is independent of the prize's location. A constant bonus can help to increase the relative expected private value of locations with a high prior probability. As a result, it can help obtain the optimal social payoff, when the reason for not having the required properties without the designer's intervention is a low-prior location having a too-high private value. For example, consider a search game with costless search (i.e., $c \equiv 0$), where there are two locations ω, ω' in the same cell of player i with priors $\mu(\omega) = 0.1$ and $\mu(\omega') = 0.2$ and with private values of $v_i^1(\omega) = 5$ and $v_i^1(\omega') = 1$, and $v_i^m \equiv \frac{1}{m} v_i^1$. The too-high private value of location ω violates solitary-search dominance because the expected private value in ω ($0.5 = 0.1 \cdot 5$) is more than twice the expected private value in ω' ($0.2 \cdot 1$). A constant bonus of 1 would restore solitary-search dominance (making the expected private value of ω and ω' to be equal to $0.6 = 0.1 \cdot (5+1)$ and $0.4 = 0.2 \cdot (1+1)$, respectively).

When the designer can offer a location-dependent and player-dependent bonus, it allows him to obtain solitary-search dominance and ordinal consistency when faced with any profile of rewards. An interesting open question is how the designer can maximize the social payoff, while minimizing the expected bonus. For example, assume that the payoffs are ordinally consistent, but they are not solitary-search dominant. Theorem 1 suggests that the designer should boost locations that have lower expected private values (which violate solitary-search dominance). Note that these locations might not coincide with the locations that are not searched by any player in the inefficient equilibrium. This is demonstrated in Example 6.

Example 6. Consider the following search game with common values (as illustrated in Figure 5): $\left(N = \{1,2\}, \Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}, \Pi, \mu, K \equiv 1, c \equiv 0, v_i^m \equiv \frac{1}{m}, v_5 \equiv 1\right)$, where the prior is $\mu(\omega_1) = 44\%$, $\mu(\omega_2) = 21\%$, $\mu(\omega_3) = 20\%$ and $\mu(\omega_4) = 15\%$, player 1 observes whether the prize's location is 1 or not, i.e., $\Pi_1 = \{\{\omega_1\}, \{\omega_2, \omega_3, \omega_4\}\}$, and player 2 observes whether the prize's location is at most 2 or not, i.e., $\Pi_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$. The game admits

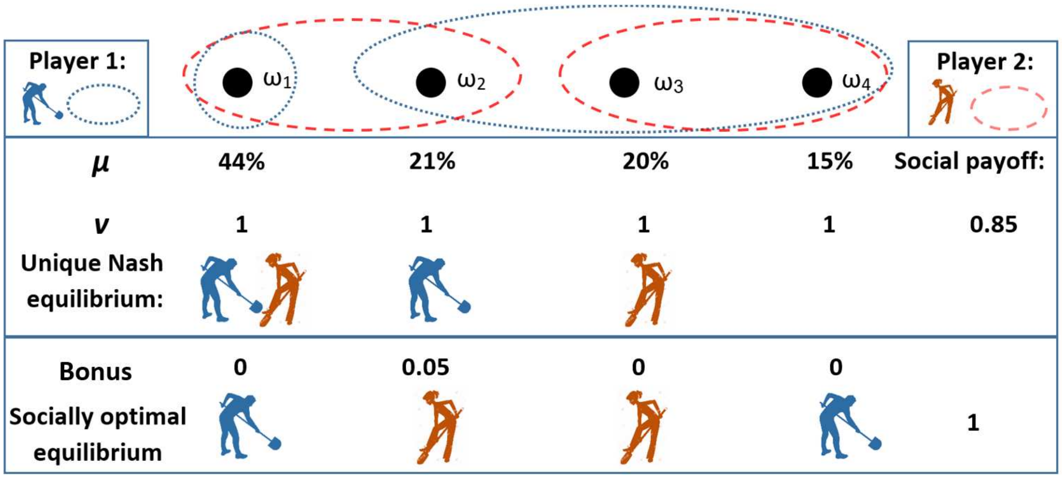

Figure 5: Illustration of Example 6: Impact of Monetary Bonuses on the Social Payoff

a unique equilibrium, in which player 1 searches in ω_1 and ω_2, while player 2 searches in locations ω_1 and ω_3. This equilibrium yields an expected social payoff of 0.85 because no player searches in ω_4. Note that solitary-search dominance is violated because of the low probability of location ω_2 (rather than a low probability of ω_4).

If the designer can offer a bonus of 0.05 that increases the private value in location ω_2 by 5% to 1.05 (which requires a modest expected bonus of $21\% \cdot 0.05 \approx 0.01$), then the modified rewards satisfy solitary-search dominance, and, as a result, the game admits a socially optimal equilibrium with a social payoff of 1 (in which player 1 searches in locations ω_1 and ω_4, while player 2 searches in locations ω_2 and ω_3).

4 Feasible Outcomes and Socially Optimal Payoff

In this section we characterize the socially optimal payoff, namely, the highest social payoff under any strategy profile, not necessarily an equilibrium. Nevertheless, note that by Theorem 1, the socially optimal payoff achieved in every result or example of this section is also yielded by some equilibrium of the game, if the payoffs are ordinally consistent and solitary-search dominant.

To study the socially optimal payoff, we study a slightly broader question: which social outcomes are feasible (where a social outcome is a specification of the locations in which the prize will be found by anyone)? We show that a social outcome is feasible iff it satisfies a condition that is formed in the spirit of Hall's marriage theorem and is an expression of the outcome's compatibility with the information structure. In particular, the socially optimal payoff of the game equals the maximal social payoff of such outcomes.

Pure outcomes A *pure outcome* is a function $f:\Omega\to\{0,1\}$ that specifies, for every location, whether that location is being searched (by anyone) or not. A pure outcome f is *feasible* if there exists a strategy profile $s\in S$ that induces f, i.e., if $f(\omega)=1_{\{m_s(\omega)>0\}}$. Let f_s denote the outcome induced by strategy profile $s\in S$. We may think of an outcome as a possible goal set by society. In a more abstract model than ours, in which society does not maintain exact values but still has (perhaps incomplete) preferences over various outcomes, a social planner would like to know which outcomes are feasible.

We say that a pure outcome is compatible with the information structure if the number of locations being searched within any subset of locations does not exceed the sum of players' capacities over all cells that intersect that subset. Formally:

Definition 5. Fix a search game G. A pure outcome f is compatible with the information structure (abbr., *compatible*) if for each subset $W\subseteq\Omega$, the following inequality holds:

$$\sum_{\omega\in W} f(\omega)\leq\sum_{i\in N} K_i\cdot\sum_{\pi_i\in\Pi_i} 1_{\pi_i\cap W\neq\emptyset}\,. \tag{3}$$

Compatibility is clearly necessary for an outcome to be feasible in a setup in which players cannot share their information, and each player decides where to search as a function of her own signal. In such a setup, each player i has $\sum_{\pi_i\in\Pi_i}1_{\pi_i\cap W\neq\emptyset}$ cells that intersect the set W, and thus she cannot search in more than $K_i\cdot\sum_{\pi_i\in\Pi_i}1_{\pi_i\cap W\neq\emptyset}$ locations within W. This implies that all players combined cannot search in more than $\sum_{i\in N}K_i\cdot\sum_{\pi_i\in\Pi_i}1_{\pi_i\cap W\neq\emptyset}$ locations within W. By representing the setup as a bipartite graph and applying Hall's marriage theorem (Hall, 1935), it follows that compatibility is also a sufficient condition for feasibility.

Proposition 2. *A pure outcome f in a search game is feasible iff it is compatible.*

Sketch of proof; formal proof is omitted because it is implied by Theorem 2. For simplicity, assume that each player has a capacity of one. Consider a bipartite undirected graph in which the left side of the graph includes the players' cells in the search game, and the right side includes the locations for which f is equal to one (as illustrated in Figure 6). The graph's edges connect each cell to the locations that are contained in that cell. A matching of all locations in this graph, i.e., a set of disjoint edges (namely, no node appears twice) such that every location belongs to some edge, corresponds to a strategy profile that induces f. Hall's theorem states that such a matching exists iff for any subset of locations W, the number of its neighbors $|N(W)|$ is at least $|W|$. The neighbors of W in this graph are the cells that intersect W; therefore, this condition is equivalent to f being compatible. $\square$

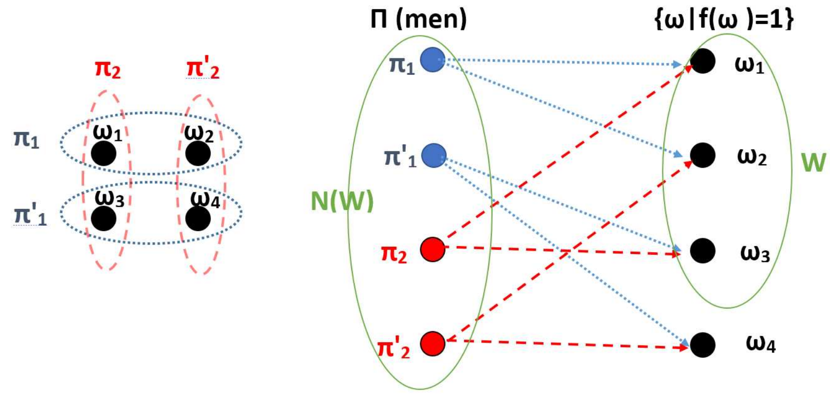

Figure 6: **Illustration of Proposition 2**. The LHS of the figure demonstrates the information partitions in a two-player search game (with capacities equal to 1). The RHS translates this into a bipartite graph, where its left part ("men") includes the cells of all players, and its right part ("women") includes all locations ω satisfying $f(\omega) = 1$. The figure further shows an example of a subset of locations W and the corresponding set of its neighbors - $N(W)$.

The following example and corollary apply Proposition 2 to obtain a simple sufficient condition for the existence of exhaustive strategy profiles, in terms of the size of the largest cell of each player.

Example 7. Suppose that there are three players, each has capacity $K_i = 1$, and every cell of every player contains exactly three locations. Consider the pure outcome $f(\omega) = 1$ for every ω. To see that f is compatible, let $W \subset \Omega$ be a subset of locations. The number of cells π_i of player i that intersect W is at least $\frac{|W|}{3}$, because the size of every cell is 3. Therefore,

$$\sum_{i \in N} \sum_{\pi_i \in \Pi_i} 1_{\pi_i \cap W \neq \emptyset} \geq 3 \cdot \frac{|W|}{3} = |W| = \sum_{\omega \in W} f(\omega),$$

i.e., f is compatible. By Proposition 2, f is feasible, i.e., this game admits an exhaustive strategy profile, and the socially optimal payoff equals $\sum_{\omega \in \Omega} \mu_{\mathfrak{s}}(\omega) \cdot v_{\mathfrak{s}}(\omega)$.

