MPRA
Munich Personal RePEc Archive

Social Welfare in Search Games with Asymmetric Information

Bavly, Gilad and Heller, Yuval and Schreiber, Amnon

Bar-Ilan University

27 February 2020

Online at https://mpra.ub.uni-muenchen.de/98838/
MPRA Paper No. 98838, posted 29 Feb 2020 15:52 UTC

Social Welfare in Search Games with Asymmetric Information

Gilad Bavly, Yuval Heller, Amnon Schreiber*

February 27, 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: search duplication, decentralized research, social welfare, incomplete information.

JEL Codes: C72, D82, D83.

1 Introduction

Many real-life situations involve decentralized research with the following four key properties: (1) agents explore different routes to making a discovery (henceforth, finding a *prize*), (2) each agent has private information about the plausibility of different routes (henceforth, the prize's potential *locations*), (3) society gains from a successful discovery, and (4) each agent gains from being the discoverer (though, the private gain might be different from the social gain), where the gain is reduced if she is not the sole discoverer. 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. We thank Sergiu Hart, Igal Milchtaich, Abraham Neyman, Ron Peretz, 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 at Bar-Ilan University, the Sir Isaac Wolfson Chair in Economics and Business Administration, and ISF grant 1626/18.

Example 1. Society faces a problem of finding a gene that induces a rare genetic disorder. There are various possible research directions that may lead to finding the culprit gene. Different research labs (or pharmaceutical R&D divisions) have heterogeneous private information about the identity of the culprit gene. Society gains if the gene is found by at least one lab. A lab that discovers the gene gains from the discovery (credit or award for the scientists, or future 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 *et al.*, 1979; Fershtman & Rubinstein, 1997; Chatterjee & Evans, 2004; Konrad, 2014; Akcigit & Liu, 2015; Letina, 2016; Liu & Wong, 2019), design of innovation contests (e.g., Che & Gale, 2003; Adamczyk *et al.*, 2012, Erat & Krishnan, 2012; Bryan & Lemus, 2017; Letina & Schmutzler, 2019; Mihm & Schlapp, 2019; Matros *et al.*, 2019), pharmaceutical research (e.g., Matros & Smirnov, 2016; de Roos *et al.*, 2018), academic research (e.g., Kleinberg & Oren, 2011), and product design within a firm (e.g., Loch *et al.*, 2001). Most of the existing literature assumes that all agents have symmetric information regarding the prize's location.[1] The main methodological innovation of the present model is the introduction of asymmetric information into search games.

The social payoff is clearly constrained by the information structure, as each player acts solely based on her private information, since we do not allow players to share their information. The social payoff may also be constrained by the fact that players' individual preferences can differ from society's, and players have strategic considerations as well. Thus, our main goal is to study 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. Player i is able to search in c_i locations (all at once). She receives some private coarse signal about the actual location of the prize, and chooses which c_i locations to search. It is a one-shot game (i.e., if the prize is not found, they do not get to search again) with simultaneous actions. We assume that each private signal is a deterministic function of the prize's location. We allow the prize's value to depend on the location. Also, the value for society and the individual values of players may all be different.

If multiple players find the prize, then each finder's value is divided by the overall number of finders (possibly, reflecting a symmetric lottery determining who is recognized as the "true"

[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 the 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.

finder). Moreover, the value may be further reduced by a fixed multiplicative factor (possibly reflecting price wars or credit wars between the finders). By contrast, the social value of the prize is unaffected by the number of finders.

First Main Result As the main question is what society can achieve in equilibrium, our answer consists of two parts. Our first main result states that under "reasonable" conditions there exists a (pure) equilibrium that yields the first-best social payoff, i.e., the highest social payoff that any strategy profile can yield. The conditions are that for any two locations ω and ω' that some player considers possible, (1) that player's preference between the two locations does not strictly contradict society's preference, i.e., it cannot be that her private value in ω is strictly higher than in ω' while it is the other way round for society, and (2) the player always prefers searching ω by herself to searching ω' with other players.

It is relatively easy to see that neither condition can be dropped. As for their sufficiency, 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. Note, however, that not all equilibria yield the first-best payoff. 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. The result is actually slightly broader than that: a strategy profile determines an *outcome*, namely, a specification of the locations in which the prize will be found (by anyone) if the prize is indeed there, and the locations in which it will not. Similarly, a probability distribution over strategy profiles (a.k.a. a correlated strategy) induces a *mixed outcome*, namely, a probability of being found for each location. We characterize the set of mixed outcomes that are induced by some correlated strategy. In particular, the first-best social payoff equals the maximal social payoff of outcomes within that set.

Our model does not allow players to coordinate partial search efforts within locations so that they do not overlap (i.e., allow players i and j to each assign an effort of 50% to location ω, and let the outcome be that ω is always searched by exactly one of them). As it turns out, this constraint does not affect the first-best payoff. To see this, suppose to the contrary that we did allow for such coordination within locations. We show that any outcome of such a setup can already be achieved in our model by some probability distribution over strategy profiles (hence, in particular, the social payoff of that outcome is at least equaled by one of those strategy profiles).[2]

[2] Notably, this is not the case with a more elaborate information structure; see Section 5.

This result is actually strongly connected to the characterization of the first-best payoff discussed above, constituting a key part of its proof. Overall, our proofs rely on representing a search game as a bipartite graph and adapting and extending various classic results from graph theory to our setup: Hall's marriage theorem (Hall, 1935), max-flow min-cut theorem (Ford & Fulkerson, 1956), and Birkhoff–von Neumann theorem (Birkhoff, 1946; Von Neumann, 1953).[3] Finally, we examine the implications of removing some assumptions of our model, such as simultaneous one-shot actions and deterministic signals.

Structure Section 2 presents our model. We study the conditions for 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 elaborate information structures, where signals are not deterministic functions of the prize's location. 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. The players search for a prize hidden in one of the locations described below. In addition to the players, we consider 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. In the motivating example described above, one can think of society as representing the regulator in the pharmaceutical industry, who cares for the welfare of future patients who may be affected by the genetic disorder.

We model the incomplete information à la Aumann (1976). Let Ω be a finite set of all states of the world. A typical state is denoted by ω. Importantly, each state of the world in our model corresponds to a different location of the prize (henceforth, we refer to the elements of Ω both as *states* and as *locations*), which formalizes our assumption that the prize's location determines the (coarse) signals of all players.[4] For each player i, let Π_i be 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,

[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., Ben-Zwi, 2017).

[4]In Section 5 we explain why deterministic signals can be modeled by having each location correspond to a single state. In addition, we analyze there a more general model in which signals are not determined by the prize's location.

let $\pi_i(\omega)$ denote the cell of player i that contains state ω. If the true state of the world is ω, then player i knows (due to the private signal that she observes) that the state of the world is one of the states in $\pi_i(\omega)$. Therefore, we can, without loss of generality, identify each signal that a player may observe with the cell that the signal induces. Figure 1 demonstrates an information structure in a two-player search game.

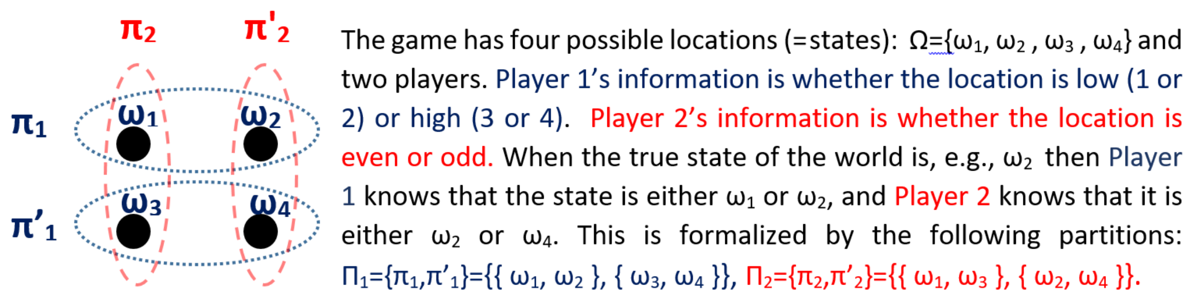

Figure 1: Illustration of information structure of a two-player search game

To make the model more general, we allow heterogeneity in the prior beliefs of the players (as well as in the players' private payoffs, as detailed below). Let $\mu_i \in \Delta(\Omega)$ (resp., $\mu_{\mathfrak{s}} \in \Delta(\Omega)$) denote the prior belief of player i (resp., society) about the state, where $\Delta(\Omega)$ denotes the set of distributions over Ω. For a subset of elements (event) $E \in \Omega$, let $\mu_i(E) = \sum_{\omega \in E} \mu_i(\omega)$ denote the prior probability that player i assigns to the event E. For simplicity, we assume that each player assigns a positive probability to each of her cells, i.e., $\mu_i(\pi_i) > 0$ for each $\pi_i \in \Pi_i$ and $i \in N$. When the true (unknown) state of the world is ω, each player i assigns a posterior belief of $\mu_i(\omega'|\pi_i(\omega)) = {\mu_i(\omega')}/{\mu_i(\pi_i(\omega))}$ to the state being $\omega' \in \pi_i(\omega)$, and 0 to any state $\omega' \notin \pi_i(\omega)$. When everyone shares the same prior we refer to it as the *common prior* and denote it simply by μ.

