

Collective Progress: Dynamics of Exit Waves*

Doruk Cetemen[†] Can Urgan[‡] Leeat Yariv[§]

June 30, 2021

Abstract. We study a model of collective search by teams. Discoveries beget discoveries and correlated search results are governed by a Brownian path. Search results' variation at any point—the search scope—is jointly controlled. Agents individually choose when to cease search and implement their best discovery. We characterize equilibrium and optimal policies. Search scope is constant and independent of search outcomes as long as no member leaves. It declines after departures. A simple drawdown stopping boundary governs each agent's search termination. We show the emergence of endogenous exit waves, whereby possibly heterogeneous agents cease search all at once.

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

JEL codes: C73, D81, D83, O35

*We thank Arjada Bardhi, Alessandro Bonatti, Francesc Dilme, Hari Govindan, Faruk Gul, Toomas Hinnosaar, Alessandro Lizzeri, Chiara Margaria, Pietro Ortoleva, Wolfgang Pesendorfer, and Heikki Rantakari for helpful comments and suggestions. We also thank seminar audiences at Central European University, Collegio Carlo Alberto, Princeton, Seoul National University, Tel-Aviv University, UCLA, University of Rochester, and VSET. Yariv gratefully acknowledges financial support from the National Science Foundation through grants SES-1629613 and SES-1949381.

[†]Collegio Carlo Alberto; doruk.cetemen@carloalberto.org

[‡]Princeton University; curgun@princeton.edu

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

1 Introduction

Discoveries are often made by teams. 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. 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. 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.

This paper offers a new framework for studying collective progress based on a process of search. We identify how the breadth of search and decisions to terminate search vary with members' characteristics, the synergies in place. We also show that exit waves, where multiple members halt search simultaneously, are an inherent feature of such processes.

Technological developments rarely occur in a vacuum and discoveries build on one another. We therefore consider environments in which search results are correlated over time and follow a Brownian path, as first modeled by [Callander \(2011\)](#). The scope of search, captured by the Brownian path's instantaneous variance, is chosen at each moment by the searching alliance. Specifically, each member of a searching alliance incurs a strictly positive cost that depends on that member's own search scope. While jointly searching, the Brownian path's instantaneous variance corresponds to the sum of members' search scopes. 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. 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 previously-observed maxima. As search progresses, members gradually terminate their search until it halts altogether.¹

¹Most of our qualitative results carry over when introducing penalties for later exits, though

We characterize equilibrium search in Markov strategies, where state variables correspond to the current search results, the attained maximum, and the active alliance. We show that, in any active alliance, search scope is constant and independent of search results as long as no member leaves. Individual search scopes increase when members depart, reflecting the more limited free-riding opportunities present. The optimal 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.

The ratio of marginal to fixed costs governs both equilibrium search scopes and drawdown sizes and serves as a proxy for the synergies present in an alliance. In particular, agents may prefer to team up with others exhibiting both higher marginal and fixed costs, provided the ratio guarantees they are more willing to contribute to the collective search.

Relative to an individual searching on her own, standard free-riding motives drive search scopes down in an alliance. This is a form of a discouragement effect, whereby members do not search as intensely when they expect others to bear some of the search costs. Nonetheless, externalities make search more valuable in a team: a member can reap the benefits of her peers' efforts. There is therefore also an encouragement effect, reminiscent of that present in experimentation settings, that leads team members to search for longer than they would have on their own.

Our equilibrium characterization allows us to identify members' patterns of exits. In general, those exhibiting high ratios of marginal to fixed costs leave earlier than those exhibiting low such ratios. We show that, even when individual costs are fully heterogenous, 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, their sequencing—who leaves first, second, etc., and with whom—does not.

Beyond its substantive implications, our equilibrium characterization offers a technical contribution. As we detail in our literature review below, extant analyses of single-agent search processes often resort to modeling short-lived agents, absent

naturally such penalties alter exit patterns. In particular, penalties for later exits can introduce exit waves mechanically—once one agent departs, others may follow suit to avoid penalties.

any controls. In contrast, we analyze the evolution of collective search by forward-looking and sophisticated agents who can utilize a costly control.

In the last part of the paper, we characterize the socially optimal search scope and stopping policies. The socially optimal search scope is also constant and independent of search results within any active alliance. Naturally, the positive externalities induced by each member's investment in search scope imply that the socially optimal level is higher than that chosen in equilibrium. Furthermore, in contrast to equilibrium search scopes, as alliance members terminate their search, the optimal scope 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.² 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. 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 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.

²As we show, allowing for non-Markovian equilibria does not eliminate some of the inefficiencies we highlight.

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 studied by [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 work that ensued considers behavior of short-lived agents as well. [Urgun and Yariv \(2021\)](#) analyze an individual-search setting similar to the one analyzed here.

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 may impede 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 \(2019\)](#) 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.³

There are also several papers illustrating patterns reminiscent of the clustered

³Dynamic contribution games without experimentation or uncertainty have also been heavily studied, see for instance [Admati and Perry \(1991\)](#), [Marx and Matthews \(2000\)](#), [Yildirim \(2006\)](#), and [Cetemen, Hwang, and Kaya \(2020\)](#).

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 as well as observes others' prior decisions. Firms' actions reveal information and can generate a wave. [Gul and Lundholm \(1995\)](#) analyze a two-agent model in which both try to predict the value of a project using their 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 the other agent 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.⁴

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

3 A Model of Collective Search

Consider a team of N agents—product developers, academic researchers, etc.—searching through a terrain of ideas in continuous time. Time is indexed by t and runs through $[0, \infty)$. Each seeks good outcomes and ultimately benefits from the maximal value they have found when they stop their search. Formally, we assume all agents are risk neutral. At each time t , agent $i = 1, 2, \dots, N$ decides on the scope

⁴There 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. Profits decrease as more firms enter. When there is a technological shock, some firms become more productive than others, potentially leading to clustered exits.

of her search $\sigma_{i,t}^A \in [\underline{\sigma}, \bar{\sigma}]$, where $\bar{\sigma} \geq \underline{\sigma} > 0$ and $A \subseteq \{1, \dots, N\}$ is the alliance of agents still searching at time t . Agents' scope choices are observed within their alliance. As we soon describe, the search scope naturally feeds into the breadth of search conducted by the alliance of active agents. For $i = 1, 2, \dots, N$, any search scope σ comes at a cost of $c_i(\sigma)$, where c_i is twice continuously differentiable, increasing, and convex, with a second derivative bounded above zero over $[\underline{\sigma}, \bar{\sigma}]$. The special case of $\underline{\sigma} = \bar{\sigma}$ corresponds to settings in which search scope is not controlled and agents' only choose when to stop search.

