Dynamic Choice of Information Sources

Tatiana Mayskaya

Higher School of Economics

North American Summer Meeting of the Econometric Society
21 June 2018

Why do we learn?

Why do we learn?

> "Physics is like sex: sure, it may give some practical results, but that's not why we do it.

Richard Feynman

Why do we learn?

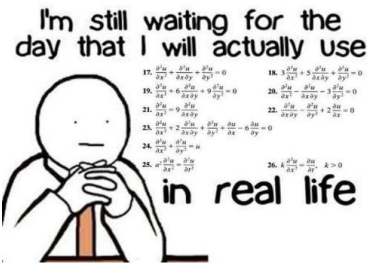

"Physics is like sex: sure, it may give some practical results, but that's not why we do it."

Richard Feynman

Why do we learn?

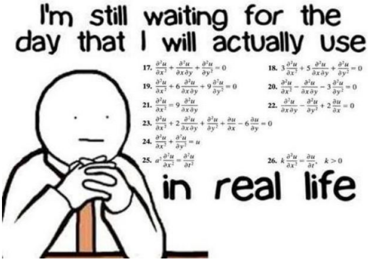

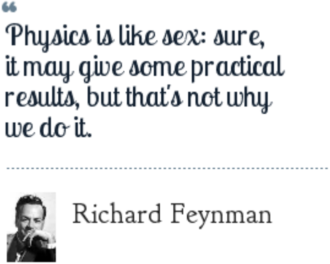

Motivation

We have common sense understanding that knowledge is important, it will pay off. But often in a distant future, and right now we might not know where and how we are going to use this knowledge.

How to choose what to learn if we don't know the goal?

Example

- A company has two research divisions
- Each division $k = 1, 2$ investigates the profitability of a certain project, which can be either success $(\theta_k = 1)$ or failure $(\theta_k = 0)$
  - θ_1 and θ_2 are not necessarily independent

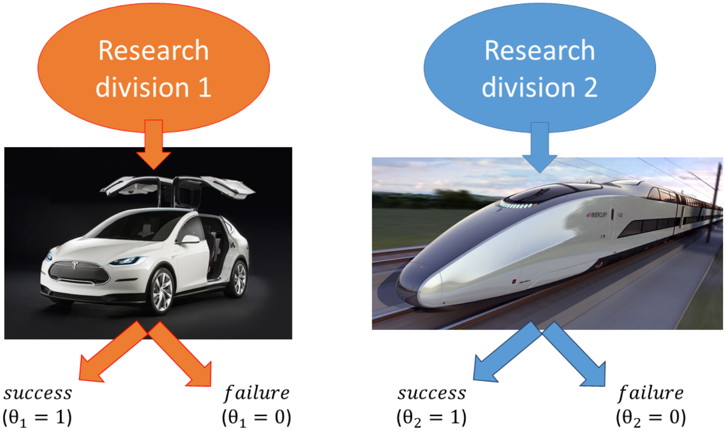

Example

- At each instance of time, the company allocates a unit of resources between two divisions
- Each division k spends its resources $x_k \in [0,1]$ to search for a proof that its project is successful $\theta_k = 1$

	$\theta_k = 1$	$\theta_k = 0$
Prob("good news")	$x_k dt$	0
Prob("no news")	$1 - x_k dt$	1