More generally, the argument in Example 7 implies the following corollary.

Corollary 3. *Let G be a search game, and let $M_i = \max(|\pi_i| : \pi_i \in \Pi_i)$ be the size of the largest cell of player i. If $\sum_{i \in N} K_i / M_i \geq 1$ then G admits an exhaustive strategy profile.*

Mixed outcomes A *mixed outcome* is a function $f : \Omega \to [0, 1]$ that assigns a probability to each location. We interpret $f(\omega)$ as the probability that the prize is found, conditional on the prize's location being ω. A mixed outcome may be a goal set by society, perhaps involving such considerations as fairness, equal opportunity, etc.

A *correlated strategy profile* $\sigma \in \Delta(S)$ is a lottery over the set of pure strategy profiles. A mixed outcome is feasible if it can be induced by a correlated strategy profile. That is, $f : \Omega \to [0, 1]$ is *feasible* if there exists a correlated strategy profile $\sigma \in \Delta(S)$ such that

$$f(\omega) = \sum_{s \in S} \sigma(s) \cdot f_s(\omega) = \sum_{s \in S} \sigma(s) \cdot 1_{\{m_s(\omega) > 0\}}.$$

Let f_σ denote the mixed outcome induced by correlated strategy profile $\sigma \in \Delta(S)$.

A mixed outcome is compatible with the information structure if the sum of the probabilities of finding the prize (henceforth, finding probabilities) of any subset of locations does not exceed the sum of players' capacities over all cells that intersect that subset. Formally:

Definition 6. Fix a search game G. A mixed outcome f is compatible with the information structure (abbr., *compatible*) if for each subset of locations $W \subseteq \Omega$, the following inequality holds:

$$\sum_{\omega \in W} f(\omega) \leq \sum_{i \in N} K_i \cdot \sum_{\pi_i \in \Pi_i} 1_{\pi_i \cap W \neq \emptyset}\ . \tag{4}$$

Let F_C be the set of compatible mixed outcomes. Clearly, compatibility is necessary for a mixed outcome to be feasible, because any feasible mixed outcome lies in the convex hull of feasible pure outcomes, all of which satisfy compatibility. In what follows we show that the converse is also true, i.e., that compatibility is sufficient for a mixed outcome to be feasible.

Theorem 2. *A mixed outcome f in a search game is feasible iff it is compatible.*

Note that we cannot directly use Hall's theorem for this result, due to the outcome being mixed, rather than pure (as Hall's theorem applies only to a "binary" matching of zeros and ones). Instead, the proof (presented in Section 4.1) includes two parts: (1) we introduce the notion of coordinated-search profiles, and apply (Prop. 3) the max-flow min-cut theorem to show that any compatible mixed outcome can be induced by a coordinated-search profile, and (2) we apply (Prop. 4) the Birkhoff–von Neumann theorem to show that coordinated-search profiles induce feasible mixed outcomes.

Before presenting the next example, we define a strategy profile *s as redundancy-free* if (1) every player always uses her entire capacity (i.e., $|s_i(\pi_i)| = K_i$ for every cell π_i of every player i), and (2) there is no search duplication (i.e., $m_s(\omega) \leq 1$ for every $\omega \in \Omega$). Since a player can search in no more than $K_i \cdot |\Pi_i|$ locations, a strategy profile is redundancy-free iff the number of locations being searched equals $\sum_{i \in N} K_i |\Pi_i|$. If a game admits redundancy-free strategy profiles, then they are exactly the location-maximizing profiles.

The following example and corollary apply Theorem 2 to obtain a simple sufficient condition for the existence of redundancy-free strategy profiles, in terms of the size of the smallest cell of each player.

Example 8. Suppose that there are three players, each has capacity $K_i = 1$, and every cell of every player contains exactly five locations. Consider the mixed outcome $f(\omega) = 0.6$ for every ω, and let $W \subset \Omega$ be a subset of locations. The number of cells π_i of player i that intersect W is at least $\frac{|W|}{5}$; therefore,

$$\sum_{i \in N} \sum_{\pi_i \in \Pi_i} 1_{\pi_i \cap W \neq \emptyset} \geq 3 \cdot \frac{|W|}{5} = 0.6 \cdot |W| = \sum_{\omega \in W} f(\omega),$$

i.e., f is compatible. By Theorem 2, f can be induced by a correlated strategy profile. Since $\sum_{\omega \in \Omega} f(\omega) = 0.6 \cdot |\Omega| = \sum_{i \in N} K_i |\Pi_i|$, this implies that the game admits a redundancy-free strategy profile.

More generally, the argument in Example 8 implies the following corollary.

Corollary 4. *Let G be a search game, and let $m_i = \min(|\pi_i| : \pi_i \in \Pi_i)$ be the size the smallest cell of player i. If $\sum_{i \in N} K_i/m_i \leq 1$ then G admits a redundancy-free profile.*

If a correlated strategy σ achieves some level of social payoff, then at least one of the pure strategies in the support of σ yields at least that much. Therefore, Theorem 2 implies a characterization of the socially optimal payoff of search games: the socially optimal payoff is the highest payoff induced by a compatible mixed outcome. Formally:

Corollary 5. *Let G be a search game. Then $U_{opt} = \max_{f \in F_C} \sum_{\omega \in \Omega} f(\omega) \mu(\omega) v_{\mathfrak{s}}(\omega)$.*

4.1 Coordinated Search

In this subsection we consider a variant of our model, and show that it does not increase the socially optimal payoff. The analysis turns out to be closely related to Theorem 2, and is helpful in deriving its proof.

The setup Coordinated search allows players to coordinate partial search efforts within a location. Specifically, we now allow each player to divide fractions of her search capacity among the different locations. This is formalized as follows. Fix a search game G. For any $k \in \mathbb{N}$ and any cell π, let $\mathcal{D}(\pi, k)$ denote the set of all functions $\eta : \pi \to [0,1]$ that satisfy $\sum_{\omega \in \pi} \eta(\omega) \leq k$. That is, an element of $\mathcal{D}(\pi, k)$ is a function that assigns a search effort to each location in π such that the total effort is at most k. A *coordinated-search profile* is a tuple $\tau = (\tau_1, ..., \tau_n)$, where each function τ_i assigns to each cell $\pi_i \in \Pi_i$ an element of $\mathcal{D}(\pi_i, K_i)$. We interpret $\tau_i(\pi_i, \omega) \equiv \tau_i(\pi_i)(\omega)$ as the (fractional) search effort player i exerts in location $\omega \in \pi_i$ (when the player observes the signal π_i). Let T be the set of all coordinated-search profiles. Observe that any (pure) strategy profile in G is a coordinated-search profile (i.e., $S \subseteq T$, where an element of $\mathcal{D}(\pi, k)$ that assigns only search efforts of zeros and ones is identified with the corresponding subset of π with at most k elements).

Importantly, we assume that fractional search efforts of different players are summed optimally from society's point of view. For example, if player i exerts an effort of 50% in location ω in cell $\pi_i(\omega)$ and player j exerts an effort of 40% in location ω in cell $\pi_j(\omega)$, then the prize is found with a total probability of 90%, conditional on the location being[10] ω.

Thus, any coordinated-search profile induces a mixed outcome, where the finding probability assigned to each location ω is the sum of the fractional search efforts exerted by each player i in location ω in cell $\pi_i(\omega)$ (bounded by the maximal finding probability of one). Formally, the mixed outcome f_τ induced by the coordinated-search profile τ is defined by $f_\tau(\omega) = \min\left(\sum_{i \in N} \tau_i(\pi_i(\omega), \omega), 1\right)$. Hence, the social payoff $U^c(\tau)$ induced by a coordinated-search profile τ is

$$U^c(\tau) = \sum_{\omega \in \Omega} \left(\min\left(\sum_{i \in N} \tau_i(\pi_i(\omega), \omega), 1 \right) \right) \cdot \mu(\omega) \cdot v_{\mathfrak{s}}(\omega).$$

First, we observe that any coordinated-search profile induces a compatible mixed outcome.

Claim 1. Fix search game G and $\tau \in T$. Then f_τ is a compatible mixed outcome.

Proof. Fix a subset of locations $W \subseteq \Omega$. Then the following inequality holds (where the last inequality is implied by $\sum_{\omega \in \pi_i} \tau_i(\pi_i, \omega) \leq K_i$):

$$\sum_{\omega \in W} f_\tau(\omega) = \sum_{\omega \in W} \min\left(\sum_i \tau_i(\pi_i(\omega), \omega), 1 \right) \leq \sum_{i \in N} \sum_{\omega \in W} \tau_i(\pi_i(\omega), \omega) \leq \sum_{i \in N} K_i \sum_{\pi_i \in \Pi_i} 1_{\pi_i \cap W \neq \emptyset}. \square$$

Next we employ the max-flow min-cut theorem to show that the converse is true as well: any compatible mixed outcome can be induced by a coordinated-search profile.

Proposition 3. *Fix a search game G and a compatible mixed outcome f. Then there exists a coordinated-search profile τ that induces f (i.e., $f = f_\tau$).*

Sketch of proof; formal proof in Appendix B.2. We construct a *flow network*: a directed graph whose edges have *flow capacities*. The graph connects every cell to the locations contained in it, with infinite flow capacity (as illustrated in Figure 7). We add a source vertex that connects to every cell, with flow capacity K_i, and a sink vertex to which every location ω is

[10]This probability is strictly higher than if it were a mixed strategy profile, where with positive probability ($20\% = 40\% \cdot 50\%$) there would be search duplication (where both players search in ω), and the total probability of finding the prize, conditional on the location being ω, would be strictly less than 90%.

connected, with flow capacity $f(\omega)$. A *cut* is a subset of edges without which there exists no path from the source to the sink. The compatibility of f implies that the minimal cut has a total capacity of $\sum_{\omega\in\Omega} f(\omega)$. Therefore, by the max-flow min-cut theorem (Ford & Fulkerson, 1956; see a textbook presentation in Cormen *et al.*, 2009, p. 723, Thm. 26.6), the network admits a flow of $\sum_{\omega\in\Omega} f(\omega)$. We define τ by letting $\tau_i(\pi_i,\omega)$ equal the flow from π_i to ω. □

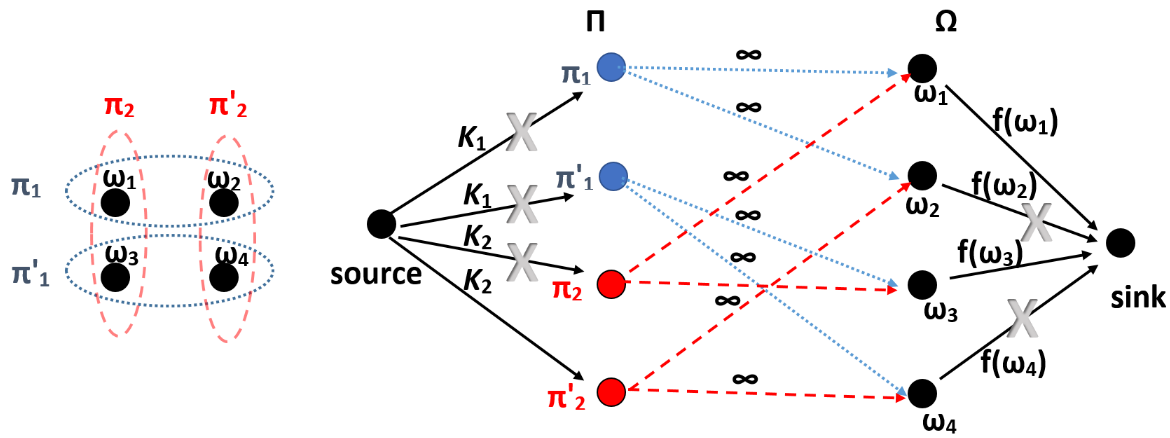

Figure 7: **Illustration of Proposition 3**. The left side of the figure demonstrates partitions in a two-player search game. The right side demonstrates the constructed directed graph in which (1) a source node is linked to every player's cells by an edge with the player's capacity, and (2) each cell is linked by unlimited edges to all the locations within that cell, and (3) each location ω is linked to a sink node by an edge with capacity $f(\omega)$. The gray X-s demonstrate an example of a cut, i.e., a subset of edges whose removal from the graph disconnects the source from the sink.

