Dynamic Choice of Information Sources

Tatiana Mayskaya

California Institute of Technology / Higher School of Economics

The 28th International Conference on Game Theory
20 July 2017

Outline

Example: Develop Cure for Disease

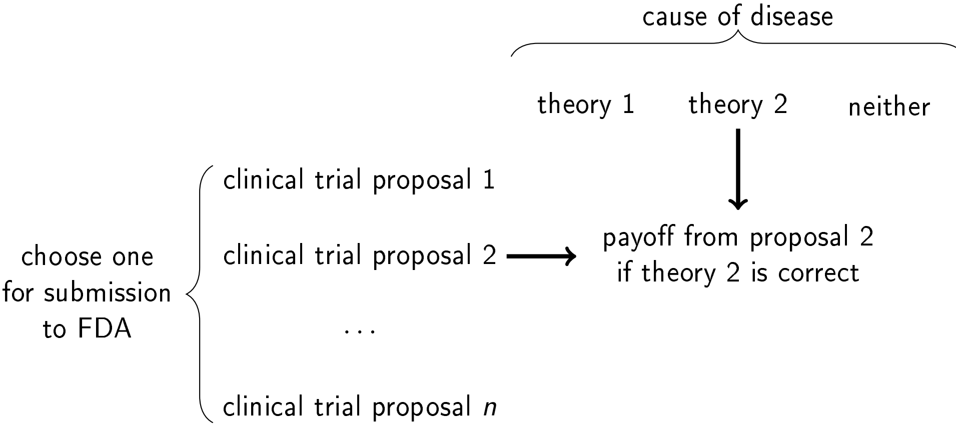

Example: Develop Cure for Disease

information collection about the true cause of the disease ⟶ choose proposal for submission to FDA ⟶ get payoff from the chosen proposal

Example: Develop Cure for Disease

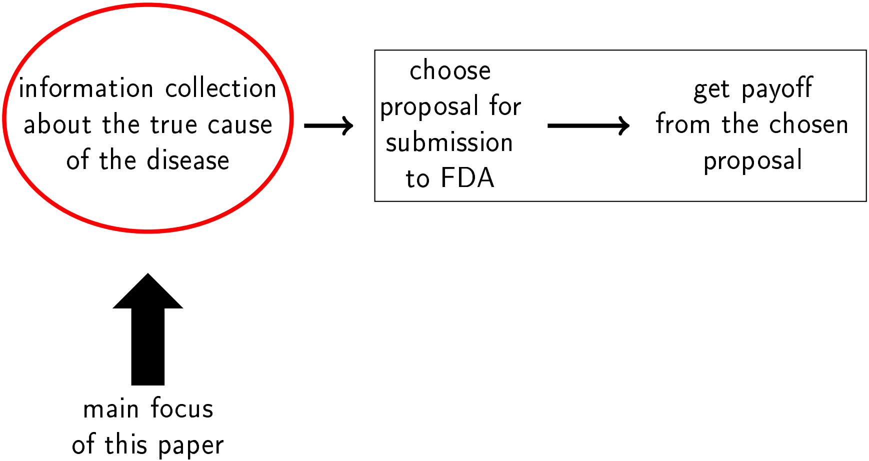

Outline

Modeling Information Collection

- one way to collect information
  ⇒ when to stop collecting information? *dynamic problem*

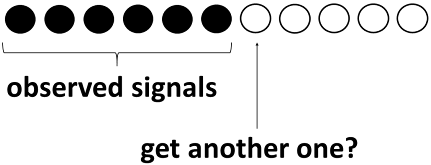

Modeling Information Collection

- one way to collect information ⇒ when to stop collecting information? *dynamic problem*

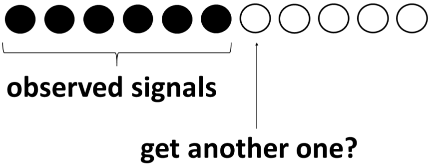

- information collection is a one-time event ⇒ what kind of information to collect? *static problem*