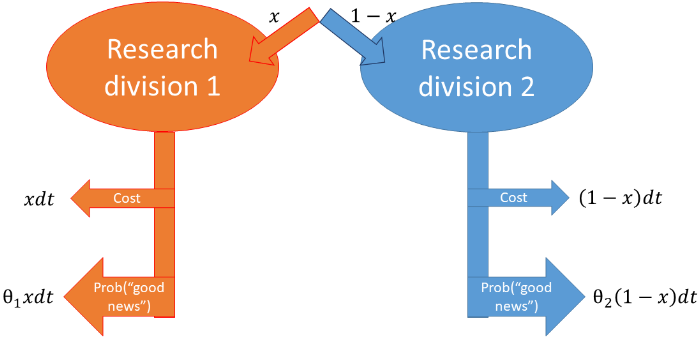

Example

- The company decides when to stop research and make the decision

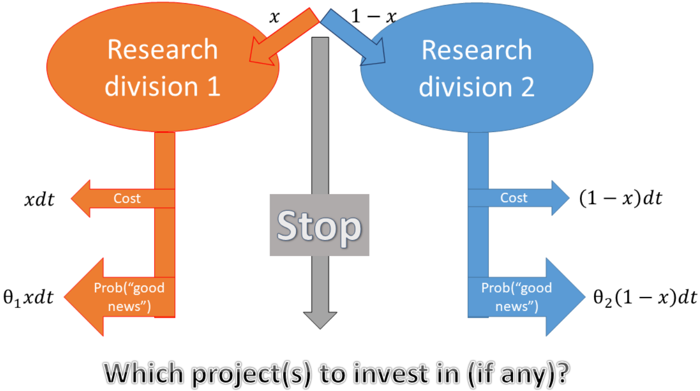

General Framework

Learning Stage

- four states of the world: $(\theta_1, \theta_2) \in \{(1,0),(0,1),(1,1),(0,0)\}$
- two information sources:
  - source 1 = search for a proof that $\theta_1 = 1$
  - source 2 = search for a proof that $\theta_2 = 1$

Decision Making Stage

- $\mathcal{A} = \{a_1, \ldots, a_n\}$: set of alternatives to choose from
- $u_i(a)$: payoff from choosing alternative a

	$(1,0)$	$(0,1)$	$(1,1)$	$(0,0)$
a_1	$u_1(a_1)$	$u_2(a_1)$	$u_3(a_1)$	$u_4(a_1)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
a_n	$u_1(a_n)$	$u_2(a_n)$	$u_3(a_n)$	$u_4(a_n)$

General Framework

Learning Stage *How to learn?*

- four states of the world: $(\theta_1, \theta_2) \in \{(1,0), (0,1), (1,1), (0,0)\}$
- two information sources:
  - source 1 = search for a proof that $\theta_1 = 1$
  - source 2 = search for a proof that $\theta_2 = 1$

Decision Making Stage *Why to learn?*

- $\mathcal{A} = \{a_1, \ldots, a_n\}$: set of alternatives to choose from
- $u_i(a)$: payoff from choosing alternative a

	$(1,0)$	$(0,1)$	$(1,1)$	$(0,0)$
a_1	$u_1(a_1)$	$u_2(a_1)$	$u_3(a_1)$	$u_4(a_1)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
a_n	$u_1(a_n)$	$u_2(a_n)$	$u_3(a_n)$	$u_4(a_n)$

Contribution

Optimal strategy balances the trade-off:

- *belief-based incentives*: use the source that leads to a higher Prob("good news")
  - depends on current **beliefs**
- *payoff-based incentives*: learn about the project that leads to a higher profit if successful
  - depends on **payoff** matrix

Contribution

This paper shows

1. Optimal strategy has two phases: in the absence of "good news"

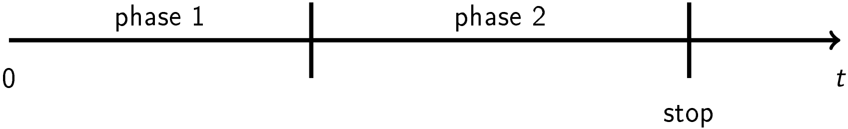

Contribution

This paper shows

1. Optimal strategy has two phases: in the absence of “good news”

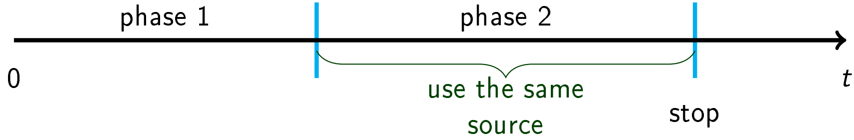

- phase 2 source depends on the *payoff matrix*
- thresholds depend on the *payoff matrix*

