Dynamic Choice of Information Sources

Tatiana Mayskaya

Higher School of Economics

May 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?

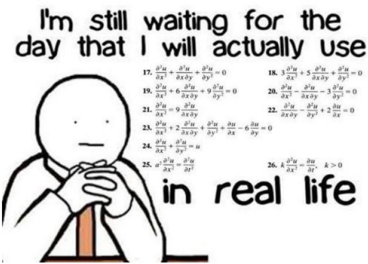

> Physics is like sex: sure, it may give some practical results, but that's not why we do it.

Richard Feynman

Why do we learn?

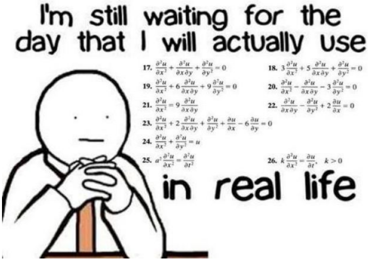

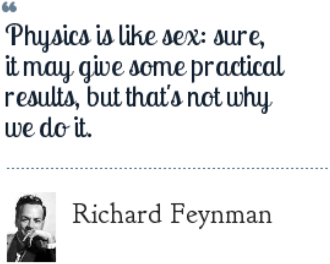

Motivation

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

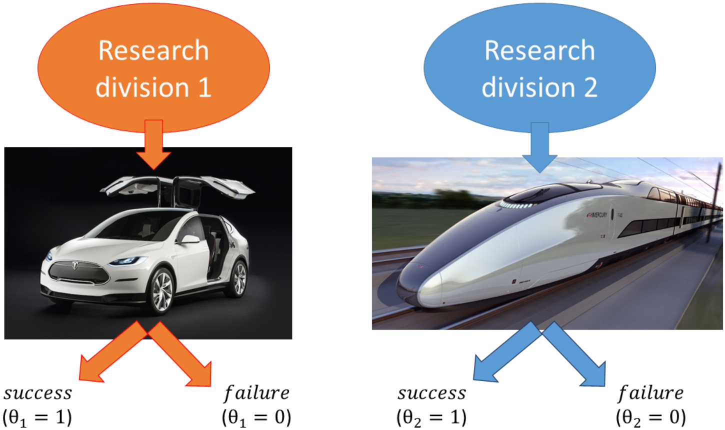

Example

- 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")	$\lambda_k x_k dt$	0
Prob("no news")	$1 - \lambda_k x_k dt$	1

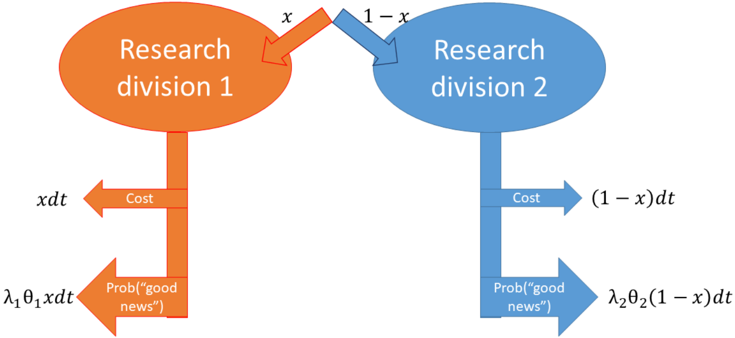

Example

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

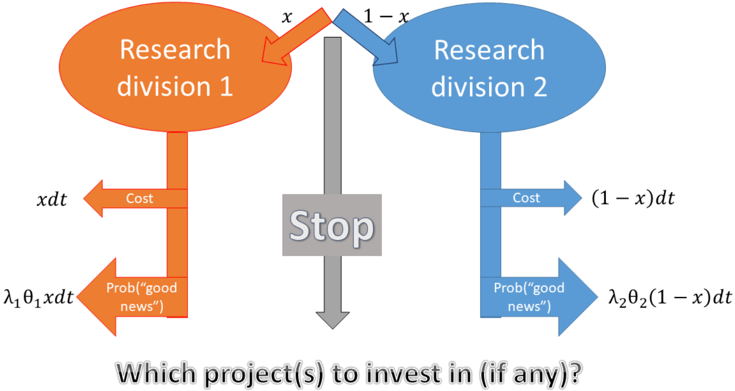

General Framework

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

- $\mathcal{A} = \{a_1, \dots, a_n\}$: 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** and information **sources**' characteristics
- *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”

