Collaborative Search for a Public Good

by

Masha Titova

UC San Diego

May 2, 2019

Motivation

- A **group of agents** often must search for possible solutions to a given problem.
- The resulting solution, as well as the information gathered during search, are often a **public good**.
- Examples of **collaborative search for a public good** are
  - consumer search,
  - search for investment opportunities,
  - adoption of new technologies,
  - research and development.

Modeling Choices

- I extend the sequential search model of Weitzman (1979) to 2 searchers.
- Each public good (project) is represented by a **box**:
  - uncertain **reward** revealed upon paying a **search cost**.
- Once the search process is over, the best uncovered project is implemented.

Questions Asked

- What is the optimal **search order** among risky alternatives?
- What are the **incentives to free ride** on colleague's search efforts?
- How does collaborative search by a group of people compare to the **(socially optimal) individual search**?

Preview of the Results

- The **search order** and the **stopping rule** are **that of a social planner**:
  - the same project is implemented at the end,
  - the same information is gathered in the same order.
- There is **delay at each stage of the search process**
  - each agent free rides in hopes that her colleague will pay the search cost.
- Overall, collaborative search is **inefficient**, but **preferred by each individual agent** to searching alone.

Literature

- Collective Experimentation:
  - Bolton and Harris (1999), Keller et al. (2005), Keller and Rady (2010), Bonatti and Rantakari (2016).

    What I do: consider a multi-armed bandit and study the *order* and *stopping rule*.

- Delegation and Approval of Experimentation:
  - Manso (2011), Lewis (2012), Halac et al. (2016), Guo (2016), McClellan (2019).

    What I do: compare the optimal search by one agent to the optimal search in teams.

- Collaboration in Teams:
  - Bonatti and Hörner (2011), Campbell et al. (2014), Georgiadis (2015).

    What I do: agents choose the order in which to search and decide when to stop.

- Dynamic Provision of Public Goods:
  - Fershtman and Nitzan (1991), Marx and Matthews (1991), Admati and Perry (1991), Compte and Jehiel (2004), Kessing (2007), Bowen et al. (2019).

    What I do: study *search* for a public good.

The Model

Setup

- 2 players:
  - risk-neutral,
  - maximize the expected present value of the best uncovered reward (**free recall**),
  - discount time at an exponential rate $\delta = e^{-r\Delta t}$.
- Each period, one player is randomly (with prob. 1/2) **chosen** to perform the search.
- The game ends if either
  - there are no options left to search among,
  - the players agree **unanimously** to terminate the search process.

Actions

- When player i is **chosen**, she can
  - open exactly one box of her choice,
  - do nothing,
  - propose to terminate the game.
- In the latter case, her **opponent** (player j) can
  - accept the offer,
  - reject it.

Public Goods

- N unopened **boxes**. Box $b_k \equiv \left(c_k, F_k(\cdot)\right)$
  - contains an uncertain **reward** $x_k \sim F_k(\cdot)$ (independent),
  - c_k is the **search cost** paid to learn the contents of the box,
  - the reward is drawn in the following period.
- Initially, there is a fallback reward $z_0 = 0$.

State Variables

- At each stage, the **state** $s = (z, \mathcal{B}^c)$ of the problem is
  - the **current best option** z,
    - e.g. at $t = 0$ it is $z_0 = 0$;
  - the **set of unopened boxes** $\mathcal{B}^c$.

Markov Perfect Equilibrium

- Let $\Phi_i^{ch}(s)$ and $\Phi_i^{op}(s)$ be the **discounted continuation payoff**, depending on player i's role in state s.
- Let $\alpha_i(s) \equiv \left(\alpha_i^{ch}(s), \alpha_i^{op}(s)\right)$ be a **stationary Markov strategy**.

A pair of strategies $\left(\alpha_1(s), \alpha_2(s)\right)$ is an **MPE** if $\forall i, \forall j \neq i, \forall s$

$$\alpha_i^{ch}(s) = \arg\max_{\hat{\alpha}_i^{ch}(s)} \Phi_i^{ch}(s), \quad \alpha_i^{op}(s) = \arg\max_{\hat{\alpha}_i^{op}(s)} \Phi_i^{op}(s)$$

given $\left(\alpha_2^{ch}(s), \alpha_2^{op}(s)\right)$ and subject to

$$\Phi_i^{ch}(z, \emptyset) = \Phi_i^{op}(z, \emptyset) = z.$$

One Box

Social Planner: Weitzman (1979)

- The **social planner** solves the **individual search problem**.
- If there is only one box *left*, the SP opens it iff

$$-c_k + \delta S(z, F_k) \geq z, \tag{SP}$$

where

$$S(z, F_k) \equiv \mathbb{E}\left[\max\{z, x_k\}\right] = z \int_{-\infty}^{z} dF_k(z) + \int_{z}^{+\infty} x dF_k(x).$$

Reservation Value of a Box

- Let $\bar{z}_k$ solve
$$-c_k + \delta S(\bar{z}_k, F_k) = \bar{z}_k.$$

- Then, it is easy to show that
$$-c_k + \delta S(z, F_k) \geq z \iff \bar{z}_k \geq z. \qquad \text{(SP)}$$

- $\bar{z}_k$ is the **reservation value** of this box b_k that **contains all relevant information about this box**.