We allow heterogeneity in the number of locations that each player can search. Specifically, each player i chooses up to $c_i \in \mathbb{N}$ locations in which she searches, where c_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 c_i elements. We interpret $s_i(\pi_i)$ as the set of up to c_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 searches in the location ω, if $\omega \in s_i(\pi_i)$ for some $\pi_i \in \Pi_i$). 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.

Values and Duplication When the prize is found it yields a private value for its finder, and a social value for society. We allow these values to depend on the state of the world. For any state ω, let $v_i(\omega) \in \mathbb{R}^+ \equiv \{x \in \mathbb{R} | x \geq 0\}$ (resp., $v_{\mathfrak{s}}(\omega) \in \mathbb{R}^+$) denote the prize's private value for player i (resp., social value for society) when player i is the sole finder of the prize in state ω. Observe that the social value does not depend on the identity of the prize's finder. The players and society are both assumed to be risk neutral with respect to their payoffs. We say that the game has *common values* if $v_i(\omega) = v_j(\omega) = v_{\mathfrak{s}}(\omega)$ for every two players $i, j \in N$ and every state $\omega \in \Omega$, and in this case we omit the subscript and write the common value as $v(\omega)$. It turns out that our results depend on the prior $\mu_i(\omega)$ and on the value $v_i(\omega)$ only through their product $\mu_i(\omega) \cdot v_i(\omega)$. We refer to this product as the *expected private value* of player i in state ω (and $\mu_{\mathfrak{s}}(\omega) \cdot v_{\mathfrak{s}}(\omega)$ is called the expected social value in state ω).

If multiple players find the prize, we assume that each finder's private value is divided by the number of simultaneous finders (possibly reflecting a symmetric lottery determining which of these finders is recognized as the "true" finder).[5] In addition, each finder's private value is multiplied by a *duplication factor* $\rho \in [0,1]$ if two or more players find the prize simultaneously. The case in which $\rho = 0$ corresponds to setups in which a price (Bertrand) competition between the pharmaceutical firms or a "credit war" between the research labs (see Example 1) destroys the finder's private value in case of a simultaneous discovery (e.g., Chatterjee & Evans, 2004; Matros & Smirnov, 2016; de Roos *et al.*, 2018). The opposite case of $\rho = 1$ 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 (e.g., Fershtman & Rubinstein, 1997; Konrad, 2014; Chen *et al.*, 2015). In the intermediate case of $\rho \in (0,1)$ some of the prize's private value is lost when there are multiple finders.[6]

We assume that the social value of the prize is not reduced when there are multiple finders, which seems plausible in many 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 gain from a new discovery is not likely to be reduced when two scientists fight over the credit.

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

[5]In order to simplify notation, our model allows each player to search in each location only once, and when there are multiple finders, each has the same probability to be the recognized "true" finder. All of our results can be adapted to a setup in which players may search multiple times in the same location, and they may choose to do so when competing with other players searching in the same location (i.e., in this modified setup, if Player 1 (resp.,Player 2) searches twice (resp., once) in the location ω, then, conditional on the true state being ω, Player 1 (resp., Player 2) is recognized as the "true" finder with probability $2/3$ (resp., $1/3$)).

[6]All of our results remain the same if one assumes that the duplication factor is decreasing in the number of finders; i.e., if the private value of the prize is $\rho(m) \cdot v_i$, where m is the number of players simultaneously finding the prize, $\rho(m)$ is weakly decreasing in m, and ρ stands for $\rho(2)/\rho(1)$.

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

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

As described above, the payoff of player i conditional on the state being ω, denoted by $u_i(s|\omega)$, is equal to $v_i(\omega)$ if player i is the only player to search in ω, and is equal to $\rho \cdot v_i(\omega) / m_s(\omega)$ if player i has searched ω together with other players, and is equal to zero if the player has not searched ω. The (ex-ante) expected payoff of player i is denoted by $u_i(s)$. Formally, the expressions for $u_i(s|\omega)$ and for $u_i(s)$ are given by

$$u_i(s|\omega) = \begin{cases} v_i(\omega) & \omega \in s_i(\pi_i(\omega)) \text{ and } m_s(\omega) = 1 \\ \rho \cdot \frac{v_i(\omega)}{m_s(\omega)} & \omega \in s_i(\pi_i(\omega)) \text{ and } m_s(\omega) > 1 \\ 0 & \omega \notin s_i(\pi_i(\omega)) \end{cases} \qquad u_i(s) = \sum_{\omega \in \Omega} \mu_i(\omega) \cdot u_i(s|\omega).$$

A strategy profile $s = (s_1, ..., s_n)$ is *a Nash equilibrium* of search game G if no player can gain from 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 search game G and a strategy profile s. Let $U(s|\omega) = v_{\mathfrak{s}}(\omega) \cdot \mathbf{1}_{m_s(\omega) \geq 1}$ denote the social payoff, conditional on the state being ω. The expected social payoff, given the social prior $\mu_{\mathfrak{s}}$, is equal to $U(s) = \sum_{\omega \in \Omega} \mu_{\mathfrak{s}}(\omega) \cdot U(s|\omega)$. Let U_S denote the *socially optimal payoff* (or the first-best payoff):

$$U_S = \max_{s \in S} U(s).$$

A strategy profile s is *socially optimal* if it achieves the socially optimal payoff, i.e., if $U(s) = U_S$. A strategy profile is *state-maximizing* if it maximizes the number of states in which the prize is found, i.e., if for any strategy profile $s' \in S$ the following inequality holds

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

Note that the set of socially optimal strategy profiles may be different from the set of state-maximizing strategy profiles. The two notions coincide if society assigns the same expected value to every location, i.e., if $v_{\mathfrak{s}}(\omega) \cdot \mu_{\mathfrak{s}}(\omega) = v_{\mathfrak{s}}(\omega') \cdot \mu_{\mathfrak{s}}(\omega')$ for any two states $\omega, \omega' \in \Omega$. 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 state-maximizing.

3 Socially Optimal Equilibrium

In this section we characterize the existence of socially optimal equilibria; that is, we study the conditions under which the strategic constraints (namely, that each player maximizes her private payoff, rather than the social payoff) do not limit the social payoff.

3.1 Search Games are Weakly Acyclic

A sequence of strategy profiles $\left(s^1, \dots, s^K\right)$ is an improvement path (Monderer & Shapley, 1996) if each strategy profile s^{k+1} differs from its preceding profile s^k by the strategy of a single player, and this player's payoff is strictly higher in s^{k+1}.

Definition 1. Fix a search game G. A sequence of strategy profiles $\left(s^1, \dots, s^K\right)$ is an *improvement path* if for every $k \in \{1, \dots, K-1\}$ there exists a player $i_k \in N$ such that: (1) $s_j^k = s_j^{k+1}$ for every player $j \neq i$, and (2) $u_i\left(s^{k+1}\right) > u_i\left(s^k\right)$.

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.

Definition 2 (Milchtaich, 1996)**.** A game is *weakly acyclic* if for any strategy profile $s_1 \in S$, there exists an improvement path $\left(s^1, \dots, s^K\right)$, such that s^K is a (pure) Nash equilibrium.

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

Sketch of proof; formal proof is in Appendix B.1. Player i has c_i units of capacity, which we index by numbers between 1 and c_i. A *cell-unit* of player i is a pair (π, j), where $\pi \in \Pi_i$ is a cell, and j a unit index. We can think of a player's strategy as being composed of many "smaller" choices, one choice of location for each of the player's cell-units.[7] Given a strategy profile, the (ex-ante) expected payoff of a player is the sum of her cell-unit payoffs (i.e., the (ex-ante) expected payoff from the location of each cell-unit). When a cell-unit of player i improves its own payoff then it also improves the payoff of player i, and, conversely, when player i can improve her payoff there must exist a cell-unit of hers that can improve its own payoff by changing its choice.

[7] The game played by cell-units is in the spirit of Selten's (1975) agent-normal form representation.

The key part is Lemma 1 in Appendix B.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 improvements that ends with all the members of $B \cup \{\alpha\}$ best-responding. To prove this, first let α move from its current location ω_0 to its best-response location ω_1. Now add a dummy player at ω_0. Now begins Phase I of the improvements, which ends when nobody has an incentive to move or when somebody has just moved into ω_0.

Note that at this point: (a) there is one location ("the plus location"; currently it is ω_1) that is chosen by one more cell-unit compared to when we started, while the number is the same for all other locations; and (b) only cell-units located at the plus location may have an incentive to move. At every stage of Phase I, a cell-unit moves from the plus location to its best-response location, making the best-response location the new plus location. By induction from one stage to the next we see that properties (a) and (b) continue to hold at every stage of the phase.

Consider a cell-unit that moves (making its new location the plus location). Afterwards, its payoff cannot drop below its current level; it may only be higher (if the plus location is somewhere else). In particular, its payoff will never drop back to the level it was at before it moved. This implies that the sequence cannot enter a cycle; therefore, the sequence must end.

When Phase I ends we remove the dummy from ω_0. If Phase I ends because someone has just moved into ω_0, then there is no longer a plus location (since the dummy is removed), and property (b) implies that everyone is now best-responding, and we are done. Otherwise, Phase I ends because everyone was best-responding *when the dummy was in* ω_0. Let s^* denote the strategy profile we have reached. Note that: (a') there is one location ("the minus location"; currently it is ω_0) chosen by one less player compared to s^* *with the dummy*, while the number is the same for all other locations; and (b') cell-units may have an incentive to move only to the minus location.