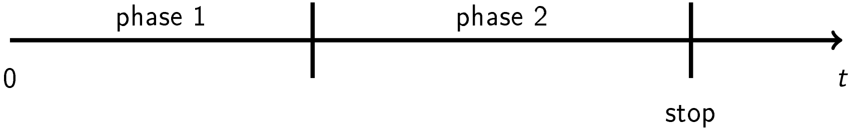

Contribution

This paper shows

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

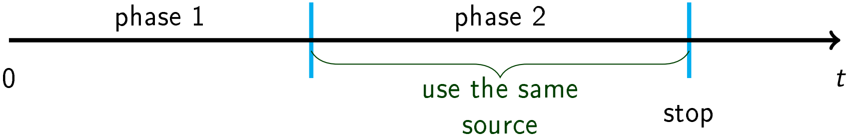

- phase 2 source depends on the *payoff matrix*
- thresholds depend on the *payoff matrix*

Contribution

This paper shows

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

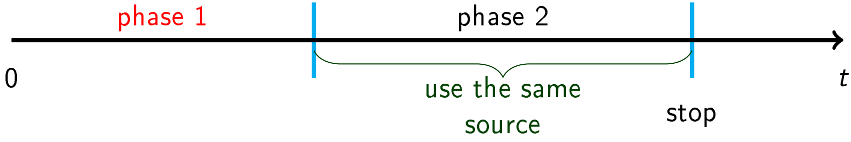

- phase 2 source depends on the *payoff matrix*
- thresholds depend on the *payoff matrix*

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

$$\boxed{\text{use the source with the highest index}}$$

Main Result

Expression for the index

Properties of the index

1. It's (almost) independent of the payoff matrix
   - independence of the payoff matrix makes the index *constant across all decision problems*
   - the index is mainly a characteristic of the information source itself

Properties of the index

1. It's (almost) independent of the payoff matrix
   - independence of the payoff matrix makes the index *constant across all decision problems*
   - the index is mainly a characteristic of the information source itself

Examples:

information sources = different diagnostic tests;
decision problem also specifies possible treatments

information sources = news topics in media;
decision problem also reflects individual characteristics, such as policy preferences

Properties of the index

1. It's (almost) independent of the payoff matrix
2. When the index is independent of the payoff matrix, it is equal to

$$\text{Prob}(\text{“good news”}) = \lambda_k \cdot \underbrace{\mathbb{P}\left(\theta_k = 1\right)}_{\text{current beliefs}}$$

- the index of source k depends *only* on the characteristics of that source k (probability of receiving a positive signal from that source); it does *not* depend on neither the intensity of the other source, nor the correlation between the states
- similarity with the Gittins index

Properties of the index

1. It's (almost) independent of the payoff matrix
2. When the index is independent of the payoff matrix, it is equal to

$$\text{Prob}(\text{“good news”}) = \lambda_k \cdot \underbrace{\mathbb{P}\left(\theta_k = 1\right)}_{\text{current beliefs}}$$

3. When the index does depend on the payoff matrix, it depends only on the payoffs at state (1,1)

Example with two investment projects:

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

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

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

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

Literature

	Che&Mierendorff (2017, R&R in AER)	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

- 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 (2018): Poisson process
- Multi-armed bandit problem: each source gives both the payoff and information about the payoff distribution
- 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)

    Fershtman&Rubinstein 1997, Matros&Smirnov 2016
- *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

Outline

Payoff Matrix and Beliefs

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

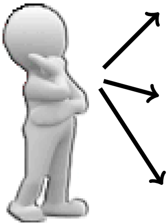

State (θ_1, θ_2):	$(1,0)$	$(0,1)$	$(0,0)$
Alternative a_1	$u_1(a_1)$	$u_2(a_1)$	$u_3(a_1)$
Alternative a_2	$u_1(a_2)$	$u_2(a_2)$	$u_3(a_2)$
$\ldots$	$\ldots$	$\ldots$	$\ldots$

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

Payoff Matrix and Beliefs

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

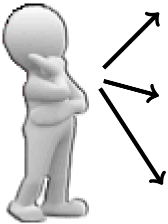

	State (θ_1, θ_2):	$(1,0)$	$(0,1)$	$(0,0)$
	Beliefs:	p_1	p_2	$1 - p_1 - p_2$
	Alternative a_1	$u_1(a_1)$	$u_2(a_1)$	$u_3(a_1)$
	Alternative a_2	$u_1(a_2)$	$u_2(a_2)$	$u_3(a_2)$
	$\ldots$	$\ldots$	$\ldots$	$\ldots$