Contribution

This paper shows

1. Optimal strategy has two phases: in the absence of "good news"

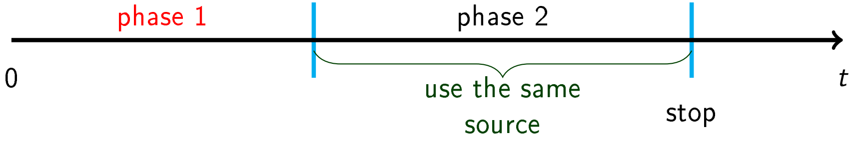

   - phase 2 source depends on the *payoff matrix*
   - thresholds depend on the *payoff matrix*

2. Phase 1 rule introduces the notion is a source *index*

use the source with the highest index

Main Result

Expression for the index

Optimal Strategy: $\theta_1 + \theta_2 \leq 1$

Optimal Strategy: $\theta_1 + \theta_2 \leq 1$

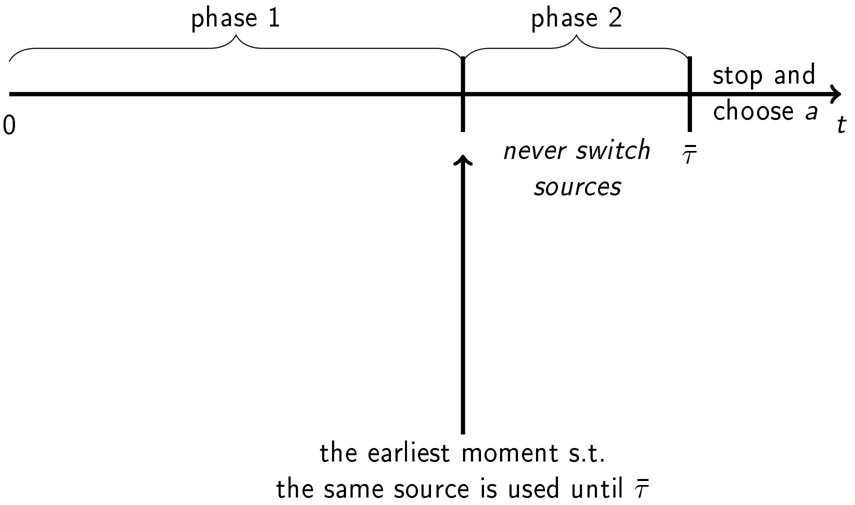

Optimal Strategy: $\theta_1 + \theta_2 \leq 1$

