Collective Progress: Dynamics of Exit Waves*

Doruk Cetemen† Can Urgun‡ Leeat Yariv§

July 25, 2022

Abstract. We introduce a framework for studying collective search by teams. Discoveries are correlated over time and governed by a Brownian path, where search speed is jointly controlled. Agents individually choose when to cease search and implement their best discovery. We characterize equilibrium and optimal policies. Search speeds are constant within active alliances and depend on complementarities between members. A drawdown stopping boundary governs each agent's search termination. The consequent exit waves, whereby possibly heterogeneous agents cease search simultaneously, exhibit deterministic sequencing but stochastic timing. We highlight environments with lower than optimal equilibrium speeds and search durations, and different exit waves.

Keywords: Retrospective Search, Optimal Stopping, Collective Action, Exit Waves

JEL codes: C73, D81, D83, O35

*We thank Arjada Bardhi, Alessandro Bonatti, Steven Callander, Sylvain Chassang, Francesc Dilme, Hari Govindan, Faruk Gul, Toomas Hinnosaar, Alessandro Lizzeri, Chiara Margaria, Pietro Ortoleva, Wolfgang Pesendorfer, Heikki Rantakari, the Editor and two anonymous reviewers, as well as seminar and conference audiences, for helpful comments and suggestions. Yariv gratefully acknowledges financial support from the National Science Foundation through grants SES-1629613 and SES-1949381.

†City, University of London; doruk.cetemen@city.ac.uk

‡Princeton University; curgun@princeton.edu

§Princeton University, CEPR, and NBER; lyariv@princeton.edu

1 Introduction

Discoveries are often made by teams. Wuchty, Jones, and Uzzi (2007) trace 19.9 million academic papers and 2.1 million patents over 5 decades; they demonstrate that teams increasingly dominate individuals in the production of knowledge. Advances in motor vehicles, communication devices, and pharmaceuticals frequently take place as joint ventures. Understanding collective progress is therefore vital for the analysis of innovation. What determines the pace of innovation? How does the composition of joint ventures affect outcomes? When does product innovation stop? These questions are at the heart of this paper.

Much of the literature on teamwork has focused on experimentation models, starting from the canonical work of Bolton and Harris (1999) and Keller, Rady, and Cripps (2005). Those models center on teams' efforts to ascertain whether one direction or project is superior to another. Nonetheless, many discovery processes follow a path of search, with numerous alternatives. Building on past discoveries, teams come up with new ones. Furthermore, there is a richness of dynamics in collective efforts not captured in prior models—alliances tend to dissolve over time, with exiting members exploiting knowledge accrued during their collaborations.[1]

This paper offers a new framework for studying collective progress based on a process of search. We identify how the search speed and decisions to terminate search vary with members' characteristics and the synergies in place. We also show that exit waves, where multiple members halt search simultaneously, are an inherent feature of such processes; while their timing is stochastic, their order is deterministic. From a design perspective, our results provide a characterization of the optimal joint venture operations.

Technological developments rarely occur in a vacuum: innovations build on one another, and alliances dissolve with new discoveries over time. For example, when developing new car chassis for improved fuel efficiency, the Partnership for a New Generation of Vehicles, which was formed in 1993 and comprised US government agencies and car manufacturers, followed a pre-prescribed path of experimentation: each step in development relying on previous insights. When the partnership dissolved in 2001, automakers used the accumulated know-how to each produce new car models. Similarly, NGO and university alliances are common for program advancement in the developing world. Experimentation pertaining to social programs frequently follows a path—e.g., altering reminders or the modes by which they are provided—and prior conclusions serve as stepping stones for new discoveries. Furthermore, at any point, NGOs or researchers can weaken their

[1] For instance, Eftekhari and Timmermans (2021) use the Danish Integrated Database for Labor Market Research to document joint ventures' dissolution. They report over 18.3 percent of original joint venture members shifting to a new, smaller, joint venture, while other members cease efforts.

involvement and issue "products" independently: implement a policy, write an academic paper, etc. Such examples abound.

In our model, search results are correlated over time and follow a Brownian path as first modeled by Callander (2011). This modeling approach captures the idea that future discoveries build on current ones. Its axiomatic foundations in the innovation context appear in Jovanovic and Rob (1990). The realized path represents the underlying "truth"—e.g., the link between car chassis composition and vehicle stability, the link between number of text reminders and adherence to a particular policy, and so on. Traversing the realized path provides information on the most promising discovery.

The search speed—the per unit of time volume of innovations attempted, or the distance traversed on the realized path of discoveries—is chosen at each moment by the searching alliance. Each alliance member incurs a strictly positive cost that depends on her chosen search speed, her current investment in the search. Individual search speeds are aggregated to generate the alliance's overall search speed. The aggregation format captures the synergies in place, allowing for both substitutability and complementarity between members' search investments.

Any member can terminate her search at any point. A member ceasing her search receives a lump-sum payoff corresponding to the maximal value the search has produced till her departure. For example, when automakers left their 1993 partnership, they each retained the ability to use the most promising technologies and expertise developed jointly.[2] As in this example, in many applications, departing members make investments to act upon their discoveries. Adjustments due to others' later innovations are costly in terms of money, time, or legal constraints when patent protections are in place. For simplicity, we assume benefits are reaped only from discoveries made during members' active search. Certainly, some alliance members may choose to continue their search even after other members have exited. These remaining members experience prolonged search costs, but benefit from any further breakthroughs, as reflected by search results that exceed the most promising discoveries previously observed. As search progresses, members gradually terminate their search until it halts altogether.[3]

We characterize equilibrium search in Markov strategies, where state variables correspond to the current search results, the attained maximum, and the active alliance.

In any active alliance, we show that individual and aggregated alliance search speeds are constant and independent of search results as long as no member leaves. When indi-

[2]General Motors developed the 80 MPG Precept, Ford designed the 72 MPG Prodigy, and Chrysler built the 72 MPG ESX-3. They utilized the jointly-developed technology and featured similar construction and performance; for details, see the US Department of Energy report from May 15, 2000.

[3]Most of our qualitative results carry over when introducing penalties for later exits; see our conclusions and Online Appendix. Naturally, penalties for later exits can induce exit waves mechanically—once one agent departs, others may follow suit to avoid arriving second to market.

vidual search speeds are substitutable, they increase when members depart, reflecting the more limited free-riding opportunities present.

Product development speed has been a major focus of study in operation management (see the meta-studies by Chen, Damanpour, and Reilly, 2010 and Cankurtaran, Langerak, and Griffin, 2013). That literature inspects the link between speed and outcomes. Our analysis sheds new light on the parameters affecting equilibrium speeds: actively searching members' investment costs and the complementarities between them. For any fixed searching alliance, there is a positive relation between speed and the value of obtained discoveries. However, equilibrium speeds adjust as members depart and alliance speeds may decline over time. Since further discoveries tend to improve product values, naively considering correlations between speeds and outcomes may yield misleading results.

The equilibrium time at which members depart and alliances shrink is governed by a simple stopping boundary, often referred to as a drawdown stopping boundary. Such boundaries are defined by one number, the drawdown size. Whenever search results fall by more than the drawdown size relative to the maximal observation achieved, a subset of members ceases search.

Our equilibrium characterization allows us to identify members' exit times. For a large class of speed aggregation formats, members exhibiting high ratios of marginal to fixed costs leave early. Even when individual costs are fully heterogeneous, clustered exits, or exit waves, may occur in equilibrium. Importantly, while the precise timing of exit waves may depend on the realized path of discoveries, we show that their sequencing does not. That is, who leaves first, second, etc., and with whom is deterministic.

The synergies between alliance members, the complementarities in their speed investments, govern both resulting equilibrium search speeds and drawdown sizes. Greater complementarities lead individual speed choices to converge and can cause overall alliance speeds to decline. They are also associated with more incremental exits, whereby agents exhibiting higher search costs depart before their lower-cost partners.

Our equilibrium characterization has clear empirical implications. It suggests the possibility of estimating costs and synergies from observed exit times and project valuations—say, revenues in startup companies. As we show, alliance size is an important statistic to control for, but historical performance is not.

Beyond its substantive implications, our characterization offers a technical contribution. As we detail in our literature review below, existent analyses of single-agent search processes with correlated observations often resort to modeling short-lived agents, absent any controls. In contrast, we analyze the evolution of collective search by forward-looking and sophisticated agents who can utilize a costly control—the search speed.

We view correlation as an important feature of discovery processes. Nonetheless, it is

useful to contrast our results with those derived from settings with independent observations, in the spirit of McCall (1970) and Mortensen (1970). One can consider a discrete-time model in which, at each period, the active alliance draws an independent value from, say, a normal distribution with expectation or variance that depend on members' investments. In such settings, agents depart whenever a sufficiently *high* value is realized, when the immediate value surpasses the option value of waiting. With correlated discoveries, agents depart when observing sufficiently *low* values: a disappointing discovery indicates that far more research is required to obtain a breakthrough. Furthermore, contrasting our setting, with independent samples, the order of exits is *stochastic*. Regardless of the cost profile, whenever a sufficiently high observation is realized, all agents stop their search at once. For moderate realizations, only a subset of agents may terminate search. As it turns out, an analogous model to ours with independent discoveries is far less tractable. We include details in the Online Appendix.[4]

In the last part of the paper, we characterize the socially optimal search speeds and stopping policies. The socially optimal search speeds are also constant and independent of search results within any active alliance. With substitutable individual speeds, the positive externalities induced by each member's investment in search speed imply that the socially optimal level is higher than that chosen in equilibrium. Furthermore, in contrast to equilibrium search speeds, as alliance members terminate their search, the optimal speed of those remaining declines. Optimal exits are governed by drawdown stopping boundaries, although the drawdown sizes corresponding to each active alliance differ from those determined in equilibrium—optimal drawdown sizes are larger, corresponding to longer search durations.[5] In terms of exit waves, clustered exits may be optimal even when individuals incur fully heterogeneous costs. As in equilibrium, the sequence of optimal exit waves is deterministic and independent of the realized search path. However, optimal exit waves may differ substantially from those induced in equilibrium.

Finding the optimal sequence of exit waves is a challenging combinatorial problem. A social planner needs to consider all possible ordered partitions of the original searching team and assess search outcomes from the corresponding exit wave sequences. The number of these ordered partitions grows exponentially fast, posing a computational challenge. We show a simple method for identifying the optimal sequencing for one class of settings, when individual search costs are proportional to one another. Similar to equilibrium, the

[4]Discrete time is inherent with independent discoveries. With a continuum of independent observations, extremely high draws occur within any infinitesimal period, and stopping is immediate. If we discretize our setting, exit waves might exhibit some stochasticity, but their pattern converges to the one we characterize as time intervals between observations shrink.

[5]As we discuss in the paper's last section, allowing for non-Markovian equilibria does not eliminate the inefficiencies we highlight. Intuitively, excessively early search termination is impossible to punish.

social planner terminates the search of those with the highest search costs first. This limits the exit wave sequences to consider. We illustrate a simple procedure, akin to a greedy algorithm (see, e.g., Papadimitriou and Steiglitz, 1998) that yields the optimal exit wave sequence. In rough terms, the social planner can use a recursive procedure, first identifying the optimal last alliance to search—the alliance that would generate the highest welfare when all members are constrained to stop jointly. Once that alliance is identified, the social planner can find the optimal penultimate alliance. And so on. The procedure allows us to highlight settings in which equilibrium exit waves differ substantially from those set optimally.

2 Literature Review

Since Weitzman (1979), much of the search with recall literature has focused on individual agents' discovery process, where the set of options is independent of one another. Our consideration of a Brownian path of discoveries, capturing intertemporal correlations, is inspired by the setting of Callander (2011). He studies short-lived agents who decide whether to choose an optimal, previously explored, result or experiment on their own. Most of the ensuing work considers short-lived agents as well. Callander and Hummel (2014) study long-run experimentation by a sequence of policymakers who operate over a two-period horizon. They show that preemption motives induce policymakers to experiment more than they would in isolation. Urgun and Yariv (2021) analyze an individual-search setting similar to the one analyzed here, where agents are long-lived. See also Décamps, Gensbittel, and Mariotti (2021) and Wong (2021). The current paper provides a full characterization of collective search by forward-looking and sophisticated agents who can utilize a costly control.

In recent years, substantial attention has been dedicated to the study of collective experimentation. Much of this literature focuses on learning spillovers between team members. For instance, the classic papers of Bolton and Harris (1999), Keller et al. (2005) extend the two-armed bandit problem to a team setting, where agents learn from others. Information is a public good. Thus, there is a free-rider problem that discourages experimentation. Nonetheless, there may also be an encouragement effect through the prospect of others' future experimentation. See Hörner and Skrzypacz (2016) for a survey.

Another strand of literature inspects settings in which stopping is determined collectively. Albrecht, Anderson, and Vroman (2010) and Strulovici (2010) consider sequential search and experimentation, respectively, where a committee votes on when to stop. They illustrate when collective dynamics impede on search or experimentation. Bonatti and Rantakari (2016) offer a model in which agents exert effort on different projects but stop

experimentation jointly. Optimally, one agent advances her preferred project quickly. Her opponent agrees to early advanced projects in order to limit effort. Deb, Kuvalekar, and Lipnowski (2020) take a design perspective—for a given deadline at which a project has to be chosen, the principal commits to a selection rule. Titova (2021) studies a public-good setting in which a team decides whether to implement a public good. Payoffs are revealed through a Pandora's box problem à la Weitzman (1979). Optimal information and projects are selected, but free-riding may generate inefficient delays.[6]

There are also several papers illustrating patterns reminiscent of the clustered exits we characterize, mostly in settings in which agents have private information. Bulow and Klemperer (1994) consider a seller who dynamically reduces the price of identical goods until demand meets supply. Agents have independent valuations and decide if and when to buy. In equilibrium, frenzies, where multiple agents buy at the same price, may occur. Caplin and Leahy (1994) study a three-period irreversible-investment game in which each firm receives private information on the aggregate state of the economy and observes others' prior decisions. Firms' actions reveal information and can generate a wave. Gul and Lundholm (1995) analyze a model with two agents who predict the value of a project using private information. Each decides when to issue a prediction, where delay entails a flow cost. The timing of decisions is then informative, and clustered predictions occur in equilibrium. Rosenberg, Solan, and Vieille (2007) study a multi-agent version of the standard real-options problem (see Dixit and Pindyck, 1994). Agents observe private signals about common returns to a risky project, as well as the actions of others. If one agent switches to a safe project—namely, exercises an option—this can lead others to immediately switch to the safe project as well. See also Murto and Välimäki (2011) and Anderson, Smith, and Park (2017). In a static information-collection setting, Bardhi and Bobkova (2021) characterize optimal subsets, or mini-publics, to be activated.[7]

The techniques we develop relate to the applied mathematics literature on optimal stopping, see Azéma and Yor (1979) and Peskir and Shiryaev (2006) for particularly relevant sources.

3 A Model of Collective Search

Consider a team of N agents—product developers, policymakers, academic researchers, etc.—searching through a terrain of ideas in continuous time. Time is indexed by t and

[6] Dynamic contribution games without experimentation or uncertainty have been studied by, e.g., Admati and Perry (1991), Marx and Matthews (2000), Yildirim (2006), and Cetemen, Hwang, and Kaya (2020).

[7] There is also a literature that tries to explain industry "shakeouts," corresponding to times at which firm numbers plummet, absent a decline in output. For example, Jovanovic and MacDonald (1994) suggest shakeouts result from exogenous technological shocks. Initially, firms enter new profitable markets. When there is a technological shock, some firms become more productive than others, potentially leading to clustered exits.

runs through $[0,\infty)$. Each seeks good outcomes and ultimately benefits from the maximal value they have found when they stop their search. We assume all agents are risk neutral.

We model the progress of discoveries using a Wiener process, where the realized sample path describes the link between new technologies and their expected value to each of the participating agents. We assume there is a natural progression of exploration. For example, in technological development, incremental increases in the number of transistors on microchips or number of pixels in digital cameras affect the plausibility of new devices. In motor vehicle technology development, there is a natural order of investigation: first, the composition and coating of the chassis might be considered, then different battery formats, followed by their combination. Similarly, in policy development, the order of experimentation is often pre-specified; particular nudges may be considered in sequence, followed by their various bundles. Such examples are ubiquitous across realms, from the development of new food recipes to the academic accumulation of techniques using textbooks in which each chapter builds on the former. Modeling the link between technologies via a Wiener process allows us to capture the correlation between expected values of similar technologies, and the impact of search speed of those who engage in search. Axiomatic foundations in the innovation context go back to Jovanovic and Rob (1990).[8]

Formally, time proxies for the sequence of ordered technologies in our model. For any time t, denote by B_t the standard Brownian motion with $B_0 = 0$. The realized sample path captures the expected value of each (ordered) discovery.

Agents can affect the speed at which the path of discoveries is traversed. In the examples above, the investment of resources—money, lab space, human capital, etc.—affects how rapidly search is conducted. At each time t, when alliance $A \subseteq \{1,...,N\}$ is actively searching, each agent $i \in A$ decides on the speed or intensity of her search $\sigma_{i,t}^A \in [\underline{\sigma}, \overline{\sigma}]$, where $\overline{\sigma} \geq \underline{\sigma} > 0$ and $A \subseteq \{1,...,N\}$. For $i \in A$, any search speed σ comes at a cost of $c_i(\sigma)$, where c_i is twice continuously differentiable, strictly increasing and convex on $[\underline{\sigma}, \overline{\sigma}]$, and $c(\underline{\sigma}) > 0$. The special case of $\underline{\sigma} = \overline{\sigma}$ corresponds to settings in which search speed is not controlled and agents only choose when to stop search.