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

Information Sources

- The decision maker can stop the learning process at any moment.
- At each moment of time during the learning process, the decision maker chooses how to split a unit of attention ($dT_{t,1} + dT_{t,2} = dt$) between two information sources
- Information source k ($k = 1, 2$):

$$X^{(k)} = \begin{cases} \text{Poisson process with intensity } \lambda_k > 0, & \theta_k = 1 \\ 0, & \theta_k = 0 \end{cases}$$

- Cost of source k is $c_k > 0$

		(1,0)	(0,1)	(0,0)
source 1	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	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

Optimization Problem

Strategy (T, τ, α):

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

$\tau \geq 0$: stopping time

α : alternative chosen at 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 corresponds to the true state

Assumption: no discounting

Optimization problem:

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

Strategy

strategy (T, τ, α)	contingency plan $(T, \bar{\tau}, a)$
T — random	T — deterministic: *conditional on not finding the state*
τ — random	$\bar{\tau} = \sup \tau$ — deterministic: *stopping time when state is not revealed by* τ
α — function	$a = \alpha(\tau = \bar{\tau}) \in \mathcal{A}$: *alternative chosen if state is not revealed by* τ *(define it as the* ***default alternative****)*

What to do if the state is not revealed?

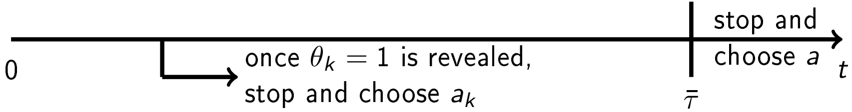

Denote

$p_{k,t}$ = belief at time t about $\theta_k = 1$, *if the state is not revealed by* t