While Phase I can be described as restabilizing after one cell-unit is added, the analogous Phase II that now follows restabilizes after one cell-unit is removed. At every stage of Phase II, a cell-unit strictly improves its payoff by moving into the minus location (if there is more than one candidate we take one whose current payoff is minimal). By induction we see that properties (a') and (b') still hold during the whole phase. The sequence cannot enter a cycle by the same argument of Phase I, and hence Phase II eventually ends. Then everyone is best-responding, and the lemma is proven.

To prove weak acyclicity, start from any profile s^1. Using the lemma inductively we add one cell-unit at a time, until eventually everyone is best-responding. □

An immediate corollary of Prop. 1 is that any search game admits a pure equilibrium.

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

3.2 Existence of a Socially Optimal Equilibrium

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

Consistency Our first property requires consistency between the private payoff and the social payoff in different locations. We say that a search game has *consistent payoffs* if for any two states ω and ω' in the same cell of player i, if the expected private value for player i is strictly lower in ω than in ω', then the expected social value is weakly lower in ω.

Definition 3. Search game $G = (N, \Omega, v, \mu, \Pi, c, \rho)$ has *consistent payoffs* if for any player i, any cell $\pi_i \in \Pi_i$, and any two states $\omega, \omega' \in \pi_i$, the following implication holds:

$$\mu_i(\omega) \cdot v_i(\omega) < \mu_i\left(\omega'\right) \cdot v_i\left(\omega'\right) \Rightarrow \mu_{\mathfrak{s}}(\omega) \cdot v_{\mathfrak{s}}(\omega) \leq \mu_{\mathfrak{s}}\left(\omega'\right) \cdot v_{\mathfrak{s}}\left(\omega'\right).$$

Observe that if everyone shares a common prior and a common value, then the search game has consistent payoffs. Further observe that if society has uniform expected values (i.e., if $\mu_{\mathfrak{s}}(\omega) \cdot v_{\mathfrak{s}}(\omega) = \mu_{\mathfrak{s}}(\omega') \cdot v_{\mathfrak{s}}(\omega')$ for any two states $\omega, \omega' \in \Omega$), then the search game has consistent payoffs regardless of what players' private payoffs are.

Balancedness Our second property requires that the expected private value does not vary too much within a player's cell. We say that a search game is r-balanced if the expected value of any player in any two states in the same cell differs by a factor of at most r.

Definition 4. Let $r \geq 1$. Search game $G = (N, \Omega, v, \mu, \Pi, c, \rho)$ has r*-balanced payoffs* if for any player i, any cell $\pi_i \in \Pi_i$, and any pair $\omega, \omega' \in \pi_i$, the following inequality holds:

$$v_i(\omega) \cdot \mu_i(\omega) \leq r \cdot v_i\left(\omega'\right) \cdot \mu_i\left(\omega'\right).$$

We say that a game has *balanced payoffs* if it has $2/\rho$-balanced payoffs (or if $\rho = 0$).

Our first main result states that any search game with consistent and balanced payoffs admits a socially optimal pure Nash equilibrium. Formally:

Theorem 1. *Let G be a search game with consistent and balanced 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 consistency and balancedness jointly imply that the social payoff cannot decrease along that sequence of unilateral improvements. Without loss of generality we can assume that each unilateral improvement consists of changing merely a single location within a single cell, since this is in fact what the proof of Proposition 1 shows. In each improvement, if the improving player leaves a location in which there were multiple searchers, then the social payoff cannot decrease. Balancedness implies that an improving player never moves from a location ω in which she was the sole searcher to an occupied location ω', due to the following inequality:

$$v_i(\omega) \cdot \mu_i(\omega) \geq \rho/2 \cdot v_i\left(\omega'\right) \cdot \mu_i\left(\omega'\right) \geq \rho/m_s(\omega') \cdot v_i\left(\omega'\right) \cdot \mu_i\left(\omega'\right).$$

Finally, consistency implies that if an improving player moves from being the sole searcher in location ω to being the sole searcher in location ω' (and this move strictly increases her payoff), then the social payoff must weakly increase. □

In Section 3.4 we discuss implications of Theorem 1 on the design of innovation contests. Our next result states that even without the consistency assumption, some efficiency is still guaranteed, in the sense that there exists an equilibrium that maximizes the number of states in which the players search. Formally:

Corollary 2. *Every search game G with balanced payoffs admits a state-maximizing equilibrium.*

Proof. Let $\hat{G} = (N, \Omega, \hat{v}, \mu, \Pi, c, \rho)$ be a search game similar to $G = (N, \Omega, v, \mu, \Pi, c, \rho)$, except that society has uniform expected payoffs, i.e., $\mu_{\mathfrak{s}}(\omega) \cdot \hat{v}_{\mathfrak{s}}(\omega) = \mu_{\mathfrak{s}}(\omega') \cdot \hat{v}_{\mathfrak{s}}(\omega')$ for any two states $\omega, \omega' \in \Omega$. Observe that $\hat{G}$ is a search game with consistent and balanced payoffs. This implies that the game $\hat{G}$ admits a socially optimal equilibrium $\hat{s}$. Observe that the definition of $\hat{v}_{\mathfrak{s}}$ implies that $\hat{s}$ is a state-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, when $\rho = 0$ (i.e., duplication destroys the prize's private values), balancedness is trivially satisfied; therefore, every search game with $\rho = 0$ admits a state-maximizing equilibrium.

Price of Anarchy Theorem 1 shows that there is an equilibrium that maximizes the social payoff (i.e., that the price of stability is 1) in any search game with consistent and balanced 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). We note, on the other hand, that in any search game with common values and a common prior, a simple argument shows that the price of anarchy cannot exceed n (the number of players).

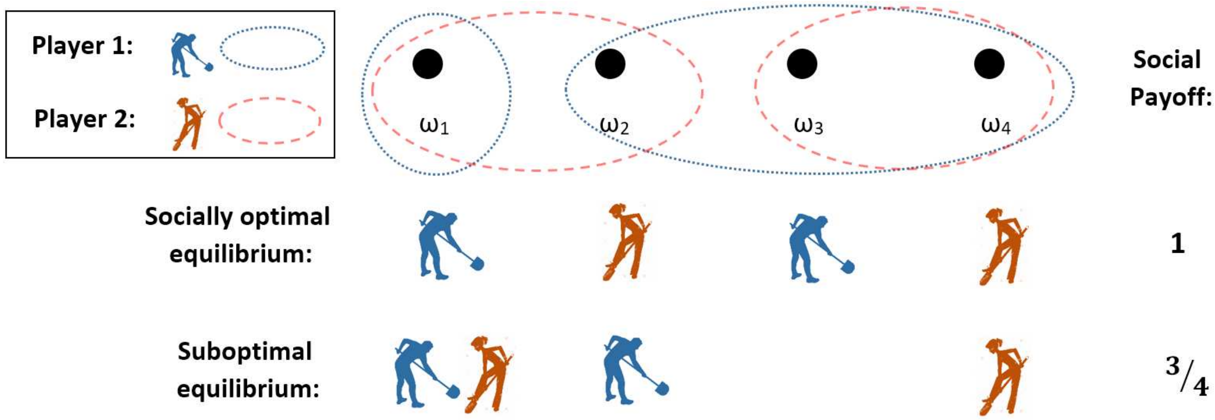

Figure 2: **Example for the price of anarchy**. The figure presents two equilibria in a two-player search game with 1-balanced consistent payoffs (the ellipses represent the partition elements) with a uniform common prior and a common value of $v \equiv 1$. The first (resp., second) equilibrium is (resp., 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 Balancedness Example 2 demonstrates that the balancedness condition is necessary for Theorem 1, in the sense that for any $r > {}^2\!/\!\rho$ there exists a search game with r-balanced consistent payoffs that does not admit a socially optimal equilibrium.

Example 2. Let $\rho \in (0,1]$ and let $r > {}^2\!/\!\rho$. Let

$$G = \left(N = \{1,2\}, \Omega = \left\{\omega, \omega'\right\}, v \equiv 1, \mu, \Pi \equiv \Omega, c \equiv 1, \rho\right)$$

be a search game with trivial information partitions ($\Pi \equiv \Omega,$), common uniform values $v \equiv 1$, and a common prior $\mu = \mu_1 = \mu_2 = \mu_{\mathbf{5}}$ defined as follows: $\mu(\omega) = \frac{r}{1+r}$ and $\mu(\omega') = \frac{1}{1+r}$. Observe that G has consistent and r-balanced payoffs. In what follows we show that 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 location ω yields a payoff of

$$\rho \cdot \frac{\mu(\omega)}{2} = \rho \cdot \frac{r}{2 \cdot (1+r)} = \rho \cdot \frac{r}{2} \cdot \mu\left(\omega'\right) > \mu\left(\omega'\right),$$

while searching in location ω' yields a strictly smaller payoff of $\mu(\omega')$. This, in turn, implies that the unique equilibrium is both players searching in ω, which is suboptimal (its payoff is $\frac{r}{1+r} < 1$, while searching in different locations yields a social payoff of 1).

Necessity of 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, the unique Nash equilibrium is not necessarily socially optimal if the consistency requirement is not satisfied.