We model the progress of discoveries using a Weiner process, which allows us to capture the correlation of new developments over time, and the impact of search scope of those who engage in search.⁵ Formally, for any time t , denote by B_t the standard Brownian motion with $B_0 = 0$, and let σ_t^A denote the controlled breadth of search, which will depend on the search scopes of all members of the active alliance A as we soon describe. The observed value at t —which can be thought of as the expected value of the discovery—is denoted by X_t , where $X_0 = 0$ and the law of motion is given by:

$$dX_t = \sigma_t^A dB_t.$$

Whenever the alliance A of agents is searching, we assume $\sigma_t^A = \sum_{i \in A} \sigma_{i,t}^A$.⁶

The search scope can be interpreted in two ways. First, it can capture search breadth. 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. Second, given our modeling of search values, the search scope can also be thought of as capturing search *speed*. Changing the search scope from 1 to σ at any small interval of time is tantamount to “speeding up” the process by a factor of σ^2 . As we soon show, search *returns* depend linearly on search scope.

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 *con-*

⁵We view correlation as an important feature of discovery processes. Nonetheless, from a purely theoretical perspective, one could analyze an analogous model with independent samples. As it turns out, such a model is far less tractable. Details appear in the Online Appendix.

⁶In the Online Appendix, we show that our analysis can be directly extended to the case in which, for any alliance A , we have $\sigma_t^A = f^A(\{\sigma_{i,t}^A\}_{i \in A})$, with f^A a differentiable function. Comparative statics would naturally depend on alliances' technologies captured by $\{f^A\}_A$.

trols drift rather than search scopes, 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.⁷ We view endogenous search scopes as natural for most applications, where investments in innovation either affect the speed at which progress is made, or entail non-trivial risks.

3.1 Payoffs

Each agent is rewarded according to the maximal project value observed up to her stopping time. Let M_t denote the maximum value observed by time t :

$$M_t = \left(\max_{0 \leq r \leq t} X_r \vee M_0 \right),$$

where we assume that $M_0 = 0$.

For any aggregate fixed search scope σ , at time t , $\mathbb{E}(M_t) = \sigma \sqrt{2t/\pi}$. Thus, the choice of search scope 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 scope of individual i , which may depend on the alliances she is active in.⁸ Any progress made after an agent stops searching does not impact her payoffs.

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 scope over time and her stopping policy. Formally, it is a pair of functions (σ_i^A, τ_i^A) , where $A \subseteq \{1, \dots, N\}$ and $i \in A$. In principle, (σ_i^A, τ_i^A) may depend on time, as well as the entire path of

⁷Taylor et al. (1975) characterize the maximal value of search with constant drift. The resulting value is far less amenable to further analysis than ours.

⁸In Section 6.1, we discuss an extension in which agents who stop later are penalized.

observed search values and corresponding maxima. 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 agents use strategies of the form (σ_i^A, τ_i^A) that depend only on the state variables X_t , M_t , and A_t . Formally, $\sigma_i^A : \mathbb{R}^2 \rightarrow [\underline{\sigma}, \bar{\sigma}]$, and τ_i^A is a random variable over \mathbb{R}_+ such that $\Pr(\tau_i^A = t | \mathcal{F}_t) = \Pr(\tau_i^A = t | X_t, M_t)$ for all i .⁹

We further assume that a continuous stopping boundary determines when each agent halts her search. Formally, for all i and all alliances A such that $i \in A$, the stopping policy takes the following form:

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

where $g_i^A(\cdot)$ is a continuous function. This formulation implicitly implies that, upon indifference, agents exit the search. 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. As we soon show, in our setting, departing agents would never benefit from continuing the search in a smaller alliance: the externalities offered by a larger alliance are always beneficial.

Given $\{(\sigma_j^A, \tau_j^A)\}_{j \neq i}$, agent i 's best-response strategy simply maximizes her expected payoff given this profile. Formally, it is determined by solving the following problem for each alliance A such that $i \in A$:

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

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 scopes and stopping boundaries. We also identify the sequencing of agents' search termination, and the patterns of equilibrium exit waves.

⁹The inefficiencies we highlight do not vanish when considering equilibria in non-Markovian strategies, see our discussion in Section 6.2.

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,

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

Since agents use continuous stopping boundaries, we can write

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

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

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

Proposition 1 (Team Search Scope). *For any agent i in an active alliance A , equilibrium search scopes are constant, $\sigma_i^A(M_t, X_t) = \sigma_i^A$. Whenever interior, search scopes satisfy the system:*

$$\frac{2c_i(\sigma_i^A)}{c'_i(\sigma_i^A)} = \sigma^A = \sum_{i \in A} \sigma_i^A \quad \forall i \in A.$$

Why are search scopes constant as long as a certain alliance of agents is active? The rough intuition is the following. Consider an agent i in an active alliance A . Suppose i believes that all other agents j in the alliance search with scope 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: $(\sum_{k \in A} \sigma_k^A)^2$ 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, $\frac{c_i(\sigma_i^A)}{(\sum_{k \in A} \sigma_k^A)^2} = \frac{c_i(\sigma_i^A)}{(\sigma^A)^2}$. The identity in the proposition reflects the corresponding first-order condition.

In general, there might be multiple solutions to the system in Proposition 1, some possibly corresponding to less efficient equilibria. Nonetheless, [Fleming and Rishel \(2012\)](#) (Theorem 6.4) guarantees that any equilibrium features continuous search scopes within any alliance. Hence, within an active alliance, agents can utilize only one of the solutions.¹⁰ Importantly, it is the *ratio* of costs to marginal

¹⁰The conditions of Theorem 6.4 in [Fleming and Rishel \(2012\)](#) follow from our assumption that

costs that govern equilibrium search scopes. In particular, in our setting, teaming up with agents who have both higher costs and marginal costs can be beneficial in terms of externalities.

Whenever interior solutions to the system in Proposition 1 are unique, comparisons of search scopes within various alliances are well defined. Uniqueness of interior solutions is guaranteed when, e.g., all scope costs are log-convex. A direct corollary of Proposition 1 is then the following.