Optimal Strategy

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

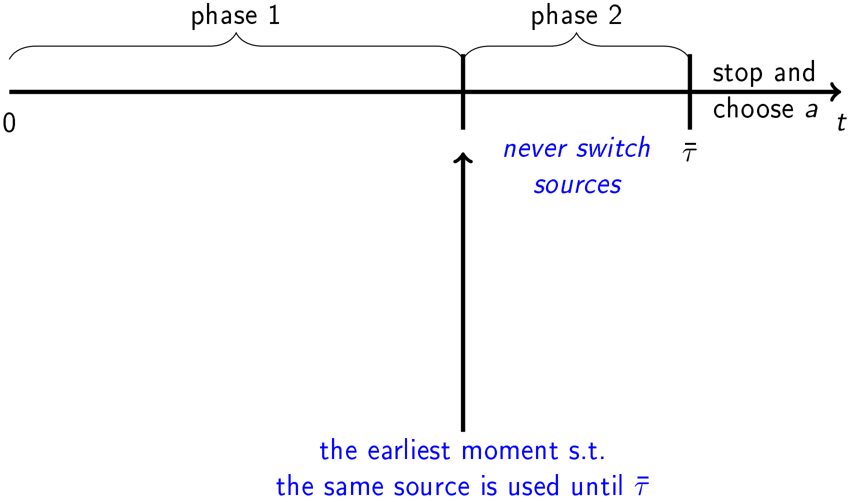

Optimal Strategy

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

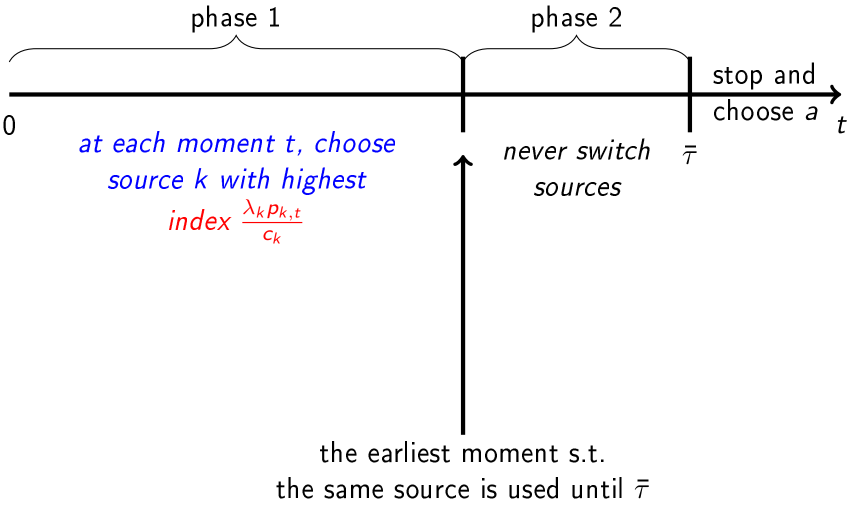

Optimal Strategy

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

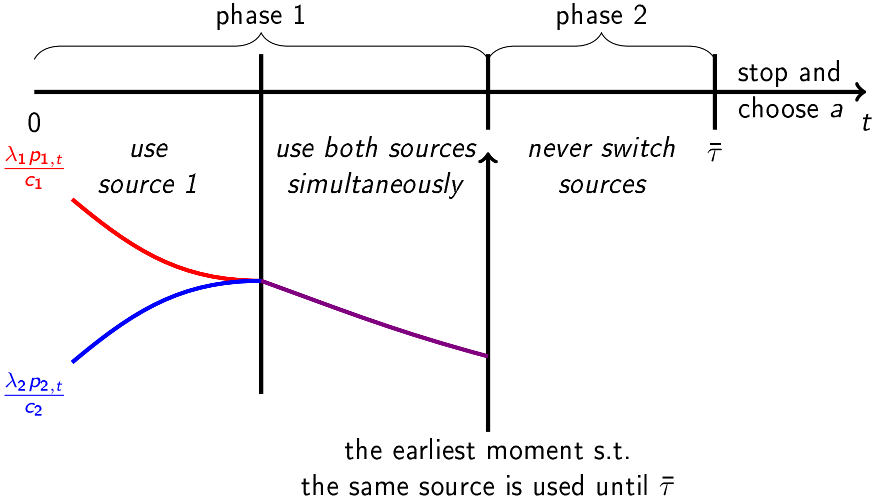

Optimal Strategy

Intuition

For simplicity, assume $\lambda_1 = \lambda_2$, $c_1 = c_2$

phase 1

phase 2

0

at each moment t, choose source k with highest *index $p_{k,t}$*

never switch sources

$\bar{\tau}$

stop and choose a

t

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

absence of evidence ⟶ deterministic search path

eventually search in *both* places ⟶ does not matter where to search first

Optimal Strategy
Intuition

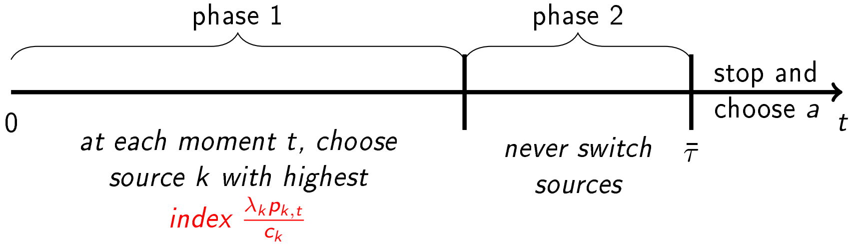

Source k is more efficient if

- intensity $\lambda_k \uparrow$
- cost $c_k \downarrow$

- We looked at the optimal strategy as a **function of time**
- Optimal strategy is Markovian: optimal action depends only on *current beliefs* (p_1, p_2) $\Rightarrow$ let's look at the optimal strategy as a **function of beliefs**

- We looked at the optimal strategy as a **function of time**
- Optimal strategy is Markovian: optimal action depends only on *current beliefs* (p_1, p_2) $\Rightarrow$ let's look at the optimal strategy as a **function of beliefs**
- Why do we need to do this?

function of time	function of beliefs
two-phase representation & introduction of index	specifies thresholds, optimal default alternative as long as the index $\Downarrow$ gives the full description

- We looked at the optimal strategy as a **function of time**
- Optimal strategy is Markovian: optimal action depends only on *current beliefs* (p_1, p_2) $\Rightarrow$ let's look at the optimal strategy as a **function of beliefs**
- Why do we need to do this?

function of time	function of beliefs
two-phase representation & introduction of index	specifies thresholds, optimal default alternative as long as the index $\Downarrow$ gives the full description

- Why haven't we started with this?
  1. focus of the paper — index
     - the notion of index it is easier to present on the time line
  2. thresholds and optimal default alternatives depend on the payoff matrix $\Rightarrow$ general case description of the optimal strategy on the belief triangle is messy (too many cases are possible)

     $\Rightarrow$ I'll present only a few examples

Optimal Strategy: Example

on belief triangle

Payoff matrix

state (θ_1, θ_2)	(1,0)	(0,1)	(0,0)
alternative a_1	1	0	0
alternative a_2	0	1	0
alternative a_3	0	0	1

Interpretation: get 1 if guess the state and 0 otherwise

Optimal Strategy: Example

on belief triangle

Payoff matrix

state (θ_1, θ_2)	(1,0)	(0,1)	(0,0)
alternative a_1	1	0	0
alternative a_2	0	1	0
alternative a_3	0	0	1