The individual search speeds determine the alliance's overall search speed. Denote by σ_t^A the vector of individual search speeds $\{\sigma_{i,t}^A\}_{i\in A}$, where entries are ordered via agents' indices. Whenever the alliance A of agents is searching, we let $\hat{\sigma}_t^A = f^A(\sigma_t^A)$, where for all alliances A, the aggregator f^A is compact-valued, bounded away from zero, and strictly increasing and differentiable in each of its arguments. The search speed of the alliance, the distance traversed on the realized path per unit of time, is given by $\left(\hat{\sigma}_t^A\right)^2$. While this

[8] Our assumption that search occurs in continuous time helps with tractability. It also approximates the observation that in many search settings, particularly research endeavors, investigations are inherently incremental. Agents are not aware of the full set of alternatives at the outset. As search progresses, more alternatives enter the menu of possibilities.

notation simplifies our presentation, we will often slightly abuse terminology and refer to $\hat{\sigma}_t^A$, rather than $\left(\hat{\sigma}_t^A\right)^2$, as the alliance's search speed.

Speeding up search is tantamount to the "scaling" of time. We utilize the fact that such scaling is equivalent to a change in the standard deviation of the original Wiener process (see, e.g., Section 8.5 in Øksendal (2003)). That is, when an alliance's speed is given by $\left(\hat{\sigma}_t^A\right)^2$, we can describe the generated values observed at time t—the expected value of the discovery—that we denote by X_t, using the following law of motion:

$$dX_t = \hat{\sigma}_t^A dB_t,$$

with $X_0 = 0$.

In general, holding their contributions fixed, smaller alliances can be associated with lower or higher aggregate speeds than larger alliances—more agents investing can be beneficial, or cause various coordination challenges that hamper the alliance's speed. For comparative statics and examples, we often focus on the special case of a (modified) Constant Elasticity of Substitution (CES) aggregator, where larger alliances generate higher speeds:

$$f^A(\sigma^A) = \left(\sum_{i \in A} w_i (\sigma_i^A)^\rho\right)^{\frac{1}{\rho}},$$

where $\rho \in (-\infty, 0) \cup (0, 1]$ and the weights satisfy $w_i > 0$ for all i and $\sum_{i=1}^n w_i = 1$.[9] This aggregator is useful for inspecting the effects of complementarities among alliance members. As $\rho \to -\infty$, individual choices become perfect complements; when $\rho = 1$, individual choices are perfect substitutes.[10]

Discovery speeds are certainly of first-order importance when it comes to research and development (see Chen et al., 2010 and Cankurtaran et al., 2013). There is, however, an additional interpretation of the choices individuals make in our model. As we described, from an ex-ante perspective, the choice of speed is equivalent to the choice of instantaneous standard deviation of the Wiener process. One could imagine individual choices corresponding to the breadth, or scope, of search. These feed into the operating alliance's search scope. Indeed, investment in development, through acquisition of instruments or expert time, often entails an increase in risk; it either leads to substantial leaps, or to more pronounced losses, naturally translating into a greater variance of outcomes.

[9] We could also define speed aggregators to take the Cobb-Douglas form when $\rho = 0$. These aggregators allow weights to not sum up to 1 in alliances that comprise a subset of agents. We could alternatively assume alliance-specific weights that sum up to 1 in each alliance. In order to ensure that larger alliances have the capacity to generate greater speeds, we could then include a productivity factor that depends on the alliance size and increases at least linearly in it. We use this version for presentation simplicity.

[10] In line with many applications, we assume search termination is irreversible. We note, however, that with CES speed aggregators, departing agents would never benefit from continuing the search in a smaller alliance: the externalities offered by a larger alliance are always beneficial.

We assume the discovery process exhibits no drift: in applications, the mere passage of time rarely improves or worsens search outcomes over standard horizons of research and development. Naturally, one could consider a team that *controls* drift rather than search speeds, which would also translate to the returns of search with recall. The analysis would follow similar lines to those we describe, although with an important loss in tractability.[11] We view endogenous search speeds as natural for most applications, where investments in innovation directly affect how quickly progress is made.

3.1 Payoffs

Each agent is rewarded according to the maximal project value observed up to her stopping time. This assumption reflects the idea that alliance members making use of their search discoveries make production investments that are difficult to alter as new innovations emerge in the market—e.g., car manufacturers may invest in factories tailored to the technologies they aim at utilizing, policymakers set policies in motion, and their academic counterparts write papers based on discoveries they took part in. In addition, patent protections can increase the costs of borrowing innovations occurring after active search has terminated. Formally, let M_t denote the maximum value observed by time t. That is, $M_t = \max_{0\leq r\leq t} X_r$, with $M_0 = X_0 = 0$.

For any aggregate fixed search speed $\hat{\sigma}$, at time t, $\mathbb{E}(M_t) = \hat{\sigma}\sqrt{2t/\pi}$. Thus, the choice of search speed translates directly to the expected returns from search.

When any agent i stops at time τ, her resulting payoff is given by

$$M_\tau - \int_0^\tau c_i(\sigma_{i,t})dt,$$

where $\sigma_{i,t}$ is the timed search speed of individual i, which may depend on the alliances she is active in.[12]

Agents observe one another's search. In particular, whenever agents stop searching, other agents realize their search will continue within a smaller alliance.

3.2 Strategies and Equilibrium

At any time t, the state of the environment is summarized by X_t, M_t, and A_t, where A_t is the active alliance of agents still searching.

A strategy for agent i dictates her chosen search speed over time and her stopping policy. Formally, it is a set of function pairs (σ_i, τ_i). In principle, (σ_i, τ_i) may depend on

[11] Taylor et al. (1975) characterize the maximal value of search with constant drift. The resulting value is far less amenable to further analysis than ours.

[12] In Section 7.1, we discuss an extension in which agents who stop later are penalized. We assume flow costs for the sake of tractability. Discounting introduces novel technical challenges, see Urgun and Yariv (2021).

time, as well as the entire path of observed search values, corresponding maxima, and active alliances. Let $\{\mathcal{F}_t\}$ denote the natural filtration induced by the governing Brownian motion. Agents' strategies are adapted to this filtration.

We restrict attention to Markov strategies. That is, we assume each agent i uses a strategy of the form (σ_i^A, τ_i^A) that depends only on the state variables X_t, M_t, and A_t.[13] Formally, $\sigma_i^A : \mathbb{R}^2 \to [\underline{\sigma}, \bar{\sigma}]$, and τ_i^A is a stopping time adapted to X_t and M_t such that $\Pr(\tau_i^A = t|\mathcal{F}_t) = \Pr(\tau_i^A = t|X_t, M_t)$ for all i. Agent i's resulting stopping time is $\tau_i = \inf\{t \geq 0 : t = \tau_i^A \text{ and } A_t = A\}$.

For any agent, a stopping time without a finite expectation is dominated: it generates unbounded negative expected payoffs. We therefore focus on strategies with finite expected stopping times. We further assume that a continuous stopping boundary determines when each agent stops within any active alliance. Formally, for all i and all alliances A such that $i \in A$, the stopping policy takes the following form:

$$\tau_i^A = \inf\{t \geq 0 : X_t \leq g_i^A(M_t)\},$$

where $g_i^A(\cdot)$ is a continuous function. This formulation implicitly implies that, upon indifference, agents exit the search.[14] Our assumption that stopping boundaries are continuous is without loss of generality as long as any agent is willing to search on her own, which we show in the Online Appendix.

A strategy for an agent i is a collection of stopping times τ_i^A and mappings that indicate the individual speed σ_i^A for each alliance A such that $i \in A$.

For a given profile $\left\{(\sigma_j^A, \tau_j^A)\right\}_{j \neq i}$, agent i's best-response strategy maximizes her expected payoff. It is determined by solving the following problem:

$$\sup_{\tau_i, \{\sigma_{i,t}^A\}_{t=0}^{\tau_i}} \mathbb{E}_{\left\{(\sigma_j^A, \tau_j^A)\right\}_{j \in A \setminus \{i\}}} \left[M_{\tau_i} - \int_0^{\tau_i} c_i\left(\sigma_{i,t}^A\right) dt \right], \text{ where } \tau_i = \inf\{t \geq 0 : t = \tau_i^A \text{ and } A_t = A\}.$$

An *equilibrium* is a profile of Markov strategies satisfying the assumptions above and constituting best responses for all agents.

4 Equilibrium Team Search

In this section, we characterize the outcomes of team search. We describe the equilibrium search speeds and stopping boundaries. We also identify the sequencing of agents' search termination, and the patterns of equilibrium exit waves.

[13] The inefficiencies we highlight do not vanish when considering equilibria in non-Markovian strategies, see our discussion in Section 7.2.

[14] This kind of stopping time τ_i^A is commonly known as an Azéma-Yor stopping time (Azéma and Yor, 1979), with the function g_i^A defining the corresponding stopping boundary.

4.1 Equilibrium Characterization

Given our restriction on agents' strategies, it follows that any alliance A gets smaller at the minimal stopping time of its members. That is, the time τ^A at which the first members of A stop search is given by $\tau^A = \min_{i \in A} \tau_i^A$. Equivalently,

$$\tau^A = \inf\{t \geq 0 : X_t \leq \max_{i \in A} g_i^A(M_t)\}.$$

Since agents use continuous stopping boundaries, we can write

$$\tau^A = \inf\{t \geq 0 : X_t \leq g^A(M_t)\},$$

where $g^A(M_t) \equiv \max_{i \in A} g_i^A(M_t)$ is continuous.

We start by identifying equilibrium search speeds. Individual search speeds depend only on the active alliance and are constant as long as no member departs.

Proposition 1 (Team Search Speed)**.** *For any agent i in an active alliance A, at any point in time, equilibrium search speeds satisfy the following system whenever interior:*

$$\frac{2c_i(\sigma_i^A)}{c_i'(\sigma_i^A)} \frac{\partial f^A(\sigma^A)}{\partial \sigma_i^A} = f^A(\sigma^A).$$

In general, there might be multiple solutions to the system in Proposition 1, some possibly corresponding to less efficient equilibria. With multiple solutions, our team search problem is compounded with a coordination problem. In principle, agents could use publicly observable Markov states—the achieved maximum or the current observation—to coordinate on different solutions. When the system has a unique solution, the proposition implies that equilibrium search speeds are constant. One restriction guaranteeing uniqueness is what we term *regularity*. An environment is *regular* if the Jacobian of the system in Proposition 1 is non-singular. In this system, marginal returns to individual speeds are adjusted by agents' cost to marginal cost ratios. In particular, regularity depends on both speed aggregators and individual costs and holds as long as complementarities in speed aggregation are not "too" strong. Indeed, regularity holds for a large class of settings, as our examples below illustrate.[15]

Corollary 1 (Constant Equilibrium Speed)**.** *In a regular environment, individual search speeds are constant within an alliance. That is, for any $i \in A$, we have $\sigma_i^A(M_t, X_t) = \sigma_i^A$.*

The intuition for Proposition 1 is the following. At almost all times throughout any agent's search, $X_t < M_t$. Thus, almost always, agents cannot affect the magnitude of the

[15] Our analysis is valid even absent the regularity assumption. If there is only a discrete set of solutions and strategies are restricted to be continuous, the constant speed conclusion continues to hold. We note that absent constraints on the environment's fundamentals, speed maximization or Pareto optimality do not always select a unique equilibrium. The Online Appendix provides a class of examples illustrating this point.

maximal value in any infinitesimal interval of time. Since we assume flow search costs, agents then choose their speed to balance the time it takes to reach a transition—either a new maximum, or search termination—and the speed cost, given the speed chosen by other alliance members. The achieved maximum or the current observed value have no bearing on this calculus.

More formally, consider an agent i in an active alliance A. Suppose i believes that all other agents j in the alliance search with speed σ_j^A. When away from agent i's stopping boundary, agent i can contemplate a small interval of time in which she is unlikely to hit her stopping boundary. For that small interval, agent i considers the induced speed of the process: $\left(f^A(\sigma^A)\right)^2$ and the cost she incurs, $c_i(\sigma_i^A)$. Ultimately, the agent aims at minimizing the cost per speed, or the overall cost to traverse any distance on the path, $\frac{c_i(\sigma_i^A)}{(f^A(\sigma^A))^2}$. The identity in the proposition reflects the corresponding first-order condition (naturally, one needs to ensure the solution is indeed a minimizer; otherwise, under regularity, the equilibrium individual speed is not interior).

Importantly, it is the *ratio* of costs to marginal costs that govern equilibrium search speeds. In particular, in our setting, teaming up with agents who have both higher costs and marginal costs can be beneficial in terms of externalities, a point we return to when discussing comparative statics in our setting.

The proposition suggests that speed is constant within any active alliance. However, the search speed adjusts as alliance members depart. This observation may play a role when assessing the relation between search speed and search outcomes. As already noted, there is a mechanical link between search speed and expected maximal values—for any aggregate *fixed* search speed $\hat{\sigma}$, at time t, $\mathbb{E}(M_t) = \hat{\sigma}\sqrt{2t/\pi}$. However, over the full path of search, alliances may change and search speeds adjust. Thus, the link between *average* speed and the observed maximum at any time need not be linear and depends on the searching team's features. This may help explain why empirical studies correlating average speeds and search outcomes yield inconclusive results (see Chen et al., 2010 and Cankurtaran et al., 2013 for relevant meta-studies in the context of product development).

When costs are log-convex, the ratio of costs to marginal costs is monotonic. When speed aggregators take the CES form, the environment is regular, and comparisons of search speeds within various alliances follow directly from the proposition.

Corollary 2 (Search Speed and Alliance Size)**.** *Suppose speed aggregators take the CES form with $\rho > 0$, costs are log-convex, and an interior solution exists for the system specified in Proposition 1. As an alliance shrinks, individual members' search speeds increase, while total search speed decreases. That is, for any $i, j \in A$, we have $\sigma_i^{A\setminus\{j\}} \geq \sigma_i^A$ while $\hat{\sigma}^A > \hat{\sigma}^{A\setminus\{j\}}$.*

The corollary highlights a form of free-riding. When speed aggregators take the CES

form with $\rho > 0$, search speed is substitutable across individuals. The more agents searching, the less each one searches. Since individual search speeds decrease within an alliance, the total search speed in any active alliance is smaller than that which would be generated by the alliance's members searching independently.[16]

The implication of the corollary is that, with substitutability, breakthroughs—new maximum values that exceed the previous maximum by a certain fixed amount—take longer and longer to achieve. This observation is in line with some evidence from industries, products, and firms showing that, over time, research effort rises while research productivity sharply declines, see Bloom, Jones, Van Reenen, and Webb (2020).

While free-riding in teams is a common phenomenon, in our setting, it occurs only in particular settings. In general, with complementarities, agents' individual search speeds can go up or down as alliance members depart. We return to the effects of complementarities in Section 5.1.

We now turn to the characterization of equilibrium stopping boundaries. We show that agents cease their search whenever search results fall by more than a set amount relative to the observed maximum. Consequently, the order in which agents terminate their search is fixed and does not depend on the realized path of search values.

Proposition 2 (Alliance Stopping Boundary)**.** *In a regular environment, there exists an equilibrium such that, for any agent i in any active alliance A,*

$$g_i^A(M) = M - \frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)}.$$

In particular, agent $i \in \arg\min_j \frac{(f^A(\sigma^A))^2}{2c_j(\sigma_j^A)}$ is the first to stop in any alliance A. Furthermore, given equilibrium search speeds, there is a unique equilibrium in which stopping boundaries are weakly undominated.

Stopping boundaries of the form $g(M) = M - d$ are often termed *drawdown stopping boundaries* with *drawdown size* of d. In equilibrium, agents stop whenever the gap between the observed maximum and the current observation exceeds their drawdown size, as identified in the proposition.

To glean some intuition for the structure of the equilibrium stopping boundary, consider some alliance A and suppose all agents believe that other members of the alliance will continue searching indefinitely with search speeds given by Proposition 1. Since the environment is regular, agents use a constant search speed. Each individual agent i's optimization problem then boils down to a solo searcher's optimization, with others' search

[16] Welfare is always lower when individuals search independently. Any agent receives a higher payoff within an alliance than she would on her own. Indeed, any agent can emulate her solo-search policy in an alliance and guarantee at least as high a payoff.

simply affecting the experienced search costs. Namely, the induced cost of implementing search speed σ is $c_i(\sigma - \sum_{j\in A, j\neq i} \sigma_j^A)$. Since agent i's optimization problem is identical when observing X and M, or $X+k$ and $M+k$ for any arbitrary constant k, her stopping boundary must coincide as well and hence takes the form of a drawdown stopping boundary, see Urgun and Yariv (2021) for further details. Denote the corresponding drawdown size by d_i^A. Suppose $d_i^A = \min_{j\in A} d_j^A$. Consider then another iteration of best responses, where all agents use the drawdown stopping boundary calculated as above. Agent i would still be best responding since, from her perspective, others in the alliance would continue searching for as long as she does. Furthermore, while other agents may want to alter their stopping boundary, intuitively, none would want to cease search before agent i since that would contradict their desire to continue searching for at least as long as agent i in the first place.