Example 3. Let $G = (N = \{1\}, \Omega = \{\omega, \omega'\}, v, \mu, \Pi \equiv \Omega, c \equiv 1, \rho = 0)$ be a search game with a common prior (i.e., $\mu_1 = \mu_{\mathfrak{s}} = \mu$) $\mu(\omega) = 1/4$, $\mu(\omega') = 3/4$, and with values of $v_{\mathfrak{s}}(\omega) = 2$, $v_{\mathfrak{s}}(\omega') = 1$, $v_1(\omega) = 4$, and $v_1(\omega') = 1$. Observe that the game's payoffs are trivially balanced due to ρ being zero. 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 Simultaneous Searches An (implicit) key assumption in our model is that all searches are done simultaneously. In what follows we demonstrate that if searches are done sequentially, then Theorem 1 is no longer true. We present two examples, one of a setup in which a player can observe the locations in which her opponents searched in the past, and the other of a setup without observability. Both examples share the following properties: (1) the prior is uniform, (2) the private and social values of the prize are equal to δ^m in all states, where m is the round in which the prize is found, for some common discount factor $\delta \in (0,1)$, and (3) there are two players who play sequentially: player 1 searches in odd rounds and player 2 searches in even rounds.

Consider first an example with observability of the opponent's past searches. Assume that the prize is hidden within a matrix of states where player 1 knows the row and player 2 knows the column. Observe that the unique socially optimal strategy profile finds the prize in at most two rounds: player 1 searches in the correct row in round 1, and player 2 searches in the prize's true location in the second round. Observe that this socially optimal strategy profile is not a Nash equilibrium, as player 1 would gain by deviating to searching in a wrong row in round 1 (in order to deceive the opponent about the prize's row) and searching in the prize's location in round 3 (after observing the column in which player 2 searches in round 2).

Next consider an example without observability of past searches. Assume that there are two locations ω, ω', and partitions are trivial. In any socially optimal strategy profile, player 1 searches deterministically in round 1 (say, in location ω), and player 2 searches in the remaining location (ω') in round 2, which guarantees that the prize is found within at most

2 rounds. Assume to the contrary that player 1 plays deterministically in round 1 in a Nash equilibrium (say, she searches in ω). This implies that player 2 searches in the remaining state ω' in round 2. This, in turn, implies that player 1 would gain by deviating to searching in ω' in round 1 (and searching in ω in round 3), as this guarantees that player 1 always finds the prize.

3.4 Insights for Innovation Contests

In what follows we discuss the implications of Theorem 1 in the setup of an innovation contest (e.g., Che & Gale, 2003; Adamczyk *et al.*, 2012; Erat & Krishnan, 2012; Bryan & Lemus, 2017; Letina & Schmutzler, 2019; Mihm & Schlapp, 2019; Matros *et al.*, 2019), in which a contest designer, who wishes to maximize the social payoff, might influence the private payoffs of players by offering a monetary reward to the prize's finder, which is added to the finder's private value.

Observe first that if the private payoffs satisfy consistency and balancedness, then Theorem 1 implies that the designer can maximize the social payoff without offering any reward: 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 either balancedness or consistency is violated in the search game (without additional monetary rewards).

Consider first a setup in which the contest designer can only offer a constant reward, which is independent of the prize's location. A constant reward 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 no depreciation in the case of duplication (i.e., $\rho=1$), where there are two states ω,ω' in the same cell of player i with priors $\mu_i(\omega)=0.1$ and $\mu_i(\omega')=0.2$ and with private values of $v_i(\omega)=5$ and $v_i(\omega')=1$. The too-high private value of state ω violates balancedness (= 2-balancedness) 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 reward of 1 would restore balancedness (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 reward, it allows him to obtain balancedness and consistency when faced with any profile of private payoffs and private priors. An interesting open question is how the designer can maximize the social payoff, while minimizing the expected reward. For example, assume that the payoffs of the search game are consistent, but they are not balanced. Theorem 1 suggests that the designer should reward locations that have lower expected private values (which violate balancedness). 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 4.

Example 4. Consider the following search game with a common prior and common values (as illustrated in Figure 3): $(N = \{1,2\}, \Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}, v \equiv 1, \mu, \Pi, c \equiv 1, \rho = 1)$, where

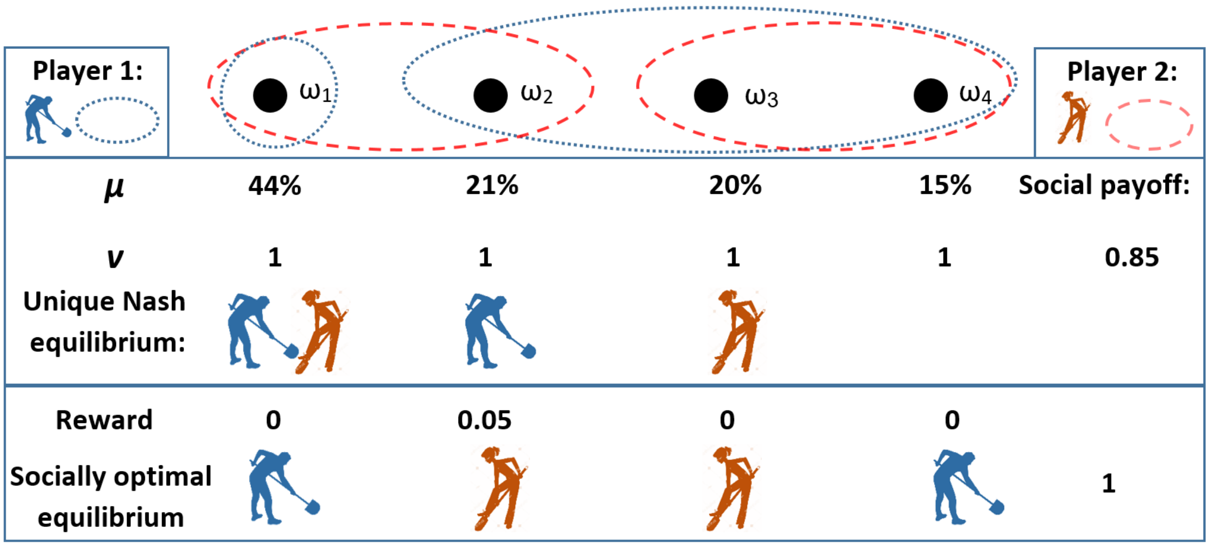

Figure 3: Illustration of Example 4: Impact of Rewards on the Social Payoff

the common prior is $\mu(\omega_1) = 44\%$, $\mu(\omega_2) = 21\%$, $\mu(\omega_3) = 20\%$ and $\mu(\omega_4) = 15\%$, player 1 observes whether the state is 1 or not, i.e., $\Pi_1 = \{\{\omega_1\}, \{\omega_2, \omega_3, \omega_4\}\}$, and player 2 observes whether the state is at most 2 or not, i.e., $\Pi_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$. The game admits a unique equilibrium, in which player 1 searches in locations 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 the cause of unbalancedness is the low expected payoff in state ω_2 (rather than a low expected payoff in ω_4).

If the social planner can offer a reward of 0.05 that increases the total private payoff in location ω_2 by 5% to 1.05 (which requires a modest expected reward of $21\% \cdot 0.05 \approx 0.01$), then the modified private payoffs satisfy balancedness, 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 the Socially Optimal Payoff

In this section we characterize the socially optimal payoff in search games.

4.1 Feasible 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 priors and 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:

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

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 $\sum_{\pi_i \in \Pi_i} c_i \cdot 1_{\pi_i \cap W \neq \emptyset}$ cells that intersect the set W, and thus she cannot search in more than $c_i \cdot \sum_{\pi_i \in \Pi_i} 1_{\pi_i \cap W \neq \emptyset}$ locations within W. This implies that all players combined cannot search in more than $\sum_{i \in N} c_i \cdot \sum_{\pi_i \in \Pi_i} 1_{\pi_i \cap W \neq \emptyset}$ locations within W. By representing the setup as a bipartite graph and applying Hall's marriage theorem (Hall, 1935), it can be shown 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 states for which f is equal to one (as illustrated in Figure 4). The graph's edges connect each cell to the states that are contained in that cell. A matching of all states in this graph, i.e., a set of disjoint edges (namely, no node appears twice) such that every state 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 states 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. □

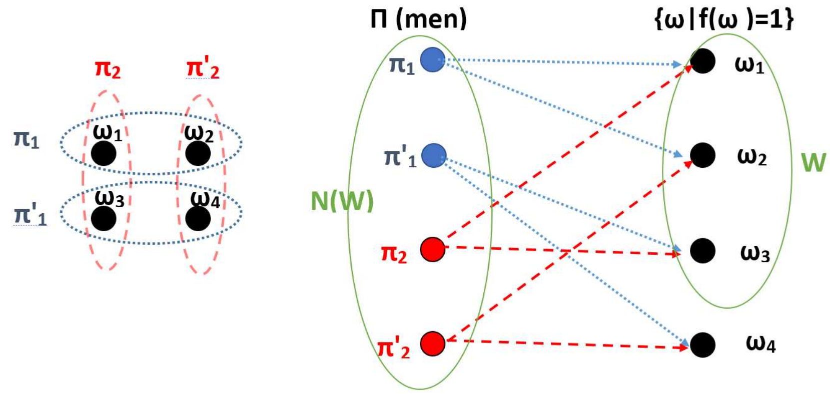

Figure 4: **Illustration of Proposition 2**. The LHS of the figure demonstrates the information partitions in a two-player search game (with capacities equal to 1). The RHS translates this into a bipartite graph, where its left part ("men") includes the cells of all players, and its right part ("women") includes all states ω satisfying $f(\omega) = 1$. The figure further shows an example of a subset of states 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 5. Suppose that there are three players, each has capacity $c_i = 1$, and every cell of every player contains exactly three states. Consider the pure outcome $f(\omega) = 1$ for every ω. To see that f is compatible, let $W \subset \Omega$ be a subset of states. 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 employed in Example 5 proves the following corollary of Proposition 2.