Interpretation: get 1 if guess the state and 0 otherwise

- It turns out that the optimal strategy representation of the belief triangle depends on c_1, c_2, λ_1 and λ_2 only through the ratios $\frac{c_1}{\lambda_1}$ and $\frac{c_2}{\lambda_2}$
- Consider symmetric sources:

$$\frac{c_1}{\lambda_1} = \frac{c_2}{\lambda_2}$$

Optimal Strategy: Example

on belief triangle

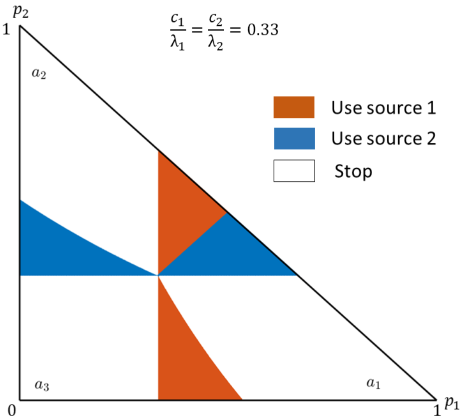

Optimal Strategy: Example

on belief triangle

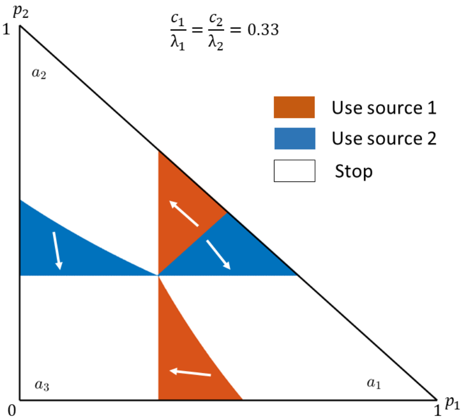

At most one source is used $\Rightarrow$ only **phase 2**

Optimal Strategy: Example

on belief triangle

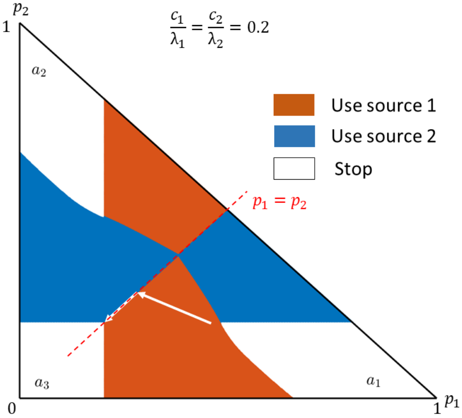

When uncertainty is the highest (middle of triangle), **phase 1** rule is optimal

Optimal Strategy: Example

on belief triangle

	(1,0)	(0,1)	(0,0)
a_1	1	0	0
a_2	0	1	0
a_3	0	0	1

	(1,0)	(0,1)	(0,0)
a_1	2	0	0
a_2	0	1	0
a_3	0	0	1

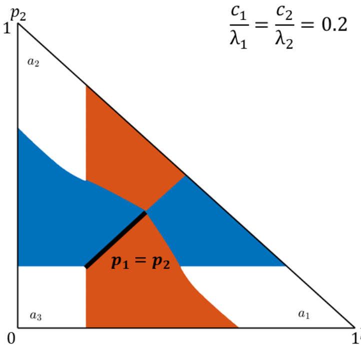

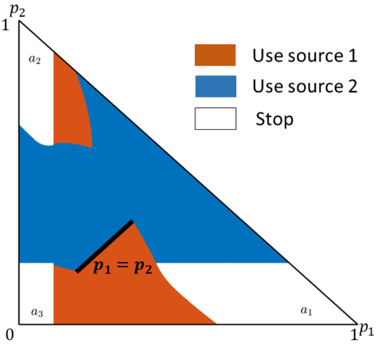

Phase 1 rule is **independent** of the payoff matrix

Outline

Payoff Matrix and Beliefs

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

State (θ_1, θ_2):	$(1, 0)$	$(0, 1)$	$(1, 1)$	$(0, 0)$
Beliefs:	p_1	p_2	p_3	$1 - p_1 - p_2 - p_3$
Alternative a_1	$u_1(a_1)$	$u_2(a_1)$	$u_3(a_1)$	$u_4(a_1)$
Alternative a_2	$u_1(a_2)$	$u_2(a_2)$	$u_3(a_2)$	$u_4(a_2)$
$\ldots$	$\ldots$	$\ldots$	$\ldots$	$\ldots$