This line of argument suggests that, given equilibrium search speeds, the stopping boundary of the first agent i to terminate search in any alliance A is determined uniquely when focusing on equilibria in which stopping boundaries are weakly undominated.[17] Multiplicity of equilibria arises from the stopping boundaries of other agents $j \in A$. Indeed, any agent j who stops strictly after agent i is indifferent across all stopping boundaries $g_j^A(\cdot)$ that satisfy $g_j^A(M) < g_i^A(M)$ for all M. Naturally, all such choices of stopping boundaries by agents other than i do not impact when the alliance first loses some of its members, nor the search speed while it is fully active. Consequently, equilibrium outcomes are unique.[18]

When search is over independent samples, as in the classical models of McCall (1970) and Mortensen (1970), agents stop when sufficiently *high* values are realized (see our Online Appendix for details of a model analogous to ours exhibiting independence). In contrast, when discoveries are correlated, as in our setting, low realizations indicate that far more research is needed to accomplish a breakthrough. Agents therefore stop when observing sufficiently *low* realized values. Nonetheless, the optimal policy has a similar threshold flavor captured by the drawdown size.

When search speeds are substitutable, Corollary 2 provides sufficient conditions for individual speeds to increase and overall speeds to decrease as alliances shrink. Proposition 2 implies that the drawdown size corresponding to any individual is increasing in the alliance's search speed and decreasing in her search speed. Therefore, under the conditions of Corollary 2, individual drawdown sizes decrease as alliances shrink. Intuitively, when

[17] The focus on weakly undominated stopping boundaries—given the equilibrium search speeds—allows us to rule out inefficient equilibria that are an artifact of coordination failures, with multiple agents stopping at an earlier time than desired since other alliance members do so.

[18] Our analysis indicates a link to other cooperative solution concepts in the spirit of the core. At any point in time, were active agents free to form any coalition to pursue search, or cease search, the externalities present in our environment would imply a unique outcome corresponding to the equilibrium outcome we identify.

search speed is substitutable across individuals, members departing from an alliance reduce free-riding opportunities and, consequently, the option value of continuing search. This leads agents to be more demanding for their search to continue.

Propositions 1 and 2 suggest an approach for estimating a joint search process' fundamentals, the relevant costs and complementarities. If one assumes a parametric family of costs and speed aggregators, the discoveries at points of exit—say, the revenues generated by members leaving an alliance and moving into production—and the times at which they occur, can allow a researcher to restrict the set of plausible parameters. While it is crucial to observe the composition of active alliances, historical features of the search would not affect such an exercise. See our discussion in Section 5.1 as well.

4.2 Equilibrium Exit Waves

When all agents have the same costs and solutions are interior, equilibrium takes a simple form. Team members choose identical search speeds, as determined by Proposition 1. They also leave in unison—there is only one exit wave. Proposition 2 suggests that joint departures may occur even when individual costs differ.

To see how those happen, consider any active alliance A. Suppose agent i is first to exit: $d_i^A = \min_{j \in A} d_j^A$. Let $Z^1 = \{i\}$. Now consider the alliance $A \setminus Z^1$ resulting from i's departure. For all remaining agents, there is then a new drawdown size that governs the decision to stop search. These new drawdown sizes are $\{d_j^{A \setminus Z^1}\}_{j \in A \setminus Z^1}$. The discrete drop in overall search speed induced by i's departure may imply that $d_j^{A \setminus Z^1} \leq d_i^A$ for some $j \in A \setminus Z^1$. Let Z^2 correspond to all these agents together with agent i. It follows that, as soon as agent i terminates her search, so will all other agents in Z^2. We can continue this process recursively to identify the clustered exits that occur in equilibrium. Their characterization depends only on the magnitudes of the drawdown sizes identified in Proposition 2. In particular, they are identified deterministically. Thus,

Corollary 3 (Equilibrium Exit Waves)**.** *In a regular environment, the order of exits is deterministic, while exit times are stochastic.*

Our description above suggests that one agent leaving may trigger the departure of multiple agents—a form of snow-balling effect. This implies that targeted interventions, subsidizing the search of only particular agents, may impact the entire path of exit waves.

The deterministic order of exits stands in stark contrast to what would occur with independent search observations à la McCall (1970) or Mortensen (1970). Consider a simple setting in which, at every discrete period, agents in an alliance observe an independent sample from a normal distribution whose expectation and variance depend on members' costly investments. Moderately high realizations could lead to different sets of agents

departing than extremely high realizations. Indeed, regardless of the cost profile, for sufficiently high realizations, all members would depart at once. That is, with independent samples, the order of exits is stochastic. See the Online Appendix for details. Discrete time is inherent with independent discoveries. Indeed, with a continuum of independent observations, extremely high draws occur within any infinitesimal period, and stopping is immediate. However, the contrast with our setting is not a pure artifact of the continuous-time setting we study. If we discretize time in our setting, exit waves might exhibit some stochasticity, but their pattern converges to the one we characterize as time intervals between observations shrink; see Whitt (1980) for related approximation results.[19]

5 Equilibrium Features

We now discuss several features of equilibria in our setting. To highlight the effects of complementarities and the structure of exit waves, we assume speed aggregators take the CES form with equal weights. That is, we assume $w_i = \frac{1}{N}$ for all i.

5.1 Impacts of Complementarities and Costs

We start by analyzing how complementarities and cost differences across the agents affect equilibrium features. For simplicity, we assume $N = 2$ and consider a special case of exponential costs: $c_1(\sigma) = e^{\sigma}$ and $c_2(\sigma) = e^{\alpha\sigma}$ with $\alpha > 1$. Thus, agent 2 has higher costs and marginal costs relative to agent 1.

For any level of complementarities, agent 1 selects a higher speed than agent 2, for whom investments are more costly. Panel (a) of Figure 1 displays agents' speed choices for different cost functions of agent 2: $\alpha = 2, 5, 10$. As complementarities increase (ρ decreases), individual choices converge, with agent 1's speed declining and agent 2's speed increasing. Intuitively, when ρ is high, agents' speed investments are substitutes. The high-cost agent 2 can then free-ride on the low-cost agent 1 and the wedge in investments is pronounced. In contrast, when ρ is low, agent 2 cannot effectively free-ride on agent 1; if agent 2 chooses a low speed, agent 1 would experience lower incentives to invest. As ρ becomes unboundedly low, the speed aggregator takes a Leontief form, where the minimum of the agents' individual speeds determines the team speed. Agents then converge to choosing the same individual speed.

[19] If the sample-generating process entails independent, memory-less, and fixed positive jumps, where the observed value might increase discontinuously at random times, we expect—using techniques introduced by Gapeev (2007)—the deterministic order of exits to hold. In innovation processes, such positive jumps can capture discrete breakthroughs. Nonetheless, with random negative jumps, where observed values might randomly decline discontinuously, the order of exits need not be deterministic: sufficiently bad jumps can induce multiple agents to cease search at the same time.

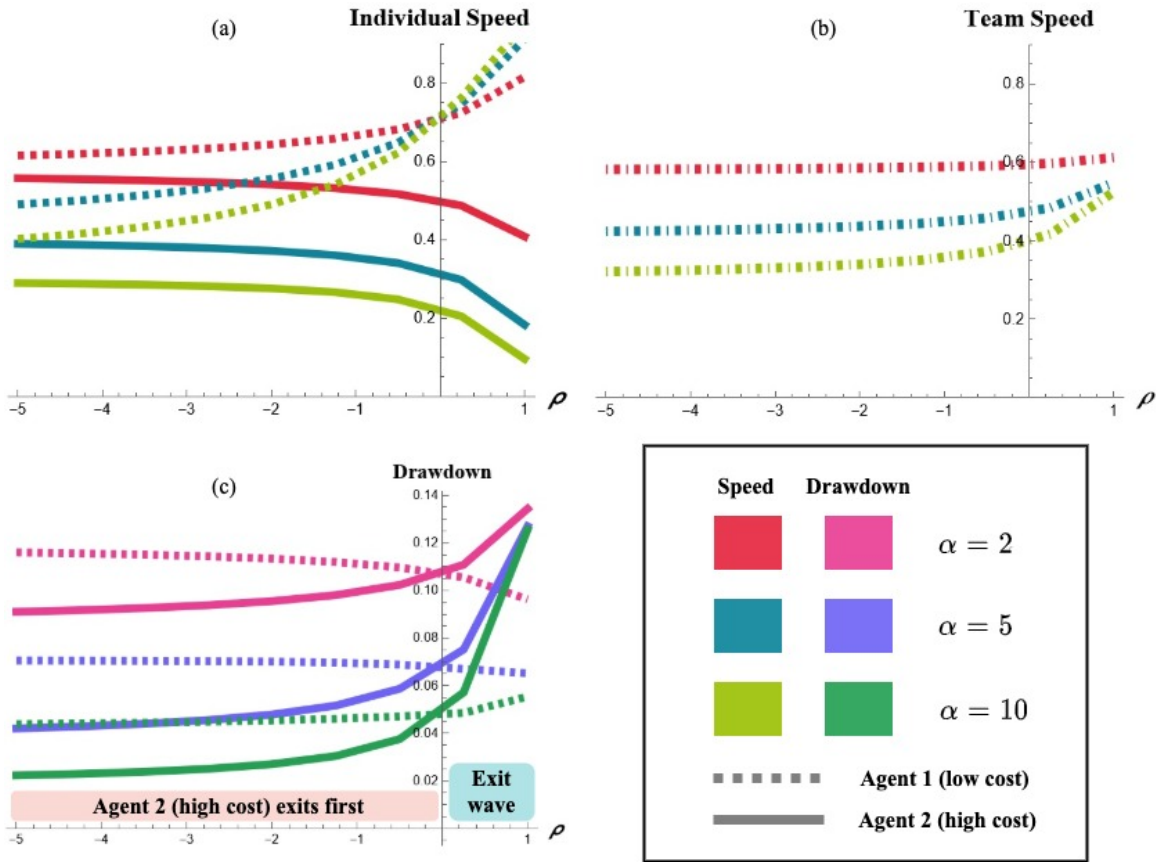

FIGURE 1: Equilibrium comparative statics with exponential costs ($c_1(\sigma) = e^{\sigma}$, $c_2(\sigma) = e^{\alpha\sigma}$)

Panel (a) of Figure 1 also illustrates the impact of increasing the costs of agent 2. As agent 2's costs increase, her chosen speeds decrease. The change in agent 1's choices depends on the complementarities in place. When speeds are substitutes ($\rho > 0$), the reduction in agent 2's speed induces agent 1 to compensate by increasing her search speed. When speeds are complements ($\rho < 0$), the reduction in agent 2's speed disincentivizes agent 1 from investing and leads to a decline in her search speed as well.

Panel (b) of Figure 1 depicts the aggregated team speed. For any level of complementarities, an increase in agent 2's costs is not helpful to the team, and overall speeds decline. Complementarities affect negatively the investment incentives of the high-cost agent 2. Consequently, the team's speed is higher as agents' investments become more substitutable (ρ increases). This observation highlights the role of complementarities in our setting: they imply a form of dependence. When alliance members' search activities feed into one another, one agent speeding up can only take the alliance so far if others are spending very little. Thus, complementarities are a limited "remedy" to free-riding effects

in the presence of substitutabilities. This is reminiscent of observations in static contests, see Kolmar and Rommeswinkel (2013).

Panel (c) of Figure 1 illustrates agents' individual drawdown sizes. As agent 2's cost increases, since the overall speed decreases as seen in panel (b), the net value of continuing search for both agents decreases and drawdown sizes decline in size—the alliance stops its joint search sooner. The comparison of the drawdown sizes of the two agents depends on complementarities. When agents' speeds are complements, agents' speed choices are similar, and the high-cost agent 2 exhibits a lower drawdown size than agent 1. In contrast, when agents' speeds are substitutes, as described in panel (a), agent 2 can free-ride on agent 1's efforts. As a consequence, agent 1 is more keen to stop search and exhibits a lower drawdown size.

Both agents benefit from the presence of another agent in our setting. The solo individual drawdown sizes are therefore higher than those identified in panel (c) of Figure 1. When speeds are substitutes, once agent 1 terminates her search, agent 2 cannot exploit another's efforts and terminates search as well: her solo-individual drawdown size is lower than that of agent 1 in the team. Thus, with $\rho > 0$, there is an exit wave with both agents leaving at once. In contrast, when speeds are complements, once agent 2 terminates her search, the loss to agent 1 is less pronounced and she is willing to continue searching: her solo-individual drawdown size is higher than that of agent 2 in the team. That is, when $\rho < 0$, agents stop their search sequentially.

The mirror image of these comparative statics emerges when considering changes in agent 1's cost relative to agent 2. Overall, increasing the costs of one agent is never helpful to the other when costs are exponential. The effects of changes in the cost parameter α are driven by the fact that increases in α increase costs and marginal costs, but decrease their ratio. In general, a point-wise increase in $c_2(\cdot)$ and $c_2'(\cdot)$ can be accompanied by a decrease or increase in their ratio. From Proposition 1, since that ratio governs agents' speeds in the alliance, and consequently their search duration, a partner with higher costs and marginal costs can, at times, be beneficial. The Online Appendix provides details on this observation.

From an empirical perspective, even with only two agents, observing the exit patterns and the values generated by each agent upon search termination restricts the set of possible parameters. An exit wave occurs only when substitutability is in place; Sequential departures occur only when there are speed complementarities. With sequential departures, the relative values at departure times—say, revenues from products produced—further restrict the set of plausible parameters. These inferences are independent of what transpired prior to the alliance searching on its own.

5.2 Exit Waves with Well-Ordered Costs

In the setting considered in Section 5.1, the low-cost agent never terminates search strictly before the other. However, the precise search speeds and drawdown sizes depend on the profile of costs. We now consider another class of costs, where the identification of exit waves and their comparative statics is particularly simple. Suppose agents' cost functions are proportional to one another: $c = c_1\beta_1 = c_2\beta_2 \cdots = c_N\beta_N$, where $\beta_1 = 1 < \beta_2 < ... < \beta_N$. That is, agent 1 has the highest search costs, while agent N has the lowest search costs. Suppose further that the environment is regular, and speed aggregators are symmetric.[20]

Proposition 1 implies that all agents in an active alliance choose the same search speed, assuming an interior solution exists.

Agents' search speed changes only when their alliance shrinks. In this special case, we can pin down the weak order by which agents stop their search without calculating their corresponding drawdown sizes, which greatly simplifies the analysis. Specifically, Proposition 2 implies that agent N exits no sooner than agent $N-1$, who exits no sooner than agent $N-2$, and so on. In equilibrium, agents with higher costs terminate search earlier. Can non-trivial exit waves occur when agents' costs are strictly ordered?

Consider any active alliance $\{j,\dots,N\}$. If

$$d_j^{\{j,\dots,N\}} \geq d_{j+1}^{\{j+1,\dots,N\}}, d_{j+2}^{\{j+2,\dots,N\}}, d_{j+k}^{\{j+k,\dots,N\}},$$

then agents $j, j+1, j+2, ..., j+k$ will all terminate their search at the same time. Figure 2 depicts an example for $N = 10$ individuals. In the figure, once agent 1 leaves, agents 2 and 3 leave as well. Similarly, once agent 4 leaves, agent 5 leaves. And so on. Ultimately, the drawdown sizes that govern agents' departures correspond to the "upper envelope" of the graph depicting $d_j^{\{j,\dots,N\}}$ as a function of j.

When costs are sufficiently close to one another, all agents exit at once. When costs are sufficiently far from one another, agents exit at different points. A decrease in β_1, keeping $c_1\beta_1$ and all other parameters fixed, increases the agent 1's search costs and leads to her earlier search termination, potentially too soon for other agents to exit. Consequently, the number of exit waves weakly increases. In contrast, a decrease in β_N, keeping $c_N\beta_N$ and all other parameters fixed, increase agent N's search costs, making her more inclined to exit when agent $N-1$ does. Consequently, the number of exit waves weakly decreases.

[20] A function f of m variables is symmetric if, for any permutation $\pi : \{1,...,m\} \to \{1,...,m\}$ and any $(\sigma_1,...\sigma_m)$ in the function's domain, $(\sigma_{\pi(1)},...,\sigma_{\pi(m)})$ is also in the function's domain and $f(\sigma_{\pi(1)},...,\sigma_{\pi(m)}) = f(\sigma_1,...\sigma_m)$.

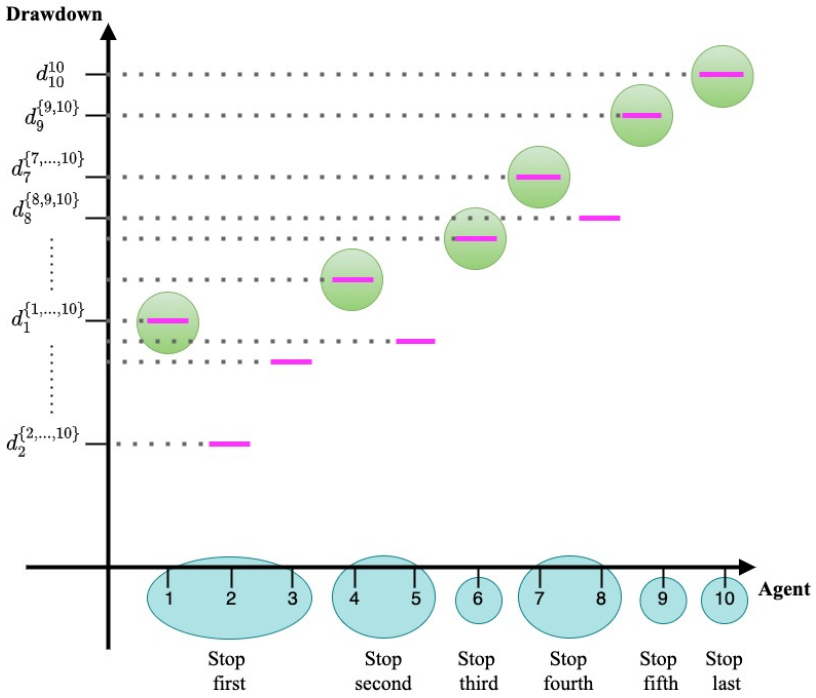

FIGURE 2: Equilibrium exit waves with well-ordered costs

6 The Social Planner's Problem

We now consider a social planner who dictates agents' search speeds and exit policies to maximize overall utilitarian efficiency of the team. This analysis highlights the type of inefficiencies that strategic forces in our joint search process imply.