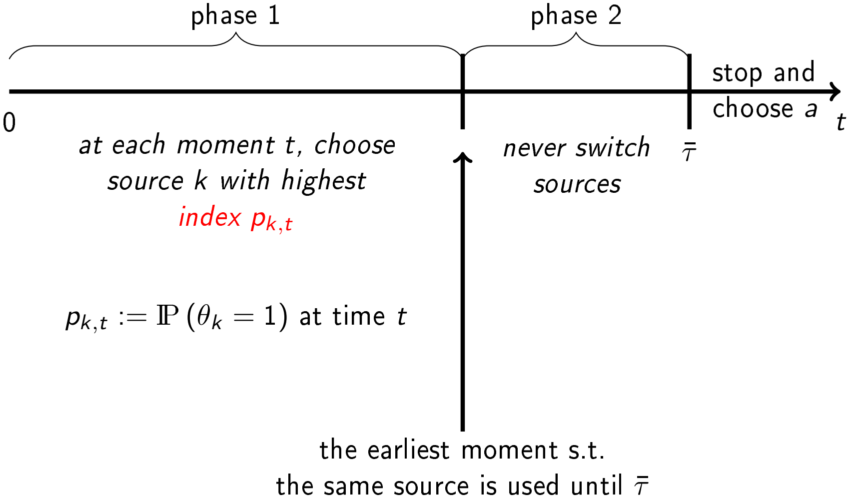

Optimal Strategy: $\theta_1 + \theta_2 \leq 1$

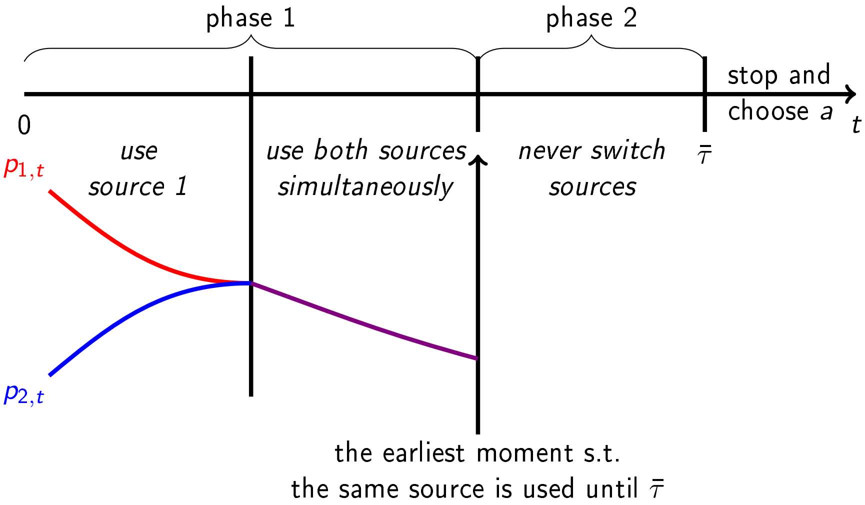

Optimal Strategy: $\theta_1 + \theta_2 \leq 1$

Intuition

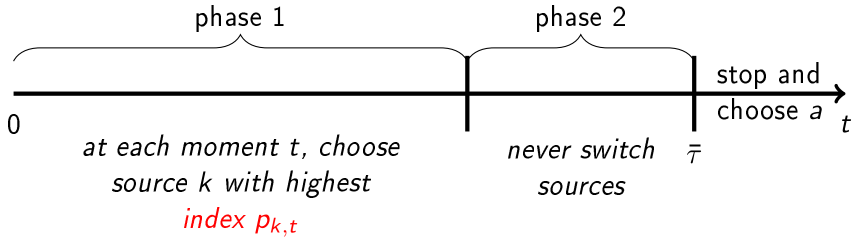

→ find the true state ⟶ check most likely theory (no discounting)

→ absence of evidence ⟶ deterministic search path | eventually use *both* sources → does not matter where to search first

Optimal Strategy: General Case

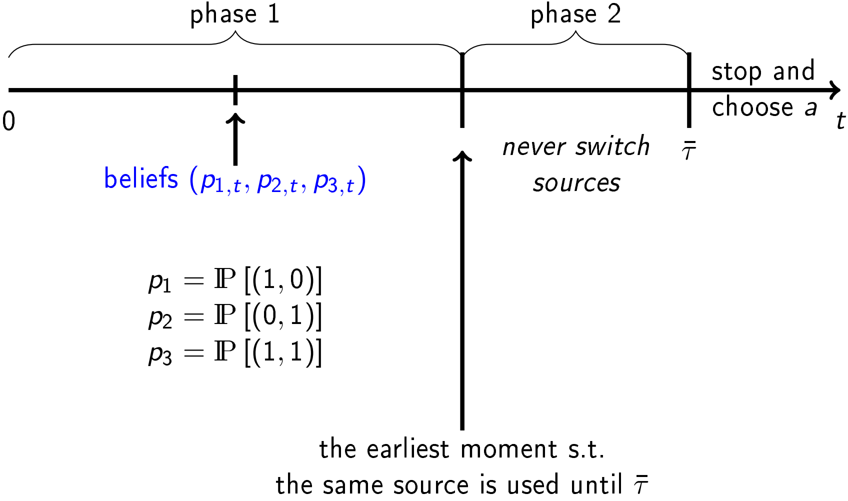

$$p_1 = \mathbb{P}[(1,0)]$$
$$p_2 = \mathbb{P}[(0,1)]$$
$$p_3 = \mathbb{P}[(1,1)]$$