Corollary 1 (Search Scope and Alliance Size). *Suppose costs are log-convex and an interior solution exists for the systems specified in Proposition 1. As an alliance shrinks, individual members' search scopes increase, while total search scope decreases. That is, for any $i, j \in A$, we have $\sigma_i^{A \setminus \{j\}} \geq \sigma_i^A$ while $\sigma^A > \sigma^{A \setminus \{j\}}$.*

The corollary highlights a form of free-riding. Search scope is substitutable across individuals. The more agents searching, the less each one searches. Since individual search scopes decrease within an alliance, the total search scope in any active alliance is smaller than that which would be generated by the alliance's members searching independently.¹¹

The corollary indicates that agents departing would never benefit from continuing search on their own, nor from switching to search in a smaller alliance than the one they have left. In particular, our assumption that agents who cease search in an alliance reap the benefits from past discoveries rather than pursue further discoveries with other newly-departed agents is without loss of generality.

We now turn to the characterization of equilibrium stopping boundaries. 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). *There exists an equilibrium such that, for any agent i in any active alliance A ,*

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

the cost function's second derivative is bounded above zero. For an alliance composed of one individual, there is a unique optimal solution due to concavity of the objective function.

¹¹Any 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.

In particular, agent $i \in \arg \min_j \frac{(\sigma^A)^2}{2c_j(\sigma_j^A)}$ is the first to stop in any alliance A . Furthermore, given equilibrium search scopes, 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 scopes given by Proposition 1. Each individual agent i 's optimization problem then boils down to a solo searcher's optimization, with others' search simply affecting the experienced search costs. Namely, the induced cost of implementing search scope 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 scopes, 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.¹² 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

¹²The focus on weakly undominated stopping boundaries—given the equilibrium search scopes—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.

not impact when the alliance first loses some of its members, nor the search scope while it is fully active. Consequently, equilibrium outcomes are unique.¹³

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 scopes, 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 that governs their decision to stop search. These new drawdowns are $\{d_j^{A \setminus Z^1}\}_{j \in A \setminus Z^1}$. The discrete drop in overall search scope 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 2 (Equilibrium Exit Waves). *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.

4.3 Well-ordered Costs

We now consider a particular setting, where the identification of exit waves and their comparative statics is particularly simple.

¹³Our 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.

Suppose agents' cost functions are proportional to one another: $c = c_1\beta_1 = c_2\beta_2 \cdots = c_N\beta_N$, where $\beta_1 = 1 < \beta_2 < \dots < \beta_N$. That is, agent 1 has the highest search costs, while agent N has the lowest search costs.

Proposition 1 implies that all agents in an active alliance choose the same search scope, assuming an interior solution exists. Suppose σ denotes the search scope all agents utilize in the full alliance. It follows that $\frac{2c(\sigma)}{c'(\sigma)} = N\sigma$.

As N increases, individual search scopes decrease. Let \hat{N} be the maximal integer such that $\frac{2c(\underline{\sigma})}{c'(\underline{\sigma})} < \hat{N}\underline{\sigma}$. For any $N > \hat{N}$, there is no interior equilibrium. Furthermore, when there are \hat{N} agents in the team, individual search scopes are initially roughly at their minimum $\underline{\sigma}$, while overall search scope is $\hat{N}\underline{\sigma}$.

Agents' search scope 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 $j, j + 1, j + 2, \dots, j + k$ will all terminate their search at the same time. Figure 1 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 drawdowns that govern agents' departures correspond to the "upper envelope" of the graph depicting $d_j^{\{j, \dots, N\}}$ as a function of j .

Despite agents' costs being strictly ordered, clustered exits are possible. In fact, when costs are close to one another, all agents might exit at once. Indeed, from Proposition 2, $d_j^{\{j, \dots, N\}} = \frac{(\sigma^{\{j, \dots, N\}})^2 \beta_j}{2c(\sigma^{\{j, \dots, N\}})}$. From Corollary 1, $\sigma^{\{j, \dots, N\}}$ decreases in j , while $\sigma_j^{\{j, \dots, N\}}$ increases in j . Therefore, for $\{\beta_j\}$ sufficiently close to one another, $d_1^{\{1, \dots, N\}} > d_2^{\{2, \dots, N\}} > \dots > d_N^{\{N\}}$ and all agents exit at once. Naturally, 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




FIGURE 1: Equilibrium exit waves with well-ordered costs

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.

5 The Social Planner's Problem

We now consider a social planner who dictates agents' search scopes 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.

5.1 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 scopes 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 scope of each member in that alliance.

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 reduces in size. Whenever $G(M, X, A) = \emptyset$, no agent is left searching and the search terminates.

The social planner's second control is the profile of search scopes within any alliance A , which can be written as $\sigma_i^A(M, X) : \mathbb{R}^2 \mapsto [\underline{\sigma}, \bar{\sigma}]$ for each $i \in A$. As before, agents that already exited cannot be induced to choose positive search scope and do not participate in any future search: exit is irreversible. We therefore write $\sigma_i^A(M, X) = 0$ for each $i \notin A$. For any active alliance A , we write:

$$\sigma^A(M, X) = \sum_{i \in A} \sigma_i^A(M, X)$$

and, 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:

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

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

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

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:

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

We assume that whenever the social planner is indifferent between maintaining a certain set of agents searching or having them exit, she chooses the latter.

Given a pair of controls (G, σ) , with slight abuse of notation, let $A_1 = \{1, \dots, N\}$ denote the first active alliance, containing all agents.¹⁴ Using (1), let $A_2 = \tilde{A}_{\tau^{A_1}}$ be the alliance that succeeds the initial alliance, 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 be induced to stop their search. 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 then have a sequence of active alliances A_1, A_2, \dots, A_K with associated stopping times $\tau^{A_1}, \tau^{A_2}, \dots, \tau^{A_K}$.

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 now 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

¹⁴We 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.

active alliance A when observing X and M . It would then also be optimal to halt the search of 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 scopes 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 scopes 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' at 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

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

where $g^A(M) = \sup\{X : G(M, X, A) \neq A\}$.¹⁵ This kind of stopping time τ^A is commonly known as an Azéma-Yor stopping time (Azéma and Yor, 1979), with the function g^A defining the corresponding stopping boundary.

For any active alliance A , we note that $g^A(M) < M$ for all M . In other words, it is never 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.¹⁶

¹⁵We 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.

¹⁶This would not hold were the social planner's objective concave in the maximum observed. Concavity introduces new challenges, see Urgun and Yariv (2021) for a discussion of its impact on single-agent decisions. Its investigation would be an interesting direction for the future.