The SP **opens** box b_k iff **the box is good enough** i.e. when the reservation value of this box is higher than the current best option.

2 agents, 1 box: the Opponent

- When the opponent receives a termination offer, he can
  - **accept**, get z immediately,
  - **reject**, eventually open the box, and get

$$\frac{\delta}{2-\delta} \cdot \left[-c_k + \delta S(z, F_k) \right];$$

$$= \frac{1}{2}\delta + \left(\frac{1}{2}\delta\right)^2 + \ldots = \frac{\delta}{2-\delta}$$

- the offer is **rejected** if and only if

$$z \leq \frac{\delta}{2-\delta} \cdot \left[-c_k + \delta S(z, F_k) \right] \iff z \leq z_k^{IR}. \qquad \text{(IR)}$$

2 agents, 1 box: the Chosen Player

- Next consider the problem of the **chosen** player.
- If $z > \bar{z}_k$, **proposing termination** is strictly dominant
  - this offer is always accepted since $z_k^{IR} < \bar{z}_k$.
- If $z \leq \bar{z}_k$, the chosen player can do better by mixing between
  - **opening the box**,
  - **doing nothing**.

The Equilibrium in Mixed Strategies

- Suppose each player, when chosen, **opens the box** with prob. π and **does nothing** with prob. $(1-\pi)$.

- In equilibrium, the chosen player must be indifferent btw
  - **opening herself**: $-c_k + \delta S(z, F_k)$,
  - **someone opening it in the future**:

$$\frac{\pi\delta}{1-(1-\pi)\delta} \cdot \left[-\frac{c_k}{2} + \delta S(z, F_k)\right]$$

$$= \pi\delta + (1-\pi)\pi \cdot \delta^2 + (1-\pi)^2\pi \cdot \delta^3 + \cdots = \frac{\pi\delta}{1-(1-\pi)\delta}$$

    – the search cost is paid $1/2$ the time in expectation

- π is obtained from the indifference condition.

The Equilibrium

- **The chosen player**
  - if $z \leq \bar{z}_k$, **opens the box** b_k with prob.

$$\pi_k = \begin{cases} \frac{2(1-\delta)}{\delta c_k}\left[-c_k + \delta S(z, F_k)\right] < 1 \text{ if } c_k > S(z, F_k) \cdot \frac{2\delta(1-\delta)}{2-\delta}, \\ 1 \text{ otherwise }, \end{cases}$$

    and **does nothing** with prob. $1 - \pi_k$;
  - if $z > \bar{z}_k$, proposes to terminate the game.
- **The opponent**
  - **accepts** the termination proposal if $z > z_k^{IR}$;
  - **rejects** the proposal if $z \leq z_k^{IR}$.

Delay and Welfare Implications

- On equilibrium path, the box is opened *eventually* if $z \leq \bar{z}_k$
  - this is the **socially optimal** cutoff.
- For *large* search costs, the box is opened with a **delay**
  - whenever $\pi_k < 1$, the chosen player is **free riding**,
  - if Δt is the time interval between periods, then the **expected delay** is $\Delta t \cdot \frac{1-\pi_k}{\pi_k}$.
- Each agent pays the search cost 1/2 of the time on average.

Properties of π_k

Higher π means less delay.

- For **very low** values of c_k, there is **no delay** because it is strictly dominant to open the box right away.
- Otherwise, $\pi_k(z)$ is **increasing** and convex in z.
-
- **Comparative statics**: $\pi_k(z)$ is **increasing** in the reservation value of the box, i.e. as
  - the search cost c_k decreases,
  - the distribution of rewards gets "better" (in terms of FOSD or MPS).

Many Boxes

Social Planner: Optimal Search Protocol

Theorem : Weitzman (1979)

- **Selection Rule**: if a box is to be opened, it should be that closed box with *highest reservation value*.
- **Stopping Rule**: terminate search whenever the best sampled reward exceeds the reservation value of every closed box.

Collaborative Search: Optimal Search Protocol

- Let $\bar{z}_k = \max_{b_l \in \mathcal{B}^c} \bar{z}_l$.
- **The chosen player**
  - if $z \leq \bar{z}_k$, **opens the box** b_k with prob. $\tilde{\pi}_k \in (0, 1]$ and **does nothing** with prob. $1 - \tilde{\pi}_k$;
  - if $z > \bar{z}_k$, proposes to terminate the game.
- **The opponent**, upon receiving a termination offer
  - **accepts** the termination proposal if $z > \tilde{z}_k^{IR}$;
  - **rejects** the proposal if $z \leq \tilde{z}_k^{IR}$.