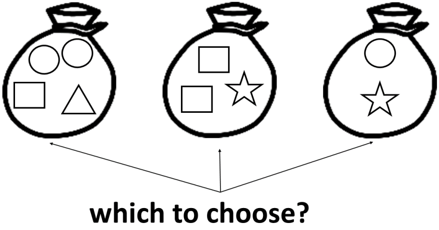

Modeling Information Collection

- one way to collect information ⇒ when to stop collecting information? *dynamic problem*

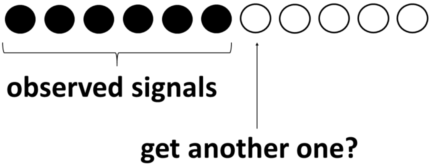

- information collection is a one-time event ⇒ what kind of information to collect? *static problem*

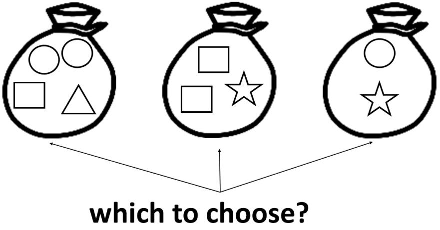

- This paper: many ways to collect information + information collection is a process ⇒ *Dynamic Choice of Information Sources*

get another one? if so, which one?

Information Sources

information source 1 = search for a proof that theory 1 is correct
information source 2 = search for a proof that theory 2 is correct

	true cause of disease		
	theory 1 *problem with head*	**theory 2** *problem with heart*	**neither** *problem with something else*
source 1 *looking in head*	at each moment, $\exists$ positive prob. to find evidence for theory 1	no chance to find anything	no chance to find anything
source 2 *looking in heart*	no chance to find anything	at each moment, $\exists$ positive prob. to find evidence for theory 2	no chance to find anything

▸ formal setup

Information Sources

information source 1 = search for a proof that theory 1 is correct
information source 2 = search for a proof that theory 2 is correct

Cost of information collection:

1. **Per unit of time cost**: search during $\Delta t \Rightarrow$ pay const$\times\Delta t$.
2. **No discounting**.

▸ formal setup

Outline

Contribution - I. Individual decision maker problem

What is the optimal strategy?

Assume: same cost of searching and same search effectiveness.

Tradeoff: search for a proof for the most likely theory (belief-based search) vs. for the most important theory (payoff-based search).

Example:

- prob(problem with head)>prob(problem with heart) $\Rightarrow$
  Belief-based search: looking in head.
- payoff from any proposal depends *only* on whether the problem is in the heart or not $\Rightarrow$
  Payoff-based search: looking in heart.

formal statement

Contribution - I. Individual decision maker problem

What is the optimal strategy?

Assume: same cost of searching and same search effectiveness.

<u>Tradeoff</u>: search for a proof for the most likely theory (belief-based search) vs. for the most important theory (payoff-based search).

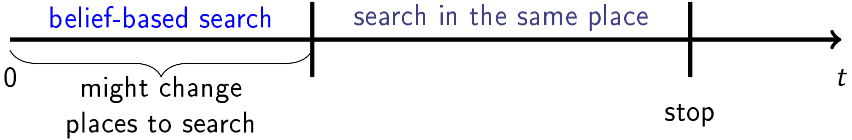

▸ formal statement

Contribution - I. Individual decision maker problem

What is the optimal strategy?

Assume: same cost of searching and same search effectiveness.

Tradeoff: search for a proof for the most likely theory (belief-based search) vs. for the most important theory (payoff-based search).