6.1 The Social Objective

The social planner aims to maximize the agents' expected utilitarian welfare. The instruments at her disposal are the times at which various agents exit—the sequence of active alliances—and the search speeds within each active alliance.

Standard arguments allow us to restrict attention to Markovian policies for the social planner, see Puterman (2014). Formally, we consider a Markov decision problem in which the state at each date t is three-dimensional and comprising (i) the set of active agents A_t, (ii) the current maximum M_t, and (iii) the current observed project value X_t. The social planner chooses a continuation alliance of agents—a subset of the current alliance A_t—and the search speed of each member in that alliance.

As before, σ_t^A denotes the vector of individual search speeds at time t, $\{\sigma_{i,t}^A\}_{i\in A}$, where entries are ordered via agents' indices. Whenever the alliance A of agents is searching, the

alliance's search speed is given by $\hat{\sigma}_t^A = f^A(\sigma_t^A)$.

The social planner has two Markovian controls. The first pertains to the selection of a continuation alliance, and denoted by $G(M, X, A) : \mathbb{R}^2 \times 2^N \mapsto 2^A$. The mapping G determines the subset of agents continuing the search as a function of the current state. In particular, if $G(M, X, A) = A$, the current alliance continues the search. If $\emptyset \neq G(M, X, A) \subsetneq A$, the alliance shrinks in size. Whenever $G(M, X, A) = \emptyset$, no agent is left searching and the search terminates.[21]

The social planner's second control is the profile of search speeds within any alliance A, which can be written as $\sigma_i^A(M, X) : \mathbb{R}^2 \mapsto [\underline{\sigma}, \bar{\sigma}]$ for each $i \in A$. We maintain the constraint that agents that already exited cannot be induced to choose positive search speed and do not participate in any future search: exit is irreversible. As a shorthand, we drop the arguments when there is no risk of confusion.

Given these controls, we can now associate a stopping time for each active alliance A. This is the first time at which the alliance shrinks in size. That is:

$$\tau^A = \inf\{t \geq 0 : G(M_t, X_t, A) \neq A\}. \tag{1}$$

If an alliance A is never reached, we set $\tau^A = 0$.

Let $\tilde{A}_t$ denote the induced process of active alliances. For any active agent i, the time at which her search stops is given by

$$\tau_i = \inf\{t \geq 0 : i \notin G(M_t, X_t, \tilde{A}_t)\}.$$

This is the first time at which agent i is not included in an active alliance.

At any time t, the welfare of individual $i \in \tilde{A}_t$, given the controls $\{G, \sigma_i\}$, is

$$W_i(M_t, X_t, \tilde{A}_t | \sigma_i, G) = \mathbb{E}\left[M_{\tau_i} - \int_t^{\tau_i} c_i(\sigma_i^{\tilde{A}_s}(M_s, X_s) ds\right].$$

For any $i \notin \tilde{A}_t$, we set $W_i(M_t, X_t, \tilde{A}_t | \sigma_i, G) = 0$. The social planner's problem is then:

$$W(M_t, X_t, \tilde{A}_t) = \sup_{\{G, \sigma_i\}} \sum_i W_i(M_t, X_t, \tilde{A}_t | \sigma_i, G).$$

Given a pair of controls (G, σ), with slight abuse of notation, let $A_1 = \{1, ..., N\}$ denote the first active alliance, containing all agents.[22] Using (1), let $A_2 = \tilde{A}_{\tau^{A_1}}$ be the alliance that succeeds A_1, the alliance resulting from the first agents halting their search. In principle, A_2 could entail some randomness—depending on the path observed, different agents may

[21] For simplicity, we restrict attention to deterministic continuation alliances. As our analysis shows, this restriction has no bearing on the social planner's welfare.

[22] We abuse notation by using subscripts to denote the alliance's order in the sequence, rather than time, in order to maintain clarity and simplified notation throughout our analysis.

be induced to exit. We then use (1) to define τ^{A_2}, the (random) time at which the second set of agents stops search and define $A_3 = \tilde{A}_{\tau^{A_2}}$ as the (potentially random) resulting alliance. We continue recursively to establish the (random) time τ^{A_k} at which the k'th set of agents stops search and define $A_{k+1} = \tilde{A}_{\tau^{A_k}}$ as the (potentially random) resulting alliance. Let K denote the (potentially random) number of different active alliances the social planner utilizes till search terminates for all. For any controls $\{G, \sigma_i\}$, we have a sequence of active alliances $A_1, A_2, ..., A_K$ with associated stopping times $\tau^{A_1}, \tau^{A_2}, ..., \tau^{A_K}$.

Suppose our team-search problem starts at the state (M, X, A). We set $\tau^{A_0} = 0$ and $A_{K+1} = \emptyset$ so that the social planner's problem can be written as:

$$W(M, X, A_1) = \sup_{\{G,\sigma_i\}} \mathbb{E}\left[\sum_{k=1}^{K}\left(|A_k \setminus A_{k+1}|M_{\tau^{A_k}} - \int_{\tau^{A_{k-1}}}^{\tau^{A_k}} \sum_{i\in A_k} c_i(\sigma_{i,t}^{A_k})dt\right)\right].$$

Equivalently, we can write the problem recursively starting from any state (M, X, A_k):

$$W(M, X, A_k) = \sup_{\{G,\sigma_i\}} \mathbb{E}\left[|A_k \setminus A_{k+1}|M_{\tau^{A_k}} - \int_{0}^{\tau^{A_k}} \sum_{i\in A_k} c_i(\sigma_{i,t}^{A_k})dt + W(M_{\tau^{A_k}}, X_{\tau^{A_k}}, A_{k+1})\right].$$

Suppose the social planner finds it optimal to halt the search of agent i in an active alliance A when observing X and M. It would then also be optimal to halt the search for this agent when observing X' and M with any $X' < X$. Intuitively, the social planner's solution would be the same were the process shifted by a constant. Therefore, her choice when observing value X' and a maximum M is the same as when observing X and maximum value $M' \equiv M + X - X' > M$. As we soon show, search speeds do not explicitly depend on the achieved maximum. Hence, when observing X and M', were the social planner to continue agent i's search for a small time interval, the optimal search speeds in the active alliance would coincide with those she would pick for the same alliance were search continued when observing X and M. However, the likelihood of surpassing M' in this small time interval is lower than the likelihood of surpassing M. Furthermore, the social planner could gain M' from releasing agent i with the current observed maximum relative to the lower M she would get from releasing that agent when observing X and M. Thus, if it is optimal to halt agent i's search when observing X and M, it is also optimal to halt that agent's search when observing X and M'. We can therefore write

$$\tau^A = \inf\{t \geq 0 : X_t \leq g^A(M_t)\},$$

where $g^A(M) = \sup\{X : G(M, X, A) \neq A\}$.[23]

For any active alliance A, we note that $g^A(M) < M$ for all M. In other words, it is never

[23] We implicitly assume, without loss of generality, that whenever the social planner is indifferent between halting the search of a subset of agents or continuing their search, she chooses the former.

optimal to stop that alliance at any t such that $M_t = X_t$. If an alliance searches for a non-trivial amount of time at its inception, say at time t_0, it must be that $M_{t_0} > X_{t_0}$. The alliance would then continue searching jointly even were the planner to observe, at some time t, the value X_t and recorded maximum of M_t with $X_t = M_t = M_{t_0}$. But then the same should hold when $M_t = X_t = y$, with arbitrary y; this corresponds to a shifted problem and does not alter welfare considerations.[24]

6.2 Optimal Team Search

Our first result illustrates that, as in equilibrium, the social planner chooses constant search speeds for each active alliance. However, these search speeds differ from those dictated by equilibrium.

Proposition 3 (Optimal Search Speed)**.** *Search speeds within an alliance are constant and depend only on the alliance's composition. Furthermore, whenever interior, search speeds satisfy the system:*

$$\frac{2\sum_{i\in A} c_i(\sigma_i^A)}{c_i'(\sigma_i^A)} \frac{\partial f^A(\sigma^A)}{\partial \sigma_i^A} = f^A(\sigma^A).$$

The intuition for this result resembles that provided for equilibrium choices. For any active alliance A, the social planner considers the induced speed of the process, given by $\left(f^A(\sigma^A)\right)^2$ and the cost she incurs, $\sum_{k\in A} c_k(\sigma_k^A)$. The social planner then aims at minimizing the cost per speed, or the overall cost to traverse any distance on the path, $\frac{\sum_{k\in A} c_k(\sigma_k^A)}{\left(f^A(\sigma^A)\right)^2}$. The identity in the proposition reflects the corresponding first-order condition (as in equilibrium, one needs to ensure this condition corresponds to a minimizer; otherwise, the optimal speeds are not interior). As we soon show, the social planner's problem has a unique solution, up to relabeling of agents. In particular, the socially optimal speed is constant even absent our regularity assumption.

When costs are log-convex and speed aggregators take the CES form with substitutes ($\rho > 0$), the proposition implies that socially optimal search speeds are higher than those prescribed in equilibrium. Furthermore, when alliance A is active, each alliance as a whole searches weakly more under the social planner's solution. Intuitively, agents' efforts exhibit two positive externalities. First, a greater search speed contributes positively to other members of the current alliance. Second, increased search efforts improve future alliances' welfare. The social planner internalizes these positive externalities and thus

[24] This would not hold were the social planner's objective concave in the maximum observed. Concavity introduces new challenges, see Urgun and Yariv (2021) for its impact on single-agent decisions. Its investigation would be an interesting direction for the future.

specifies greater overall search investments. Corollary 2 indicates that, as alliances shrink, remaining agents increase their search speed. The impacts of agents departing are quite different in the social planner's solution. As members depart, the externalities of each remaining agent decline: there are fewer others their search speed helps. Consequently, the socially optimal search speed of each individual agent *declines*. That is:

Corollary 4 (Optimal Speed and Alliance Size)**.** *Suppose speed aggregators take the CES form with $\rho > 0$, costs are log-convex, and the social planner's search speeds are interior. Then, in any alliance, an agent's equilibrium search speed is lower than that agent's search speed in the social planner's solution. Furthermore, in the social planner's solution, each agent's search speed decreases as her alliance shrinks in size.*

The sequencing of alliances and their search duration also differ between the social planner's solution and the corresponding equilibrium:

Proposition 4 (Optimal Alliance Sequencing)**.** *The socially optimal sequence of alliances is deterministic, and unique up to agents' relabeling. For any deterministic sequence of alliances $A_1,...,A_k$ exerting optimal search speeds, the socially optimal stopping boundaries are drawdown stopping boundaries. That is, for each alliance A_k, $g^{A_k}(M) = M - d_{A_k}$ with $d_{A_k} \in R_+$. Furthermore, the drawdown sizes $\{d_{A_k}\}$ exhibit a recursive structure: for any k,*

$$d_{A_k} = \frac{|A_k \setminus A_{k+1}|}{2\left(\frac{\sum_{i\in A_k} c_i(\sigma_i^{A_k})}{\left(f^{A_k}(\sigma^{A_k})\right)^2} - \frac{\sum_{i\in A_{k+1}} c_i(\sigma_i^{A_{k+1}})}{\left(f^{A_{k+1}}(\sigma^{A_{k+1}})\right)^2}\right)}.$$

Why does the social planner use drawdown stopping boundaries for various alliances? Intuitively, for any active alliance A_k, the social planner considers the marginal group of agents $A_k \setminus A_{k+1}$ whose search will be terminated next. The relevant marginal added cost per speed for that group is then:

$$\frac{\sum_{i\in A_k} c_i(\sigma_i^{A_k})}{\left(f^{A_k}(\sigma^{A_k})\right)^2} - \frac{\sum_{i\in A_{k+1}} c_i(\sigma_i^{A_{k+1}})}{\left(f^{A_{k+1}}(\sigma^{A_{k+1}})\right)^2}.$$

Each of these agents would receive the established maximum once they depart, thereby generating a multiplier of $|A_k \setminus A_{k+1}|$ of the maximum in the social planner's objective. The resulting stopping boundary then emulates that of a single decision maker, a special case of Proposition 2, with scaled up returns to each maximum established when the alliance shrinks, and adjusted costs as above.

To glean some intuition into the deterministic nature of the sequence of alliances, suppose that the social planner, starting with some active alliance A, proceeds to either alliance A' or alliance A'', depending on the realized path, with $A', A'' \subset A$. Following our

discussion above, both transitions—from A to A' and from A to A''—are associated with a drawdown stopping boundary, with drawdown sizes of d' and d'', respectively. If $d' < d''$, starting from alliance A, the social planner would always shrink the alliance to A' as the relevant stopping boundary would always be reached first. Similarly, if $d'' < d'$, the social planner would always reduce the alliance to A''. In other words, different drawdown stopping boundaries never cross one another, and so the path of alliances is deterministic.

To see how multiplicity might emerge, consider the following simple example. Suppose $N = 3$ and that agents 2 and 3 are identical, with equal cost functions and interchangeable effects on speed aggregators. Suppose further that agent 1's cost is substantially lower than that of agents 2 and 3 and that speed aggregators take much larger values for alliances of two agents than for any other alliance. The social planner would then terminate the search of either agent 2 or agent 3 at the outset of the search—since these agents are identical, which of the two does not matter. In this case, there are clearly multiple solutions, but their resulting speed choices, alliance paths, and payoffs are equivalent up to the relabeling of agents 2 and 3. The proposition shows that this is the only type of multiplicity possible.

Propositions 3 and 4 suggest that the general structure of efficient search is similar to that conducted in equilibrium. Agents depart the search process in a pre-specified order and do so using drawdown stopping boundaries. Furthermore, within each active alliance, search speeds are constant over time. Nonetheless, the optimal sequence of active alliances, their corresponding drawdown sizes, and the search speeds do not generally coincide with those prescribed by equilibrium. In equilibrium, Corollary 3 implies that targeted subsidization of certain agents can dramatically alter the structure of exit waves. Proposition 4 provides guidance on which agents should ideally be subsidized.[25]

Certainly, agents who search exert positive externalities on others searching. Naturally, then, the social planner exploits these externalities by extending the time individuals spend searching. In fact, the expressions derived for the optimal and equilibrium alliance drawdown sizes imply directly the following.

Corollary 5 (Longer Optimal Search)**.** *Suppose speed aggregators take the CES form with $\rho > 0$, costs are log-convex, and the equilibrium and social planner's search speeds are interior. Consider any alliance that is active on path in both the social planner's solution and in equilibrium. Then, the drawdown chosen by the social planner for that alliance is weakly larger than the equilibrium drawdown of the same alliance.*

The results of this section provide some features of the optimal solution.[26] However,

[25] Outside the scope of the current paper, it would be interesting to analyze how a limited budget should be utilized to subsidize searching agents efficiently.

[26] In addition, in the Online Appendix, we show a recursive formulation of the social planner's objective—

they do not offer a general characterization of the optimal sequence of alliances, which is the result of a challenging combinatorial optimization problem—in principle, the planner needs to consider all possible exit patterns, corresponding to ordered partitions of the team. A sharper characterization requires more structure on the environment's fundamentals. In the next subsection, we impose such a structure and solve the social planner's problem completely, illustrating the optimal sequence of alliances and contrasting it with that emerging in equilibrium.

6.3 Optimal Team Search with Well-Ordered Costs

Suppose, as in Section 5.2, that agents' cost functions are proportional to one another and point-wise ordered: $c_1\beta_1 = c_2\beta_2 = \cdots = c_N\beta_N$, where $\beta_1 = 1 < \beta_2 < ... < \beta_N$. Suppose further that all speed aggregators are symmetric and that, for any alliance A, the speed aggregator f^A depends only on the alliance's cardinality $|A|$.[27]

We start by showing that the social planner uses a similar sequencing of active alliances to that used in equilibrium.

Lemma 1 (Optimal and Equilibrium Alliance Sequence)**.** *In the social planner's solution, agent i never terminates search before agent j if $i > j$. In particular, whenever agent i terminates search before agent j in equilibrium, the social planner terminates agent i's search either with, or before, agent j's.*

Intuitively, the social planner optimally terminates the search of agents with the highest search costs first, so agent 1's search is terminated no later than agent 2's search, which is terminated no later than agent 3's, etc. This mimics, "weakly," the order governed by equilibrium. Nonetheless, the social planner's sequencing need not echo that prescribed by equilibrium since clustered exits can differ dramatically, as we soon show.

It will be useful to introduce the following notation for our characterization of the socially optimal sequence of alliances. Let $B_k = \{k, k+1, ..., N\}$ for all $k = 1, ..., N$. Lemma 1 and our equilibrium characterization imply that the optimal sequence of active alliances has to correspond to a subset of $\{B_k\}_{k=1}^N$. This already suggests the computational simplicity well-ordered costs allow. For instance, instead of considering $2^N - 1$ alliances that could conceivably be the last ones active, we need to consider only N.