5.2 Optimal Team Search

Our first result illustrates that the social planner chooses constant search scopes for each active alliance. However, the specification of these search scopes differs from that dictated by equilibrium.

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

$$\frac{2 \sum_{i \in A} c_i(\sigma_i^A)}{c'_i(\sigma_i^A)} = \sum_{i \in A} \sigma_i^A \quad \forall i \in A.$$

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 $(\sum_{k \in A} \sigma_k^A)^2$ 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,

$$\frac{\sum_{k \in A} c_k(\sigma_k^A)}{(\sum_{k \in A} \sigma_k^A)^2} = \frac{\sum_{k \in A} c_k(\sigma_k^A)}{(\sigma^A)^2}.$$

The identity in the proposition reflects the corresponding first-order condition.

When costs are log-convex, the proposition implies that socially optimal search scopes 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, the social planner internalizes the positive externalities entailed by agents' contributions to the scope of search and thus specifies greater overall search investments. Comparative statics of the socially optimal search scopes resemble those described for equilibrium choices in Section 4.

In equilibrium, Corollary 1 indicated that, as alliances shrink, remaining agents increase their search scope. 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 scope helps. Consequently, the socially optimal search scope of each individual agent *declines*. That is,

Corollary 3 (Optimal Scope and Alliance Size). *Suppose costs are log-convex and the equilibrium and social planner's search scopes are interior. Then, in any alliance,*

an agent's equilibrium search scope is lower than that agent's search scope in the social planner's solution. Furthermore, in the social planner's solution, each agent's search scope 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. For any deterministic sequence of alliances A_1, \dots, A_k exerting optimal search scopes, 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 $d_{A_k} \in \mathbb{R}_+$. Furthermore, the drawdown sizes $\{d_{A_k}\}$ exhibit a recursive structure: for any k ,*

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

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

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

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.

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 scopes are constant over time. Nonetheless, the optimal sequence of active alliances, their corresponding drawdown sizes, and the search scopes do not generally coincide with those prescribed by equilibrium.

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 4 (Longer Optimal Search). *Suppose costs are log-convex and the equilibrium and social planner’s search scopes 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.¹⁷ However, 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.

5.3 Optimal Team Search with Well-ordered Costs

Suppose, as in Section 4.3, that agents’ cost functions are proportional to one another and point-wise ordered: $c_1\beta_1 = c_2\beta_2 = \dots = c_N\beta_N$, where $\beta_1 = 1 < \beta_2 < \dots < \beta_N$.

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

¹⁷In addition, in the Online Appendix, we show a recursive formulation of the social planner’s objective—the resulting welfare—in terms of the optimal drawdowns and search scopes.

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 $B_k = \{k, k + 1, \dots, N\}$ for all $k = 1, \dots, N$. 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 $d_{B \rightarrow B'}$ the socially optimal drawdown size associated with alliance B , when it is followed by alliance B' , as described in Proposition 4. In particular, $d_{B \rightarrow \emptyset}$ 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 $\{d_{B_k \rightarrow \emptyset}\}_{k=1}^N$. Let $L_1 = \arg \max_{k=1, \dots, N} d_{B_k \rightarrow \emptyset}$. 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 $\{d_{B_k \rightarrow B_{L_1}}\}_{k=1}^{L_1-1}$. Let $L_2 = \arg \max_{k=1, \dots, L_1-1} d_{B_k \rightarrow B_{L_1}}$. 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 $B_1, B_{L_1}, \dots, B_{L_{n-1}}$.*




FIGURE 2: Socially optimal exit waves

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 .¹⁸ This allows us to determine the last active alliance chosen by the social planner, B_{L_1} , as in panel (a) of Figure 2.

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 2. 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 2.

5.4 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 $1 < \beta_2 < \beta_3$. 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 3 focuses on the case in which the social planner would cluster all agents' exits (each tick on the axes corresponds to one unit of the corresponding multiplier, so that both β_2 and β_3 range from 0 to 24). The figure depicts the different regions of β_2 and β_3 combinations that generate the four possible structures of equilibrium exit waves. Since $\beta_3 > \beta_2$, all regions are above the gray 45-degree diagonal 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

¹⁸As mentioned, this simplifies the computation problem substantially. Instead of considering $2^N - 1$ alliances, we need to consider only N .




FIGURE 3: Equilibrium exit wave patterns when the optimal policy entails one exit wave including all three agents

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.

6 Conclusions and Discussion

This paper analyzes team search patterns. We show that the equilibrium and socially optimal search scopes are constant within an alliance. However, as alliance members depart, individual search scopes increase in equilibrium and decrease under the optimal policy. We also characterize the deterministic path of exit waves generated in equilibrium. In particular, even when team members are fully heterogeneous, clustered exits may occur. 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, explicit rewards for

innovating early and the utilization of non-Markovian strategies in equilibrium. In the Online Appendix, we also analyze the limitations introduced by a fixed search scope that cannot be altered, and our model’s implications for settings with independent search observations.

6.1 Equilibrium with Penalties for Later Innovations

Suppose stopping earlier grants one an advantage. 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. 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 .¹⁹ As we show in the Online Appendix, the order of exits remains deterministic. Furthermore, as long as both agents are searching, the search scope 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 $g_L(\cdot)$ remains her equilibrium stopping boundary regardless of α and is governed by the drawdown identified in Proposition 2. The follower’s stopping boundary, however, may change.

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 scope when searching within the full team, σ_T denote the total search scope in the full team, and σ_F denote the follower’s optimal solo search scope. Similar calculations to those underlying Proposition 2 yield the follower’s stopping boundary $g_F(\cdot)$:

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

¹⁹The analysis naturally extends to N agents via a decreasing sequence of discounts: $\alpha_0 = 1 \geq \alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_N$. 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.

where

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

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 scope 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 $\alpha < 1$.²⁰ For sufficiently high M , search continuation would again not be profitable. As α increases, the threshold level \bar{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.

6.2 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.²¹ To see why, consider a team of two agents and suppose the optimal search scope can be implemented in equilibrium—say, when there is only one viable scope, $\underline{\sigma} = \bar{\sigma}$. Our results show that the social planner would like agents to

²⁰Specifically, 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.

²¹This contrasts insights on collective experimentation, see [Hörner, Klein, and Rady \(2020\)](#).

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 scopes. 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. 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 will derive. Proofs of this result can be found in various sources, including [Dubins, Shepp, and Shiryaev \(1994\)](#), [Urgun and Yariv \(2021\)](#) and [Peskir \(1998\)](#) among others and hence omitted.