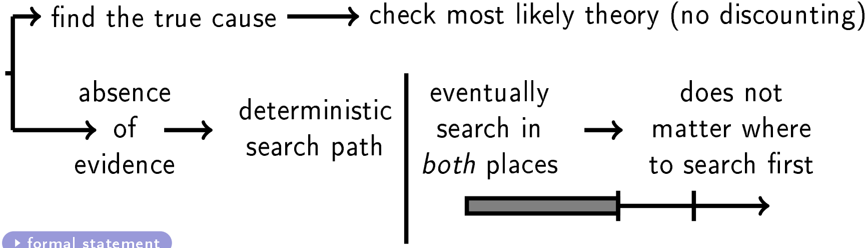

formal statement

Contribution - I. Individual decision maker problem

What is the optimal strategy?

Assume: same cost of searching and same search effectiveness.

Tradeoff: search for a proof for the most likely theory (belief-based search) vs. for the most important theory (payoff-based search).

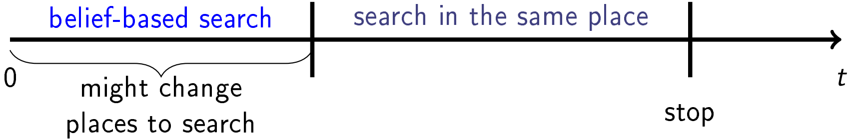

If Prob(neither theory is correct)=0, then

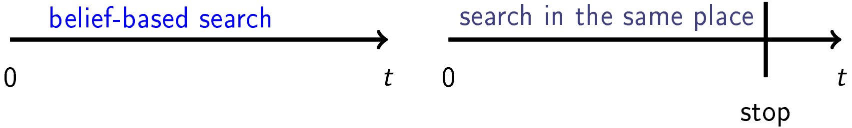

▸ formal statement

Outline

Example

Setup

Result 1: Optimal Learning Strategy

Result 2: Information Sources as Complements and Substitutes

Literature Review

Conclusion

Contribution - II. Market for information

The optimal strategy shapes the demand for information.
How does this strategy characterize the sources as information *goods*?
Do they complement or substitute each other?

▸ formal statement

Contribution - II. Market for information

The optimal strategy shapes the demand for information.
How does this strategy characterize the sources as information *goods*?
Do they complement or substitute each other?

- depends on current beliefs $\Rightarrow$ change over time

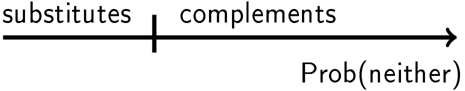

- depends on per unit of time cost (price) of information

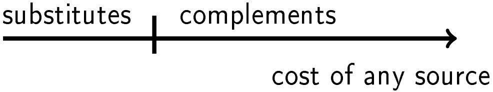

▸ formal statement

Literature Review

1. Closest paper is Che and Mierendorff (2017): Prob(neither)=0, discounting
2. Drift-diffusion model: info about alternatives directly, not through a state of the world; Fudenberg, Strack, and Strzalecki (2015) and Ke, Shen, and Villas-Boas (2016): Brownian motion; Pancs and Nikandrova (2017): Poisson process
3. Rational inattention theory: general information structures; static models except Zhong (2017) who shows that Poisson process that seeks most likely state is the optimal information structure
4. Multi-armed bandit problem: each source gives both the payoff and information about the payoff distribution (clinical trials)
5. Sequential optimal experimental design: approximate solution except Liang, Mu, and Syrgkanis (2017): learning from a *finite* set of *Gaussian* signals
6. Search problem: minimize cost of learning conditional on finding the state (Ahlswede&Wegener 1987), or finding the state has a direct consequences to the payoff (Fershtman&Rubinstein 1997, Matros&Smirnov 2016)
7. Complements & Substitutes: Gul and Pesendorfer (2012), Chen and Waggoner (2016)

Conclusion

1. In a long-term research, the only goal should be to search for the truth. Why we want to know the truth matters only on the last stage of the research.
2. Whether two information goods are complements or substitutes can change over time and with the price.