For $B' \subsetneq B$, we denote by $d_{B \to B'}$ the socially optimal drawdown size associated with alliance B, when it is followed by alliance B', as described in Proposition 4. In particular,

the resulting welfare—in terms of the optimal drawdown sizes and search speeds.

[27] The assumption that speed aggregators depends only on the alliance's cardinality ensures that the social planner does not terminate the search of certain agents only to increase the productivity of those remaining, an externality that would trivially not be internalized in equilibrium.

$d_{B\to\emptyset}$ denotes the optimal drawdown of an alliance B when it is the last active alliance. We now characterize the optimal sequence of alliances.

Proposition 5 (Optimal Alliance Sequence with Well-Ordered Costs)**.** *The optimal sequence of alliances is identified as follows:*

- *There is a unique maximizer of* $\{d_{B_k\to\emptyset}\}_{k=1}^{N}$. *Let* $L_1 = \underset{k=1,\ldots,N}{\arg\max} d_{B_k\to\emptyset}$. *The last active alliance is* B_{L_1}, *with* $L_1 \leq N$. *If* $L_1 = 1$, *all agents optimally terminate their search at the same time. Otherwise,*
- *There is a unique maximizer of* $\{d_{B_k\to B_{L_1}}\}_{k=1}^{L_1-1}$. *Let* $L_2 = \underset{k=1,\ldots,L_1-1}{\arg\max} d_{B_k\to B_{L_1}}$. *The penultimate active alliance is* B_{L_2}, *with* $L_2 < L_1$. *If* $L_2 = 1$, *there are optimally only two active alliances:* B_1 *followed by* B_{L_1}. *Otherwise,*
- *Proceed iteratively until reach* L_n, *where* $L_n = 1$. *The socially optimal order of alliances is given by* $B_1, B_{L_{n-1}}, \ldots, B_{L_1}$.

The optimal sequence of alliances is constructed recursively. Consider first the case in which an alliance's search is terminated jointly. That is, once search terminates for one of the alliance's members, it is terminated for all others. Our analysis in the previous section suggests that, restricted in this way, the social planner would optimally determine the stopping time using a drawdown stopping boundary. Naturally, any possible alliance would be associated with a different optimal drawdown size. Higher drawdown sizes correspond to alliances the planner would prefer to have searching for longer periods. It is therefore natural to suspect that the alliance corresponding to the highest such drawdown size is the last active alliance. Since we already determined that optimal search exits occur in "weak" order, with agent i never exiting after agent $i+1$, it suffices to consider drawdown sizes corresponding to each alliance B_k.[28] This allows us to determine the last active alliance chosen by the social planner, B_{L_1}, as in panel (a) of Figure 3.

Once B_{L_1} is identified, we proceed to the penultimate active alliance. Namely, we consider all plausible super-sets of B_{L_1} and assess drawdown sizes when the social planner is constrained to transition directly to B_{L_1}, see panel (b) of Figure 3. The alliance generating the maximal such drawdown size is the one the planner would want to keep searching the longest, foreseeing her optimal utilization of the next alliance B_{L_1}. That is the penultimate alliance. We continue recursively until reaching the maximal active alliance B_1, see panel (c) of Figure 3.

[28] As mentioned, this simplifies the computation problem substantially. Instead of considering $2^N - 1$ alliances, we need to consider only N.

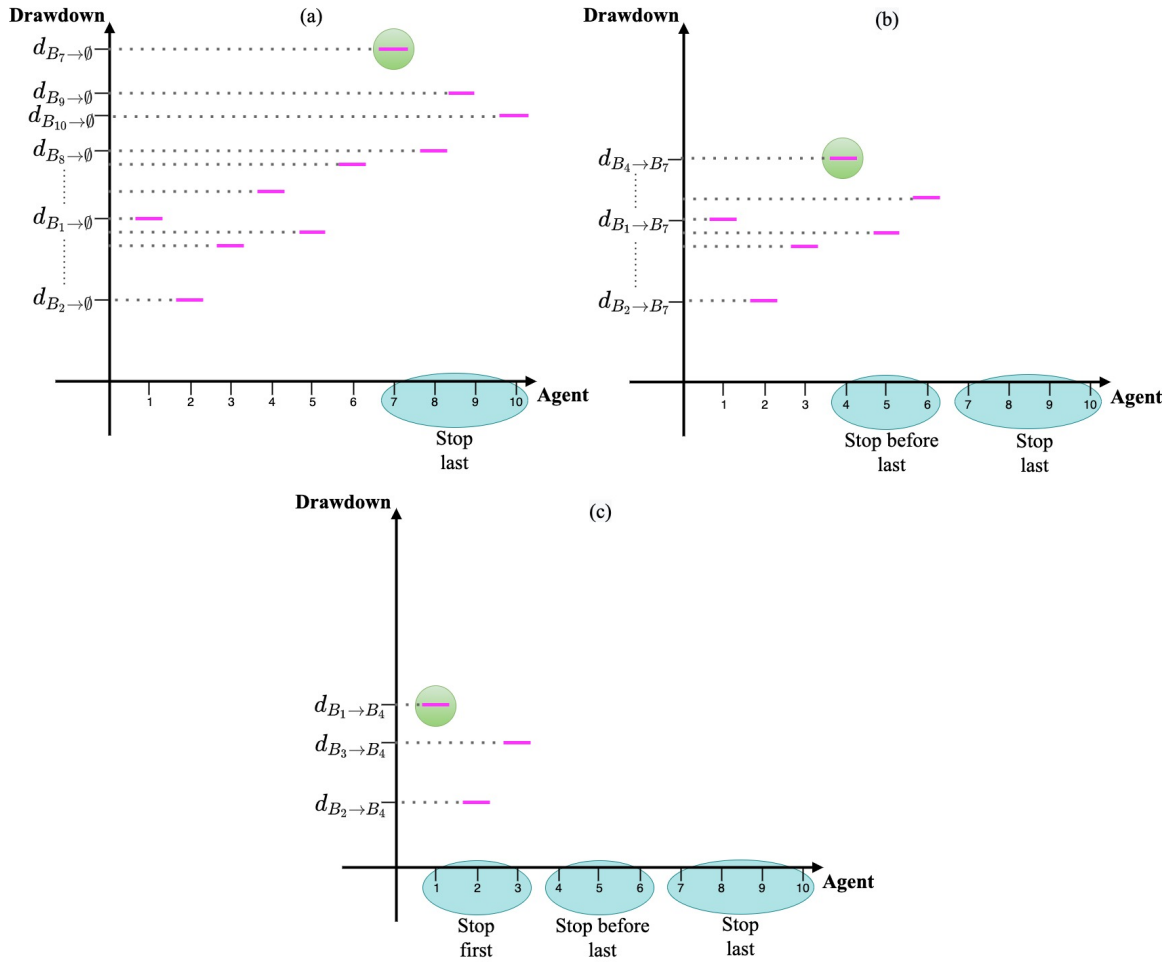

FIGURE 3: Socially optimal exit waves

6.4 Comparing Exit Waves in an Exponential World

In order to contrast the structure of equilibrium and socially optimal exit waves, we now consider a particular example. Suppose the team comprises three agents, $N = 3$, and assume cost functions are exponential and well-ordered: $c(\sigma) = c_1(\sigma) = e^{b\sigma} = \beta_2 c_2(\sigma) = \beta_3 c_3(\sigma)$, where $1 < \beta_2 < \beta_3$. Suppose further that speed aggregators take the CES form, with perfect substitutes ($\rho = 1$) and equal weights ($w_i = \frac{1}{3}$ for all i). There are four possible exit wave structures: all agents can leave at once; agent 1 might leave first, followed by the clustered exit of the lower-cost agents 2 and 3; agents 1 and 2 might leave together, followed by agent 3; or agents may exit at different points.

Figure 4 focuses on parameters under which the social planner clusters all agents' exits and, in equilibrium, agents implement a symmetric speed profile. The figure depicts the different regions of β_2 and β_3 combinations that generate the four possible structures of

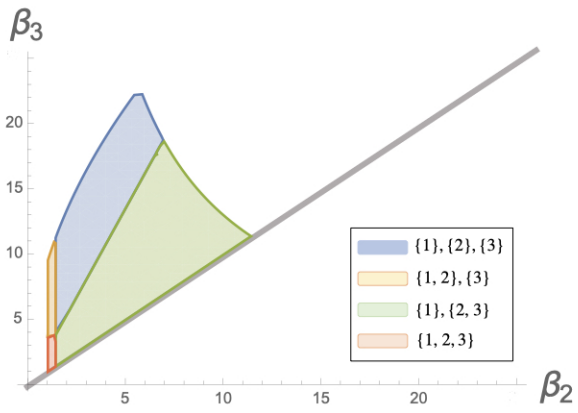

FIGURE 4: Equilibrium exit wave patterns when the optimal policy entails one exit wave including all three agents

equilibrium exit waves. Since $\beta_3 > \beta_2$, all regions are above the gray 45-degree line. We use $\{1,2,3\}$ to denote one clustered exit wave including all agents; $\{1,2\},\{3\}$ to denote an exit wave consisting of agents 1 and 2, followed by the exit of agent 3; and so on.

When the cost multipliers are sufficiently close to one another, agents exit in unison, even in equilibrium. When β_2 is sufficiently close to 1, but β_3 is sufficiently higher, agent 3 has substantially lower search costs. Since agents 1 and 2 do not internalize their externalities on agent 3, they prefer to leave early on, generating two exit waves. Similarly, when β_2 and β_3 are sufficiently high but close to one another, two exit waves occur in equilibrium. Last, when agents' costs are sufficiently different, equilibrium dictates agents exiting at different points, resulting in three exit waves, even when externalities are sufficiently strong so that the social planner would prefer to have the agents search together till they all exit. Naturally, for sufficiently high β_2 and β_3, the wedge in costs is big, and even the social planner would prefer to split agents' exits. The Online Appendix contains detailed characterization of the equilibrium and social planner's solutions, and displays similar figures for other exit-wave structures chosen by the social planner.

7 Conclusions and Discussion

This paper analyzes team search patterns in a setting with long-lived and sophisticated agents. We show that the equilibrium and socially optimal search speeds are constant within an alliance. However, as alliance members depart, individual search speeds increase in equilibrium and decrease under the optimal policy. We also characterize the

deterministic path of exit waves generated in equilibrium. The optimal path of exit waves shares features with the equilibrium path in terms of the structure of stopping boundaries that govern departures. However, search externalities naturally prolong optimal search in teams and alter resulting exit waves.

In what follows, we consider two extensions of our model, allowing for competitive forces that generate explicit rewards for early innovation, and the utilization of non-Markovian equilibrium strategies. In the Online Appendix, we also analyze the limitations introduced by a fixed, non-alterable search speed, and our model's implications for settings with independent search observations.

7.1 Market Competition through Penalties for Later Innovations

Market competition could yield an advantage for those who stop their search early. For example, a firm that produces the first product of its type might capture a market segment that is later more challenging to capture. Similarly, researchers arguably get additional credit for being the first to suggest a modeling framework or a measurement technique.

For simplicity, consider a team of two agents with speed aggregators that take the perfect-substitutes CES form ($\rho = 1$) with equal weights ($w_i = \frac{1}{2}$ for $i = 1, 2$). Assume that the first agent to stop, say at time t, receives M_t. The second agent to stop, say at time $s > t$, receives αM_s, with $\alpha \leq 1$. If both agents stop at the same time t, they both receive M_t.[29] As we show in the Online Appendix, the order of exits remains deterministic. Furthermore, as long as both agents are searching, the search speed and the initial stopping boundary are identical to those in our benchmark setting, where $\alpha = 1$. Thus, if there is a unique exit wave when $\alpha = 1$, that is still the case when $\alpha < 1$.

Suppose there are two distinct exit waves with $\alpha = 1$. Then, there is a *leader*—the agent who exits early—and a *follower*—the agent who exits later. The leader's stopping boundary $g_L(\cdot)$ is governed by the equilibrium drawdown identified in Proposition 2 regardless of α. The follower's stopping boundary, however, may change with α.

To characterize the follower's stopping boundary, denote the costs of the leader by $c_L(\cdot)$ and those of the follower by $c_F(\cdot)$. Let σ_L denote the leader's search speed when searching within the full team, σ_T denote the total search speed in the full team, and σ_F denote the follower's optimal solo search speed. Similar calculations to those underlying Proposition

[29] The analysis naturally extends to N agents via a decreasing sequence of discounts: $\alpha_0 = 1 \geq \alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_N$. In addition, one could consider a continuous version of this setup, where the second agent who stops at time $s > t$ receives $M_t + \alpha(M_s - M_t)$. That model generates qualitatively similar results, but is more cumbersome to analyze.

2 yield the follower's stopping boundary $g_F(\cdot)$:

$$g_F(M) = \begin{cases} M - \frac{\alpha\sigma_F^2}{2c_F(\sigma_F)} \text{ if } M < \overline{M} \text{ and } \frac{\alpha\sigma_F^2}{2c_F(\sigma_F)} > \frac{\sigma_T^2}{2c_L(\sigma_L)}, \\ g_L(M) \text{ otherwise,} \end{cases}$$

where

$$\overline{M} = \frac{1}{1-\alpha}\frac{c_F(\sigma_F)}{\sigma_F^2}\left(\frac{\alpha\sigma_F^2}{2c_F(\sigma_F)} - \frac{\sigma_T^2}{2c_L(\sigma_L)}\right)^2.$$

To glean some intuition, consider the follower's problem after the leader's departure. The follower faces a similar problem to the individual agent's problem, with identical search costs and rewards scaled down by α. This case falls within the analysis of Urgun and Yariv (2021). The search speed is unaffected by the attenuated rewards, but the drawdown size is scaled linearly by α—as α declines, the rewards from search become less meaningful, and the follower ceases search more willingly. Naturally, for sufficiently low α, search continuation would not be worthwhile altogether, regardless of the maximal observation achieved when the leader exits. That corresponds to the drawdown used by the follower alone, $\frac{\alpha\sigma_F^2}{2c_F(\sigma_F)}$, being smaller than the full alliance's drawdown, $\frac{\sigma_T^2}{2c_L(\sigma_L)}$. In that case, the stopping boundary of the leader governs the exit of both. In addition, when the maximal observation M achieved when the leader exits is high enough, the loss from leaving at a later point, $(1-\alpha)M$ is substantial for any $\alpha < 1$.[30] For sufficiently high M, search continuation would again not be profitable. As α increases, the threshold level $\overline{M}$ increases. To summarize, for the follower to continue search after the leader, α needs to be sufficiently high and the current maximum sufficiently small.

Importantly, when later innovations are penalized, there are no preemption motives. The main impact is on later innovators, who face weakened incentives to search. Mechanically, larger exit waves occur for a larger set of parameters. Nonetheless, the main messages of the paper extend directly to such settings.

Certainly, market competition could affect search interactions in a variety of other ways, through entry considerations, dynamic pricing schemes of products, and so on. Incorporating these would be an interesting avenue for future research.

7.2 Non-Markovian Strategies

Our equilibrium analysis restricts attention to Markovian strategies. In our setting, the use of non-Markovian strategies cannot yield the socially optimal solution in general.[31]

[30] Specifically, the gain from continuation for the follower is given by $(1-\alpha)M + (d_L - d_F)^2\frac{c_F}{\sigma_F^2}$, where d_L and d_F are the drawdown sizes for the leader and the follower, respectively.

[31] This contrasts insights on collective experimentation, see Hörner, Klein, and Rady (2021).

To see why, consider a team of two agents and suppose the optimal search speed can be implemented in equilibrium—say, when there is only one viable speed, $\underline{\sigma} = \overline{\sigma}$. Our results show that, in some settings, the social planner would like agents to search for a longer time than the (Markovian) equilibrium we identify would prescribe. Suppose agent 1 is the first to exit in such an equilibrium, where stopping strategies are not weakly dominated given the search speeds. As long as agent 2 is searching, agent 1 has a unique best response. She would like to use a drawdown size d_1, while the social planner would like her to use a drawdown size $d_1' > d_1$. However, regardless of the space of strategies, there is no way to punish agent 1 for leaving early, and no way to foretell that she will do so. When $\underline{\sigma} < \overline{\sigma}$, agents have room to punish one another for not using the efficient speed by reducing their own. Nonetheless, there is still no way to punish agents for departing too soon. A full analysis of equilibria in non-Markovian strategies is left for the future.

A Appendix

Corollary proofs are immediate and, for completeness, available in the Online Appendix. In what follows, we provide proofs for the paper's main results.

A.1 Proofs for Equilibrium Team Search

First, we note a useful lemma, commonly known as "reflection on the diagonal". This lemma allows us to omit the partial derivatives pertaining to M in the control problem in the various Hamilton-Jacobi-Bellman (HJB) equations that we soon derive. Proofs of this result can be found in various sources, including Dubins, Shepp, and Shiryaev (1994), Urgun and Yariv (2021), and Peskir (1998), and hence omitted.