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

Note that if a game admits an exhaustive profile and has balanced payoffs, then it admits an exhaustive equilibrium, by Corollary 2.

Mixed outcomes A *mixed outcome* is a function $f:\Omega\rightarrow[0,1]$ that assigns a probability to each state. We interpret $f(\omega)$ as the probability that the prize is found, conditional on the state of the world 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, $f:\Omega\rightarrow[0,1]$ is *feasible* if there exists a correlated strategy profile $\sigma\in\Delta(S)$ such that

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

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 states is bounded by the sum of all relevant capacities, i.e., the 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 states $W\subseteq\Omega$, the following inequality holds:

$$\sum_{\omega\in W}f(\omega)\leq\sum_{i\in N}\sum_{\pi_i\in\Pi_i}c_i\cdot 1_{\pi_i\cap W\neq\emptyset}\,. \tag{2}$$

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 "binary" matching of zeros and ones). Instead, the proof (presented in Section 4.2) includes two parts: (1) we introduce the notion of coordinated-search profiles, and apply (Proposition 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 (Proposition 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., $|s_i(\pi_i)\cap\pi_i|=c_i$ 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 $c_i |\Pi_i|$ states, a strategy profile is redundancy-free iff the number of states being searched equals $\sum_{i\in N} c_i |\Pi_i|$. If a game admits redundancy-free strategy profiles, then they are exactly the state-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 6. Suppose that there are three players, each has capacity $c_i = 1$, and every cell of every player contains exactly five states. Consider the mixed outcome $f(\omega) = 0.6$ for every ω, and let $W \subset \Omega$ be a subset of states. 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 $\sum_{\omega\in\Omega} f(\omega) = 0.6 \cdot |\Omega| = \sum_{i\in N} c_i |\Pi_i|$, this implies that the game admits a redundancy-free strategy profile (and if the game is balanced then it admits a redundancy-free equilibrium).

More generally, the argument employed in Example 6 proves 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 $\sum_{i\in N} c_i/m_i \leq 1$ then G admits a redundancy-free strategy profile.*

Theorem 2 implies a characterization of the socially optimal payoff of search games (which is an equilibrium payoff if the game has consistent and balanced payoffs): The socially optimal payoff is the highest payoff induced by a compatible mixed outcome. Formally (where the "moreover" part is implied by Theorem 1):

Corollary 5. *Let G be a search game. Then*

$$U_S = \max_{f\in F_C} \sum_{\omega\in\Omega} f(\omega)\, \mu_{\mathfrak{s}}(\omega)\, v_{\mathfrak{s}}(\omega).$$

Moreover, if G has consistent and balanced payoffs, then U_S is an equilibrium payoff.

4.2 Coordinated Search

Alternative setup with coordinated search Consider an alternative setup that allows players to coordinate partial search efforts within a state. Specifically, we now allow each player to divide fractions of her search capacity among the different states. 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 state 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 $\mathcal{D}(\pi_i, c_i)$. We interpret $\tau_i(\pi_i, \omega) \equiv \tau_i(\pi_i)(\omega)$ as the (fractional) search effort player i exerts in state $\omega \in \pi_i$ (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 element of $\mathcal{P}(\pi, k)$).

Importantly, we assume that fractional search efforts of different players are summed optimally from society's point of view. For example, if player i assigns 50% search capacity to location ω in cell $\pi_i(\omega)$ and player j assigns 40% search capacity to location ω in cell $\pi_j(\omega)$, then the prize is found with a total probability of 90%, conditional on the state being[8] ω. Specifically, the social payoff $U^c(\tau)$ induced by a coordinated-search profile τ is

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

Observe that any coordinated-search profile induces a mixed outcome, where the finding probability assigned to each state ω is the sum of the fractional search efforts exerted by each player i in state 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 $f_\tau(\omega) = \min\left(\sum_{i \in N} \tau_i(\pi_i(\omega), \omega), 1\right)$.

Observe that any coordinated-search profile induces a compatible mixed outcome.

Claim 1. Fix $G = (N, \Omega, v, \mu, \Pi, c, \rho)$ and $\tau \in T$. Then f_τ is a compatible mixed outcome.

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

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

Our next result employs the max-flow min-cut theorem to show that the converse claim is true as well: any compatible mixed outcome can be induced by a coordinated-search profile.

[8] 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 state being ω, would be strictly less than 90%.

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 states contained in it, with infinite flow capacity (as illustrated in Figure 5). We add a source vertex that connects to every cell, with flow capacity c_i, and a sink vertex to which every state ω is 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. □

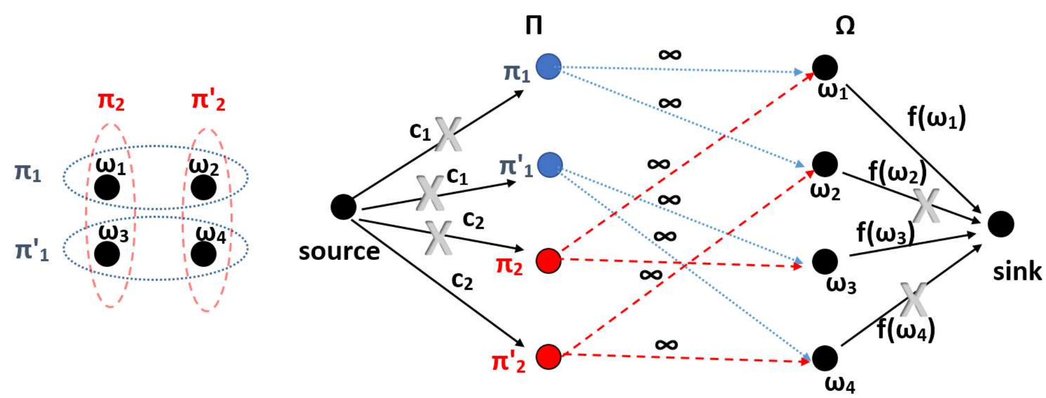

Figure 5: **Illustration of Proposition 3**. The left side of the figure demonstrates partitions in a two-player search game. The right side demonstrates the constructed directed graph in which (1) a source node is linked to every player's cells by an edge with the player's capacity, and (2) each cell is linked by unlimited edges to all the states within that cell, and (3) each state 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. Another potential constraint on the social payoff in our main model is that players are not allowed to efficiently coordinate fractional search efforts within a location as in a coordinated-search profile. 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 strategy 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 7. Let the set of states $\Omega = A \cup B_1 \cup B_2 \cup B_3$ be a union of four disjoint sets of equal size. There are three players, each with capacity $c_i = 1$. 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 6. 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 state in π_i. Thus, for any $a \in A$, $\sum_{i\in N} \tau_i(\pi_i(a),a) = 1/3 + 1/3 + 1/3 = 1$, and 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 (and if the game is balanced, it also admits an exhaustive equilibrium).

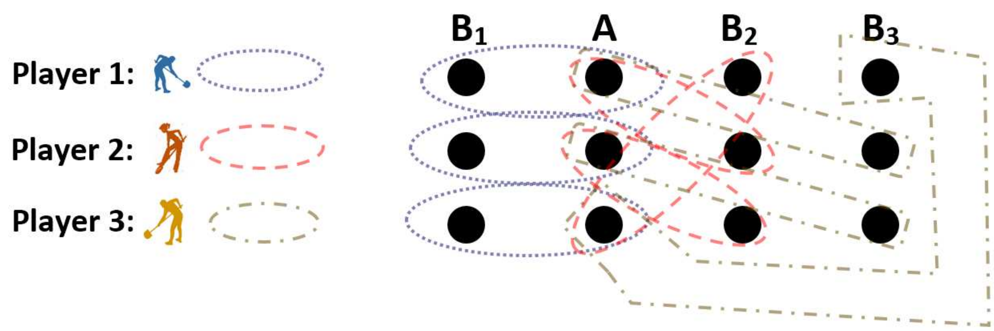

Figure 6: Illustration of Example 7 with $|\Omega| = 12$ states. The figure shows the players' partitions (where the cell containing the remaining six states of each player is not drawn to make the figure less crowded).

Note that $|\Pi_i| = |\Omega|/3$, implying that $|\Omega| = \sum_{i\in N} c_i |\Pi_i|$; 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 state ω.

Proposition 3 and Proposition 4 jointly imply Theorem 2. Moreover, Proposition 4 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 $G = (N,\Omega,v,\mu,\Pi,c,\rho)$ 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 k states, and derive tight conditions for the existence of equilibria with appealing properties.