Properties of the Equilibrium

- The **search order** and **termination rule** are myopic
  - only depend on the highest reservation value $\bar{z}_k$,
  - **socially optimal** on the equilibrium path.
- **The prob. of opening the box** $\tilde{\pi}_k(s)$ is NOT myopic
  - can only be estimated numerically,
  - known lower bound π_k (from the one box case),
  - less than one for large enough search costs $\Longrightarrow$ **delay at each stage of the learning process**.

Dynamics of the Delay

- How does the delay change as they search?
  - The more boxes are opened, **the better the uncovered reward**, so

    $$z \uparrow \Longrightarrow \pi \uparrow \text{ so the delay decreases.}$$

  - The more they search, **the worse boxes are left** so

    $$\bar{z}_k \downarrow \Longrightarrow \pi \downarrow \text{ so the delay increases.}$$

Discussion

- All results (probably) still **hold** if
  - there are *N* players,
  - players alternate or are chosen with unequal probability,
  - there is no option to do nothing.
- The results **do not hold** if players value boxes differently:
  - the best uncovered reward is not a public good,
  - they have different discount factors,
  - players have diferent costs of opening the same box.

Conclusion

- This paper examines a model of **sequential search** for a **public good** by a **group of agents**.
- I find that
  - the **search order** and **stopping rule** are **socially optimal**;
  - **delay** occurs **at every stage of the search process** because agents free ride;
  - each agent prefers to search in group rather than by herself.

Thank You!

Bellman Equation for the Social Planner

- The **Bellman equation** is

$$\Phi(s) = \max \left\{ z, \max_{b_k \in \mathcal{B}^c} \left\{ -c_k + \delta \Phi\left(s^{-b_k}\right) \right\} \right\}$$

where

- payoff if she quits
- cont. value of opening the *best* box
  - the new state $s^{-b_k} \equiv \left(\mathbb{E}\left[\max\{z, x_k\} \right], \mathcal{B}^c \setminus \{b_k\} \right)$

Properties of π_k

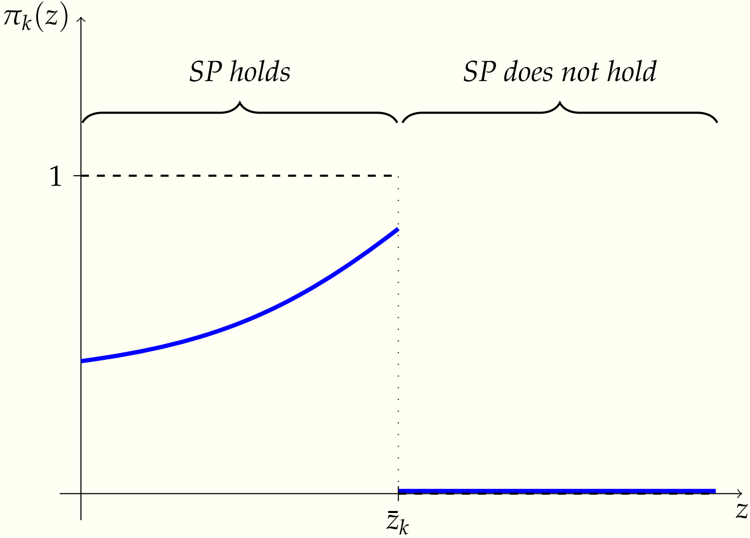

go back

Bellman Equations for 2 Searchers

- Let $\bar{\Phi}_i = 1/2\Phi_i^{ch}(s) + 1/2\Phi_i^{op}(s)$ be the **average discounted continuation payoff**.
- When player i is **chosen**, her Bellman equation is

$$\Phi_i^{ch}(s) = \max_{\alpha_i^{ch}} \left\{ \alpha_j^{op}(s) \cdot z, \delta\bar{\Phi}_i(s), \max_{b_k \in \mathcal{B}^c} \left\{ -c_k + \delta\bar{\Phi}_i\big(s^{-b_k}\big) \right\} \right\},$$

- When player i is the **opponent**, her Bellman equation is

$$\Phi_i^{op}(s) = \max_{\alpha_i^{op}} \left\{ \mathbb{1}_{\{\alpha_j^{ch}(s)=T\}} \cdot r_i \cdot z,\ \delta\bar{\Phi}_i(s') \right\},$$

$$\text{s.t. } s' = \begin{cases} s & \text{if } \alpha_j^{ch}(s) = T, r_i = 0 \text{ or } \alpha_j^{ch}(s) = \varnothing, \\ s^{-b_k} & \text{if } \alpha_j^{ch}(s) = b_k. \end{cases}$$