Lemma A.1. *The infinitesimal generator of the two dimensional process $Z = (M, X)$ satisfies the following:*

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 $\frac{\partial V}{\partial M} = 0$.

Proof of Proposition 1. For any agent i in an alliance A , the value function takes the following form

$$V_i^A(M, X) = \mathbb{E} \left[\int_0^{\tau^A} c_i(\sigma_{i,t} | M, X) dt + \mathbb{E} \left[V_i^A(M_{\tau^A}, g^A(M_{\tau^A})) \right] \right], \quad (2)$$

where

$$\mathbb{E}\left[V_i^A(M_{\tau^A}, g^A(M_{\tau^A}))\right] = \begin{cases} M_{\tau^A} & \text{if } g_i^A(M_{\tau^A}) = g^A(M_{\tau^A}) \\ V_i^{A \setminus \{j\}} & \text{if } g_j^A(M_{\tau^A}) = g^A(M_{\tau^A}) > g_i^A(M_{\tau^A}), j \in A \end{cases}.$$

In words, with Markov strategies, agent i 's expected value is derived from two components: the cost accrued until her alliance shrinks, and the continuation value once that happens. If the alliance shrinks with agent i 's departure, her continuation value is simply the maximum value when she exits.

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 $g_i^A(M) \leq X \leq M$. 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 the following recursive fashion:

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

Rearranging terms, we get:

$$\begin{aligned} V_i^A(M, M) - M &= \frac{M - g_i^A(M)}{X - g_i^A(M)} \left[V_i^A(M, X) - M \right. \\ &\quad \left. + \int_{g_i^A(M)}^M G_{g_i^A(M), M}(x, y) c_i(\sigma_i^A(M, y)) \frac{2}{(\sigma^A(M, y))^2} dy \right]. \end{aligned}$$

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 $X \rightarrow g_i^A(M)$ can be written as $\lim_{X \rightarrow g_i^A(M)} \frac{V_i^A(M, X) - M}{X - g_i^A(M)}$. By smooth pasting, this must be equal the derivative of the value from stopping, $\frac{\partial}{\partial x} M = 0$.

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

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

This, in turn, implies that

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

Now, taking the second derivative with respect to X and simplifying yields:

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

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

$$\frac{2c_i(\sigma_i^A(M, X))}{c'_i(\sigma_i^A(M, X))} = \sigma_i^A(M, X) + \sum_{j \in A \setminus \{i\}} \sigma_j^A(M, X).$$

Suppose now that 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:

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

Rearranging terms, we get:

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

Again, taking the limit as $X \rightarrow g(M)$ from above, and letting $\frac{\partial_D V_i^A(M, g(M))}{\partial X}$ denote

the upper Dini derivative of $V_i^A(M, g(M))$ at $g(M)$, we have:²²

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

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

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

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

$$\frac{2c_i(\sigma_i^A(M, X))}{c_i'(\sigma_i^A(M, X))} = \sigma_i^A(M, X) + \sum_{k \in A \setminus \{i\}} \sigma_k^A(M, X).$$

Our assumption that the cost function's second derivative is bounded above zero implies that the conditions of Theorem 6.4 of [Fleming and Rishel \(2012\)](#) are satisfied. Hence, any equilibrium features continuous search scopes within any alliance. In particular, within an active alliance, agents can utilize only one of the solutions for the system above. ■

Proof of Proposition 2. The statement of Proposition 2 is a combination of 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 $g_i^A(M) = \max_{j \in A} g_j^A(M)$ 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

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

²²Since σ_i 's are bounded, V is Lipschitz, hence the Dini derivative is finite.

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 $g_i^A(M)$:

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

which leads to the following solution:

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

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

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

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

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

Claim A.2. Suppose that for some i in an active alliance of A , $\frac{(\sigma^A)^2}{2c_i(\sigma_i^A)} \leq \frac{(\sigma^A)^2}{2c_j(\sigma_j^A)}$ for all $j \in A$. Then i is the first to exit alliance A .²³

Proof of Claim. Suppose $\frac{(\sigma^A)^2}{2c_i(\sigma_i^A)} \leq \frac{(\sigma^A)^2}{2c_j(\sigma_j^A)}$ 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 $A \setminus K$. 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

²³If there are multiple agents who satisfy the condition, all exhibiting the same drawdown size, they all exit jointly, weakly before others.

boundary is characterized by a drawdown. However, from the Claim's restriction,

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

For each $k \in A$, the stopping boundary $g_k^A(M) = M - \frac{(\sigma^A)^2}{2c_k(\sigma_k^A)}$ 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 $0 < \epsilon < \frac{(\sigma^A)^2}{2c_j(\sigma_j^A)} - \frac{(\sigma^A)^2}{2c_i(\sigma_i^A)}$. 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

Proof of Proposition 3. Let $\{\sigma_i^A(M, X, A)\}$ 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 dictated by this optimal solution. Optimality implies that the induced search scopes with A_k should solve:

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

Following similar steps to the proof of proposition 1 in the equilibrium analysis, the continuation HJB for the social planner can be written as:

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

It follows that

$$\frac{2 \sum_{i \in A_k} c_i(\sigma_i^{A_k}(X))}{\sum_{i \in A_k} \sigma_i^{A_k}(X)} = c'_j(\sigma_j^{A_k}(X)) \quad \forall j \in A_k.$$

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

$$W_K(M, X) = \sup_{\tau^K, \{\sigma_i\}_{i \in A_K}} \mathbb{E} \left[|A_K| M_{\tau^K} - \int_0^{\tau^K} \sum_{i \in A_K} c_i(\sigma_i^{A_K}) dt \mid M, X \right].$$

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:

$$g^{A_K}(M) = M - d_{A_K},$$

where $d_{A_k} = \frac{|A_k| \sigma^{A_k}}{2 \sum_{i \in A_k} c_i(\sigma_i^{A_k})}$.