5 Extension with Stochastic Signals

In this section we present a general model of signals, dropping the assumption that every location corresponds to a single state. 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. A state of the world 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. With a slight abuse of notation, we denote by ω a location that includes only the state ω. For a cell $\pi_i \in \Pi_i$, let $\ell(\pi_i)$ be the set of locations that are consistent with the state being in π_i, i.e., $\ell(\pi_i) = \{\ell(\omega) \,|\, \omega \in \pi_i\}$.

A *generalized search game* is a tuple $\tilde{G} = (N, \Omega, v, \mu, \Pi, L, c, \rho)$, where L is the partition of locations, and 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 (in which each location corresponds to a single state).

A *strategy* of player i is a function s_i that assigns to each cell π_i a subset of locations with at most c_i elements that satisfies $s_i(\pi_i) \subseteq \ell(\pi_i)$. We interpret $s_i(\pi_i)$ as the set of up to c_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

$$u_i(s|\omega) = \begin{cases} v_i(\omega) & \ell(\omega) \in s_i(\pi_i(\omega)) \text{ and } m_s(\omega) = 1 \\ \rho \cdot \frac{v_i(\omega)}{m_s(\omega)} & \ell(\omega) \in s_i(\pi_i(\omega)) \text{ and } m_s(\omega) > 1 \\ 0 & \ell(\omega) \notin s_i(\pi_i(\omega)). \end{cases}$$

Consistency is redefined as follows. A search game has *consistent payoffs* if for any two locations ℓ and ℓ', if the prize's interim expected private value for player i (when she is the sole finder) is strictly lower in ℓ than in ℓ', then the expected social value of the prize is weakly lower in ℓ.

Definition 7. Generalized Search game $\tilde{G}$ has *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:

$$\sum_{\omega \in \ell \cap \pi_i} \mu_i(\omega) \cdot v_i(\omega) < \sum_{\omega \in \ell' \cap \pi_i} \mu_i(\omega) \cdot v_i(\omega) \Rightarrow \sum_{\omega \in \ell \cap \pi_i} \mu_{\mathfrak{s}}(\omega) \cdot v_{\mathfrak{s}}(\omega) \leq \sum_{\omega \in \ell' \cap \pi_i} \mu_{\mathfrak{s}}\left(\omega'\right) \cdot v_{\mathfrak{s}}\left(\omega'\right).$$

Balancedness is redefined as follows. We say that a search game is r-balanced if the prize's interim expected private value for player i (when she is the sole finder), in any two locations which are possible given her posterior information, differs by a factor of at most r. Formally:

Definition 8. Let $r \geq 1$. Generalized search game $\tilde{G}$ has *r-balanced 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 inequality holds:

$$\sum_{\omega \in \ell \cap \pi_i} v_i(\omega) \cdot \mu_i(\omega) \leq r \cdot \sum_{\omega \in \ell' \cap \pi_i} v_i(\omega) \cdot \mu_i(\omega).$$

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 two states $\omega, \omega' \in \Omega$ and any two 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. This implies that the prize's location and the signal of one of the players jointly determine the signal of all players.

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 keeps the payoff of all players the same.

Definition 10. Simple search game G and generalized search game $\tilde{G}$ with the same set of players, the same capacities, and the same duplication factor are *equivalent* if there exists a bijection $f : S\left(\tilde{G}\right) \to S(G)$, 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.

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, which implies that all our results can be extended to generalized search games with weakly deterministic signals. Formally:

Claim 2. Let $\tilde{G}$ be a generalized search game with weakly deterministic signals. Then there exists an equivalent simple search game $G = (N, \Omega, v, \mu, \Pi, c, \rho)$. Moreover, if $\tilde{G}$ is consistent or balanced, then so is G.

Sketch of proof; the straightforward long 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 value of each merged state is defined as the weighted average of the values of 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 consistency and balancedness of the payoffs, are straightforward. The simple process of constructing the equivalent simple game is demonstrated in Figure 7. □

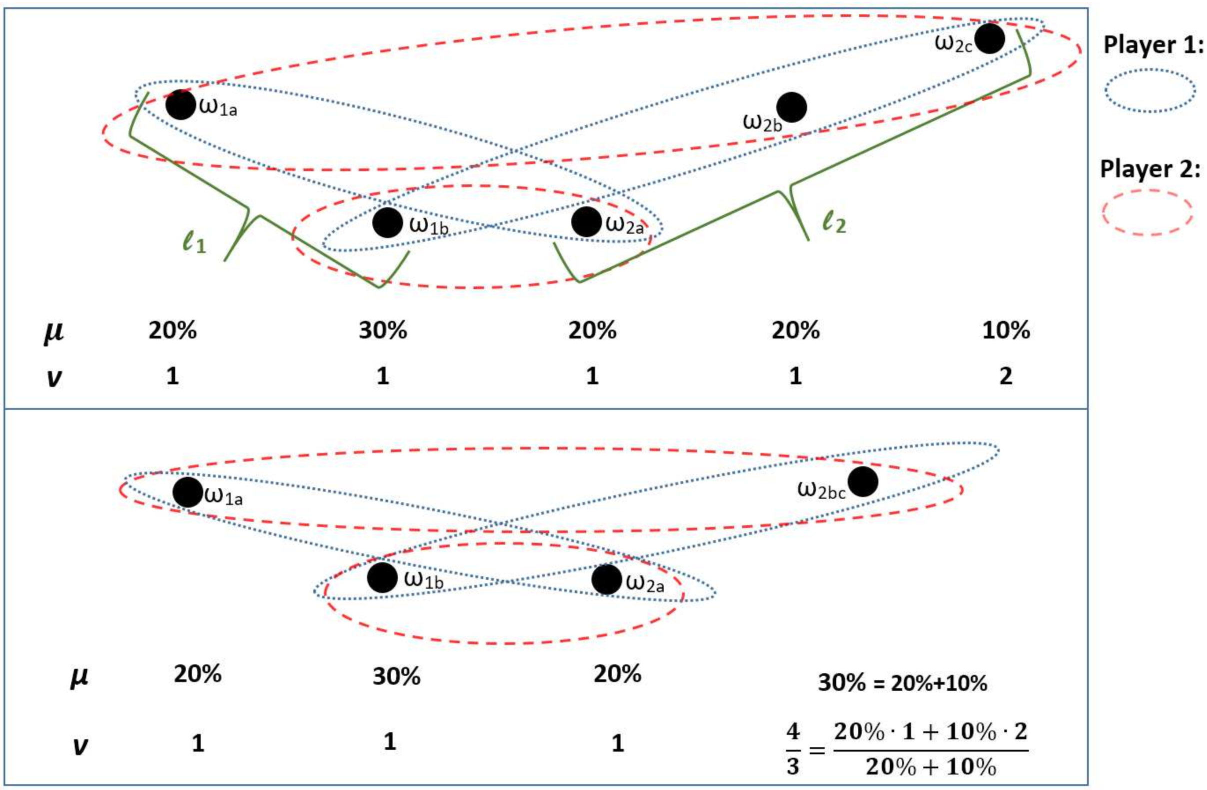

Figure 7: **Illustration of Claim 2**. The upper panel presents a generalized two-player search game $\tilde{G}$ with weakly deterministic signals. The lower panel describes the equivalent simple game G in which (1) states ω_{2b} and ω_{2c} are merged to ω_{2bc}, and (2) each location has been divided into singletons.

5.3 Counterexamples without Weakly Deterministic Signals

Figure 8 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 Proposition 4 (which implies a failure of Theorem 2).

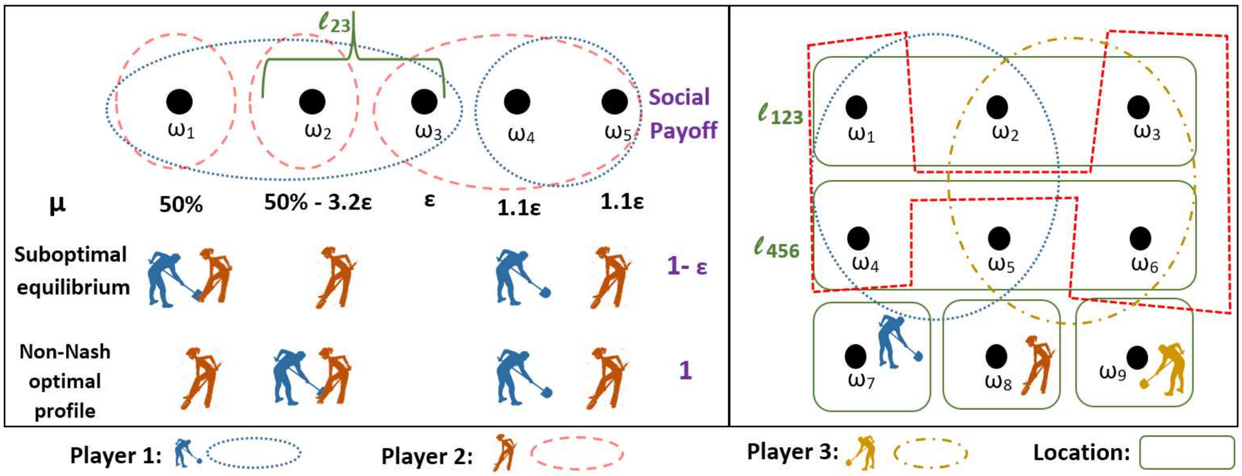

Figure 8: **The left panel presents a counterexample to Theorem 1** (without weakly deterministic signals). It shows the essentially unique suboptimal equilibrium and the non-Nash socially optimal profile in a generalized two-player search game with 1.1-consistent and balanced payoffs and $c \equiv 1$. **The right panel presents a counterexample to Proposition 4**. It shows a three-player generalized search game with $c \equiv 1$ 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 ℓ_{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 of the world.