General Case vs Special Case: $\theta_1 + \theta_2 \leq 1$

In the absence of a positive signal, follow a contingency plan:

General Case vs Special Case: $\theta_1 + \theta_2 \leq 1$

Key difference

Once a positive signal is received (state $\theta_k = 1$ is revealed), there might still be uncertainty about the state

In the absence of a positive signal, follow a contingency plan:

General Case vs Special Case: $\theta_1 + \theta_2 \leq 1$

Key difference

Once a positive signal is received (state $\theta_k = 1$ is revealed), there might still be uncertainty about the state

$\Rightarrow$

1. It might be optimal to learn even after a positive signal

In the absence of a positive signal, follow a contingency plan:

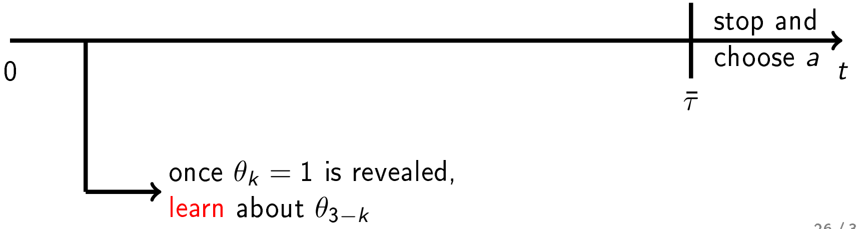

General Case vs Special Case: $\theta_1 + \theta_2 \leq 1$

Key difference

Once a positive signal is received (state $\theta_k = 1$ is revealed), there might still be uncertainty about the state

$\Rightarrow$

1. It might be optimal to learn even after a positive signal
2. Even if it is optimal to stop after receiving a positive signal, the optimal alternative after receiving a positive signal depends on beliefs

In the absence of a positive signal, follow a contingency plan:

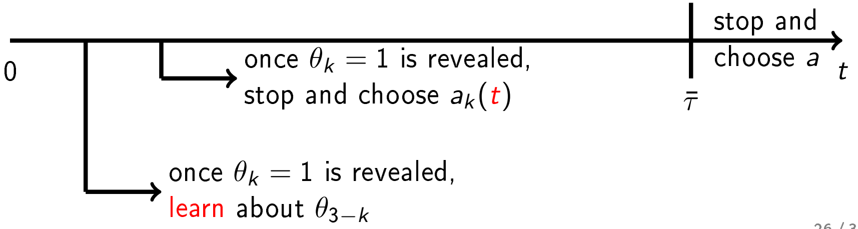

In the absence of a positive signal, follow a contingency plan:

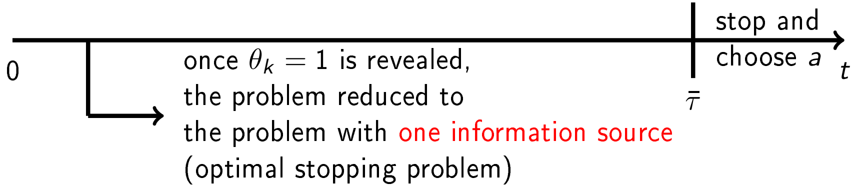

In the absence of a positive signal, follow a contingency plan:

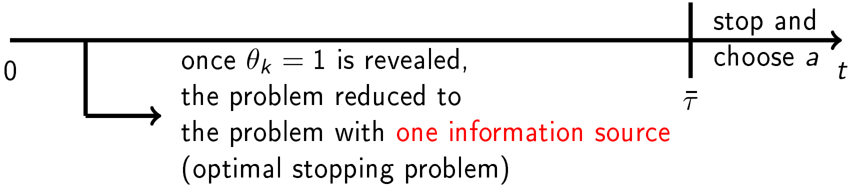

1. When $\theta_1 = 1$ is revealed: only source 2 might be used

$$
\begin{array}{cccc}
(1,0) & \cancel{(0,1)} & (1,1) & \cancel{(0,0)} \\
p_1 & p_2 & p_3 & 1-p_1-p_2-p_3 \\
\downarrow & \downarrow & \downarrow & \downarrow \\
p_1' = \frac{p_1}{p_1+p_3} & 0 & 1-p_1' = \frac{p_3}{p_1+p_3} & 0
\end{array}
$$

For example:

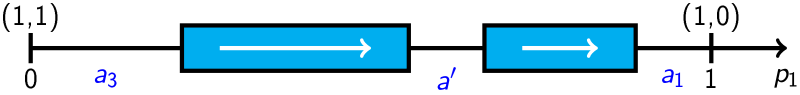

In the absence of a positive signal, follow a contingency plan:

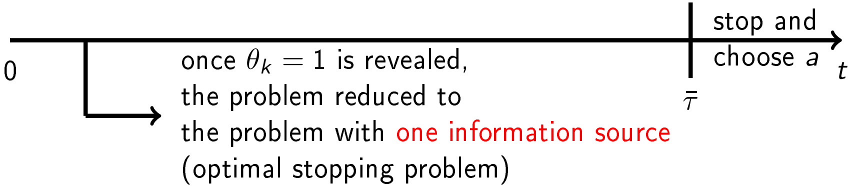

1. When $\theta_1 = 1$ is revealed:
   - problem with only source 2

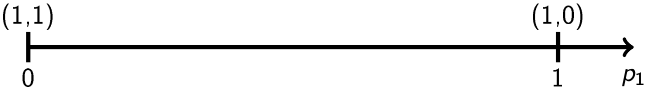

2. When $\theta_2 = 1$ is revealed:
   - problem with only source 1

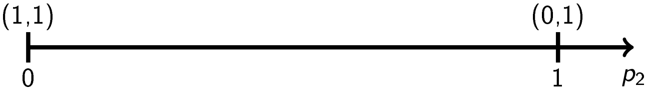

To find the optimal strategy, it only remains to find the optimal contingency plan

Optimal Strategy

Theorem: The optimal contingency plan has two phases.

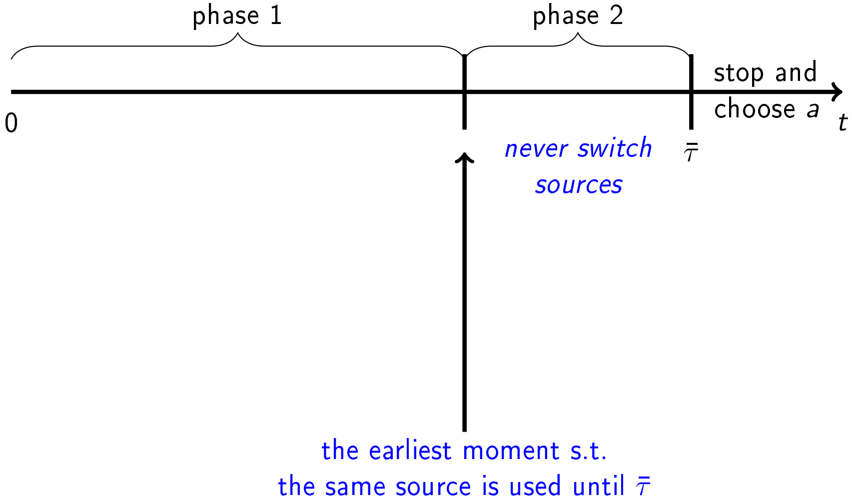

Optimal Strategy

Theorem: The optimal contingency plan has two phases.

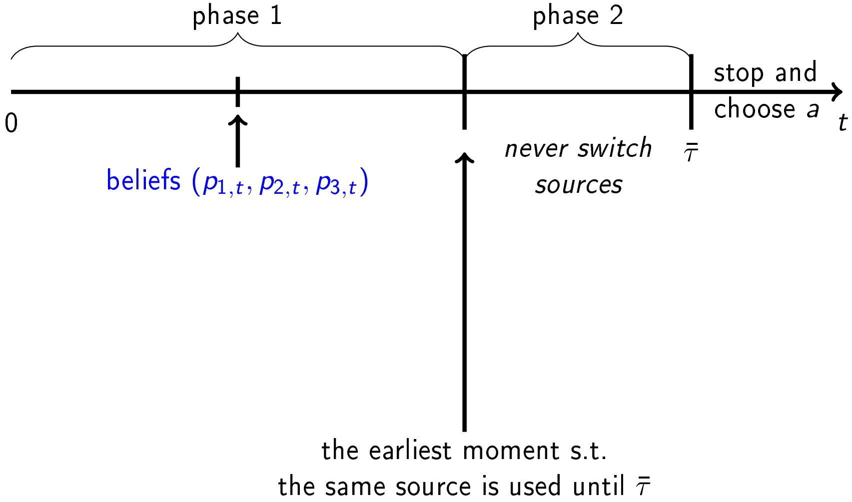

Phase 1

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

Phase 1: Index

Assumption

$\theta_1 = 1 \Rightarrow$ stop and choose a_1; $\theta_2 = 1 \Rightarrow$ stop and choose a_2