Coordination does not add new outcomes The social payoff is constrained by the fact that the players are not allowed to share their private signals. This constraint is captured by the compatibility condition presented above. The inability, in our main model (as opposed to the coordinated-search variant), to efficiently coordinate players' fractional search efforts within a location is another potential constraint on the social payoff. Nevertheless, in what follows we show that this constraint does not limit the social payoff. Specifically, we apply the Birkhoff–von Neumann theorem (Birkhoff, 1946; Von Neumann, 1953; see Berman & Plemmons, 1994, p. 50, for a textbook presentation) to show that any mixed outcome that can be induced by a coordinated-search profile is feasible (i.e., it can be induced by a correlated strategy profile with no coordination of fractional efforts).

Proposition 4. *Fix search game G and a coordinated-search profile $\tau \in T$. Then there exists a correlated strategy profile $\sigma \in \Delta(S)$ such that $f_\tau = f_\sigma$.*

Sketch of proof (see Appendix B.3 for the formal proof). To simplify the sketch of the proof assume that all capacities are equal to one. Let τ be a coordinated-search profile. We can represent the profile τ as a matrix $(A^{\tau}_{\pi\omega})_{\omega\in\Omega,\pi\,\text{is a cell}}$, where

$$A^{\tau}_{\pi\omega} = \begin{cases} \tau_i(\pi,\omega) & \omega\in\pi,\,\pi\in\Pi_i \\ 0 & \text{elsewhere}. \end{cases}$$

Observe that $A^{\tau}_{\pi\omega}$ is a nonnegative matrix, and that the sum of each row π is at most one. Let $B^{\tau}_{\pi\omega}$ be a matrix derived from $A^{\tau}_{\pi\omega}$ by decreasing elements of the matrix such that the sum of each column ω that exceeded one in $A^{\tau}_{\pi\omega}$ is equal to one in $B^{\tau}_{\pi\omega}$. Observe that $B^{\tau}_{\pi\omega}$ is a doubly substochastic matrix; i.e., it is a nonnegative matrix for which the sum of each column and each row is at most one. A simple adaptation of the Birkhoff–von Neumann theorem shows that $B^{\tau}_{\pi\omega}$ can be represented as a convex combination of matrices $C^{1}_{\pi\omega},...,C^{K}_{\pi\omega}$ (i.e., $B^{\tau}_{\pi\omega}=\sum w_k\cdot C^{k}_{\pi\omega}$ where $\sum w_k=1$ and $w_k\geq 0$), where each matrix $C^{k}_{\pi\omega}$: (1) contains only zeros and ones, and (2) contains in each row and in each column at most a single value of one. Observe that each such matrix $C^{k}_{\pi\omega}$ corresponds to a pure profile s^k in the search game, and that the outcome f_τ is a weighted sum of the outcomes induced by the profiles s^k. This implies that f_τ is feasible because it is induced by the correlated strategy profile $\sigma=\sum w_k\cdot s^k$. □

In the following example, Proposition 4 is used to prove that the game admits an exhaustive strategy profile.

Example 9. Let the set of locations $\Omega = A\uplus B_1\uplus B_2\uplus B_3$ be a union of four disjoint sets of equal size. There are three players, each with capacity $K_i=1$. The partition Π_i of player i consists of cells of size two $\{a,b\}$, where $a\in A$ and $b\in B_i$, and of cells of size six, whose members come from $B_j\cup B_k$ $(j,k\neq i)$. The partitions are illustrated in Figure 8. Define a coordinated-search profile τ as follows. For $\pi_i=\{a,b\}$, $\tau_i(\pi_i,a)={}^1\!/\!_3$ and $\tau_i(\pi_i,b)={}^2\!/\!_3$, and for π_i of size six, τ_i assigns ${}^1\!/\!_6$ to every location in π_i. Thus, for any $a\in A$, $\sum_{i\in N}\tau_i(\pi_i(a),a)={}^1\!/\!_3+{}^1\!/\!_3+{}^1\!/\!_3=1$, for any $b_1\in B_1$, $\sum_{i\in N}\tau_i(\pi_i(b_1),b_1)={}^2\!/\!_3+{}^1\!/\!_6+{}^1\!/\!_6=1$, and similarly for B_2 and B_3; therefore, $f_\tau(\omega)=1$ for every $\omega\in\Omega$. Proposition 4 implies that the game admits an exhaustive strategy profile.

Note that $|\Pi_i|=|\Omega|/3$, implying that $|\Omega|=\sum_{i\in N}K_i|\Pi_i|$; therefore, a strategy profile in this game is exhaustive iff it is redundancy-free. Redundancy-freeness can, alternatively, be deduced from the fact that under τ players always use their entire capacity and the sum of fractional search efforts $\sum_{i\in N}\tau_i(\pi_i(\omega),\omega)$ does not exceed one in any location ω.

Proposition 3 and Proposition 4 jointly imply Theorem 2. Moreover, Proposition 4

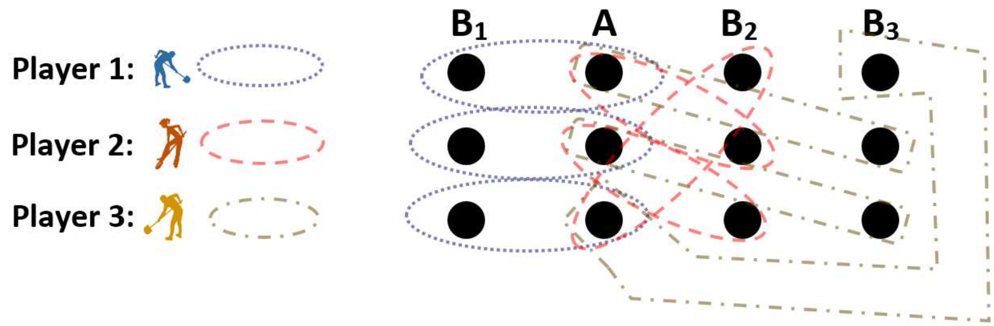

Figure 8: Illustration of Example 9 with $|\Omega| = 12$ locations. The figure shows the players' partitions (where the cell containing the remaining six locations of each player is not drawn to make the figure less crowded).

implies that if some level of social payoff is yielded by a coordinated-search profile, then the same or higher social payoff can be yielded by a pure strategy profile. That is, the ability to coordinate search efforts does not improve the social payoff. Formally:

Corollary 6. *Fix a search game G and a coordinated-search profile $\tau \in T$. Then, there exists a pure strategy profile $s \in S$ such that $U(s) \geq U^c(\tau)$.*

Proof. Proposition 4 implies that $f_\tau = \sum w_k \cdot f_{s_k}$, where $\sum w_k = 1$, $w_k \geq 0$, and $s_k \in S$ for each k. This implies that $U^c(\tau) = \sum w_k \cdot U(s_k)$, which, in turn, implies that $U^c(\tau) \leq U(s_k)$ for some k. □

In Appendix A we apply our results to search games in which the intersection of every profile of cells includes at least κ locations, and derive tight conditions for the existence of equilibria with appealing properties.

5 General Signals

In this section we present a general model of signals, dropping the assumption that every location corresponds to a single state of the world. We also discuss a weaker assumption, under which all our results still hold.

5.1 Adaptations to the Model

To extend our model to the general case, we let the set of locations and the set of states be different objects. Let Ω denote the set of states of the world (abbr., *states*). A state determines the location of the prize, and we let $\ell(\omega)$ denote the prize's location when the state of the world is ω. Thus, the different locations induce a partition of Ω, and without loss of generality we let L, the set of locations, be that partition. That is, a location $\ell \in L$ is an element of the partition, namely, a subset of states such that $\ell(\omega)$ is the same for every ω in the subset.

A *generalized search game* is a tuple $\tilde{G} = (N, \Omega, \Pi, \mu, L, K, c, v)$, where L is the partition of locations. All other components are as defined in the baseline model. To prevent confusion we use the term *simple search game* to refer to the search games of the baseline model. For a cell $\pi_i \in \Pi_i$, let $\ell(\pi_i)$ be the set of locations that are ordinally consistent with the state being in π_i, i.e., $\ell(\pi_i) = \{\ell(\omega) \,|\, \omega \in \pi_i\}$.

A *strategy* of player i is a function s_i that assigns to each cell π_i a subset of locations with at most K_i elements that satisfies $s_i(\pi_i) \subseteq \ell(\pi_i)$. We interpret $s_i(\pi_i)$ as the set of up to K_i locations in which the player searches when she observes the signal π_i. When the state of the world is ω, player i finds the prize if she searches in the location $\ell(\omega)$, i.e., if $\ell(\omega) \in s_i(\pi_i(\omega))$. We redefine the number of players who search in the prize's location when the state is ω as follows: $m_s(\omega) = \sum_{i \in N} \mathbf{1}_{\ell(\omega) \in s_i(\pi_i(\omega))}$.

A player's payoff, conditional on the state being ω, is then redefined by

$$u_i(s|\omega) = 1_{\ell(\omega) \in s_i(\pi_i(\omega))} v_i^{m_s(\omega)}(\omega) - c_i(|s_i(\pi_i(\omega))|).$$

Ordinal consistency is redefined as follows. A search game has *ordinally consistent payoffs* if for any two locations ℓ and ℓ', if the expected (interim) private value of player i is strictly lower in ℓ than in ℓ', then the expected (interim) social value is weakly lower in ℓ.