Payoff Matrix and Beliefs

$\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$ — set of alternatives

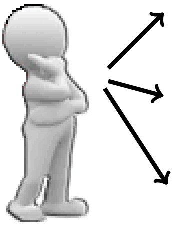

	State 1	State 2	State 3
Beliefs:	p_1	p_2	$1 - p_1 - p_2$
Alternative a_1	$u_1(a_1)$	$u_2(a_1)$	$u_3(a_1)$
Alternative a_2	$u_1(a_2)$	$u_2(a_2)$	$u_3(a_2)$
$\ldots$	$\ldots$	$\ldots$	$\ldots$

WLOG: $u_1(a_1) = \max_{a \in \mathcal{A}} u_1(a)$, $u_2(a_2) = \max_{a \in \mathcal{A}} u_2(a)$

Back to main slides

Information Collection Process

at each moment of time,
choose how to split attention
between two sources

Information source k $(k = 1, 2)$:

$$X^{(k)} = \begin{cases} \text{Poisson process with intensity } \lambda_k, & \text{state } k, \\ 0, & \text{otherwise.} \end{cases}$$

Use source k during $[t_1, t_2]$ =

observe $X^{(k)}_{t_2} - X^{(k)}_{t_1}$ and pay $c_k(t_2 - t_1)$.

Back to ▸ main slides

Information Collection Process

Consider $[t, t + dt]$.

		state 1 theory 1 *problem w/ head*	**state 2** theory 2 *problem w/ heart*	**state 3** neither
source 1 *looking in head*	pay $c_1 dt$	$\begin{cases} 1, & \text{w/pr } \lambda_1 dt \\ 0, & \text{o/w} \end{cases}$	0	0
source 2 *looking in heart*	pay $c_2 dt$	0	$\begin{cases} 1, & \text{w/pr } \lambda_2 dt \\ 0, & \text{o/w} \end{cases}$	0

Back to main slides

Information Collection Process

Consider $[t, t + dt]$.

		state 1 theory 1 *problem w/ head*	**state 2** theory 2 *problem w/ heart*	**state 3** neither
source 1 *looking in head*	pay $c_1 dT_{t,1}$	$\begin{cases} 1, & \text{w/pr } \lambda_1 dT_{t,1} \\ 0, & \text{o/w} \end{cases}$	0	0
source 2 *looking in heart*	pay $c_2 dT_{t,2}$	0	$\begin{cases} 1, & \text{w/pr } \lambda_2 dT_{t,2} \\ 0, & \text{o/w} \end{cases}$	0

- split attention: $dT_{t,1} + dT_{t,2} = dt$, $dT_{t,1} \geq 0$, $dT_{t,2} \geq 0$

Back to ▸ main slides

Optimization Problem

Strategy (T, τ, a^F):

$T = (T_1, T_2)$: attention allocation plan, $T_{t,k}$ is the total amount of attention the agent paid to source k by time t (note: $T_{t,1} + T_{t,2} = t$)

$\Rightarrow$ by time t, agent knows $X^{(1)}_{T_{t,1}}$ and $X^{(2)}_{T_{t,2}}$

$\tau \geq 0$: stopping time

α : alternative chosen at time τ *(function: info by $\tau \mapsto \mathcal{A}$)*

Total ex post payoff:

$u_j(\alpha) - c_1 T_{\tau,1} - c_2 T_{\tau,2}$, where j is the true state

Assumption: no discounting

Optimization problem:

$$\sup_{(T,\tau,\alpha)} \mathbb{E}\left[u_j(\alpha) - c_1 T_{\tau,1} - c_2 T_{\tau,2}\right]$$

Back to main slides

Optimal Strategy: Two Phases

Theorem 1: The optimal contingency plan $(T, \bar{\tau}, a)$ has two phases.