Theorem

Phase 1 rule: use the source with the highest index

$$\frac{\lambda_k \, \mathbb{P}\left(\theta_k = 1\right)}{c_k}$$

if $a_1 = a_2$

Phase 1: Index

Assumption

$\theta_1 = 1 \Rightarrow$ stop and choose a_1; $\theta_2 = 1 \Rightarrow$ stop and choose a_2

Theorem

Phase 1 rule: use the source with the highest index

$$\frac{\lambda_k(p_k + p_3)}{c_k}$$

if $a_1 = a_2$

Phase 1: Index

Assumption

$\theta_1 = 1 \Rightarrow$ stop and choose a_1; $\theta_2 = 1 \Rightarrow$ stop and choose a_2

Theorem

Phase 1 rule: use the source with the highest index

$$\frac{\lambda_k(p_k + p_3)}{c_k} + p_3 \times \frac{\lambda_1 \lambda_2}{c_1 c_2} \times \underbrace{u_3(a_k)}_{\substack{\text{payoff in} \\ \text{state } (1,1)}}$$

Theorem

Phase 1 rule: use the source with the highest index

$$p_k + p_3 u_3(a_k)$$

if $c_1 = c_2 = 1$ *and* $\lambda_1 = \lambda_2 = 1$

Theorem

Phase 1 rule: use the source with the highest index

$$p_k + p_3 u_3(a_k)$$

if $c_1 = c_2 = 1$ *and* $\lambda_1 = \lambda_2 = 1$

→ get positive signal → (1,0)

→ get positive signal → (0,1)

→ get positive signal → (1,1)

→ absence of evidence → deterministic search path | eventually search in *both* places → does not matter where to search first

Theorem

Phase 1 rule: use the source with the highest index

source 1: p_1 source 2: p_2

if $c_1 = c_2 = 1$ *and* $\lambda_1 = \lambda_2 = 1$

→ get positive signal →
- p_1 (1,0) → use source 1
- p_2 (0,1) → use source 2
- (1,1)

→ absence of evidence → deterministic search path | eventually search in *both* places → does not matter where to search first

Theorem

Phase 1 rule: use the source with the highest index

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

if $c_1 = c_2 = 1$ *and* $\lambda_1 = \lambda_2 = 1$

→ get positive signal →
- p_1 (1,0) ⟶ use source 1
- p_2 (0,1) ⟶ use source 2
- p_3 (1,1) ⟶ use the source with higher $u_3(a_k)$

→ absence of evidence → deterministic search path | eventually search in *both* places → does not matter where to search first

Properties of the Index

Theorem

Phase 1 rule: use the source with the highest index

$$\boxed{\frac{\lambda_k \, \mathbb{P}\left(\theta_k = 1\right)}{c_k} + \mathbb{P}\left(\theta_1 = \theta_2 = 1\right) \frac{\lambda_1 \lambda_2}{c_1 \, c_2} u_3(a_k)}$$

- Payoffs in states (0,0), (1,0) and (0,1) do not matter

Properties of the Index

Theorem

Phase 1 rule: use the source with the highest index

$$\boxed{\frac{\lambda_k \, \mathbb{P}\left(\theta_k = 1\right)}{c_k} + \mathbb{P}\left(\theta_1 = \theta_2 = 1\right) \frac{\lambda_1 \lambda_2}{c_1 c_2} u_3(a_k)}$$

- Payoffs in states (0,0), (1,0) and (0,1) do not matter
- As we increase the correlation between θ_1 and θ_2, the difference $u_3(a_1) - u_3(a_2)$ plays more important role

Conclusion

Question

How to choose the research direction **if we don't know the goal**?

Answer

In a long-term research, **you don't have to know the goal** *fully* to behave optimally:

Conclusion

Question

How to choose the research direction **if we don't know the goal**?

Answer

In a long-term research, **you don't have to know the goal** *fully* to behave optimally: **search for the truth**

1. Current beliefs & difficulty of research

Conclusion

Question

How to choose the research direction **if we don't know the goal**?

Answer

In a long-term research, **you don't have to know the goal** *fully* to behave optimally: **search for the truth**, taking into account *lost opportunities*:

1. Current beliefs & difficulty of research
2. Lost opportunities:

	true state		
research direction 1	$\boxed{\checkmark}$	$\longleftarrow$	how much you gain if another research direction was successful
research direction 2	$\boxed{\checkmark}$	$\longleftarrow$	discover only this one