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:

$$W_{K-1}(M, X) = \sup_{\tau^{K-1}, \{\sigma_i\}_{i \in A_{K-1}}} \mathbb{E} \left[|A_{K-1} \setminus A_K| M_{\tau^{K-1}} + W_K(M_{\tau^{K-1}}, g^{A_{K-1}}(M_{\tau^{K-1}})) \mid M, X \right] \\ - \mathbb{E} \left[\int_0^{\tau^{K-1}} \sum_{i \in A_{K-1}} c_i(\sigma_i^{A_{K-1}}) dt \mid M, X \right].$$

By optimality of the stopping time τ^{K-1} , we have smooth pasting of $W_{K-1}(M, X)$ and $W_K(M, X)$. Therefore,

$$W_{K-1}(M, g^{A_{K-1}}(M)) = |A_{K-1} \setminus A_K| M + W_K(M, g^{A_{K-1}}(M)), \\ \frac{\partial W_{K-1}(M, g^{A_{K-1}}(M))}{\partial X} \Big|_{X=g^{A_{K-1}}} = \frac{\partial (|A_{K-1} \setminus A_K| M + W_K(M, g^{A_{K-1}}(M)))}{\partial X} \Big|_{X=g^{A_{K-1}}(M)}.$$

Similar to our equilibrium analysis, and using the notation for the Green function

introduced there, we can write the welfare maximization problem as

$$W_{K-1}(M, X) = |A_{K-1} \setminus A_K| M + W_K(M, g^{A_{K-1}}(M)) \frac{M - X}{M - g^{A_{K-1}}(M)} \\ + W_{K-1}(M, M) \frac{X - g^{A_{K-1}}(M)}{M - g^{A_{K-1}}(M)} - \int_{g^{A_{K-1}}}^M G_{g^{A_{K-1}}(M), M}(X, y) \frac{2 \sum_{i \in A_{K-1}} c_i(\sigma_i^{A_{K-1}})}{(\sigma^{A_{K-1}})^2} dy.$$

Letting X approach $g^{A_{K-1}}(M)$, smooth pasting and rearranging yields:²⁴

$$W_{K-1}(M, X) = |A_{K-1} \setminus A_K| M + W_K(M, g^{A_{K-1}}(M)) \\ + (X - g^{A_{K-1}}(M)) \int_{g^{A_K}(M)}^{g^{A_{K-1}}(M)} \frac{2 \sum_{i \in A_K} c_i(\sigma_i^{A_K})}{(\sigma^{A_K})^2} dx + \int_{g^{A_{K-1}}(M)}^X (X - y) \frac{2 \sum_{i \in A_{K-1}} c_i(\sigma_i^{A_{K-1}})}{(\sigma^{A_{K-1}})^2} dy,$$

Using the closed-form representation of W_K leads to:

$$W_{K-1}(M, X) = |A_{K-1}| M + \frac{1}{2} (g^{A_{K-1}}(M) - g^{A_K}(M))^2 \frac{2 \sum_{i \in A_K} c_i(\sigma_i^{A_K})}{(\sigma^{A_K})^2} \\ + (X - g^{A_{K-1}}(M)) (g^{A_{K-1}}(M) - g^{A_K}(M)) \frac{2 \sum_{i \in A_K} c_i(\sigma_i^{A_K})}{(\sigma^{A_K})^2} \\ + \frac{1}{2} (X - g^{A_{K-1}}(M))^2 \frac{2 \sum_{i \in A_{K-1}} c_i(\sigma_i^{A_{K-1}})}{(\sigma^{A_{K-1}})^2}.$$

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

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

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

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

²⁴The equality follows from

$$\frac{\partial (|A_{K-1} \setminus A_K| M + W_K(M, g^{A_{K-1}}(M)))}{\partial X} \Big|_{X=g^{A_{K-1}}(M)} = (X - g^{A_{K-1}}(M)) \int_{g^{A_K}(M)}^{g^{A_{K-1}}(M)} \frac{2 \sum_{i \in A_K} c_i(\sigma_i^{A_K})}{(\sigma^{A_K})^2} dx.$$

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

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

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

We can then 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

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

■

Lemma A.4. *The set of agents dropping from an alliance is deterministic. That is for any alliance A , for all pairs $(M, X), (M', X')$ such that $G(M, X, A) \neq A$, we have $G(M, X, A) = G(M', X', A)$.*

Proof of Lemma A.4. We prove this result by induction on the size of the initial team N , regardless of the starting values of the maximum and the current value. The claim follows immediately for $N = 1$. In that case, the agent uses a drawdown stopping boundary and the only way for the singleton alliance to change is for the agent to terminate her search.

For the inductive step, assume that for any initial team of size $N - 1$ or less the optimal alliance sequence is deterministic. By Lemma A.3, each of these alliances is associated with a drawdown stopping boundary. Let A_1 be an alliance of size N . The continuation value when M and X are observed is:

$$W_1(M, X) = \mathbb{E} \left[M_{\tau^1} + \max_{A_2 \subsetneq A_1} \{ |A_1 \setminus A_2| M_{\tau^1} + W_{A_2}(M_{\tau^1}, g^{A_1}(M_{\tau^1})) \} - \int_0^{\tau^1} \sum_{i \in A_k} c_i(\sigma_i^{A_1}) dt \right].$$

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. We can therefore write the continuation value as:

$$\begin{aligned} W_1(M, X) = & |A_1 \setminus A_2| M + W_2(M, g^{A_1}(M)) \\ & + \left(X - g^{A_1}(M) \right) \sum_{m=2}^{K-1} \left(\int_{g^{A_{m+1}}(M)}^{g^{A_m}(M)} \frac{2 \sum_{i \in A_{m+1}} c_i(\sigma_i^{A_{m+1}})}{(\sigma^{A_{m+1}})^2} dx \right) \\ & + \int_{g^{A_1}(M)}^X (X - y) \frac{2 \sum_{i \in A_1} c_i(\sigma_i^{A_1})}{(\sigma^{A_1})^2} dy. \end{aligned}$$

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 $d_{A_1} = \frac{|A_1 \setminus A_2|}{2 \left(\frac{\sum_{i \in A_1} c_i(\sigma_i^{A_1})}{(\sigma^{A_1})^2} - \frac{\sum_{i \in A_2} c_i(\sigma_i^{A_2})}{(\sigma^{A_2})^2} \right)}$. Towards a contra-

dition, suppose that at some other path a different alliance is optimally chosen to follow the full alliance A_1 . Call that alliance $\hat{A}_2 \neq A_2$. Similar argument would then imply that the stopping boundary for A_1 is given by $g^{A_1}(M) = M - \hat{d}_{A_1}$, where $\hat{d}_{A_1} = \frac{|A_1 \setminus \hat{A}_2|}{2 \left(\frac{\sum_{i \in A_1} c_i(\sigma_i^{A_1})}{(\sigma^{A_1})^2} - \frac{\sum_{i \in \hat{A}_2} c_i(\sigma_i^{\hat{A}_2})}{(\sigma^{\hat{A}_2})^2} \right)}$. Suppose $d_{A_1} = \hat{d}_{A_1}$. Without loss of generality, as-