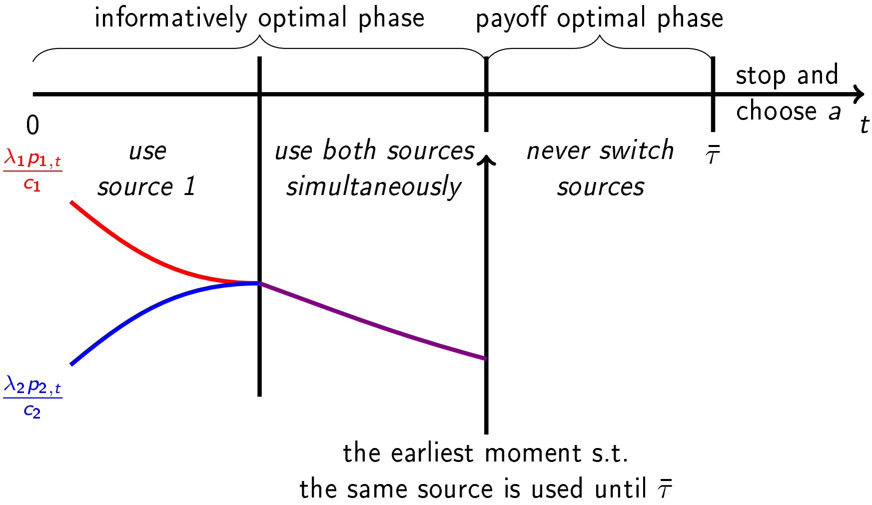

Back to main slides

Optimal Strategy: Two Phases

Theorem 1: $\exists\ t^* \in [0, \bar{\tau}],\ k^* \in \{1, 2\}$, such that

$$
\begin{cases}
0 \le t \le \hat{t}: & T_{t,\hat{k}} = t, \quad T_{t,3-\hat{k}} = 0 \\
\hat{t} \le t \le t^*: & T_{t,k} - T_{\hat{t},k} = \frac{\lambda_{3-k}}{\lambda_1+\lambda_2}(t - \hat{t}), \quad k = 1, 2 \\
t^* \le t \le \bar{\tau}: & T_{t,k^*} - T_{t^*,k^*} = t - t^*, \quad T_{t,3-k^*} - T_{t^*,3-k^*} = 0
\end{cases}
$$

where $\hat{k} = \arg\max_{i=1,2} \frac{\lambda_i p_{i,0}}{c_i}$, $\hat{t} = \arg\min_{t\in[0,t^*]} \left| \frac{\lambda_{\hat{k}} p_{\hat{k},t}}{c_{\hat{k}}} - \frac{\lambda_{3-\hat{k}} p_{3-\hat{k},t}}{c_{3-\hat{k}}} \right|$.

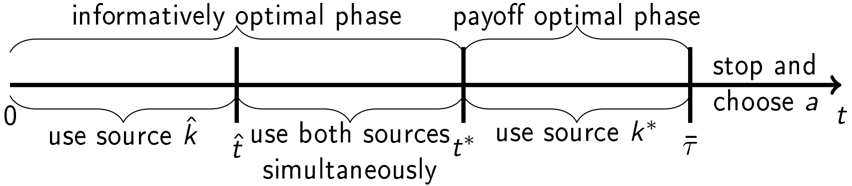

$\frac{\lambda_{\hat{k}} p_{\hat{k},t}}{c_{\hat{k}}} > \frac{\lambda_{3-\hat{k}} p_{3-\hat{k},t}}{c_{3-\hat{k}}}$ $\quad$ $\frac{\lambda_1 p_{1,t}}{c_1} = \frac{\lambda_2 p_{2,t}}{c_2}$

Back to main slides

Definition

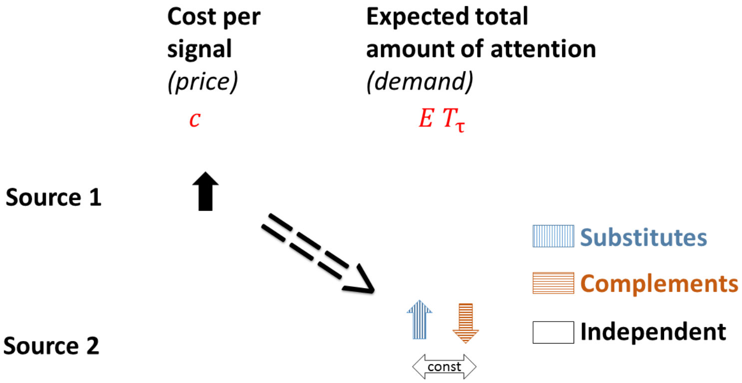

Back to ▸ main slides

Complements and Substitutes

Assumption 1 (no structural change)

The optimal default alternative does not change.