Phase 1

Fix current beliefs (p_1, p_2, p_3).
If a positive signal occurs ***now****, what is the optimal strategy?*

Phase 1

Fix current beliefs (p_1, p_2, p_3).
*If a positive signal occurs **now**, what is the optimal strategy?*

Four possibilities:

1. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ learn from source 1
2. $\theta_1 = 1 \Rightarrow$ stop ; $\theta_2 = 1 \Rightarrow$ learn from source 1
3. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ stop
4. $\theta_1 = 1 \Rightarrow$ stop ; $\theta_2 = 1 \Rightarrow$ stop

Phase 1

Fix current beliefs (p_1, p_2, p_3).
If a positive signal occurs ***now****, what is the optimal strategy?*

Four possibilities:

1. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ learn from source 1
   - phase 1 rule: indifferent between both sources
2. $\theta_1 = 1 \Rightarrow$ stop ; $\theta_2 = 1 \Rightarrow$ learn from source 1
3. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ stop
4. $\theta_1 = 1 \Rightarrow$ stop ; $\theta_2 = 1 \Rightarrow$ stop

Phase 1

Fix current beliefs (p_1, p_2, p_3).
*If a positive signal occurs **now**, what is the optimal strategy?*

Four possibilities:

1. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ learn from source 1
   - phase 1 rule: indifferent between both sources
2. $\theta_1 = 1 \Rightarrow$ stop ; $\theta_2 = 1 \Rightarrow$ learn from source 1
   - phase 1 rule: use source 1
3. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ stop
   - phase 1 rule: use source 2
4. $\theta_1 = 1 \Rightarrow$ stop ; $\theta_2 = 1 \Rightarrow$ stop

Phase 1

Fix current beliefs (p_1, p_2, p_3).
*If a positive signal occurs **now**, what is the optimal strategy?*

Four possibilities:

1. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ learn from source 1
   - phase 1 rule: indifferent between both sources
2. $\theta_1 = 1 \Rightarrow$ stop ; $\theta_2 = 1 \Rightarrow$ learn from source 1
   - phase 1 rule: use source 1
3. $\theta_1 = 1 \Rightarrow$ learn from source 2; $\theta_2 = 1 \Rightarrow$ stop
   - phase 1 rule: use source 2
4. $\theta_1 = 1 \Rightarrow$ stop and choose a_1; $\theta_2 = 1 \Rightarrow$ stop and choose a_2
   - phase 1 rule: use the source with the highest index