Lemma A.1. *The infinitesimal generator of the two-dimensional process $Z = (M, X)$ satisfies the following for any C^2 function W:*

1. *If $M_t > X_t$, then $\mathcal{A}_Z^{\sigma_t} = \mathcal{A}_X^{\sigma_t} = \frac{1}{2}(\sigma_t)^2 \frac{\partial^2}{\partial X^2}$.*
2. *If $M_t = X_t$, then $\frac{\partial W}{\partial M} = 0$.*

Proof of Proposition 1. Let $V_i^A(M, X)$ denote the continuation value of agent i in an active alliance A when the observed maximum is M and the current observation is X. Given the Markov structure of the problem and Lemma A.1, the **Hamilton-Jacobi-Bellman (HJB) equation** for such an agent i in an active alliance A, before that alliance shrinks in size, takes the following form:

$$\sup_{\sigma_i}\{\frac{1}{2}(f^A(\sigma^A(M,X))^2 \frac{\partial^2 V_i^A(M,X)}{\partial X^2} - c_i(\sigma_i)\} = 0.$$

The corresponding first-order condition then yields:

$$f^A(\sigma^A(M,X)\frac{\partial f^A(\sigma^A(M,X)}{\partial \sigma_i^A(M,X)}\frac{\partial^2 V_i^A(M,X)}{\partial X^2} = c_i'(\sigma_i).$$

The next step is to characterize the second derivative of the value function. For any agent i in an alliance A, the value function takes the following form:

$$V_i^A(M,X) = \mathbb{E}\left[-\int_0^{\tau^A} c_i(\sigma_{i,t}^A|M,X)dt + V_i^A(M_{\tau^A}, g^A(M_{\tau^A}))\right].$$

The Green function on the interval $[a,b]$ is defined as follows:

$$G_{a,b}(x,y) = \begin{cases} \frac{(b-x)(y-a)}{b-a} & \text{if } a < y < x < b \\ \frac{(b-y)(x-a)}{b-a} & \text{if } a < x < y < b \end{cases}.$$

Following standard techniques, we can write the equilibrium value function of agent i in an alliance A as follows:

$$\begin{aligned} V_i^A(M,X) = {} & V_i^A(M,g^A(M))\frac{M-X}{M-g^A(M)} + V_i^A(M,M)\frac{X-g^A(M)}{M-g^A(M)} \\ & - \int_{g^A(M)}^{M} G_{g^A(M),M}(X,y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy. \end{aligned}$$

Intuitively, there are two possible transitions agent i needs to contemplate that are reflected in the above formulation. The first term corresponds to the stopping boundary being reached with the current maximum. The second term corresponds to a new maximum being achieved. Reaching either of these two states entails a flow of costs, which corresponds to the third term.[32]

For a given observed maximum M, there are two cases to consider for an active agent i in A: either her stopping boundary is the highest within the active alliance, or not. We discuss these in sequence.

Suppose first that $g_i^A(M) = \max_{j\in A} g_j^A(M)$. Consider any observed value X such that $g_i^A(M) \leq X \leq M$. By value matching, agent i's value from reaching the stopping boundary is the observed maximum, since she is the first to stop search. As above, we can write the equilibrium value function of agent i as follows (regardless of whether other agents

[32]This way of writing the accrued costs is the consequence of a change of variables: instead of integrating over time, we integrate over the states, adjusted by the measure of time spent in each state, which is captured by the Green function.

terminate search at the same time):

$$V_i^A(M,X) = M\frac{M-X}{M-g_i^A(M)} + V_i^A(M,M)\frac{X-g_i^A(M)}{M-g_i^A(M)} - \int_{g_i^A(M)}^{M} G_{g_i^A(M),M}(X,y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy.$$

Rearranging terms, we get:

$$V_i^A(M,M) - M = \frac{M-g_i^A(M)}{X-g_i^A(M)}\Big[V_i^A(M,X) - M + \int_{g_i^A(M)}^{M} G_{g_i^A(M),M}(X,y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy\Big].$$

Since agent i optimally terminates her search at $g_i^A(M)$, smooth pasting must hold at $g_i^A(M)$. The derivative of the continuation value as X approaches $g_i^A(M)$ can be written as $\lim_{X\to g_i^A(M)}\frac{V_i^A(M,X)-M}{X-g_i^A(M)}$. By smooth pasting, it must equal the derivative of the value from stopping, $\frac{\partial}{\partial X}M = 0$.

Consider the above equality for $V_i^A(M,M)$. Taking the limit as $X \to g_i^A(M)$,

$$V_i^A(M,M) = M + \int_{g_i^A(M)}^{M}(M-y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy.$$

This, in turn, implies that

$$V_i^A(M,X) = M + \int_{g_i^A(M)}^{X}(X-y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy.$$

Taking the second derivative with respect to X and simplifying yields:

$$\frac{\partial^2 V_i^A(M,X)}{\partial X^2} = \frac{2c_i(\sigma_i^A(M,X))}{(f^A(\sigma_i^A(M,X)))^2}.$$

Plugging the second derivative into the HJB for agent i and simplifying further generates:

$$\frac{2c_i(\sigma_i^A(M,X))}{c_i'(\sigma_i^A(M,X))}\frac{\partial f^A(\sigma^A(M,X))}{\partial \sigma_i^A(M,X)} = f^A(\sigma^A(M,X)).$$

Suppose, instead, that agent i does not have the highest stopping boundary, $\max_{k\in A} g_k^A(M) > g_i^A(M)$. Let $F(M) = \{j \in N : \max_{k\in A} g_k^A(M) = g_j^A(M)\}$. Choose an arbitrary agent $j \in F(M)$.

As above, we can write the continuation payoff of i as follows:

$$V_i^A(M,X) = V_i^A(M, g_j^A(M))\frac{M-X}{M-g_j^A(M)} + V_i^A(M,M)\frac{X-g_j^A(M)}{M-g_j^A(M)} - \int_{g_j^A(M)}^{M} G_{g_j^A(M),M}(X,y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy.$$

This expression echoes the one above. In particular, the first term corresponds to the case in which agent j's stopping boundary is reached before a new maximum. Certainly, in this case, there might be a set $B \supseteq \{j\}$ of agents, possibly random, that terminate search once g_j^A is reached. In this case, $V_i^A(M, g_j^A(M)) = V_i^{A\setminus B}(M, g_j^A(M))$. Rearranging terms, we get:

$$V_i^A(M,M) - V_i^A(M, g_j^A(M)) = \frac{M-g_j^A(M)}{X-g_j^A(M)}\left[V_i^A(M,X) - V_i^A(M,g_j^A(M)) + \int_{g_j^A(M)}^{M} G_{g_j^A(M),M}(X,y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy\right].$$

Again, taking the limit as $X \to g_j(M)$ from above, and letting $\frac{\partial_D V_i^A(M,g_j(M))}{\partial X}$ denote the upper Dini derivative of $V_i^A(M, g_j(M))$ at $g_j(M)$, we have:[33]

$$V_i^A(M,M) = \frac{\partial_D V_i^A(M,g_j(M))}{\partial X}(M-g(M)) + V_i^A(M,g_j^A(M)) + \int_{g_j^A(M)}^{M}(M-y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy.$$

Plugging this identity in $V_i^A(M,X)$'s expression and taking the second derivative:

$$\frac{\partial^2 V_i^A(M,X)}{\partial X^2} = \frac{2c_i(\sigma_i^A(M,X))}{(f^A(\sigma_i^A(M,X)))^2}.$$

Plugging this back into the HJB for agent i and simplifying further generates:

$$\frac{2c_i(\sigma_i^A(M,X))}{c_i'(\sigma_i^A(M,X))}\frac{\partial f^A(\sigma^A(M,X))}{\partial \sigma_i^A(M,X)} = f^A(\sigma^A(M,X)).$$

■

Proof of Proposition 2. The statement of Proposition 2 follows from the following claims.

Claim A.1. *For any given alliance A with $i \in A$, if $g_i^A(M^*) = \max_{j\in A} g_j^A(M^*)$ for some M^*, then*

[33] Since speeds and f are bounded, V is Lipschitz continuous, hence the Dini derivative is finite.

$g_i^A(M) = \max_{j\in A} g_j^A(M)$ *for all* M.

Proof of Claim. The proof of the claim relies on the following lemma.

Lemma A.2. *Suppose agent* $i \in A$ *has the highest stopping boundary at a given observed* M, X. *Then* $g_i^A(M)$ *is a drawdown stopping boundary.*

Proof of Lemma A.2. Suppose $\max_{j\in A} g_j^A(M) = g_i^A(M)$. As shown in the proof of Proposition 1, we have

$$V_i^A(M,X) = M + \int_{g_i^A(M)}^{X} (X-y)c_i(\sigma_i^A(M,y))\frac{2}{(f^A(\sigma^A(M,y)))^2}dy.$$

Furthermore, using Proposition 1, we know that $\sigma^A(M,X) = \sigma^A$ for all M, X and $\sigma_i^A(M,X) = \sigma_i^A$ for all M, X.

Now, differentiating $V_i^A(M,X)$ with respect to M and evaluating the derivative at $X = M$ yields the following ordinary differential equation (ODE) for $g_i^A(M)$:

$$\frac{dg_i^A(M)}{dM} = \frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)\left(M - g_i^A(M)\right)}$$

which leads to the following solution:

$$g_i^A(M) = M - \frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)}.$$

This is a drawdown stopping boundary with drawdown size $d_i^A := \frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)}$. ■

We can now proceed with the claim's proof. Suppose that $g_i^A(M^*) = \max_{j\in A} g_j^A(M^*)$ for some M^* and let $M' = \inf_{\hat{M}>M^*}\{\hat{M} | i \notin \arg\max_{j\in A} g_j^A(\hat{M})\}$. Toward a contradiction, assume that for some $\varepsilon > 0$ and $k \neq i$, for any $\hat{M} \in (M', M'+\varepsilon)$, we have $\max_{j\in A} g_j^A(\hat{M}) = g_k^A(\hat{M}) > g_i^A(\hat{M})$. From continuity of the stopping boundary and Lemma A.2,

$$g_i^A(M) = M - \frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)} \quad \text{and} \quad g_k^A(\hat{M}) = \hat{M} - \frac{(f^A(\sigma^A))^2}{2c_k(\sigma_k^A)}.$$

Our choice of i and k yields $\frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)} \leq \frac{(f^A(\sigma^A))^2}{2c_k(\sigma_k^A)}$ and $\frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)} > \frac{(f^A(\sigma^A))^2}{2c_k(\sigma_k^A)}$, in contradiction. ■

Claim A.2. *Suppose that for some* i *in an active alliance of* A, $\frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)} \leq \frac{(f^A(\sigma^A))^2}{2c_j(\sigma_j^A)}$ *for all* $j \in A$. *Then* i *is the first to exit alliance* A.[34]

[34] If there are multiple agents who satisfy the condition, all exhibiting the same drawdown size, they all exit jointly, weakly before others.

Proof of Claim. Suppose $\frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)} \leq \frac{(f^A(\sigma^A))^2}{2c_j(\sigma_j^A)}$ for all $j \in A$ but that agent i is not one of the first agents to exit from alliance A for some path of observed values. For that path, agent i ceases her search when active at a smaller alliance $A \setminus B$. Without loss of generality, suppose agent j exits alliance A first (if there are multiple such agents, pick any) when observing M and X. From Lemma A.2, agent j's stopping boundary is characterized by a drawdown. However, from the Claim's restriction,

$$M - \frac{(f^A(\sigma^A))^2}{2c_j(\sigma_j^A)} \leq M - \frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)}.$$

For each $k \in A$, the stopping boundary $g_k^A(M) = M - \frac{f^A((\sigma^A))^2}{2c_k(\sigma_k^A)}$ is identified by value matching and smooth pasting. In particular, we have $V_k^A(M, g_k^A(M)) = M$. If $g_i^A(M) > g_j^A(M)$, this implies that $V_i^A(M, g_j^A(M) + \epsilon) < M$ for $0 < \epsilon < \frac{(f^A(\sigma^A))^2}{2c_j(\sigma_j^A)} - \frac{(f^A(\sigma^A))^2}{2c_i(\sigma_i^A)}$. Therefore, agent i would prefer to stop strictly before agent j. ■

The two claims and Lemma A.2' s characterization yield the proposition's proof.

■

A.2 Proofs for the Social Planner's Solution

As in the main text, we denote the social planner's (possibly random) sequence of active alliances by $A_1, A_2, ...$, with $A_1 = \{1, ..., N\}$.

Proof of Proposition 3. Let $\left\{\sigma_i^A(M, X, A)\right\}$ and $G(M, X, A)$ correspond to a solution to the social planner's problem. Consider any alliance A_k at some observed values and let A_{k+1} denote the potentially empty *random* alliance that optimally follows it. Optimality implies that the induced search speeds with A_k should solve:

$$\sup_{\{\sigma_{i,t}\}_{i \in A_k}} \mathbb{E}\left[|A_k \setminus A_{k+1}| M_{\tau^{A_k}} - \int_0^{\tau^{A_k}} \sum_{i \in A_k} c_i(\sigma_{i,s}) ds\right].$$

Using Lemma A.1, the continuation HJB for the social planner can be written as:

$$\sup_{\{\sigma_i\}_{i \in A_k}} \left[\frac{1}{2}\left(f^{A_k}(\sigma^{A_k})\right)^2 \frac{\partial^2 W(M, X, A_k)}{\partial X^2} - \sum_{i \in A_k} c_i(\sigma_i)\right] = 0.$$

Replacing the sup with the appropriate first-order condition yields:

$$f^{A_k}(\sigma^{A_k}) \frac{\partial f^{A_k}(\sigma^{A_k})}{\partial \sigma_i^{A_k}} \frac{\partial^2 W(M, X, A_k)}{\partial X^2} = c_i'(\sigma_i^{A_k}) \quad \forall i \in A_k.$$

Analogous to Proposition 1's proof, the social planner's problem can be written as:

$$W(M,X,A_k) = \sup_{\tau^{A_k},\{\sigma_{i,s}\}_{i\in A_k}} \mathbb{E}\Big[|A_k \setminus A_{k+1}|M_{\tau^{A_k}} + W(M_{\tau^{A_k}}, g^{A_k}(M_{\tau^{A_k}}), A_{k+1}) \mid M,X\Big] - \mathbb{E}\left[\int_0^{\tau^{A_k}} \sum_{i\in A_k} c_i(\sigma_{i,s}^{A_k})ds \mid M,X\right].$$

Since the solution is Markovian, by smooth pasting, we reach the following equation:

$$\begin{aligned} W(M,X,A_k) =& |A_k \setminus A_{k+1}|M + W(M, g^{A_k}(M), A_{k+1}) \\ &+ \left(X - g^{A_k}(M)\right) \sum_{j=k}^{K-1} \left(\int_{g^{j+1}(M)}^{g^j(M)} \frac{2\sum_{i\in A_{j+1}} c_i(\sigma_i^{A_{j+1}}(M,u))}{(f^{A_{j+1}}(\sigma^{A_{j+1}}(M,u)))^2} du \right) \\ &- \int_{g^{A_k}(M)}^{X} (X-y) \frac{2\sum_{i\in A_k} c_i(\sigma_i^{A_k}(M,y))}{(f^{A_k}(\sigma^{A_k}(M,y)))^2} dy. \end{aligned}$$

Taking the second derivative of the value function with respect to X, we have the following:

$$\frac{\partial^2 W(M,X,A_k)}{\partial X^2} = \frac{\partial^2 \mathbb{E}\left[\int_0^{\tau^{A_k}} \sum_{i\in A_k} c_i(\sigma_i^{A_k}(M,X))ds | M,X\right]}{\partial X^2}.$$

Therefore,

$$\frac{2\sum_{i\in A_k} c_i(\sigma_i^{A_k}(M,X))}{c_j'(\sigma_i^{A_k}(M,X))} \frac{\partial f^{A_k}(\sigma^{A_k}(M,X))}{\partial \sigma_j^{A_k}(M,X)} = f^{A_k}(\sigma^{A_k}(M,X)) \quad \forall j \in A_k.$$

Since there is no direct dependence on X on either side, optimal search speeds are independent of observed values and constant over time for each active alliance.

■

As discussed in the text, there are two types of externalities agents exert on one another. The first is an externality within the active alliance in which they operate. This externality is reflected in the optimal search speed. The second externality is on ensuing alliances. This externality is captured by the third term (on the second line) in the expression for $W(M,X,A_k)$ above. This externality is not reflected in the optimal search speed—technically, it appears as a linear term in X and thus vanishes when we take the second derivative. As we show below, this externality does impact the stopping boundary.

In what follows, we suppress the dependence of search speeds on the current maximum and observed value, since we have shown they are constant within any active alliance.