5.4 Non-monotone Value of Information

Our final example (illustrated in Figure 9) demonstrates that in generalized search games the maximal equilibrium social payoff might be non-monotone with respect to refinement of the information of the players. Specifically, it presents a search game in which there is an exhaustive equilibrium if the players' private signal is either non-informative or fully-informative, while the equilibrium payoff is strictly smaller with intermediate informativeness.

Example 8. Let $(N = \{1,2\}, \Omega = \{\omega_1,\omega_2,\omega_3,\omega_4\}, v \equiv 1, \mu, \Pi, L = \{\{\omega_1,\omega_2\},\{\omega_3,\omega_4\}\}, c \equiv 1,$ $\rho = 1)$ be a generalized search game, where states $\ell_{12} \equiv \{\omega_1,\omega_2\}$ correspond to one prize's location, and states $\ell_{34} = \{\omega_3, \omega_4\}$ correspond to another prize's location. The common prior distribution 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 $\epsilon = 10\%$. 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 the true prize's location). Finally, consider the case of symmetric partially-informative signals that induce the symmetric partitions $\Pi \equiv \{\{\omega_1, \omega_4\}, \{\omega_2, \omega_3\}\}$. Observe that in this case (in which balancedness is violated) the unique equilibrium is both players searching in ω_3 (resp., ω_1) after observing the signal $\{\omega_2, \omega_3\}$ (resp., $\{\omega_1, \omega_4\}$), which implies that the prize is not found when the state of the world is either ω_1 or ω_3.

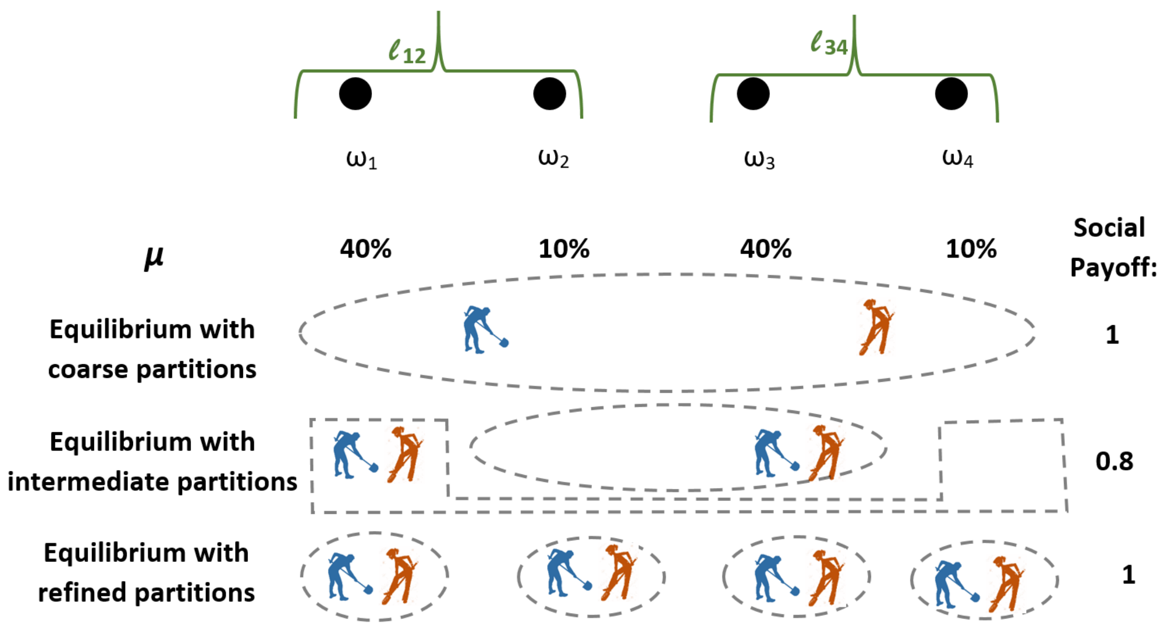

Figure 9: **Illustration of Example 8**. Non-monotone value of information in a symmetric two-player generalized search game.

A Application: Games with Intersecting Signals

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

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

Therefore, k-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 k-intersecting signals seems plausible (especially for $k = 1$) 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 k-intersecting signals and capacities of at most k have appealing efficiency properties. Namely, such games admit a redundancy-free strategy profile iff $\sum_{i\in N} c_i \cdot |\Pi_i| \leq |\Omega|$, and they admit an exhaustive strategy profile iff $\sum_{i\in N} c_i \cdot |\Pi_i| \geq |\Omega|$. Moreover, this strategy profile is an equilibrium if the payoffs are balanced. Formally:

Proposition 5. *Let G be a search game with capacity $c_i \leq k$ for every player $i \in N$ and with k-intersecting signals (resp., and with balanced payoffs). Then the game G admits*

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

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 k-intersecting, each cell of each player must contain at least $k \cdot 2^{|M|-1}$ elements. Since this number is at least $k \cdot |M|$, Corollary 4 implies that the smaller game admits a redundancy-free strategy profile. If $\sum_{i\in N} c_i \cdot |\Pi_i| \leq |\Omega|$, 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 $\sum_{i\in N} c_i \cdot |\Pi_i| \geq |\Omega|$, 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)

Player i has c_i units of capacity, which we index by numbers between 1 and c_i. A *cell-unit* of player i is a pair (π, j) where $\pi \in \Pi_i$ is a cell of player i, and $1 \leq j \leq c_i$ is a unit index. A strategy of i chooses a location for every cell-unit of i. Let us think of player i as being composed of many "smaller" decision makers, one decision maker for every cell-unit of hers. The payoff g_α of cell-unit α of player i is the (ex-ante) expected payoff that α contributes to the overall expected payoff of player i; i.e., if α chooses location ω and overall there are m cell-units choosing ω, then $g_\alpha(\omega|m)$ equals $\mu_i(\omega) \cdot v_i(\omega)$ if $m = 1$ and equals $\rho \cdot \mu_i(\omega) \cdot v_i(\omega) / m$ if $m > 1$.

Given a strategy profile s, suppose that a single cell-unit α of player i changes its choice from ω to ω'. If such a change improves the payoff of α then this is of course also an improvement for player i. Conversely, suppose that some change of strategy for player i is an improvement for her. Since her expected payoff is the sum of the expected payoffs within each cell, there must be at least one cell $\pi \in \Pi_i$ such that the expected payoff within π has improved by this change. The improvement within π consists of switching from choosing one subset of π (of size at most c_i) to choosing another subset, and there must be at least one unit index j such that the payoff of the cell-unit (π, j) has improved by switching from some ω to another ω' that was not chosen by player i under s. Hence, the cell-unit (π, j) can improve its payoff by itself in the profile s.

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

Proof. Suppose that α is not best-responding in s; otherwise we are done. Let α move from its current choice ω_0 to another location that is a best-response for α. The new strategy profile is s^2. Now we add a dummy player at ω_0, and Phase I begins: at every stage of Phase I, one member of $B \cup \{\alpha\}$ who is currently not best-responding moves to a best-response location. This continues as long as there are such non-best-responding members, unless someone moves into ω_0, in which case Phase I immediately terminates.

We claim that under any strategy profile t encountered during Phase I, it is the case that:
(a) there exists exactly one location ω that is chosen by one more cell-unit than under s, i.e., $m_t(\omega) = m_s(\omega) + 1$, while for every other location ω', $m_t(\omega') = m_s(\omega')$ (we call ω "the plus location"); and
(b) for any member of $B \cup \{\alpha\}$ whose current location is some ω and who can also choose another location ω', if there were $m_s(\omega)$ cell-units at ω (including itself) and $m_s(\omega')$ cell-units at ω', then this member would have no incentive to move from ω to ω'.

When Phase I starts, in s^2, (a) holds and α has just moved to the plus location. (b) also holds because members of B were best-responding under s, and α is currently best-responding when α's current location is the plus location, let alone when it is not the plus location. The claim is proved by induction from one stage to the next: suppose that cell-unit β improves on stage τ by moving from ω to ω'. Since β could improve on stage τ, (b) implies that ω must have been the plus location in that stage. Therefore, the plus will move with β from ω to ω', and hence (a) will still hold in stage $\tau+1$. Property (b) will still hold in stage $\tau+1$, since β will be best-responding at ω' although ω' will be the plus location by then, and the situation will not change for the other cell-units, as they did not move.

By our definition of strategies, a cell-unit cannot move to a location occupied by another cell-unit of the same player. Thus, we should note that since β moves from ω to ω', there may be another cell-unit γ of the same player that will be able to move to ω in stage $\tau+1$. Nevertheless, by (b) we know that γ has no incentive to move to ω' in stage τ, which would yield γ the same payoff $g_\beta(\omega'|m_t(\omega')+1)$ that β will receive in stage $\tau+1$, which is more (since the move is an improvement) than the payoff of β in stage τ, $g_\beta(\omega|m_t(\omega))$, which is what γ would receive by moving to ω on stage $\tau+1$. Hence, γ will have no incentive to move to ω in stage $\tau+1$, and (b) will not be violated, and the claim is proven.