Definition 7. Generalized Search game $\tilde{G}$ has *ordinally consistent payoffs* if for any player i, any cell $\pi_i \in \Pi_i$, and any two locations $\ell, \ell' \in L$, the following implication holds:

$$\sum_{\omega \in \ell \cap \pi_i} \mu(\omega) \cdot v_i^1(\omega) < \sum_{\omega' \in \ell' \cap \pi_i} \mu\left(\omega'\right) \cdot v_i^1\left(\omega'\right) \Rightarrow$$

$$\sum_{\omega \in \ell \cap \pi_i} \mu(\omega) \cdot v_{\mathfrak{s}}(\omega) \leq \sum_{\omega' \in \ell' \cap \pi_i} \mu\left(\omega'\right) \cdot v_{\mathfrak{s}}\left(\omega'\right).$$

Similarly, solitary-search dominance is redefined as follows.

Definition 8. Generalized search game $\tilde{G}$ has *solitary-search dominant payoffs* if for any player i, cell $\pi_i \in \Pi_i$, and pair of locations $\ell, \ell' \in L$ such that $\ell \cap \pi_i \neq \emptyset$, the following inequalities hold:

$$\sum_{\omega \in \ell \cap \pi_i} \mu(\omega|\pi_i) \cdot v_i^1(\omega) \geq \sum_{\omega' \in \ell' \cap \pi_i} \mu\left(\omega'|\pi_i\right) \cdot v_i^2\left(\omega'\right), \text{ and}$$

$$\sum_{\omega \in \ell \cap \pi_i} \mu(\omega|\pi_i) \cdot v_i^1(\omega) \geq c_i(K_i) - c_i(K_i - 1).$$

All other parts of the baseline model remain the same.

5.2 Equivalence Result with Weakly Deterministic Signals

Our results about simple search games hold in the current setup if we assume that the signal of any player is determined by the prize's location and the signal of another player. Formally:

Definition 9. Generalized search game $\tilde{G}$ has *weakly deterministic signals* if $\ell(\omega) = \ell(\omega')$ and $\pi_i(\omega) = \pi_i(\omega') \Rightarrow \pi_j(\omega) = \pi_j(\omega')$, for any $\omega, \omega' \in \Omega$ and players i, j.

In words, if two states ω, ω' share the same location and are indistinguishable to one of the players, then these states must be indistinguishable to all players.

The set of generalized search games with weakly deterministic signals is broader than the set of simple search games. In particular: (1) generalized search games allow a player to hold different positive posterior beliefs on the prize being in location ℓ depending on the different signals that she may observe (which is impossible in simple search games), and (2) any generalized search game in which players have the same information partitions (i.e., $\Pi_i = \Pi_j$ for any pair of players) has weakly deterministic signals. Nevertheless, we show that any generalized search game with weakly deterministic signals is strategically equivalent to a simple game (and, thus, our results hold for the broader set of games).

A generalized search game is equivalent to a simple search game if there exists a bijection between the sets of strategies of each game that preserves the expected payoff of all players.

Definition 10. Simple search game G and generalized search game $\tilde{G}$ with the same set of players and the same capacities are *equivalent* if there exist bijections $f_i : S_i\left(\tilde{G}\right) \to S_i(G)$, such that $\tilde{u}_i(\tilde{s}) = u_i(f(\tilde{s}))$ and $\tilde{U}(\tilde{s}) = U(f(\tilde{s}))$ for every strategy profile $\tilde{s} \in S\left(\tilde{G}\right)$ and every player i (where $f = (f_1, \ldots, f_n)$).

It is immediate that equivalent games have equivalent sets of Nash equilibria; i.e., $\tilde{s}$ is a Nash equilibrium of $\tilde{G}$ iff $f(\tilde{s})$ is a Nash equilibrium of G, and both equilibria yield the same payoffs to all players and to society. Next we show that any generalized search game with weakly deterministic signals is equivalent to a simple game. Formally:

Claim 2. Let $\tilde{G}$ be a generalized search game with weakly deterministic signals. Then there exists an equivalent simple search game G. Moreover, if $\tilde{G}$ is ordinally consistent or solitary-search dominant, then so is G.

Sketch of proof; formal proof is omitted for brevity. We say that two states in $\tilde{G}$ are equivalent if they have the same location and no player can distinguish between the two states (i.e., the states are elements of the same cell, for any player). We construct the equivalent simple game G by the following two steps: (1) merge equivalent states into a single state, and (2) extend the set of locations, such that each location corresponds to a single (possibly merged) state. The reward in each merged state is defined as the weighted average of the rewards in the corresponding equivalent states. The prior of each merged state is defined as the sum of the priors of the corresponding states. The equivalence of the two games, and the invariance of the ordinal consistency and solitary-search dominance of the payoffs, are straightforward. The simple process of constructing the equivalent simple game is demonstrated in Figure 9. □

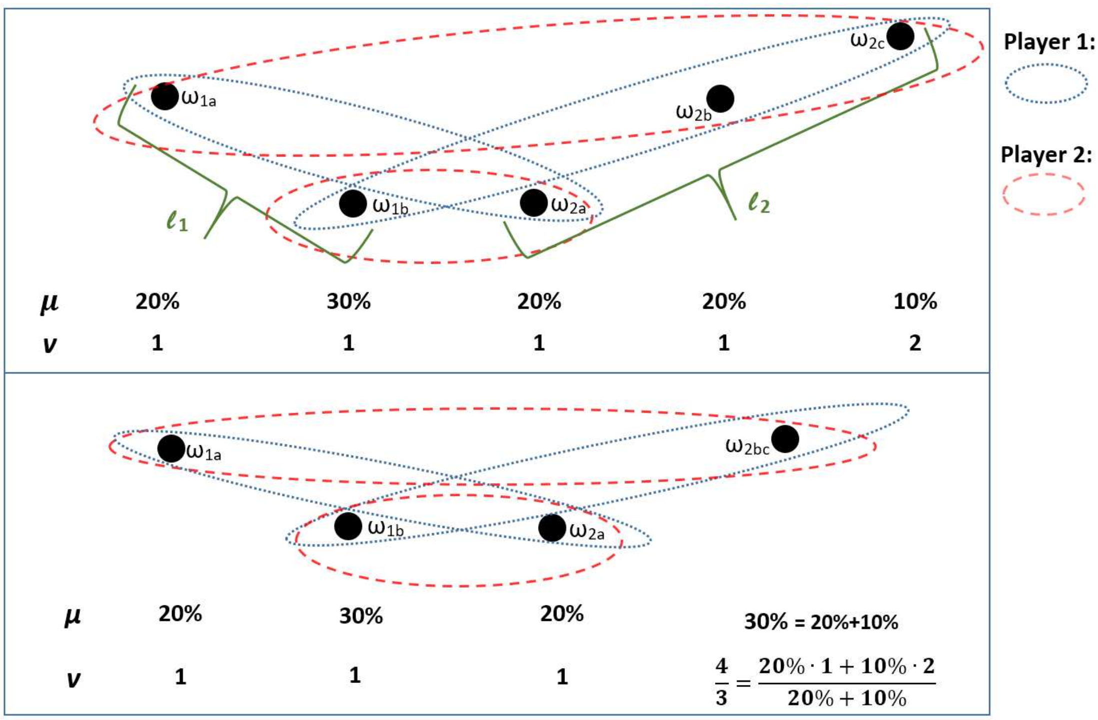

Figure 9: **Illustration of Claim 2.** The upper panel presents a generalized two-player search game $\tilde{G}$ with weakly deterministic signals. The lower panel describes the equivalent simple game G in which (1) states ω_{2b} and ω_{2c} are merged to ω_{2bc}, and (2) each location has been divided into singletons.

5.3 Non-monotone Value of Information

Refining the players' partitions always weakly increases the socially optimal payoff. Therefore, if a refinement maintains the properties of weakly deterministic signals, ordinal consistency, and solitary-search dominance, then the maximal social payoff yielded by an equilibrium also increases, by Theorem 1 and Claim 2. However, even if we restrict ourselves to games with weakly deterministic signals, a refinement may still break ordinal consistency or solitary-search dominance (by contrast, this cannot happen in simple search games). Consequently, the value of information may be negative, in the sense that the refinement decreases the equilibrium social payoff.

The following example exhibits a refinement with a negative value, and a further refinement with a positive one. More specifically, it presents a game that admits an exhaustive equilibrium if the players' signal is either uninformative or fully informative, but not if it is partially informative.

Example 10. Let $(N=\{1,2\},\Omega=\{\omega_1,\omega_2,\omega_3,\omega_4\},\Pi,\mu,L=\{\{\omega_1,\omega_2\},\{\omega_3,\omega_4\}\},K\equiv 1,c\equiv 0,v_i^m\equiv \frac{1}{m},v_{\mathfrak{s}}\equiv 1)$ be a generalized search game, where states $\ell_{12}\equiv\{\omega_1,\omega_2\}$ comprise one location, and states $\ell_{34}=\{\omega_3,\omega_4\}$ comprise another location. The prior belief μ is: $\mu(\omega_1)=\mu(\omega_3)=0.5-\epsilon$ and $\mu(\omega_2)=\mu(\omega_4)=\epsilon$, and assume that ϵ is small, say $\epsilon=10\%$, as illustrated in Figure 10.

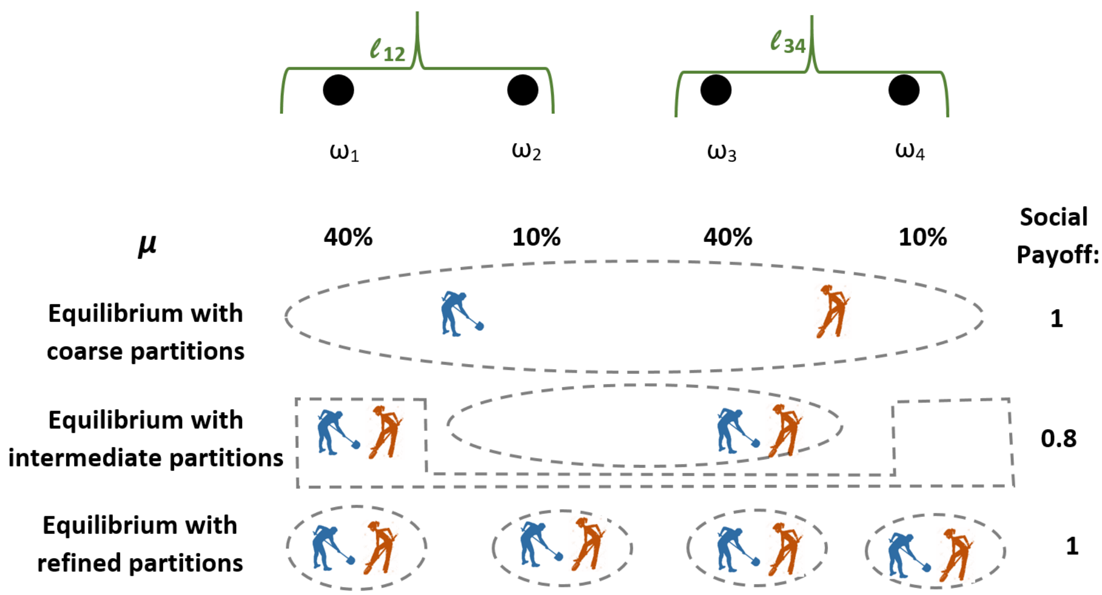

Figure 10: **Illustration of Example 10**. Non-monotone value of information in a symmetric two-player generalized search game.