Proof of Proposition 4. The proof follows from two lemmas:

Lemma A.3. *If the set of agents exiting an alliance is independent of the observed path, each alliance has a drawdown stopping boundary identified by a drawdown size d_{A_k}.*

Proof of Lemma A.3. Let A_K be the final alliance in the social planner's problem, with cardinality $|A_K|$. The social planner's problem when left with alliance A_K, and observing maximum M and current value X, takes the following form:

$$W(M,X,A_K) = \sup_{\tau^{A_K}} \mathbb{E}\left[|A_K| M_{\tau^{A_K}} - \int_0^{\tau^{A_K}} \sum_{i \in A_K} c_i(\sigma_i^{A_K}) ds \;\middle|\; M, X \right].$$

This is tantamount to a single-searcher problem, where search rewards are scaled by $|A_K|$. From Urgun and Yariv (2021), the stopping boundary is given by:

$$g^{A_K}(M) = M - d_{A_k}, \quad \text{where } d_{A_k} = \frac{|A_K|(f^{A_K}(\sigma^{A_K}))^2}{2\sum_{i \in A_K} c_i(\sigma_i^{A_K})}.$$

Consider the social planner's problem when the penultimate alliance A_{K-1} is active and the observed maximum and value are M and X, respectively:

$$\begin{aligned} W(M,X,A_{K-1}) = \sup_{\tau^{A_{K-1}}} \mathbb{E}\Big[|A_{K-1} \setminus A_K| M_{\tau^{A_{K-1}}} + W(M_{\tau^{A_{K-1}}}, g^{A_{K-1}}(M_{\tau^{A_{K-1}}}), A_K) \;\Big|\; M, X \Big] \\ - \mathbb{E}\left[\int_0^{\tau^{A_{K-1}}} \sum_{i \in A_{K-1}} c_i(\sigma_i^{A_{K-1}}) ds \;\middle|\; M, X \right]. \end{aligned}$$

By optimality of the stopping time $\tau^{A_{K-1}}$, we have value matching and smooth pasting of $W(M,X,A_{K-1})$ and $W(M,X,A_K)$. Therefore,

$$\begin{aligned} W(M, g^{A_{K-1}}(M), A_{K-1}) &= |A_{K-1} \setminus A_K| M + W(M, g^{A_{K-1}}(M), A_K), \\ \frac{\partial W(M, g^{A_{K-1}}(M), A_{K-1})}{\partial X}\Big|_{X = g^{A_{k-1}}(M)} &= \frac{\partial (|A_{K-1} \setminus A_K| M + W(M, g^{A_{K-1}}(M), A_K))}{\partial X}\Big|_{X = g^{A_{K-1}}(M)}. \end{aligned}$$

Similar to our equilibrium analysis, and using the notation for the Green function introduced there, we can write the welfare maximization problem as:

$$\begin{aligned} W(M,X,A_{K-1}) &= |A_{K-1} \setminus A_K| M + W(M, g^{A_{K-1}}(M), A_K) \frac{M - X}{M - g^{A_{K-1}}(M)} \\ &+ W(M,M,A_{K-1}) \frac{X - g_{A^{K-1}}(M)}{M - g^{A_{K-1}}(M)} - \int_{g^{A_{K-1}}}^{M} G_{g^{A_{K-1}}(M), M}(X,y) \frac{2\sum_{i \in A_{K-1}} c_i(\sigma_i^{A_{K-1}})}{(f^{A_{K-1}}(\sigma^{A_{K-1}}))^2} dy. \end{aligned}$$

Letting X approach $g^{A_{K-1}}(M)$, following the analogous steps in our equilibrium analysis, smooth pasting and rearrangement yield:

$$W(M,X,A_{K-1}) = |A_{K-1} \setminus A_K| M + W(M, g^{A_{K-1}}(M), A_K)$$

$+\Big(X-g^{A_{K-1}}(M)\Big)\int_{g^{A_K}(M)}^{g^{A_{K-1}}(M)}\frac{2\sum_{i\in A_K}c_i(\sigma_i^{A_K})}{(f^{A_K}(\sigma^{A_K}))^2}dx+\int_{g^{A_{K-1}}(M)}^{X}(X-y)\frac{2\sum_{i\in A_{K-1}}c_i(\sigma_i^{A_{K-1}})}{(f^{A_{K-1}}(\sigma^{A_{K-1}}))^2}dy.$
Using the closed-form representation of the value function leads to:

$$\begin{aligned}W(M,X,A_{K-1}) &= |A_{K-1}|M+\frac{1}{2}(g^{A_{K-1}}(M)-g^{A_K}(M))^2\frac{2\sum_{i\in A_K}c_i(\sigma_i^{A_K})}{(f^{A_K}(\sigma^{A_K}))^2}\\ &+\Big(X-g^{A_{K-1}}(M)\Big)(g^{A_{K-1}}(M)-g^{A_K}(M))\frac{2\sum_{i\in A_K}c_i(\sigma_i^{A_K})}{(f^{A_K}(\sigma^{A_K}))^2}\\ &+\frac{1}{2}(X-g^{A_{K-1}}(M))^2\frac{2\sum_{i\in A_{K-1}}c_i(\sigma_i^{A_{K-1}})}{(f^{A_{K-1}}(\sigma^{A_{K-1}}))^2}.\end{aligned}$$

To generate an ODE that identifies $g^{A_{K-1}}(M)$, we take the derivative with respect to M that, evaluated at $X=M$, equals 0. After algebraic manipulations, this ODE takes the form:

$$\frac{dg^{A_{K-1}}(M)}{dM}=\frac{|A_{K-1}\setminus A_K|}{2\Big(M-g^{A_{K-1}}(M)\Big)\left(\frac{\sum_{i\in A_{K-1}}c_i(\sigma_i^{A_{K-1}})}{(f^{A_{K-1}}(\sigma^{A_{K-1}}))^2}-\frac{\sum_{i\in A_K}c_i(\sigma_i^{A_K})}{(f^{A_K}(\sigma^{A_K}))^2}\right)}.$$

It is straightforward to verify that the unique solution for this ODE satisfying the value-matching condition takes the form $g^{A_{K-1}}(M)=M-d_{A_{K-1}}$, where

$$d_{A_{K-1}}=\frac{|A_{K-1}\setminus A_K|}{2\left(\frac{\sum_{i\in A_{K-1}}c_i(\sigma_i^{A_{K-1}})}{(f^{A_{K-1}}(\sigma^{A_{K-1}}))^2}-\frac{\sum_{i\in A_K}c_i(\sigma_i^{A_K})}{(f^{A_K}(\sigma^{A_K}))^2}\right)}.$$

In particular, the optimal stopping boundary is a drawdown stopping boundary.

Proceeding inductively, for any alliance indexed by $m\le K$, the continuation value when M and X are observed can be written as:

$$\begin{aligned}W(M,X,A_m) =&|A_m\setminus A_{m+1}|M+W(M,g^{A_m}(M),A_{m+1})\\ &+\Big(X-g^{A_m}(M)\Big)\sum_{k=m}^{K-1}\left(\int_{g^{k+1}(M)}^{g^k(M)}\frac{2\sum_{i\in A_{k+1}}c_i(\sigma_i^{A_{k+1}})}{(f^{A_{k+1}}(\sigma^{A_{k+1}}))^2}dx\right)\\ &-\int_{g^{A_m}(M)}^{X}(X-y)\frac{2\sum_{i\in A_m}c_i(\sigma_i^{A_m})}{(f^{A_m}(\sigma^{A_m}))^2}dy.\end{aligned}$$

We repeat the steps above to generate an analogous ODE for $g^{A_m}(M)$ and verify that it is uniquely identified as a drawdown stopping boundary. Namely, $g^{A_m}(M)=M-d_{A_m}$, where

$$d_{A_m}=\frac{|A_m\setminus A_{m+1}|}{2\left(\frac{\sum_{i\in A_m}c_i(\sigma_i^{A_m})}{(f^{A_m}(\sigma^{A_m}))^2}-\frac{\sum_{i\in A_{m+1}}c_i(\sigma_i^{A_{m+1}})}{(f^{A_{m+1}}(\sigma^{A_{m+1}}))^2}\right)}.$$

■

Lemma A.4. *The optimal alliance sequence is deterministic. Furthermore, the realized path of alliances is unique up to agents' relabeling.*

Proof of Lemma A.4. The claim follows immediately for $N = 1$. In that case, the solo active agent uses a drawdown stopping boundary, uniquely determining when the agent terminates her search. This is the only possible alliance sequence.

Consider a team of size $N > 1$. Suppose alliances $A1, A2, \ldots$ are implemented. By Lemma A.3, each of these alliances is associated with a drawdown stopping boundary. With the entire team searching, the continuation value when M and X are observed is:

$$W(M, X, A_1) = \mathbb{E}\left[\max_{A_2 \subsetneq A_1} \{|A_1 \setminus A_2| M_{\tau^{A_1}} + W(M_{\tau^{A_1}}, g^{A_1}(M_{\tau^{A_1}}), A_2)\} - \int_0^{\tau^{A_1}} \sum_{i \in A_k} c_i(\sigma_i^{A_1}) dt\right].$$

Suppose that, for some path, the social planner optimally transitions from alliance A_1 to a strictly smaller alliance $A_2 \neq \emptyset$. In particular, alliance A_2 contains fewer than N agents. By the inductive hypothesis, the sequence that ensues is path independent and unique up to agents' relabeling. We can therefore write the continuation value as:

$$\begin{aligned} W(M, X, A_1) =& |A_1 \setminus A_2| M + W(M, g^{A_1}(M), A_2) \\ &+ \left(X - g^{A_1}(M)\right) \sum_{m=1}^{K-1} \left(\int_{g^{A_{m+1}}(M)}^{g^{A_m}(M)} \frac{2 \sum_{i \in A_{m+1}} c_i(\sigma_i^{A_{m+1}})}{(f^{A_{m+1}}(\sigma^{A_{m+1}}))^2} dx \right) \\ &+ \int_{g^{A_1}(M)}^{X} (X - y) \frac{2 \sum_{i \in A_1} c_i(\sigma_i^{A_1})}{(f^{A_1}(\sigma^{A_1}))^2} dy. \end{aligned}$$

As before, this yields an ODE characterizing $g^{A_1}(M)$ and a unique solution of the form $g^{A_1}(M) = M - d_{A_1}$, where $d_{A_1} = \frac{|A_1 \setminus A_2|}{2\left(\frac{\sum_{i \in A_1} c_i(\sigma_i^{A_1})}{(f^{A_1}(\sigma^{A_1}))^2} - \frac{\sum_{i \in A_2} c_i(\sigma_i^{A_2})}{(f^{A_2}(\sigma^{A_2}))^2}\right)}$.

Towards a contradiction, suppose that, on some other path, a different alliance is optimally chosen to follow the initial alliance A_1. Call that alliance $\hat{A}_2 \neq A_2$. Similar arguments would then imply that the stopping boundary for A_1 is given by $g^{A_1}(M) = M - \hat{d}_{A_1}$, where $\hat{d}_{A_1} = \frac{|A_1 \setminus \hat{A}_2|}{2\left(\frac{\sum_{i \in A_1} c_i(\sigma_i^{A_1})}{(f^{A_1}(\sigma^{A_1}))^2} - \frac{\sum_{i \in \hat{A}_2} c_i(\sigma_i^{\hat{A}_2})}{(f^{\hat{A}_2}(\sigma^{\hat{A}_2}))^2}\right)}$.

We consider three cases in turn. First, suppose $d_{A_1} \neq d_{\hat{A}_1}$. In this case, the two stopping boundaries identified above, $M - d_{A_1}$ and $M - d_{\hat{A}_1}$ never intersect, in contradiction.

Second, suppose that $d_{A_1} = \hat{d}_{A_1}$ and $|\hat{A}_2| \neq |A_2|$. In this case, since both $\hat{A}_2$ and A_2 are optimal continuation alliances, $W(M, g^{A_1}(M), A_2) = W(M, g^{A_1}(M), \hat{A}_2)$. Furthermore, since $d_{A_1} = \hat{d}_{A_1}$, the stopping boundary g^{A_1} is identical for either continuation alliances. Therefore, the last term in the expression for the continuation value above, $\int_{g^{A_1}(M)}^{X} (X -$

$y)\frac{2\sum_{i\in A_1} c_i(\sigma_i^{A_1})}{(f^{A_1}(\sigma^{A_1}))^2}dy$, is independent of whether the initial alliance is followed by $\hat{A}_2$ or A_2. Optimality of both sequences then implies that

$$|A_1 \setminus A_2|M + \left(M - g^{A_1}(M)\right)\sum_{m=1}^{K-1}\left(\int_{g^{A_{m+1}}(M)}^{g^{A_m}(M)} \frac{2\sum_{i\in A_{m+1}} c_i(\sigma_i^{A_{m+1}})}{(f^{A_{m+1}}(\sigma^{A_{m+1}}))^2}dx\right) =$$
$$|A_1 \setminus \hat{A}_2|M + \left(M - g^{A_1}(M)\right)\sum_{m=1}^{K-1}\left(\int_{g^{\hat{A}_{m+1}}(M)}^{g^{\hat{A}_m}(M)} \frac{2\sum_{i\in \hat{A}_{m+1}} c_i(\sigma_i^{\hat{A}_{m+1}})}{(f^{\hat{A}_{m+1}}(\sigma^{\hat{A}_{m+1}}))^2}dx\right).$$

Therefore, for $X \in (g^{A_1}(M), M)$, the continuation payoffs when either A_2 or $\hat{A}_2$ ensue do not coincide. In particular, one is higher than the other, contradicting the optimality of both.

Last, suppose that $d_{A_1} = \hat{d}_{A_1}$ and $|\hat{A}_2| = |A_2|$. Then, from optimality,

$$\sum_{m=1}^{K-1}\left(\int_{g^{\hat{A}_{m+1}}(M)}^{g^{\hat{A}_m}(M)} \frac{2\sum_{i\in \hat{A}_{m+1}} c_i(\sigma_i^{\hat{A}_{m+1}})}{(f^{\hat{A}_{m+1}}(\sigma^{\hat{A}_{m+1}}))^2}dx\right) = \sum_{m=1}^{K-1}\left(\int_{g^{A_{m+1}}(M)}^{g^{A_m}(M)} \frac{2\sum_{i\in A_{m+1}} c_i(\sigma_i^{A_{m+1}})}{(f^{A_{m+1}}(\sigma^{A_{m+1}}))^2}dx\right).$$

Therefore, the two candidate alliance sequences remain payoff equivalent for all X until $g^{A_1}(M)$ is reached and differ only in the identity, but not the volume, of agents departing from A_1. We can then recursively follow the above line of arguments starting from alliances A_2 or $\hat{A}_2$ to establish the claim.

■

Combining the two lemmas leads to the conclusion of the proposition. ■

A.3 Proofs for Optimal Sequencing with Well-ordered Costs

Proof of Lemma 1. We use the superscripts eq and sp to denote the equilibrium and social planner's solution, respectively. When costs are well-ordered and speed aggregators are symmetric, in equilibrium, in any alliance, all agents utilize the same search speed. In particular, for any active alliance A and any $i, j \in A$, we have $\sigma_i^{A,eq} = \sigma_j^{A,eq}$. This implies that, in equilibrium, each agent k exits no later than agent $k-1$, for all $k = 2, ..., N$. Indeed, in any active alliance A, the equilibrium stopping boundary is governed by drawdown size

$$d_A^{eq} = \max_{i\in A} \frac{(f^A(\sigma^{A,eq}))^2}{2c_i(\sigma_i^{A,eq})}.$$

This maximum is achieved for the lowest agent index in A.

Suppose, towards a contradiction, that there exists a pair i, j such that $i > j$, so that $\beta_i > \beta_j$, and the social planner has agent i terminate her search strictly before agent j. There are then indices k and m, $k < m$, such that in the social planner's solution, $i, j \in A_k$ but $i \notin A_{k+1}$ and $j \in A_m$ but $j \notin A_{m+1}$.

As we showed, the social planner's solution associates a drawdown stopping boundary with each alliance. Denote the corresponding drawdown sizes d_k^{sp} and d_m^{sp} for A_k and A_m, respectively. Suppose that, instead, the social planner swaps the exits of agents i and j, exiting agent j from A_k whenever agent i was to cease her search and exit from A_k and exiting agent i from A_m whenever agent j was to cease her search and exit from A_m. Furthermore, the social planner can have agent i use the same search speed as agent j had originally in the alliances that follow A_k. Since the speed aggregators depend only on the cardinality of active alliances, the overall search speed in any alliance does not change after this modification. Consequently, expected search outcomes are unaltered. However, the overall cost decreases weakly in every alliance and strictly in all alliances $A_{k+1}, ..., A_m$, contradicting the optimality of the proposed solution. ■

Proof of Proposition 5. Recall that our results so far imply that the social planner can restrict attention to the choice between deterministic alliance sequences. Furthermore, given a deterministic sequence of alliances, Proposition 4 identifies the optimal drawdown stopping boundaries associated with that alliance sequence. For any feasible sequence of alliances (not necessarily optimal), our characterization implies that, whenever the drawdown size associated with two consecutive alliances is negative or zero, the larger alliance is utilized for no length of time. In contrast, when the drawdown size is strictly positive, the social planner gains positive welfare from maintaining the larger alliance active for a non-trivial amount of time. This observation helps us to identify the optimal sequence. The proof of Proposition 5 follows from several lemmas. For any alliance B_k, regardless of whether it is on the social planner's optimal alliance sequence, we denote the optimal overall search speed within the alliance by s^k. That is, when σ^k is the vector of individual speeds in alliance B_k, then $s^k = \left(f^{B_k}(\sigma^k)\right)^2$. The consequent overall search cost within that alliance is denoted $\tilde{c}^k$.