Theorem

Phase 1 rule: use the source with the highest index

$$p_k + p_3 u_3(a_k)$$

Theorem

Phase 1 rule: use the source with the highest index

$$p_k + p_3 u_3(a_k)$$

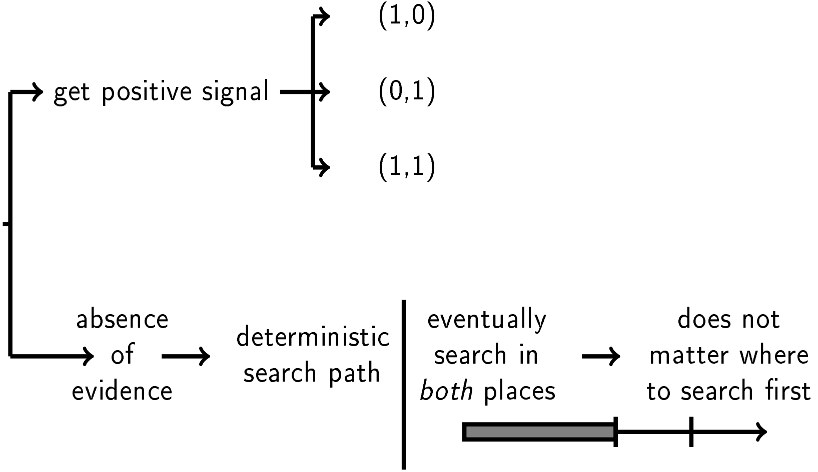

Theorem

Phase 1 rule: use the source with the highest index

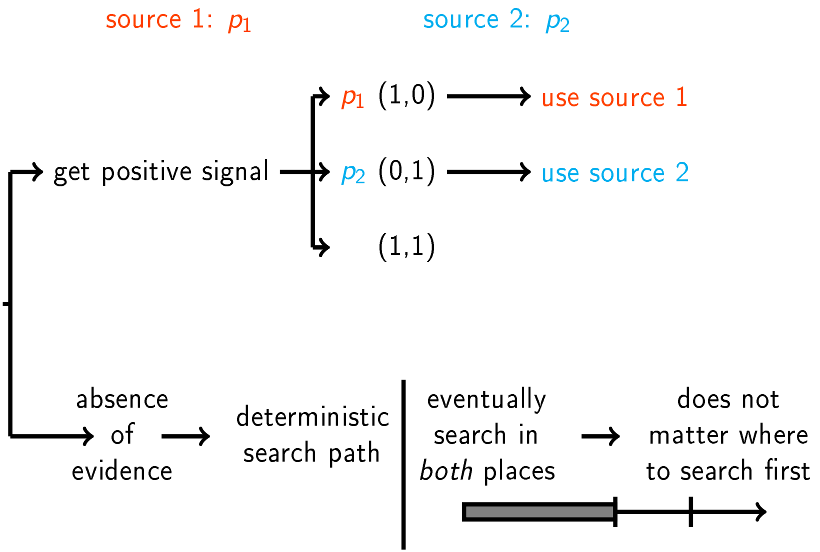

Theorem

Phase 1 rule: use the source with the highest index

source 1: $p_1 + p_3 u_3(a_1)$ source 2: $p_2 + p_3 u_3(a_2)$

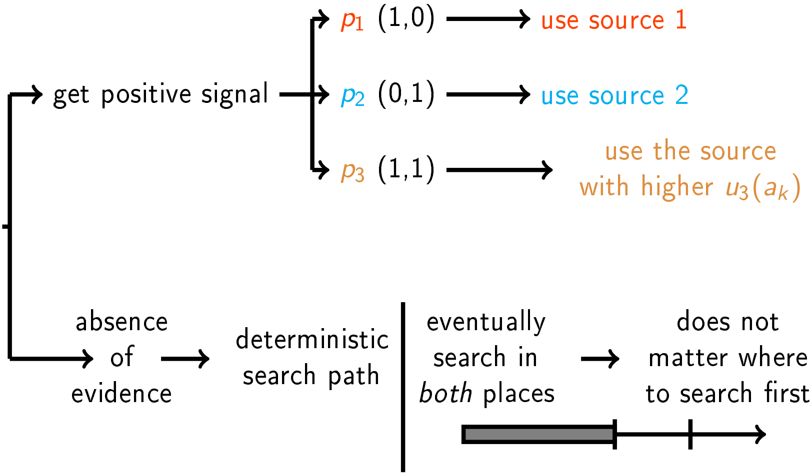

Discussion

Result

The index depends only on the relative payoffs at state (1,1). In other words, its dependence on the learning goal is very limited

Discussion

Result

The index depends only on the relative payoffs at state (1,1). In other words, its dependence on the learning goal is very limited

Example with two investment projects:

- in monopoly market, only the difference in payoffs in case of success matters for the index

$$\underbrace{v_1}_{\text{payoff from successful project 1}} - \underbrace{v_2}_{\text{payoff from successful project 2}}$$

Discussion

Result

The index depends only on the relative payoffs at state (1,1). In other words, its dependence on the learning goal is very limited

Example with two investment projects:

- in monopoly market, only the difference in payoffs in case of success matters for the index

$$\underbrace{v_1}_{\text{payoff from successful project 1}} - \underbrace{v_2}_{\text{payoff from successful project 2}}$$

- when another firm can compete with the other project, only the payoffs when both projects are successful matter for the index

	$(1,0)$	$(0,1)$	$(1,1)$	$(0,0)$
invest in project 1	v_1	0	$v_1 - d_1$	0
invest in project 2	0	v_2	$v_2 - d_2$	0

Discussion

Result

The index depends only on the relative payoffs at state (1,1). In other words, its dependence on the learning goal is very limited

Examples:

information sources = different diagnostic tests;
decision problem also specifies possible treatments
Why much do we need to know about possible treatments to know which test to apply?

Discussion

Result

The index depends only on the relative payoffs at state (1,1). In other words, its dependence on the learning goal is very limited

Examples:

information sources = different diagnostic tests;
decision problem also specifies possible treatments
Why much do we need to know about possible treatments to know which test to apply?

information sources = news topics in media;
decision problem also reflects individual characteristics, such as policy preferences
Why much do we need to know about individual characteristics to know the demand for information sources?

Literature

	Che&Mierendorff (2017, working paper)	Nikandrova&Pancs (2018, TE)	This paper
states	(1,0) and (0,1)	(1,0), (0,1), (1,1), (0,0)	
correlation	$\theta_1 + \theta_2 = 1$	$\theta_1 \perp\!\!\!\perp \theta_2$	arbitrary
alternatives	2 ("match the state")	2 ("choose successful project")	any number
focus	on a given decision problem $\Downarrow$ optimal strategy (including chosen alternative)		on the general form of the learning process $\Downarrow$ index

Literature

Existing literature: alternatives and states are tightly connected

- Multi-armed bandit problem: alternatives $\sim$ states (info about alternatives directly, not through a state of the world)
  Pancs and Nikandrova 2018, Che and Mierendorff 2017, Klein and Rady 2011
- 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 (treasure hunt)
    Matros&Smirnov 2016, Fershtman&Rubinstein 1997
- *Liang, Mu, and Syrgkanis (2017)* came to the "opposite" conclusion that myoptic learning is (almost) optional
  - completely different structure of the decision problem; in particular, only one state component, θ_1, is payoff relevant

Optimal Strategy

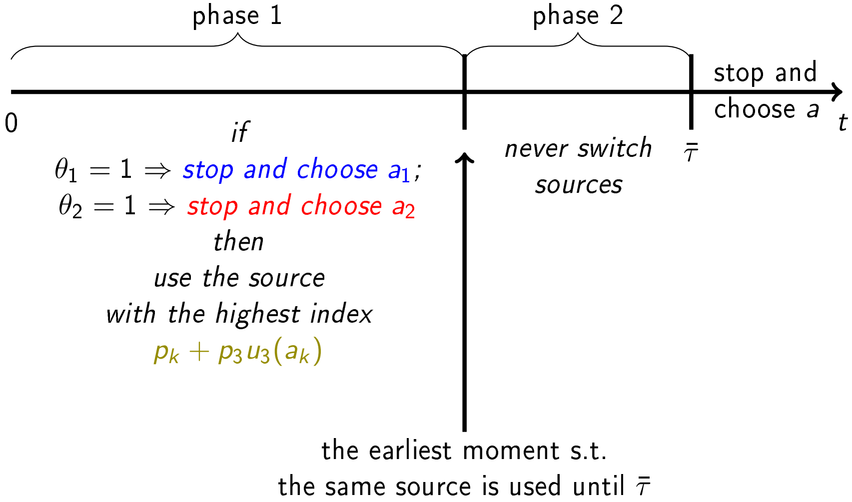

Reminder: $p_1 = \mathbb{P}[(1,0)]$, $p_2 = \mathbb{P}[(0,1)]$, $p_3 = \mathbb{P}[(1,1)]$, u_3: payoff at (1,1)