With the trivial partitions ($\Pi\equiv\{\Omega\}$), all (pure) Nash equilibria are exhaustive (the prize is always found), and are characterized by one player searching in location ℓ_{12} and the other player searching in location ℓ_{34}. Similarly, with the full-information partitions ($\Pi\equiv\{\{\omega_1\},\{\omega_2\},\{\omega_3\},\{\omega_4\}\}$), each player knows the prize's location, and the unique equilibrium is exhaustive (each player searches in the true prize's location). Finally, consider the case of symmetric partially-informative signals induced by the symmetric partitions $\Pi\equiv\{\{\omega_1,\omega_4\},\{\omega_2,\omega_3\}\}$. Observe that in this case (in which solitary-search dominance is violated) the unique equilibrium is both players searching in ℓ_{34} (resp., ℓ_{12}) after observing the signal $\{\omega_2,\omega_3\}$ (resp., $\{\omega_1,\omega_4\}$), and, therefore, the prize is not found when the state of the world is either ω_2 or ω_4.

5.4 Counterexamples without Weakly Deterministic Signals

Figure 11 demonstrates a failure of each of our main results without the assumption of weakly deterministic signals: the left panel demonstrates the failure of Theorem 1, and the right panel demonstrates the failure of Prop. 4 (and thus of Theorem 2).

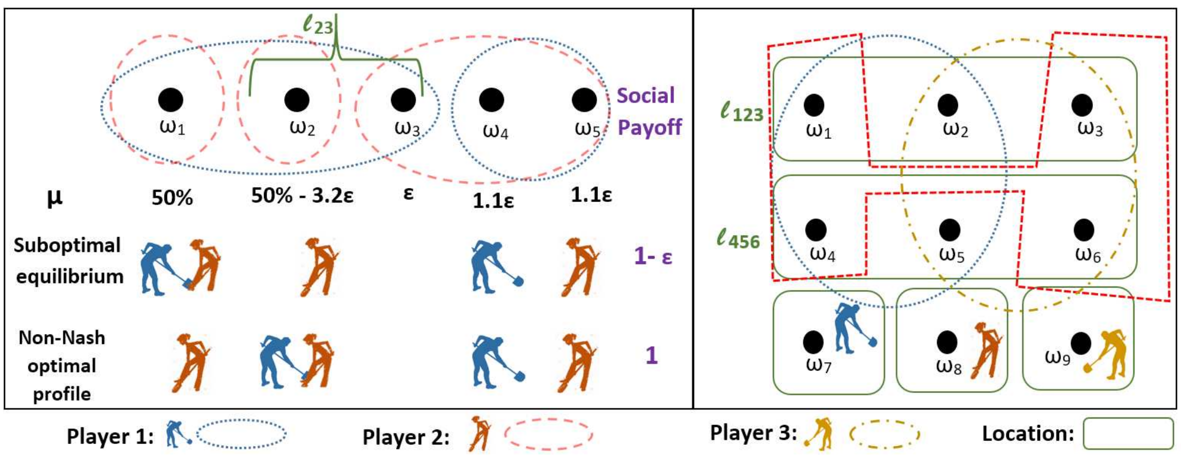

Figure 11: **The left panel presents a counterexample to Theorem 1** (without weakly deterministic signals). It shows the essentially unique suboptimal equilibrium and the non-Nash socially optimal profile in a generalized two-player search game with ordinally consistent and solitary-search dominant payoffs, $K \equiv 1$, $c \equiv 0$, $v_i^m \equiv \frac{1}{m}$, $v_{\mathfrak{s}} \equiv 1$.
The right panel presents a counterexample to Proposition 4. It shows a three-player generalized search game with $K \equiv 1$ in which coordinated search allows the players to always find the prize, whereas this is not possible without coordination. Each player's partition has two cells: one with 4 states (which is drawn in the figure), and another with the remaining 5 states (which is not drawn, to make the figure less crowded). The figure shows a coordinated-search profile that always finds the prize: each player i divides her search effort equally between locations ℓ_{123} and ℓ_{456} in her four-state cell, and exerts all of her effort to location $\{\omega_{i+6}\}$ in the other cell. By contrast, in any pure strategy profile in which ω_7, ω_8, and ω_9 are all searched by some player, either ℓ_{123} or ℓ_{456} is not searched by any player in at least one state.

6 Conclusion

Our paper studies search games in which agents explore different routes to making a discovery that would benefit both society and the discoverer (although the private gain may differ from the social gain). Our main departure from the related literature is that we introduce asymmetric information to this setup. That is, we allow each agent to have private information about the plausibility of different routes, while almost all of the existing literature assumes that all agents have the same information. We believe that this is a natural development, as asymmetric information is a key component in many real-life decentralized research situations. In addition, we allow substantial heterogeneity between the different routes (i.e., different expected values of finding the prize in different locations). We also allow heterogeneity in the rewards and costs of different players.

Our model is simplified in one aspect, as we assume that the search is a one-shot game, while the dynamic aspects of the search interaction are a key component in many of the existing models (see, e.g., Chatterjee & Evans, 2004; Akcigit & Liu, 2015; Bryan & Lemus, 2017). While a one-shot game might model reasonably well situations with severe time constraints, such as the motivating example of developing a COVID-19 vaccine as soon as possible, we think that incorporating asymmetric information in a dynamic search game is an important avenue for future research.

Our first main result (Theorem 1) states that a search game admits a (pure) equilibrium that yields the first-best social payoff if for any two locations within a player's cell: (1) the player and society have the same ordinal ranking over these two locations (ordinal consistency), and (2) the player always prefers searching in one of these locations alone to searching in the other location with other players, or to not searching at all (solitary-search dominance). Taylor (1995); Fullerton & McAfee (1999); Che & Gale (2003); Koh (2017) present setups of innovation contests in which it is socially optimal to restrict the number of participating players, because adding a player decreases the incentive of others to exert costly effort. By contrast, Theorem 1 implies that adding players to a search game with ordinally consistent and solitary-search dominant payoffs always improves the maximal social payoff that can be yielded by an equilibrium. This is so because the first-best social payoff is (weakly) increasing when players are added. Thus, when the payoffs are ordinally consistent and solitary-search dominant, it is socially optimal to allow all players to participate. It is an open question whether this property holds in our setup when we relax the assumptions of ordinal consistency and solitary-search dominance.

Our second main result shows that the first-best payoff is constrained only by compatibility with the information structure. Any outcome in which the number of locations being searched within any subset of locations does not exceed the sum of players' capacities over all cells that intersect that subset, can be induced by a mixture of pure strategy profiles. A surprising implication of this result is that an alternative setup, in which players can coordinate fractional search efforts within a location, does not increase the social payoff.

Appendix

A Application: Games with Intersecting Signals

We say that a search game has *κ-intersecting signals* if the intersection of any profile of cells (one for each player) includes at least κ locations. Formally:

Definition 11. Search game G has *κ-intersecting signals* if for any profile of cells $(\pi_1, .., \pi_n) \in \Pi_1 \times ... \times \Pi_n$ there are at least κ different locations in $\pi_1 \cap ... \cap \pi_n$.

Therefore, κ-intersecting signals have the property that each signal of player i has a positive probability conditional on any profile of signals observed by the other players. Observe that having 1-intersecting signals is substantially weaker than having independent signals (i.e., than requiring that the probability that player i observes a signal is independent of the signals observed by others). The assumption of κ-intersecting signals seems plausible (especially for $\kappa = 1$) in situations like Example 1, if each research lab has a unique expertise that is, in some sense, separate from all the information that can be obtained by the other labs.

Our final result states that search games with κ-intersecting signals and capacities of at most κ have appealing efficiency properties. Namely, such games admit a redundancy-free strategy profile iff $\sum_{i \in N} K_i \cdot |\Pi_i| \leq |\Omega|$, and they admit an exhaustive strategy profile iff $\sum_{i \in N} K_i \cdot |\Pi_i| \geq |\Omega|$. Moreover, this strategy profile is an equilibrium if the payoffs are solitary-search dominant. Formally:

Proposition 5. *Let G be a search game with capacity $K_i \leq \kappa$ for every player $i \in N$ and with κ-intersecting signals (resp., and with solitary-search dominant payoffs). Then G admits*

1. *a redundancy-free strategy profile (resp., equilibrium) iff $\sum_{i \in N} K_i \cdot |\Pi_i| \leq |\Omega|$;*

2. *an exhaustive strategy profile (resp., equilibrium) iff $\sum_{i \in N} K_i \cdot |\Pi_i| \geq |\Omega|$.*

Sketch of proof; see Appendix B.4 for the formal proof. Let M be the set of players whose partitions are not trivial. Consider a smaller auxiliary search game created by omitting all players in $N \setminus M$. Since the signals are κ-intersecting, each cell of each player must contain at least $\kappa \cdot 2^{|M|-1}$ locations. Since this number is at least $\kappa \cdot |M|$, Corollary 4 implies that the smaller game admits a redundancy-free strategy profile. If $\sum_{i \in N} K_i \cdot |\Pi_i| \leq |\Omega|$, then we can let the remaining players (with trivial partitions) choose one by one a location that has not been chosen by other players yet. The resulting strategy profile is redundancy-free in G. If $\sum_{i \in N} K_i \cdot |\Pi_i| \geq |\Omega|$, then there are sufficiently many remaining players (with trivial partitions) to cover all locations, and hence the resulting strategy profile is exhaustive. □

B Formal Proofs

B.1 Proof of Proposition 1 (Search Games are Weakly Acyclic)

Given a strategy profile s, we define, for any cell $\pi_i \in \Pi_i$ of player i, the payoff of π_i as the (interim) expected payoff of player i given that her signal is π_i, i.e., $u_i(s|\pi_i) = \sum_{\omega\in\pi_i} \mu(\omega|\pi_i) u_i(s|\omega)$. Note that player i is best-responding iff every cell of hers is best-responding.

