Collaborative Search for a Public Good

by

Maria (Masha) Titova

UC San Diego

May 2, 2019

MOTIVATION

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

- 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**?

- **search order** and **stopping rule** are **that of a social planner**:
  - same project is implemented in the end
  - same information is gathered in the same order
- there is **delay at each stage of 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 multi-armed bandit and study the *order* and *stopping rule*

- **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), Bowen et al. (2019)

    what I do: study *search* for a public good

MODEL

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

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

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

- let $\Phi_i^{ch}(s)$ and $\Phi_i^{op}(s)$ be **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**

Theorem

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** solves **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

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

- let $\bar{z}_k$ solve

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

- then,

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

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

Theorem

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

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

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

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

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

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

2 agents, 1 box: chosen player

- consider problem of **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 $z \leq \bar{z}_k$, chosen player can do better by mixing between
  - **opening the box**
  - **doing nothing**

- suppose each player, when chosen, **opens box** with prob. π and **does nothing** with prob. $(1-\pi)$.
- in equilibrium, chosen player must be indifferent btw
  - **opening herself**: $-c_k + \delta S(z, F_k)$
  - **someone opening it in the future**:

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

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

    - the search cost is paid half the time in expectation

- π is obtained from the indifference condition

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

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

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

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

Corollary

Higher π means less delay

- for **very low** values of c_k, there is **no delay** because it is strictly dominant to open 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
  - search cost c_k decreases
  - distribution of rewards gets "better" (in terms of FOSD or MPS)

MANY BOXES

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 best sampled reward exceeds reservation value of every closed box

- let $\bar{z}_k = \max_{b_l \in \mathcal{B}^c} \bar{z}_l$
- **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
- **opponent**, upon receiving a termination offer
  - **accepts** termination proposal if $z > \tilde{z}_k^{IR}$
  - **rejects** proposal if $z \leq \tilde{z}_k^{IR}$

- **search order** and **termination rule** are myopic
  - only depend on highest reservation value $\bar{z}_k$
  - **socially optimal** on equilibrium path
- **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**

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

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

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

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

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

- this paper examines a model of **sequential search** for a **public good** by a **group of agents**
- I find that
  - **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** is

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

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

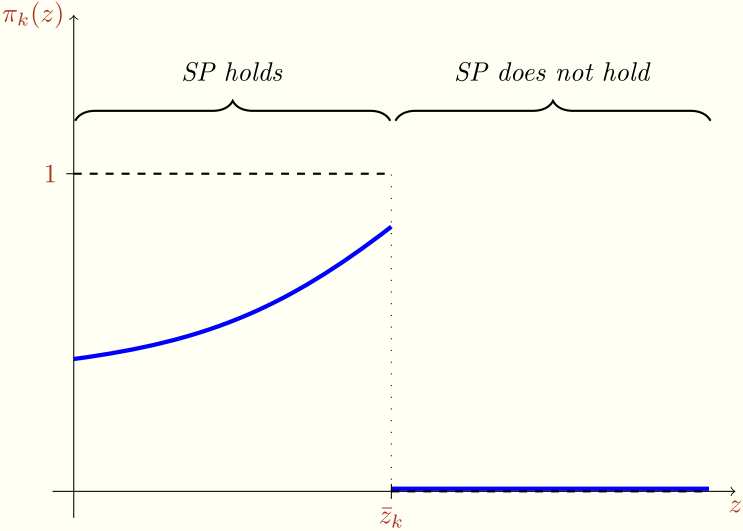

go back

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

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

- when player i is **opponent**, her Bellman equation is

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

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