sume $\sigma^{A_2} \geq \sigma^{\hat{A}_2}$. We cannot have $\hat{A}_2 \subseteq A_2$. Indeed, if that were the case, the social planner would be indifferent between keeping the agents in $A_2 \setminus \hat{A}_2$ at points at which either \hat{A}_2 or A_2 are chosen to continue. However, our tie-breaking rule implies that such indifferences are broken in favor of stopping; that is, in such cases, the smaller set \hat{A}_2 would be chosen. Therefore, $\hat{A}_2 \not\subseteq A_2$ and there exists an agent $i \in \hat{A}_2 \setminus A_2$. Suppose that whenever the social planner transitions from A_1 to A_2 , she instead transitions to $A_2 \cup \{i\}$, maintaining the search scopes of members of A_2 as before and having agent i search with the lowest scope $\underline{\sigma}$ for a sufficiently small interval of time. In that interval of time, everyone in A_2 benefits. When alliance \hat{A}_2 is picked, agent i uses at least as high a search scope, while benefitting from lower overall search scope. In particular, agent i benefits as well from this change.

Suppose $d_{A_1} \neq \hat{d}_{A_1}$. In this case, the two stopping boundaries identified above, $M - d_{A_1}$ and $M - \hat{d}_{A_1}$ never intersect, in contradiction. ■

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. As introduced in the proof of Corollary 3, we use here the superscripts eq and sp to denote the equilibrium and social planner's solution, respectively. When costs are well-ordered, in equilibrium, in any alliance, all agents utilize the same search scope. 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 $k = 2, \dots, N$. Indeed, in any active alliance A , the equilibrium stopping boundary is governed by drawdown size

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

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 two distinct alliances in the social planner's solution, A_k and A_m , with $k < m$, where $i, j \in A_k$ 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 scope as agent j had originally in the alliances that follow A_k . The overall search scope 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}, \dots, 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, Lemma A.3 identifies the optimal drawdown stopping boundaries associated with that alliance sequence. If the chosen sequence is suboptimal, some of its associated drawdown sizes might be negative or zero, implying the corresponding alliance is utilized for no length 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 scope within the alliance by $\tilde{\sigma}^k$ and the consequent overall search cost within that alliance by \tilde{c}^k .

Lemma A.5. *For any m, j, k such that $m < j < k$, 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 $d_{B_m \rightarrow B_k} > d_{B_j \rightarrow B_k}$.*

Proof of Lemma A.5. from the characterization of drawdowns in the well-ordered settings, $d_{B_m \rightarrow B_k} \neq d_{B_j \rightarrow B_k}$. Suppose that $d_{B_m \rightarrow B_k} < d_{B_j \rightarrow B_k}$. Since B_k is preceded by B_m in the optimal sequence, $d_{B_m \rightarrow B_k} > 0$. It then follows that $d_{B_j \rightarrow B_k} > 0$. 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 $m < k$, $d_{B_m \rightarrow \emptyset} > d_{B_k \rightarrow \emptyset}$ implies $d_{B_m \rightarrow B_k} > d_{B_k \rightarrow \emptyset}$.*

Proof of Lemma A.6. $d_{B_m \rightarrow \emptyset} > d_{B_k \rightarrow \emptyset}$ implies

$$\frac{1}{N-m} \frac{\tilde{c}^m}{\tilde{\sigma}^m} < \frac{1}{N-k} \frac{\tilde{c}^k}{\tilde{\sigma}^k} \implies \frac{1}{k-m} \left(\frac{\tilde{c}^m}{\tilde{\sigma}^m} - \frac{\tilde{c}^k}{\tilde{\sigma}^k} \right) < \frac{1}{N-k} \frac{\tilde{c}^k}{\tilde{\sigma}^k},$$

illustrating the claim. ■

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

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

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

Simply summing the inequalities and reorganizing yields the implied statement. ■

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

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

$$\frac{1}{N-k} \frac{\tilde{c}^k}{\tilde{\sigma}^k} < \frac{\tilde{c}^{k-1}}{\tilde{\sigma}^{k-1}} - \frac{\tilde{c}^k}{\tilde{\sigma}^k} \implies \frac{N-(k-1)}{N-k} \frac{\tilde{c}^k}{\tilde{\sigma}^k} < \frac{\tilde{c}^{k-1}}{\tilde{\sigma}^{k-1}} \implies d_{B_k \rightarrow \emptyset} > d_{B_{k-1} \rightarrow \emptyset}.$$

But this inequality implies that

$$\frac{1}{N-k} \frac{\bar{c}^k}{\bar{\sigma}^k} < \frac{1}{N-(k-1)} \frac{\bar{c}^{k-1}}{\bar{\sigma}^{k-1}} \implies \frac{\bar{c}^{k-1}}{\bar{\sigma}^{k-1}} - \frac{\bar{c}^k}{\bar{\sigma}^k} > \frac{1}{(N-(k-1))} \frac{\bar{c}^{k-1}}{\bar{\sigma}^{k-1}} \implies d_{B_{k-1} \rightarrow \emptyset} > d_{B_{k-1} \rightarrow B_k}.$$

■

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

Proof of Lemma A.9. Suppose not, so that, from 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 settings, $d_{B_k \rightarrow \emptyset} \neq d_{B_l \rightarrow \emptyset}$. The social planner would, then, benefit from transitioning from B_l to B_k instead of exiting all members of B_l since $d_{B_k \rightarrow \emptyset} > d_{B_l \rightarrow \emptyset} > 0$, in contradiction. ■

Lemma A.10. *If k satisfies $\max_j d_{B_j \rightarrow \emptyset} = d_{B_k \rightarrow \emptyset}$, 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 $B_k = B_N = \{N\}$.

For the proof, it is useful to notice that our entire analysis does not hinge on the range of viable search scopes coinciding across agents. In fact, the analysis would go through in its entirety if each agent i had an individual range of scope $[\underline{\sigma}_i, \bar{\sigma}_i]$, as long as all solutions remain interior.