Player i has K_i units of capacity, which we index by numbers between 1 and K_i. A *cell-unit* of player i is a pair (π_i, j) where $\pi_i \in \Pi_i$ is a cell, and $1 \leq j \leq K_i$ is a unit index. We can assume w.l.o.g. that a strategy specifies not only in which locations within π_i to search, but also which specific cell-unit is assigned to each of these locations. Thus, for every cell-unit α of i, a strategy of i chooses a location or chooses that α be inactive. We can think of a player as being composed of many "smaller" decision makers, one for each cell of hers, and of each cell as being composed of even smaller decision makers, one for every cell-unit of that cell (with the restriction that two cell-units of the same cell cannot search in the same location). We define the payoff of cell-unit $\alpha = (\pi_i, j)$ of player i as the payoff of its cell $u_i(s|\pi_i)$. Thus, every cell-unit of π_i gets the same payoff. Note that the expected reward of a cell-unit located at ω is $\mu(\omega|\pi_i) \cdot v_i^{m_s(\omega)}(\omega)$, and $u_i(s|\pi_i)$ equals the sum of the expected rewards of the active cell-units of π_i minus the cost c_i of the number of active cell-units of π_i.

Observe that if an inactive cell-unit (π_i, j) switches to searching in location ω, it makes the activation of another cell-unit (π_i, j') at ω' (weakly) less attractive than it previously was, because of increasing marginal costs (namely, the convexity of costs). Similarly, deactivating a cell-unit makes a second deactivation weakly less attractive.

Given a strategy profile s, if there exists a deviation of a single cell-unit that improves its own payoff then, by definition, it is also an improvement for its cell. Conversely, let us show that the existence of an improvement for a cell implies the existence of an improvement for some cell-unit. Suppose first that the cell improvement consists merely of changing the location of a few cell-units, without changing the number of active units. Then the cost remains unchanged, but the overall expected reward has increased. Hence, there must be at least one cell-unit α whose expected reward has increased by switching from its location ω to another location ω' that was not chosen by player i under s. Therefore, switching the location of α from ω to ω' is an improvement for α. Next, suppose that activating multiple cell-units is an improvement. Then there must also exist an improvement consisting of activating only one of these cell-units, due to the above observation about convex costs. Similarly, if the cell can improve by deactivating multiple cell-units then one of them can improve by deactivating itself.

Overall, we got that a player is best-responding iff all her cells are best-responding iff all her cell-units are best-responding.

Lemma 1. *Suppose that B is a set of cell-units (of various players), $\alpha \notin B$ is another cell-unit, and s^1 is a strategy profile under which every member of B is best-responding. Then there exists a finite sequence of cell-unit improvements $s^1, \ldots, s^T$ such that every member of $B \cup \{\alpha\}$ is best-responding under s^T.*

Proof. For convenience of description, let us imagine that all the inactive cell-unit of all players stay in some place that we denote by λ. The set $\Omega \cup \{\lambda\}$ of the locations plus λ will be called the set of *sites*. With this terminology, we can say that a strategy of player i chooses a site for every cell-unit of i (and λ is the only site where more than one cell-unit of the same player can be placed).

In what follows, whenever we mention cell-units, it only refers to members of $B \cup \{\alpha\}$. Note that the site of all other cell-units will remain fixed along the sequence.

Suppose that α is not best-responding in s^1; otherwise we are done. Let α switch from its current site θ^1 to another site θ^2 that is a best-response. The new strategy profile is s^2. Now α is best-responding, and we claim that any other cell-unit β of the same cell is still best-responding. We note first that β is not placed in θ^1 (since β was best-responding under s^1), and w.l.o.g. it is also not in θ^2 (otherwise, it is currently best-responding, since α is). Next we note that β cannot improve by switching to θ^1; otherwise, simply switching β to θ^2 would have been an improvement earlier, in s^1.

Suppose first that θ^1 and θ^2 are both locations. Then, since θ^2 is now occupied, and the preference relation between sites other than θ^1 and θ^2 has not changed (as the cost has not changed), β is indeed still best-responding. Next suppose that $\theta^2 = \lambda$. Then, by the convex costs observation, the attractiveness of λ has decreased by the switch from θ^1 to $\theta^2 = \lambda$, hence β still cannot improve by switching to λ. And although the cost has changed, the relation between *locations* other than θ^1 has not changed; hence, β is best-responding. Finally, suppose that $\theta^1 = \lambda$. Then the relation between locations other than θ^2 has not changed; hence, β is best-responding.

Phase I In s^2, we add a dummy player at the site θ^1, denoting the resulting strategy profile by σ^2 (for a profile σ^t that includes the dummy player, s^t will denote the same profile without the dummy). Then Phase I begins: at every stage of Phase I, one cell-unit who is currently not best-responding switches to a best-response site. This continues as long as there are non-best-responding cell-units, unless someone switches to θ^1, in which case Phase I immediately terminates.

As we will see, Phase I begins by some cell-unit switching from θ^2 to another site θ^3, then another cell-unit switching from θ^3 to another site, etc. More specifically, we claim that under any strategy profile $\sigma = \sigma^t$ encountered during Phase I,
(a) there exists exactly one site θ that is chosen by one more cell-unit than under s^1, i.e., $m_\sigma(\theta) = m_{s^1}(\theta) + 1$, while for every other site θ', $m_\sigma(\theta') = m_{s^1}(\theta')$ (we call θ "the plus site"); and
(b) for any cell-unit β whose current site is some θ with $m_\sigma(\theta) > 0$, and who can also choose another site θ', if there were $m_{s^1}(\theta)$ cell-units at θ (including β itself) and $m_{s^1}(\theta')$ cell-units at θ', then β would weakly prefer θ to θ'.
Property (b) roughly says that if β is not best-responding, it is only because β is in the plus site.

When Phase I starts, in σ^2, property (a) holds and α has just switched to the plus site θ^2. Since all cell-units of the cell of α were best-responding in s^2 (i.e., without the dummy) they obey property (b) in σ^2 (in particular, α is currently best-responding when α's current site is the plus site, let alone when it is not the plus site). As for cell-units of other cells, they obeyed (b) in s^1 and, therefore, they still do, as the switch of α or the addition of the dummy do not affect that.

The claim is proved by induction from one stage to the next: suppose that cell-unit β improves on stage t by switching from θ to θ'. Since β could improve, (b) implies that θ must have been the plus site in stage t. Therefore, the plus will move with β from θ to θ', and (a) will still hold in stage $t+1$. Note that, importantly, if the plus site is λ in some stage then every cell-unit is best-responding, since the number of partners does not affect the reward in λ, which simply equals 0; hence, Phase I will end on that stage.

As for property (b), β best-responds in θ', and other cell-units of β's cell obeyed (b) on stage t, implying that they were best-responding on that stage. It follows, by the same argument we used above for $t = 1$ (i.e., the transition from s^1 to s^2), that they also best-respond on stage $t+1$. Therefore, they obey (b); and cell-units of other cells still obey (b), as they were not affected by β's switch.

To see that Phase I cannot go on forever, recall first that it ends if the plus site is λ. Otherwise, on each stage of Phase I some cell-unit β switches from location ω to location ω', and the costs always remain fixed. Since this switch is an improvement, it strictly increases the expected reward of β. Now ω' becomes the plus site, and afterwards the expected reward of β when placed in ω' can never be lower than it is now, while it can be higher if the plus is somewhere else (or if β improves again).[11] Thus, the expected reward of β will never go

[11] One can verify that, in fact, β will not switch again during Phase I. We employ a different argument here, in order to strengthen the analogy with Phase II.

down to the level it was at before the switch. Hence, Phase I cannot turn into a cycle, and since there are only finitely many strategy profiles, Phase I must eventually end.

Recall that Phase I terminates once someone switches to θ^1. Therefore, the plus site cannot be θ^1 during this phase except maybe at the end, and hence nobody switches from θ^1 during Phase I. Therefore, all the switches are improvements not only in the game with the dummy added at θ^1, but also in the original game.

Denote the strategy profile at the end of Phase I by σ^*. Now we remove the dummy from θ^1. Suppose first that Phase I ended because somebody has just switched to θ^1. Then, the removal of the dummy means that now there is no plus site at all, and (b) implies that every cell-unit best-responds under s^* (recall that s^* is σ^* without the dummy player), and we are done.

Phase II Otherwise, Phase I ended at σ^* because everyone was best-responding (when the dummy was still at θ^1). Starting from s^*, we define Phase II analogously to Phase I (while Phase I more or less described a process of restabilizing the system after one cell-unit is added, Phase II describes restabilizing it after one cell-unit is removed), as follows. At every stage, as long as there are cell-units who are not best-responding, choose a cell, and choose a switch of a single cell-unit that would yield the highest increase in that cell's payoff.

As we will see, Phase II begins by some cell-unit switching to θ^1 from some site θ', then another cell-unit switching to θ' from another site, etc. More specifically, we claim that under any strategy profile $s = s^t$ encountered during Phase II,
(a') there exists exactly one site θ with $m_s(\theta) = m_{\sigma^*}(\theta) - 1$, while for every other site θ', $m_s(\theta') = \hat{m}_{\sigma^*}(\theta')$ (we call θ "the minus site"); and
(b') for any cell-unit β whose current site is some θ and who can also choose site θ', if there were $m_{\sigma^*}(\theta)$ cell-units at θ (including β) and $m_{\sigma^*}(\theta')$ cell-units at θ', then β would weakly prefer θ to θ'.

The analysis is almost analogous to that of Phase I. When Phase II starts, in the profile s^*, (a') holds and θ^1 is the minus site. (b') also holds because everyone was best-responding under σ^*. The claim is proved by induction from one stage to the next: suppose that cell-unit β improves on stage t by switching from θ to θ'. Improvement implies, by (b'), that θ' must have been the minus site in stage t. Therefore, the minus will move from θ' to θ, and (a') will still hold in stage $t+1$. Note that if the minus site is λ in some stage, then every cell-unit is best-responding and Phase II will end on that stage.

Let π_i be the cell of β. Any cell-unit of another cell still obeys (b'), as it was not affected by β's switch. Since the switch of β from θ to θ' was, by definition of Phase II, a best cell-unit switch for π_i, β cannot improve again. Therefore, β obeys (b'), since β is not in the minus site. Let γ be another cell-unit of π_i. Note first that γ cannot improve by switching to the current minus site θ; otherwise, switching γ to θ' earlier, on stage t, would have been a better switch than the one chosen.

We have seen that $\theta' \neq \lambda$. If also $\theta \neq \lambda$, then the preference relation between sites other than θ and θ' has not changed, and, therefore, γ still obeys (b'). Otherwise, $\theta = \lambda$. Then if γ is placed in λ it still obeys (b'), because the attractiveness of λ has increased, by the convex costs observation; and if γ is placed in some location then γ still obeys (b'), because the relation between locations other than θ' has not changed.