Assumption 2 (no independence)

The sources are not independent. Equivalently, the optimal contingency plan involves using both sources.

Back to ▸ main slides

Complements and Substitutes: Beliefs

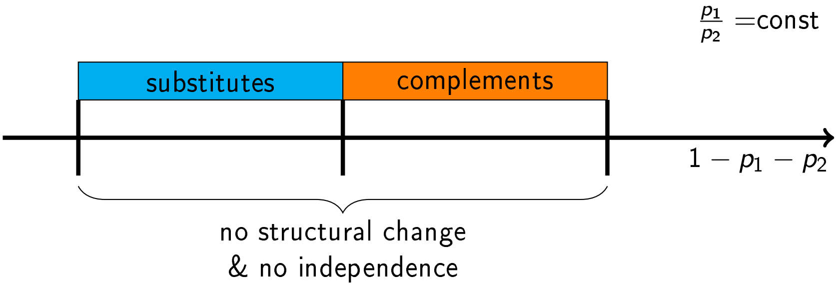

$\Rightarrow$ Qualitative nature of information market might change over time, as more information consumed

Back to ▸ main slides

Complements and Substitutes: Cost

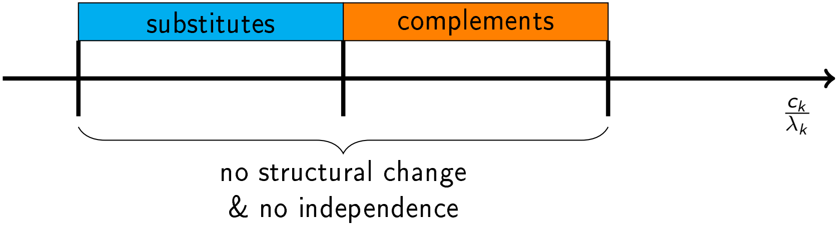

- c_k = price of source k
- the sign of $\frac{d(\text{demand for source 2})}{dc_1} = \frac{d(\text{demand for source 1})}{dc_2}$ depends on c_k and λ_k only through $\frac{c_k}{\lambda_k}$

Back to ▸ main slides

Complements and Substitutes: Payoffs

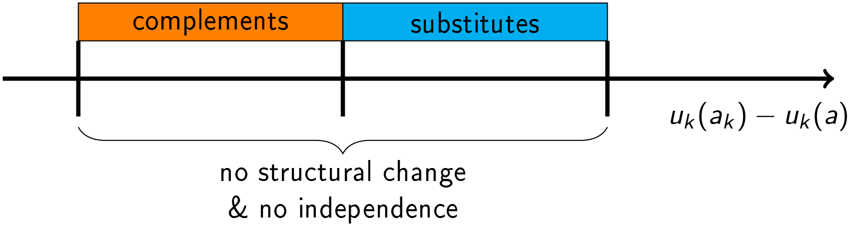

- $u_k(a_k) - u_k(a)$ = benefit from the correct choice in state k:
  - $u_k(a_k)$ = maximum payoff in state k
  - $u_k(a)$ = payoff from the default alternative in state k
- the sign of $\frac{d(\text{demand for source 2})}{dc_1} = \frac{d(\text{demand for source 1})}{dc_2}$ depends on $u_k(a_k)$ and $u_k(a)$ only through $u_k(a_k) - u_k(a)$

Back to main slides