Assume the statement is true for sets up to cardinality n . We show the statement holds for $|B_k| = n + 1$ (so that $k = N - n + 1$). 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 $d_{B_m \rightarrow \emptyset} > d_{B_l \rightarrow \emptyset}$ 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 $\hat{c}(\cdot)$ defined so that $\hat{c}(\sigma)$ is the minimal overall cost in B_m required for implementing an overall search scope σ . That is, if

$$\sigma_m^{B_m}, \dots, \sigma_N^{B_m} \in \arg \min_{\hat{\sigma}_j^{B_m} \in [\underline{\sigma}, \bar{\sigma}] \forall j, \sum_{j=m}^N \hat{\sigma}_j^{B_m} = \sigma} \sum_{j=m}^N c(\hat{\sigma}_j^{B_m}),$$

then $\hat{c}(\sigma) = \sum_{j=m}^N c(\hat{\sigma}_j^{B_m})$. Under this definition, $\sigma \in [(N - m + 1)\underline{\sigma}, (N - m + 1)\bar{\sigma}]$.

In the equivalent problem, we have m agents $1, 2, \dots, m-1, M$.²⁵ From our construction, in the optimal solution, for any $j = 1, \dots, m-1$, the corresponding draw-down size $d_{\{j, \dots, M\} \rightarrow \emptyset}$ coincides with the optimally-set drawdown size $d_{B_j \rightarrow \emptyset}$ in our original problem. Furthermore, $d_{\{M\} \rightarrow \emptyset}$ coincides with $d_{B_m \rightarrow \emptyset}$ in our original problem. Therefore, $\max_{j \in \{1, \dots, m-1, M\}} d_{\{j, \dots, M\} \rightarrow \emptyset} = d_{\{k, \dots, M\} \rightarrow \emptyset}$. By our induction hypothesis, $\{j, \dots, M\}$ 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 $p \in \{k, \dots, N-1\}$, 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 $d_{B_N \rightarrow \emptyset} > d_{B_{N-1} \rightarrow B_N} > d_{B_{N-2} \rightarrow B_{N-1}} > \dots > d_{B_k \rightarrow B_{k+1}}$. Applying Lemma A.8 repeatedly implies that $d_{B_N \rightarrow \emptyset} > d_{B_{N-1} \rightarrow \emptyset} > d_{B_{N-2} \rightarrow \emptyset} \dots > d_{B_{k+1} \rightarrow \emptyset}$. The assumed maximality of $d_{B_k \rightarrow \emptyset}$ implies, in particular, that $d_{B_k \rightarrow \emptyset} > d_{B_N \rightarrow \emptyset}$ that, combined with the above, yields $d_{B_k \rightarrow \emptyset} > d_{B_N \rightarrow \emptyset} > d_{B_{k+1} \rightarrow \emptyset}$. By Lemma A.7, we then have that $d_{B_k \rightarrow B_{k+1}} > d_{B_N \rightarrow \emptyset}$. 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 $B_{k+1}, B_{k+2}, \dots, B_N$, in contradiction.

Case 2: There does not exist any $p \in \{k, \dots, N-1\}$ 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 $d_{B_k \rightarrow \emptyset}$ implies that $d_{B_k \rightarrow \emptyset} > d_{B_N \rightarrow \emptyset}$. By Lemma A.6, $d_{B_k \rightarrow B_N} > d_{B_N \rightarrow \emptyset}$ and by Lemma A.5, $d_{B_l \rightarrow B_N} > d_{B_k \rightarrow B_N} > d_{B_N \rightarrow \emptyset}$. 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 $p, q \in \{k, \dots, N-1\}$ such that B_p is part of the optimal sequence but B_q is not. Here we have two subcases:

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

²⁵Our assumption that all alliance optimally have members using an interior search scope guarantees that this fictitious agent would choose an interior search scope as well.

$d_{B_l \rightarrow B_{N-1}} > d_{B_k \rightarrow B_{N-1}} > d_{B_{N-1} \rightarrow \emptyset}$. Finally, $d_{B_N \rightarrow \emptyset} > d_{B_{N-1} \rightarrow \emptyset}$ implies that

$$\frac{1}{2} \frac{\tilde{c}^{N-1}}{\tilde{\sigma}^{N-1}} > \frac{\tilde{c}^N}{\tilde{\sigma}^N} \implies \frac{1}{2} \frac{\tilde{c}^{N-1}}{\tilde{\sigma}^{N-1}} < \frac{\tilde{c}^{N-1}}{\tilde{\sigma}^{N-1}} - \frac{\tilde{c}^N}{\tilde{\sigma}^N} \implies d_{B_{N-1} \rightarrow \emptyset} > d_{B_{N-1} \rightarrow B_N}.$$

Thus, $d_{B_l \rightarrow B_{N-1}} > d_{B_k \rightarrow B_{N-1}} > d_{B_{N-1} \rightarrow \emptyset} > d_{B_{N-1} \rightarrow B_N}$. 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 $p \in \{k, \dots, N-2\}$. We can now emulate the argument above pertaining to the construction of an equivalent problem in which agents $\{p, \dots, N-1\}$ 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 $\max_j d_{B_j \rightarrow \emptyset} = d_{B_k \rightarrow \emptyset}$.

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 omitted.

Lemma A.11. *Consider B_k where k is such that $d_{B_k \rightarrow B_{L_1}} > d_{B_j \rightarrow B_{L_1}}$ for all $j < L_1$ and B_{L_1} as 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 $d_{B_k \rightarrow B_{L_1}} > d_{B_j \rightarrow B_{L_1}}$ for all $j < L_1$ and B_{L_1} as 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 $d_{B_k \rightarrow B_{L_1}} > d_{B_j \rightarrow B_{L_1}}$ 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*.
- 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.
- 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.
- Deb, J., A. Kuvalekar, and E. Lipnowski (2020). Fostering collaboration. *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 besse processes. *Theory of Probability & Its Applications* 38(2), 226–261.
- Fleming, W. H. and R. W. Rishel (2012). *Deterministic and stochastic optimal control*, Volume 1. Springer Science & Business Media.
- 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 (2020). 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.

- Keller, G., S. Rady, and M. Cripps (2005). Strategic experimentation with exponential bandits. *Econometrica* 73(1), 39–68.
- Marx, L. M. and S. A. Matthews (2000). Dynamic voluntary contribution to a public project. *The Review of Economic Studies* 67(2), 327–358.
- Murto, P. and J. Välimäki (2011). Learning and information aggregation in an exit game. *The Review of Economic Studies* 78(4), 1426–1461.
- 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. (2019). Collaborative search for a public good.
- 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: Journal of the Econometric Society*, 641–654.
- Yildirim, H. (2006). Getting the ball rolling: Voluntary contributions to a large-scale public project. *Journal of Public Economic Theory* 8(4), 503–528.