When Phase II ends, every cell-unit will be best-responding. To see that Phase II must eventually end, we employ the same argument as in Phase I, noting that right after some cell-unit β switches to location ω', ω' is not the minus site, and, therefore, the expected reward of β when placed in ω' can never be lower than it is now. □

To prove weak acyclicity, start from any strategy profile. By applying Lemma 1 inductively we obtain a sequence of cell-unit improvements that lead to a profile under which one cell-unit is best-responding, then two, and so on. Eventually we get a profile under which every cell-unit is best-responding, hence an equilibrium.

B.2 Proof of Prop. 3 (Any $f \in F_C$ Can Be Induced by $\tau \in T$)

Denote $\hat{\Pi} = \{(i;\pi_i) : i \in N, \pi_i \in \Pi_i\}$. We construct a *flow network*, namely, a directed graph $D = (V, E)$ with vertices V and edges $E \subset V \times V$, and a flow capacity $\kappa(v_1, v_2) \geq 0$ for every edge (v_1, v_2) (illustrated beside the sketch of this proof, in Figure 7). There are two special vertices, a *source* s and a *sink* t. The other vertices in our network are the locations Ω and the cells $\hat{\Pi}$ of the game. There is an edge from s to every $(i;\pi_i) \in \hat{\Pi}$, where $\kappa(s,(i;\pi_i)) = K_i$, and an edge from every location $\omega \in \Omega$ to t, where $\kappa(\omega, t) = f(\omega)$. Also, there is an edge from a cell $(i;\pi_i) \in \hat{\Pi}$ to a location ω iff π_i contains ω, and the flow capacity κ of such edges is infinite (for a textbook presentation of flow networks; see, e.g., Cormen *et al.*, 2009, Ch. 26).

A *cut* of D is a subset of edges $C \subset E$, such that if all the edges of C are removed then there exists no path between s and t. Suppose that C is a *minimal cut,* i.e., a cut whose sum of capacities is minimal. Then C certainly does not include any edge between a cell and a location, as those edges have an infinite flow capacity. Let $Q = \{\omega \in \Omega : (\omega, t) \in C\}$ denote the locations that the cut separates from t. Denote $W = \Omega \setminus Q$. Then C must include all the edges $\{(s,(i;\pi_i)) : i \in N, \pi_i \cap W \neq \emptyset\}$; otherwise there would still exist a path from s to

t. Hence, the total capacity of C equals

$$\sum_{\omega\in Q}\kappa(\omega,t)+\sum_{i\in N}\sum_{\pi_i\cap W\neq\emptyset}\kappa(s,(i;\pi_i))=\sum_{\omega\in Q}f(\omega)+\sum_{i\in N}\sum_{\pi_i\cap W\neq\emptyset}K_i=$$

$$\sum_{\omega\in Q}f(\omega)+\sum_{i\in N}K_i\cdot|\{\pi_i\in\Pi_i:\pi_i\cap W\neq\emptyset\}|\geq\sum_{\omega\in Q}f(\omega)+\sum_{\omega\in W}f(\omega)=\sum_{\omega\in\Omega}f(\omega).$$

Therefore, the cut that consists of all edges of type (ω,t), whose total capacity equals $\sum_{\omega\in\Omega}f(\omega)$, is minimal.

A *flow* in D is a function $\varphi:E\to\mathbb{R}^+$ such that: (i) the flow never exceeds the capacity, i.e., $\varphi(e)\leq\kappa(e)$, and (ii) the overall flow outgoing from s, namely, the sum of flows on edges outgoing from s, equals the overall flow incoming to t, namely, the sum of flows on edges incoming to t (call this quantity the *value* of the flow), and for any other vertex the incoming flow equals the outgoing flow. The max-flow min-cut theorem (Cormen *et al.*, 2009, p. 723, Theorem 26.6) states that the value of the maximal flow equals the total capacity of the minimal cut; therefore, D admits a flow φ of value $\sum_{\omega\in\Omega}f(\omega)$, and so it must be the case that $\varphi(\omega,t)=f(\omega)$ for every $\omega\in\Omega$.

Now define a coordinated-search profile τ by letting $\tau_i(\pi_i,\omega)=\varphi((i;\pi_i),\omega)$ for every $i\in N,\pi_i\in\Pi_i$, and $\omega\in\pi_i$. To see that this is a coordinated-search profile we verify that for any π_i, $\sum_{\omega\in\pi_i}\tau_i(\pi_i,\omega)=\sum_{\omega\in\pi_i}\varphi((i;\pi_i),\omega)=\varphi(s,(i;\pi_i))\leq\kappa(s,(i;\pi_i))=K_i$ (where the second equality is due to the equality of the outgoing and the incoming flow). To see that τ induces f, we verify that for any ω, it is the case that $\sum_{i\in N}\tau_i(\pi_i(\omega),\omega)=\sum_{i\in N}\varphi((i;\pi_i(\omega)),\omega)=\varphi(\omega,t)=f(\omega)$.

B.3 Proof of Proposition 4 (Any f_τ is Feasible)

A nonnegative matrix $\boldsymbol{A}$ is *doubly stochastic* (resp., *doubly substochastic*) if the sum of the elements in each row and in each column is equal to (resp., at most) one, i.e., if $\sum_j A_{ij}=1$ (resp., $\sum_j A_{ij}\leq 1$) for each row i and $\sum_i A_{ij}=1$ (resp., $\sum_i A_{ij}\leq 1$) for each column j. Note that any doubly stochastic matrix must be a square matrix (but this is not the case for a doubly substochastic matrix). A doubly stochastic (resp., doubly substochastic) matrix is a permutation (resp., subpermutation) matrix if it includes only zeros and ones, i.e., if $A_{ij}\in\{0,1\}$ for any i,j. Note that a permutation (resp., subpermutation) matrix includes exactly (resp., at most) one non-zero value in each row and in each column, and this value is equal to one. The Birkhoff–von Neumann theorem states that any doubly stochastic matrix can be written as a convex combination of permutation matrices. Formally:

Theorem 3 (Birkhoff–von Neumann Theorem)**.** *Let $\boldsymbol{A}$ be a doubly stochastic matrix. Then* *there exists a finite set of permutation matrices* $\boldsymbol{P}^1, \ldots, \boldsymbol{P}^K$ *such that* $\boldsymbol{A} = \sum_k w_k \cdot \boldsymbol{P}^k$, *where* $w_k \geq 0$ *for each* k *and* $\sum_k w_k = 1$.

We present a simple extension of Thm. 3 that states that any doubly substochastic matrix can be written as a convex combination of subpermutation matrices.[12]

Lemma 2. *Let* $\boldsymbol{A}$ *be a doubly substochastic matrix. Then there exists a finite set of subpermutation matrices* $\boldsymbol{Q}^1, \ldots, \boldsymbol{Q}^K$ *s.t.* $\boldsymbol{A} = \sum_k w_k \cdot \boldsymbol{Q}^k$, *where* $w_k \geq 0$ *for each* k *and* $\sum_k w_k = 1$.

Proof. Let I (resp., J) be the number of rows (resp., columns) in the matrix $\boldsymbol{A}$. We construct a square doubly stochastic matrix $\boldsymbol{B}$ with $I+J$ rows and columns by merging 4 submatrices (as illustrated in Figure 12): (1) the matrix $\boldsymbol{A}$ (with I rows and J columns) in the top-left part of $\boldsymbol{B}$, (2) a $J \times J$ diagonal matrix in the bottom-left part of $\boldsymbol{B}$, where each diagonal cell completes the values in each column of $\boldsymbol{A}$ to one, (3) an $I \times I$ diagonal matrix in the top-right part of $\boldsymbol{B}$, where each diagonal cell completes the values in each row of $\boldsymbol{A}$ to one, and (4) the $J \times I$ matrix $\boldsymbol{A}^T$ (the transpose of $\boldsymbol{A}$) in the bottom-right part of $\boldsymbol{B}$. It is immediate that $\boldsymbol{B}$ is a doubly stochastic matrix. By Theorem 3 there exists a finite set of permutation matrices $\boldsymbol{P}^1, \ldots, \boldsymbol{P}^K$ (with $I+J$ rows and columns) such that $\boldsymbol{B} = \sum_k w_k \cdot \boldsymbol{P}^k$, where $w_k \geq 0$ for each k and $\sum_k w_k = 1$. Let $\boldsymbol{Q}^k$ be a submatrix of $\boldsymbol{P}^k$ with the first I rows and J columns. Then it is immediate that each $\boldsymbol{Q}^k$ is a subpermutation matrix and that $\boldsymbol{A} = \sum_k w_k \cdot \boldsymbol{Q}^k$. □

Figure 12: Illustration of How to Construct the Square Matrix $\boldsymbol{B}$

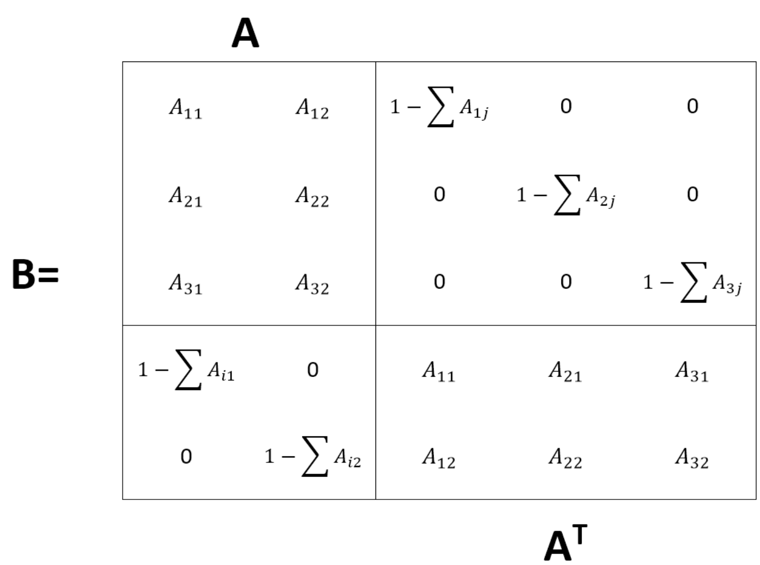

[12] One can show that Lemma 2 is implied by the extension of the Birkhoff–von Neumann Theorem presented in Budish *et al.* (2013). For completeness, we provide a self-contained proof of the lemma.

Next we rely on Lemma 2 to prove Proposition 4. Let τ be a coordinated-search profile. Similarly to the proof of Proposition 1, we define a *cell-unit* as a tuple (i, j, π_i), where $i \in N$ is a player, $j \in \{1, ..., K_i\}$ is an index corresponding to one unit of capacity of player i, and $\pi_i \in \Pi_i$ is a cell of player i. Let $\hat{\Pi}$ denote the set of all cell-units with a typical element $\hat{\pi}$, let $\hat{\Pi}_i$ denote the subset of cell-units that correspond to player i, and let $\hat{\Pi}_{i,j}$ denote the subset of cell-units that correspond to capacity unit $j \in \{1, ..., K_i\}$ of player i. We write $\omega \in \hat{\pi} = (i, j, \pi)$ if $\omega \in \pi$.