Suppose cell-unit β moves to some ω and thus improves its payoff. Then ω becomes the plus location, and afterwards the payoff of β at ω can never be worse than it is now, while it can be better if the plus is somewhere else. Thus, the payoff of β will never go down to the level before the movement. Hence, Phase I cannot turn into a cycle, and since there are only finitely many strategy profiles, Phase I must eventually end. Also note that ω_0 is never the plus location during Phase I (since Phase I terminates once someone moves to ω_0), and hence nobody moves from ω_0 to another location. Therefore, all the movements are improvements not only in the game with the dummy added at ω_0, but also in the original game.

Denote the strategy profile at the end of Phase I by s^*. Now we remove the dummy from ω_0. Suppose first that Phase I ended because somebody moved into ω_0. Since the dummy has been removed, there is no plus location at all, and (b) implies that every member of $B\cup\{\alpha\}$ is best-responding under s^*, and we are done.

Otherwise, Phase I ended at s^* because nobody had an incentive to move when the dummy was still at ω_0. 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, out of the members of $B\cup\{\alpha\}$ who are currently not best-responding, choose one whose current payoff is minimal, and let it move to a better location. This continues as long as there are non-best-responding members.

Let $\hat{m}_{s^*}(\omega)$ denote the number of cell-units *including the dummy* who chose ω at s^*. We claim that under any strategy profile t encountered during Phase II, it is the case that:
(a') there exists exactly one location ω for which $m_t(\omega)=\hat{m}_{s^*}(\omega)-1$, while for every other

location ω', $m_t(\omega') = \hat{m}_{s^*}(\omega)$ (we call ω "the minus location"); and

(b') for any member of $B \cup \{\alpha\}$ whose current location is ω and who can also choose another location ω', if there were $\hat{m}_{s^*}(\omega)$ cell-units at ω (including itself) and $\hat{m}_{s^*}(\omega')$ cell-units at ω', then this member would have no incentive to move from ω to ω'.

The analysis is almost analogous to that of Phase I. When Phase II starts, in the profile s^*, (a') holds and ω_0 is the minus location. (b') also holds because everyone was best-responding under s^* when the dummy was in place. The claim is proved by induction from one stage to the next: suppose that cell-unit β improves on stage τ by moving from ω to ω'. Improvement implies, by (b'), that ω' must have been the minus location in that stage. Therefore, the minus will move from ω' to ω, and hence (a') will still hold in stage $\tau+1$. Property (b') will also still hold in stage $\tau+1$, since β will be best-responding at ω' when the minus will be somewhere else, let alone when there is no minus anywhere, and the situation does not change for the other cell-units, as they did not move.

Since β moves from ω to ω', there may be another cell-unit γ of the same player that will be able to move to ω in stage $\tau+1$. Nevertheless, the payoff of β at ω in stage τ, $g_\beta(\omega | m_t(\omega))$, does not exceed the payoff of γ in that stage (because a cell-unit that is chosen to move has a minimal payoff), and exactly equals what γ will receive by moving to ω in stage $\tau+1$. Hence, γ will have no incentive to move to ω, (b') will not be violated, and the claim is proven.

Suppose that cell-unit β moves to some ω and thus improves its payoff. Then ω ceases to be the minus location, and afterwards the payoff of β at ω can never be worse than it is now, while it can be better if the minus is at ω. Therefore, the payoff of β will never go down to the level it was at before the movement, and will of course improve when β moves again. Hence, Phase II cannot turn into a cycle, and therefore it must eventually end. When it ends, every member of $B \cup \{\alpha\}$ will be best-responding. □

To prove weak acyclicity, start from any strategy profile s^1. 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 s^K under which every cell-unit is best-responding. These cell-unit improvements are also improvements in the original sense, and s^K is an equilibrium.

B.2 Proof of Proposition 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 5). There are two special vertices, a *source* s and a *sink* t. The other vertices in our network are the states Ω and the cells $\hat{\Pi}$ of the game. There is an edge from s to every $(i;\pi_i) \in \hat{\Pi}$, where $\kappa(s,(i;\pi_i)) = c_i$, and an edge from every state $\omega \in \Omega$ to t, where $\kappa(\omega,t) = f(\omega)$. Also, there is an edge from a cell $(i;\pi_i) \in \hat{\Pi}$ to a state 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 state, as those edges have an infinite flow capacity. Suppose that C includes, out of the edges from states to t, exactly the edges $\{(\omega,t) : \omega \in Q\}$ for some subset of states $Q \subset \Omega$. 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 exist a path from s to t. Hence, the total capacity of C equals

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

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

Therefore, the cut that consists of the 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, $\sum_{\omega \in \pi_i} \tau_i(\pi_i,\omega) = \sum_{\omega \in \pi_i} \varphi((i;\pi_i),\omega) = \varphi(s,(i;\pi_i)) \leq \kappa(s,(i;\pi_i)) = c_i$ (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 must be 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 $\boldsymbol{P}^1$, ...,$\boldsymbol{P}^K$ 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 Theorem 3 that states that any doubly substochastic matrix can be written as a convex combination of subpermutation matrices.[9]

Lemma 2. *Let $\boldsymbol{A}$ be a doubly substochastic matrix. Then there exists a finite set of subpermutation matrices $\boldsymbol{Q}^1$, ...,$\boldsymbol{Q}^K$ 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: (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 $\boldsymbol{B}$. This is illustrated in Figure 10 below. It is immediate that $\boldsymbol{B}$ is a doubly stochastic matrix. By Theorem 3 there exists a finite set of permutation matrices $\boldsymbol{P}^1$, ...,$\boldsymbol{P}^K$ (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 $\boldsymbol{A} = \sum_k w_k \cdot \boldsymbol{Q}^k$. □

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 $i \in N$

[9] 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.

Figure 10: Illustration of How to Construct the Square Matrix $\boldsymbol{B}$

$$\mathbf{B}= \begin{array}{cc|ccc} \multicolumn{2}{c|}{\mathbf{A}} & \multicolumn{3}{c}{\mathbf{D}} \\ \hline A_{11} & A_{12} & 1-\sum A_{1j} & 0 & 0 \\ A_{21} & A_{22} & 0 & 1-\sum A_{2j} & 0 \\ A_{31} & A_{32} & 0 & 0 & 1-\sum A_{3j} \\ \hline 1-\sum A_{i1} & 0 & A_{11} & A_{21} & A_{31} \\ 0 & 1-\sum A_{i2} & A_{12} & A_{22} & A_{32} \\ \hline \multicolumn{2}{c|}{\mathbf{C}} & \multicolumn{3}{c}{\mathbf{A}^{\mathbf{T}}} \end{array}$$

is a player, $j \in \{1, ..., c_i\}$ 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 $j \in \{1, ..., c_i\}$ 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 $\sum_{j=1}^{c_i} \tau'((i, j, \pi_i), \omega) = \tau_i(\pi_i, \omega)$ 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 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., having at least two cells in their partition), and denote $m = |M|$. For any player $i \in N$, state $\omega \in \Omega$, 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 k locations in its intersection. Thus, $\pi_i(\omega)$ contains at least $k \cdot \prod_{j \in N \setminus i} |\Pi_j|$ such intersection points, and

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

Hence, $|\pi_i(\omega)| \geq k \cdot 2^{m-1}$. 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 $k \cdot 2^{m-1} \geq k \cdot m$ and $c_i \leq k$, 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, $\sum_{i \in M} c_i \cdot |\Pi_i|$ 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 c_i locations that have not been chosen by other players yet, as long as there are such.

Case 1: Assume that $\sum_{i \in N} c_i \cdot |\Pi_i| \leq |\Omega|$. 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 $\sum_{i \in N} c_i \cdot |\Pi_i| = |\Omega|$, and that if $\sum_{i \in N} c_i \cdot |\Pi_i| < |\Omega|$ then the game does not admit an exhaustive strategy profile (as the players can search in at most $\sum_{i \in N} c_i \cdot |\Pi_i|$ locations).

Case 2: We are left with the case of $\sum_{i \in N} c_i \cdot |\Pi_i| > |\Omega|$ 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 balanced then, either when s is redundancy-free or when s is exhaustive, Corollary 2 implies that G also admits such an equilibrium.

References

Adamczyk, Sabrina, Bullinger, Angelika C, & Möslein, Kathrin M. 2012. Innovation contests: A review, classification and outlook. *Creativity and Innovation Management*, **21**(4), 335–360.

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.

de Roos, Nicolas, Matros, Alexander, Smirnov, Vladimir, & Wait, Andrew. 2018. Shipwrecks and treasure hunters. *Journal of Economic Dynamics and Control*, **90**, 259–283.

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.

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.*

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.*

Loch, Christoph H, Terwiesch, Christian, & Thomke, Stefan. 2001. Parallel and sequential testing of design alternatives. *Management Science*, **47**(5), 663–678.

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

Matros, Alexander, & Smirnov, Vladimir. 2016. Duplicative search. *Games and Economic Behavior*, **99**, 1–22.

Matros, Alexander, Ponomareva, Natalia, Smirnov, Vladimir, & Wait, Andrew. 2019. Search without looking. *mimeo*.

Mihm, Jürgen, & Schlapp, Jochen. 2019. Sourcing innovation: On feedback in contests. *Management science*, **65**(2), 559–576.

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.

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.