Lemma A.5. *For any m, j, k with $m < j < k$, if the welfare-maximizing sequence is such that B_k is preceded by B_m, then for any sequence where B_k is preceded by B_j, we have $d_{B_m \to B_k} > d_{B_j \to B_k}$.*

Proof of Lemma A.5. From the characterization of drawdown sizes in the well-ordered setting, $d_{B_m \to B_k} \neq d_{B_j \to B_k}$. Suppose that $d_{B_m \to B_k} \leq d_{B_j \to B_k}$. Since B_k is preceded by B_m in the optimal sequence, $d_{B_m \to B_k} > 0$. It then follows that $d_{B_j \to B_k} > 0$. Furthermore, since $m < j$,

$$\frac{|B_m| - |B_k|}{2(\tilde{c}^m/s^m - \tilde{c}^k/s^k)} < \frac{|B_j| - |B_k|}{2(\tilde{c}^j/s^j - \tilde{c}^k/s^k)} < \frac{|B_m| - |B_k|}{2(\tilde{c}^j/s^j - \tilde{c}^k/s^k)} \Rightarrow d_{B_m \to B_j} > 0.$$

This implies that it would be beneficial for the planner to have alliance B_m first transition to alliance B_j, and only then transition to alliance B_k. ■

Lemma A.6. *If $m < k$, $d_{B_m \to \emptyset} > d_{B_k \to \emptyset}$ implies $d_{B_m \to B_k} > d_{B_k \to \emptyset}$.*

Proof of Lemma A.6. Since $|B_i| = N - i + 1$ for any i, $d_{B_m \to \emptyset} > d_{B_k \to \emptyset}$ implies:

$$\frac{1}{N-m+1}\frac{\tilde{c}^m}{s^m} < \frac{1}{N-k+1}\frac{\tilde{c}^k}{s^k} \Longrightarrow \frac{1}{k-m}\left(\frac{\tilde{c}^m}{s^m} - \frac{\tilde{c}^k}{s^k}\right) < \frac{1}{N-k+1}\frac{\tilde{c}^k}{s^k},$$

illustrating the claim. ■

Lemma A.7. *For any k such that $d_{B_k \to \emptyset} > d_{B_N \to \emptyset} > d_{B_{k+1} \to \emptyset}$, we have $d_{B_k \to B_{k+1}} > d_{B_N \to \emptyset}$.*

Proof of Lemma A.7. Observe that $d_{B_k \to \emptyset} > d_{B_N \to \emptyset} > d_{B_{k+1} \to \emptyset}$ implies:

$$(N-k)\frac{\tilde{c}^N}{s^N} > \frac{\tilde{c}^{k+1}}{s^{k+1}} \quad \text{and} \quad \frac{\tilde{c}^k}{s^k} > (N-k+1)\frac{\tilde{c}^N}{s^N}.$$

Summing the inequalities and reorganizing yields the implied statement. ■

Lemma A.8. *If $d_{B_k \to \emptyset} > d_{B_{k-1} \to B_k}$, then $d_{B_k \to \emptyset} > d_{B_{k-1} \to \emptyset} > d_{B_{k-1} \to B_k}$.*

Proof of Lemma A.8. From the first inequality, $d_{B_k \to \emptyset} > d_{B_{k-1} \to B_k}$, we have,

$$\frac{1}{N-k+1}\frac{\tilde{c}^k}{s^k} < \frac{\tilde{c}^{k-1}}{s^{k-1}} - \frac{\tilde{c}^k}{s^k} \Longrightarrow \frac{N-k+2}{N-k+1}\frac{\tilde{c}^k}{s^k} < \frac{\tilde{c}^{k-1}}{s^{k-1}} \Longrightarrow d_{B_k \to \emptyset} > d_{B_{k-1} \to \emptyset}.$$

But this inequality implies that

$$\frac{1}{N-k+1}\frac{\tilde{c}^k}{s^k} < \frac{1}{N-k+2}\frac{\tilde{c}^{k-1}}{s^{k-1}} \Longrightarrow \frac{\tilde{c}^{k-1}}{s^{k-1}} - \frac{\tilde{c}^k}{s^k} > \frac{1}{N-k+2}\frac{\tilde{c}^{k-1}}{s^{k-1}} \Longrightarrow d_{B_{k-1} \to \emptyset} > d_{B_{k-1} \to B_k}.$$

■

Lemma A.9. *If $d_{B_k \to \emptyset} = \max_j d_{B_j \to \emptyset}$, then any alliance B_l with $l < k$ cannot be the welfare maximizing last alliance.*

Proof of Lemma A.9. Suppose not, so that, form some $l < k$, alliance B_l is the last. Since B_k is strictly contained in B_l, from the characterization of drawdowns in the well-ordered setting, $d_{B_k \to \emptyset} \neq d_{B_l \to \emptyset}$. Thus, $d_{B_k \to \emptyset} > d_{B_l \to \emptyset} > 0$. Following similar arguments to those in the proof of Lemma A.5, the social planner would benefit from transitioning from B_l to B_k instead of exiting all members of B_l, in contradiction. ■

Lemma A.10. *If $d_{B_k \to \emptyset} = \max_j d_{B_j \to \emptyset}$, then any alliance B_l with $l > k$ cannot be the last.*

Proof of Lemma A.10. We use induction on the cardinality of the set B_k. The claim certainly holds when $|B_k| = 1$, so that $B_k = B_N = \{N\}$.

Assume the statement is true for sets up to cardinality n. We show the statement holds for $|B_k| = n + 1$ (so that $k = N - n$). By Lemma A.9, the last alliance cannot be B_j with $j < k$. Towards a contradiction, suppose that a smaller set B_m, with $m > k$, is the last alliance.

From the inductive hypothesis, we must have $d_{B_m \to \emptyset} > d_{B_l \to \emptyset}$ for all $l > m$, as otherwise the social planner would benefit by inducing B_l to continue search instead of terminating it for all agents in B_m.

Suppose that $m < N$. Consider an equivalent problem, where alliance B_m is replaced with a single individual M that has cost function $\beta_M c(\cdot)$, where $\beta_M > \beta_{m-1}$ arbitrarily chosen. For each $j \leq m$, define $C_j = \{j, \ldots, m-1, M\}$, $\hat{f}^{C_j}$, and an interior $\sigma_M^{C_j}$ so that the social planner's optimal speed vector for any set C_j is given by $(\sigma_j^{B_j}, \ldots, \sigma_{m-1}^{B_j}, \sigma_M^{C_j})$ and $\hat{f}^{C_j}$ satisfies all our original assumptions on speed aggregators.[35]

We now face an equivalent problem with m agents $1, 2, ..., m-1, M$. From our construction, in the optimal solution, for any $j = 1, ..., m-1$, the corresponding drawdown sizes $d_{\{j,...,M\} \to \emptyset}$ and $d_{\{M\} \to \emptyset}$ coincide with the optimally-set drawdown sizes $d_{B_{j \to \emptyset}}$ and $d_{B_{m \to \emptyset}}$ in our original problem. Therefore, $\max_{j \in \{1,...,m-1,M\}} d_{\{j,...,M\} \to \emptyset} = d_{\{k,...,M\} \to \emptyset}$. By our induction hypothesis, $\{j, ..., M\}$ with $j > k$ cannot optimally be the last alliance, in contradiction.

Suppose now that $m = N$ and, towards a contradiction, assume B_N is the welfare maximizing last alliance. Now consider the sequence of welfare maximizing alliances B_p such that $B_p \subset B_k$. There are three cases to consider.

Case 1: For all $p \in \{k, ..., N-1\}$, the alliance B_p is part of the welfare-maximizing sequence. That is, agents terminate their search one by one starting from B_k onwards. Since B_N is the last alliance, we must have that $d_{B_N \to \emptyset} > d_{B_{N-1} \to B_N} > d_{B_{N-2} \to B_{N-1}} > \ldots > d_{B_k \to B_{k+1}}$. Applying Lemma A.8 repeatedly implies that $d_{B_N \to \emptyset} > d_{B_{N-1} \to \emptyset} > d_{B_{N-2} \to \emptyset} \ldots > d_{B_{k+1} \to \emptyset}$.

The assumed maximality of $d_{B_k \to \emptyset}$ implies, in particular, that $d_{B_k \to \emptyset} > d_{B_N \to \emptyset}$ that, combined with the above, yields $d_{B_k \to \emptyset} > d_{B_N \to \emptyset} > d_{B_{k+1} \to \emptyset}$. By Lemma A.7, we then have that $d_{B_k \to B_{k+1}} > d_{B_N \to \emptyset}$. It follows that whenever agents in the active alliance B_k optimally stop searching, the social planner would benefit from halting all agents' search instead of proceeding with $B_{k+1}, B_{k+2}, ..., B_N$, in contradiction.

Case 2: There does not exist any $p \in \{k, ..., N-1\}$ such that B_p is part of the optimal sequence. Thus, the penultimate alliance in the optimal sequence is B_l with $l < k$. Maximality of $d_{B_k \to \emptyset}$ implies that $d_{B_k \to \emptyset} > d_{B_N \to \emptyset}$. By Lemma A.6, $d_{B_k \to B_N} > d_{B_N \to \emptyset}$ and by Lemma A.5, $d_{B_l \to B_N} > d_{B_k \to B_N} > d_{B_N \to \emptyset}$. Thus, whenever agents in active alliance B_l optimally stop searching, the social planner would benefit from halting all agents' search instead of proceeding with B_N, in contradiction.

Case 3: There exist $p, q \in \{k, ..., N-1\}$ such that B_p is part of the optimal sequence but B_q is not. Here we have two subcases:

[35] We only restrict $\hat{f}^{C_j}$ to lead to the original optimal speed choice by agents $j, \ldots, m-1$. This restriction leaves us with a lot of freedom to select $\hat{f}^{C_j}$ in this way.

Subcase 1: B_{N-1} is the penultimate alliance. We must have $d_{B_{N-1}\to\emptyset} < d_{B_N\to\emptyset}$; otherwise, by Lemma A.6, we would have $d_{B_{N-1}\to B_N} > d_{B_N\to\emptyset}$ and it would be suboptimal to utilize alliance B_N as the last alliance. From the maximality of $d_{B_k\to\emptyset}$ and Lemma A.5, for any $l < k$ such that B_l precedes B_{N-1} on the optimal path, $d_{B_l\to B_{N-1}} > d_{B_k\to B_{N-1}} > d_{B_{N-1}\to\emptyset}$. Finally, $d_{B_N\to\emptyset} > d_{B_{N-1}\to\emptyset}$ implies that

$$\frac{1}{2}\frac{\tilde{c}^{N-1}}{s^{N-1}} > \frac{\tilde{c}^{N}}{s^{N}} \Longrightarrow \frac{1}{2}\frac{\tilde{c}^{N-1}}{s^{N-1}} < \frac{\tilde{c}^{N-1}}{s^{N-1}} - \frac{\tilde{c}^{N}}{s^{N}} \Longrightarrow d_{B_{N-1}\to\emptyset} > d_{B_{N-1}\to B_N}.$$

Thus, $d_{B_l\to B_{N-1}} > d_{B_k\to B_{N-1}} > d_{B_{N-1}\to\emptyset} > d_{B_{N-1}\to B_N}$. Therefore, whenever agents in the active alliance B_l optimally stop searching, the social planner would benefit from transitioning to B_N directly, thereby terminating the search of agent $N-1$ as well, instead of transitioning to B_{N-1} first, in contradiction.

Subcase 2: The penultimate alliance is B_p with $p \in \{k, ..., N-2\}$. We can now emulate the argument above pertaining to the construction of an equivalent problem in which agents $\{p, ..., N-1\}$ are viewed as one agent with appropriately induced search costs. We can then consider an equivalent problem with fewer agents to achieve a contradiction through our induction hypothesis. ■

It follows that the last alliance is given by B_k with $\max_j d_{B_j\to\emptyset} = d_{B_k\to\emptyset}$.

The proofs of the following Lemmas are a consequence of identical arguments to those in of Lemmas A.9 and A.10 and are therefore ommitted.

Lemma A.11. *Consider B_k, where k is such that $d_{B_k\to B_{L_1}} > d_{B_j\to B_{L_1}}$ for all $j < L_1$ and B_{L_1} is the last alliance as identified above. Then any alliance with $l < k$ cannot be the welfare maximizing second to last alliance.*

Lemma A.12. *Consider B_k, where k is such that $d_{B_k\to B_{L_1}} > d_{B_j\to B_{L_1}}$ for all $j < L_1$ and B_{L_1} is the last alliance as identified above. Then any alliance B_l with $L_1 > l > k$ cannot be the welfare maximizing second to last alliance.*

The proof of Proposition 5 then follows. Using the proposition's notation, B_{L_1} is the last alliance on the social planner's optimal path. Similarly, the penultimate alliance is given by B_k where k is such that $d_{B_k\to B_{L_1}} > d_{B_j\to B_{L_1}}$ for all $j < L_1$. We can continue recursively to establish the proposition's claim. ■

References

Admati, A. R. and M. Perry (1991). Joint projects without commitment. *The Review of Economic Studies 58*(2), 259–276.

Albrecht, J., A. Anderson, and S. Vroman (2010). Search by committee. *Journal of Economic Theory 145*(4), 1386–1407.

Anderson, A., L. Smith, and A. Park (2017). Rushes in large timing games. *Econometrica 85*(3), 871–913.

Azéma, J. and M. Yor (1979). Une solution simple au problème de skorokhod. *Séminaire de probabilités de Strasbourg 13*, 90–115.

Bardhi, A. and N. Bobkova (2021). Local evidence and diversity in minipublics. *mimeo*.

Bloom, N., C. I. Jones, J. Van Reenen, and M. Webb (2020). Are ideas getting harder to find? *American Economic Review 110*(4), 1104–44.

Bolton, P. and C. Harris (1999). Strategic experimentation. *Econometrica 67*(2), 349–374.

Bonatti, A. and H. Rantakari (2016). The politics of compromise. *American Economic Review 106*(2), 229–59.

Bulow, J. and P. Klemperer (1994). Rational frenzies and crashes. *Journal of political Economy 102*(1), 1–23.

Callander, S. (2011). Searching and learning by trial and error. *American Economic Review 101*(6), 2277–2308.

Callander, S. and P. Hummel (2014). Preemptive policy experimentation. *Econometrica 82*(4), 1509–1528.

Cankurtaran, P., F. Langerak, and A. Griffin (2013). Consequences of new product development speed: A meta-analysis. *Journal of Product Innovation Management 30*(3), 465–486.

Caplin, A. and J. Leahy (1994). Business as usual, market crashes, and wisdom after the fact. *American Economic Review*, 548–565.

Cetemen, D., I. Hwang, and A. Kaya (2020). Uncertainty-driven cooperation. *Theoretical Economics 15*, 1023–1058.

Chen, J., F. Damanpour, and R. R. Reilly (2010). Understanding antecedents of new product development speed: A meta-analysis. *Journal of Operations Management 28*(1), 17–33.

Deb, J., A. Kuvalekar, and E. Lipnowski (2020). Fostering collaboration. *mimeo*.

Décamps, J.-P., F. Gensbittel, and T. Mariotti (2021). Investment timing and technological breakthrough. *mimeo*.

Dixit, A. K. and R. S. Pindyck (1994). *Investment under uncertainty*. Princeton university press.

Dubins, L. E., L. A. Shepp, and A. N. Shiryaev (1994). Optimal stopping rules and maximal

inequalities for bessel processes. *Theory of Probability & Its Applications 38*(2), 226–261.

Eftekhari, N. and B. Timmermans (2021). New venture dissolution and the comobility of new venture teams. *Small Business Economics*, 1–20.

Gapeev, P. V. (2007). Discounted optimal stopping for maxima of some jump-diffusion processes. *Journal of Applied Probability 44*(3), 713–731.

Gul, F. and R. Lundholm (1995). Endogenous timing and the clustering of agents' decisions. *Journal of Political Economy 103*(5), 1039–1066.

Hörner, J., N. Klein, and S. Rady (2021). Overcoming free-riding in bandit games. *mimeo*.

Hörner, J. and A. Skrzypacz (2016). Learning, experimentation and information design. *mimeo*.

Jovanovic, B. and G. M. MacDonald (1994). The life cycle of a competitive industry. *Journal of Political Economy 102*(2), 322–347.

Jovanovic, B. and R. Rob (1990). Long waves and short waves: Growth through intensive and extensive search. *Econometrica: Journal of the Econometric Society*, 1391–1409.

Keller, G., S. Rady, and M. Cripps (2005). Strategic experimentation with exponential bandits. *Econometrica 73*(1), 39–68.

Kolmar, M. and H. Rommeswinkel (2013). Contests with group-specific public goods and complementarities in efforts. *Journal of Economic Behavior & Organization 89*, 9–22.

Marx, L. M. and S. A. Matthews (2000). Dynamic voluntary contribution to a public project. *The Review of Economic Studies 67*(2), 327–358.

McCall, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 113–126.

Mortensen, D. T. (1970). Job search, the duration of unemployment, and the phillips curve. *The American Economic Review 60*(5), 847–862.

Murto, P. and J. Välimäki (2011). Learning and information aggregation in an exit game. *The Review of Economic Studies 78*(4), 1426–1461.

Øksendal, B. (2003). *Stochastic differential equations*. Springer.

Papadimitriou, C. H. and K. Steiglitz (1998). *Combinatorial optimization: algorithms and complexity*. Courier Corporation.

Peskir, G. (1998). Optimal stopping of the maximum process: The maximality principle. *Annals of Probability*, 1614–1640.

Peskir, G. and A. Shiryaev (2006). *Optimal stopping and free-boundary problems*. Springer.

Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

Rosenberg, D., E. Solan, and N. Vieille (2007). Social learning in one-arm bandit problems. *Econometrica 75*(6), 1591–1611.

Strulovici, B. (2010). Learning while voting: Determinants of collective experimentation.

Econometrica 78(3), 933–971.

Taylor, H. M. et al. (1975). A stopped brownian motion formula. *The Annals of Probability 3*(2), 234–246.

Titova, M. (2021). Collaborative search for a public good. *mimeo*.

Urgun, C. and L. Yariv (2021). Retrospective search: Exploration and ambition on uncharted terrain. *mimeo*.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica*, 641–654.

Whitt, W. (1980). Some useful functions for functional limit theorems. *Mathematics of Operations Research 5*(1), 67–85.

Wong, Y. F. (2021). Spatial experimentation. *mimeo*.

Wuchty, S., B. F. Jones, and B. Uzzi (2007). The increasing dominance of teams in production of knowledge. *Science 316*(5827), 1036–1039.

Yildirim, H. (2006). Getting the ball rolling: Voluntary contributions to a large-scale public project. *Journal of Public Economic Theory 8*(4), 503–528.