A coordinated-search action profile (of the cell-units) is a function τ' that assigns to each cell-unit (i, j, π) an element of $\mathcal{D}(\pi, 1)$ (recall that an element of $\mathcal{D}(\pi, 1)$ is a function $\eta : \pi \to [0, 1]$ such that $\sum_{\omega \in \pi} \eta(\omega) \leq 1$). The coordinated-search profile τ can be represented as an equivalent coordinated-search action profile (of the cell-units) τ' that satisfies $\sum_{j=1}^{K_i} \tau'((i, j, \pi_i), \omega) = \tau_i(\pi_i, \omega)$ for each $\pi_i \in \Pi_i$ and $\omega \in \Omega$. The equivalent coordinated-search action profile τ' can be represented as a $\left|\hat{\Pi}\right| \times |\Omega|$ nonnegative matrix $\boldsymbol{C}$ as follows:

$$C_{(i,j,\pi_i),\omega} = \begin{cases} \tau'((i, j, \pi_i), \omega) & \omega \in \pi_i \in \Pi_i \\ 0 & \text{otherwise.} \end{cases}$$

Observe that the sum of each row in $\boldsymbol{C}$ is at most one, i.e., $\sum_{\omega \in \Omega} C_{\hat{\pi},\omega} \leq 1$, but the sum of a column might be greater than one. Let $\boldsymbol{A}$ be the matrix derived from $\boldsymbol{C}$ by decreasing the values of the lower cells within columns whose sum is greater than one, such that the sum of each column is at most one. Formally (where we write $\hat{\pi}' < \hat{\pi}$ if the row of $\hat{\pi}'$ is higher than the row of $\hat{\pi}$ in the matrix $\boldsymbol{C}$):

$$A_{\hat{\pi},\omega} = \begin{cases} C_{\hat{\pi},\omega} & \sum_{\hat{\pi}' \leq \hat{\pi}} C_{\hat{\pi}',\omega} \leq 1 \\ 1 - \sum_{\hat{\pi}' < \hat{\pi}} C_{\hat{\pi}',\omega} & \sum_{\hat{\pi}' < \hat{\pi}} C_{\hat{\pi}',\omega} \leq 1 < \sum_{\hat{\pi}' \leq \hat{\pi}} C_{\hat{\pi}',\omega} \\ 0 & \sum_{\hat{\pi}' < \hat{\pi}} C_{\hat{\pi}',\omega} > 1. \end{cases}$$

Observe that $\boldsymbol{A}$ is a doubly substochastic matrix (i.e., the sum of each row and of each column is at most one), and that the coordinated-search action profile corresponding to $\boldsymbol{A}$ induces the same mixed outcome as τ. By Lemma 2, there exists a finite set of subpermutation matrices $\boldsymbol{Q}^1, \ldots, \boldsymbol{Q}^K$ such that $\boldsymbol{A} = \sum_k w_k \cdot \boldsymbol{Q}^k$, where $w_k \geq 0$ for each k and $\sum_k w_k = 1$. Further observe that each subpermutation matrix $\boldsymbol{Q}^k$ corresponds to the cell-unit representation of a pure strategy profile s^k, which implies that τ induces the same mixed outcome as the correlated strategy profile $\sigma = \sum_k w_k \cdot s^k$.

B.4 Proof of Proposition 5 (Intersecting Signals)

Let $M = \{i \in N : |\Pi_i| \geq 2\}$ be the set of players whose partitions are not trivial (i.e., players that have at least two cells in their partition), and denote $m = |M|$. For any player $i \in N$, location $\omega \in \Omega$, and any tuple of cells $(\pi_j)_{j \neq i}$, the profile of n cells $(\pi_i(\omega), (\pi_j)_{j \neq i})$ has at least κ locations in its intersection. Thus, $\pi_i(\omega)$ contains at least $\kappa \cdot \prod_{j \in N \setminus i} |\Pi_j|$ such intersection points, and

$$\kappa \cdot \prod_{j \in N \setminus i} |\Pi_j| = \kappa \cdot \prod_{j \in M \setminus i} |\Pi_j| \geq \kappa \cdot 2^{m-1}.$$

Hence, $|\pi_i(\omega)| \geq \kappa \cdot 2^{m-1}$. Suppose first that $M \neq \emptyset$. Now consider a smaller, auxiliary search game created by omitting all players in $N \setminus M$, leaving only members of M to play. Since $\kappa \cdot 2^{m-1} \geq \kappa \cdot m$ and $K_i \leq \kappa$, Corollary 4 implies that the smaller game admits a redundancy-free strategy s_M. If $M = \emptyset$ then s_M is empty and the proof proceeds the same.

Under s_M, $\sum_{i \in M} K_i \cdot |\Pi_i|$ distinct locations are searched. Going back to the original game, we define a strategy profile s by complementing s_M with strategies of the members of $N \setminus M$ as follows. We let them choose, one by one (within their single cell, namely, the whole Ω), any K_i locations that have not been chosen by other players yet, as long as there are such.

Case 1: Assume that $\sum_{i \in N} K_i \cdot |\Pi_i| \leq |\Omega|$. This implies that this procedure continues until all members of $N \setminus M$ have chosen. We end up with a strategy profile s that is redundancy-free. Observe that s is also exhaustive iff $\sum_{i \in N} K_i \cdot |\Pi_i| = |\Omega|$, and that if $\sum_{i \in N} K_i \cdot |\Pi_i| < |\Omega|$ then the game does not admit an exhaustive strategy profile (as the players can search in at most $\sum_{i \in N} K_i \cdot |\Pi_i|$ locations).

Case 2: We are left with the case of $\sum_{i \in N} K_i \cdot |\Pi_i| > |\Omega|$ in which the procedure cannot be completed, as at some stage all the locations will already have been chosen before all the players have been able to choose. Let the remaining players choose arbitrarily. We end up with an exhaustive strategy profile s.

If the payoffs are solitary-search dominant then, either when s is redundancy-free or when s is exhaustive, Corollary 2 implies that G also admits such an equilibrium.

References

Ackermann, Heiner, Röglin, Heiko, & Vöcking, Berthold. 2009. Pure nash equilibria in player-specific and weighted congestion games. *Theoretical Computer Science*, **410**(17), 1552–1563.

Akcigit, Ufuk, & Liu, Qingmin. 2015. The role of information in innovation and competition. *Journal of the European Economic Association*, **14**(4), 828–870.

Aumann, Robert J. 1976. Agreeing to disagree. *The Annals of Statistics*, **4**(6), 1236–1239.

Ben-Zwi, Oren. 2017. Walrasian's characterization and a universal ascending auction. *Games and Economic Behavior*, **104**, 456–467.

Berman, Abraham, & Plemmons, Robert J. 1994. *Nonnegative Matrices in the Mathematical Sciences*. Vol. 9. SIAM.

Birkhoff, Garrett. 1946. Tres observaciones sobre el algebra lineal. *Universidad Nacional de Tucumán*, **5**, 147–154.

Blonski, Matthias. 2005. The women of cairo: Equilibria in large anonymous games. *Journal of Mathematical Economics*, **41**(3), 253–264.

Bronfman, Slava, Alon, Noga, Hassidim, Avinatan, & Romm, Assaf. 2018. Redesigning the israeli medical internship match. *ACM Transactions on Economics and Computation*, **6**(3–4), 1–18.

Bryan, Kevin A., & Lemus, Jorge. 2017. The direction of innovation. *Journal of Economic Theory*, **172**, 247–272.

Budish, Eric, Che, Yeon-Koo, Kojima, Fuhito, & Milgrom, Paul. 2013. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, **103**(2), 585–623.

Chatterjee, Kalyan, & Evans, Robert. 2004. Rivals' search for buried treasure: Competition and duplication in r&d. *RAND Journal of Economics*, **35**(1), 160–183.

Che, Yeon-Koo, & Gale, Ian. 2003. Optimal design of research contests. *American Economic Review*, **93**(3), 646–671.

Chen, Yiling, Nissim, Kobbi, & Waggoner, Bo. 2015. Fair information sharing for treasure hunting. *Pages 851–857 of: Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Cormen, Thomas H, Leiserson, Charles E, Rivest, Ronald L, & Stein, Clifford. 2009. *Introduction to Algorithms*. MIT Press: Cambridge, MA.

Erat, Sanjiv, & Krishnan, Vish. 2012. Managing delegated search over design spaces. *Management Science*, **58**(3), 606–623.

Fershtman, Chaim, & Rubinstein, Ariel. 1997. A simple model of equilibrium in search procedures. *Journal of Economic Theory*, **72**(2), 432–441.

Ford, LR, & Fulkerson, DR. 1956. Maximal flow through a network. *Canadian Journal of Mathematics*, **8**, 399–404.

Fullerton, Richard L, & McAfee, R Preston. 1999. Auctioning entry into tournaments. *Journal of Political Economy*, **107**(3), 573–605.

Hall, P. 1935. On representatives of subsets. *Journal of the London Mathematical Society*, **1**(1), 26–30.

Kleinberg, Jon, & Oren, Sigal. 2011. Mechanisms for (mis)allocating scientific credit. *Pages 529–538 of: Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing.*

Koh, Youngwoo. 2017. Incentive and sampling effects in procurement auctions with endogenous number of bidders. *International Journal of Industrial Organization*, **52**, 393–426.

Konrad, Kai A. 2014. Search duplication in research and design spaces: Exploring the role of local competition. *International Journal of Industrial Organization*, **37**, 222–228.

Letina, Igor. 2016. The road not taken: competition and the R&D portfolio. *The RAND Journal of Economics*, **47**(2), 433–460.

Letina, Igor, & Schmutzler, Armin. 2019. Inducing variety: A theory of innovation contests. *International Economic Review*, **60**(4), 1757–1780.

Liu, Qingmin, & Wong, Yu Fu. 2019. Strategic exploration. *mimeo.*

Loury, Glenn C. 1979. Market structure and innovation. *Quarterly Journal of Economics*, **93**(3), 395–410.

Milchtaich, Igal. 1996. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, **13**(1), 111–124.

Monderer, Dov, & Shapley, Lloyd S. 1996. Potential games. *Games and Economic Behavior*, **14**, 124–143.

Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, **4**(1), 25–55.

Taylor, Curtis R. 1995. Digging for golden carrots: An analysis of research tournaments. *The American Economic Review*, **85**(4), 872–890.

Tierney, Ryan. 2019. The problem of multiple commons: A market design approach. *Games and Economic Behavior*, **114**, 1–27.

Von Neumann, John. 1953. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games*, **2**(0), 5–12.