Complementary Information and Learning Traps

Annie Liang (UPenn)

Xiaosheng Mu (Columbia)

Introduction

- Information is produced within a society by many individuals over time.
- What helps a society efficiently refine its knowledge over time vs. get trapped in slower, less productive paths?

Introduction

- Information is produced within a society by many individuals over time.
- What helps a society efficiently refine its knowledge over time vs. get trapped in slower, less productive paths?
- Key mechanism we explore: informational complementarities.

Introduction

- Information is produced within a society by many individuals over time.
- What helps a society efficiently refine its knowledge over time vs. get trapped in slower, less productive paths?
- Key mechanism we explore: informational complementarities.
- Build a model in which agents sequentially acquire information to learn about an unknown parameter (e.g. research)
- Path-dependencies: past research affects what information is most valuable to acquire today.

Introduction

- Information is produced within a society by many individuals over time.
- What helps a society efficiently refine its knowledge over time vs. get trapped in slower, less productive paths?
- Key mechanism we explore: informational complementarities.
- Build a model in which agents sequentially acquire information to learn about an unknown parameter (e.g. research)
- Path-dependencies: past research affects what information is most valuable to acquire today.
- Study decentralized learning:
  - agents acquire best information for now
  - don't take into account externality on path of research
- Compare the outcome of this decentralized information acquisition process with the optimal one.

Model

Informational Environment

K Unknowns: $(\omega, b_1, \ldots, b_{K-1}) \sim \mathcal{N}(\mu, \Sigma)$

ω	payoff-relevant
$b_1, \ldots, b_{K-1}$	"confounding variables" or "biases"

Informational Environment

K Unknowns: $(\omega, b_1, \ldots, b_{K-1}) \sim \mathcal{N}(\mu, \Sigma)$

ω	payoff-relevant
$b_1, \ldots, b_{K-1}$	"confounding variables" or "biases"

N Sources/Signals: $X_i^t = c_i^1 \omega + c_i^2 b_1 + \cdots + c_i^K b_{K-1} + \varepsilon_i^t$

$\varepsilon_i^t \sim \mathcal{N}(0, 1)$ i.i.d. across time and sources

Interpretations of Payoff-Irrelevant Terms

$$X_i = \underbrace{\omega} + b_1 + \cdots + b_{K-1} + \varepsilon_i$$

level of dopamine

Interpretations of Payoff-Irrelevant Terms

$$\underbrace{X_i}_{} = \omega + b_1 + \cdots + b_{K-1} + \varepsilon_i$$

measurement of dopamine using sensor

Interpretations of Payoff-Irrelevant Terms

$$X_i = \omega + \underbrace{b_1 + \cdots + b_{K-1}}_{} + \varepsilon_i$$

confounding chemicals also picked up by sensor

Interpretations of Payoff-Irrelevant Terms

$$X_i = \underbrace{\omega}\, c_i^1 + b_1 c_i^2 + \ldots b_{K-1} c_i^K + \varepsilon_i$$

effect of micro-lending on poverty

Interpretations of Payoff-Irrelevant Terms

$$X_i = \omega \underbrace{c_i^1}_{} + b_1 c_i^2 + \ldots b_{K-1} c_i^K + \varepsilon_i$$

loan amount

Interpretations of Payoff-Irrelevant Terms

$$\underbrace{X_i} = \omega c_i^1 + b_1 c_i^2 + \ldots b_{K-1} c_i^K + \varepsilon_i$$

reduction in poverty rate

Interpretations of Payoff-Irrelevant Terms

$$X_i = \omega c_i^1 + \underbrace{b_1 c_i^2 + \dots b_{K-1} c_i^K}_{} + \varepsilon_i$$

controls: age, gender, etc.

Decision Environment

Agents indexed by $t \in \mathbb{Z}_+$ move sequentially. Each agent:

1. chooses a source $i \in \{1, \ldots, N\}$
2. observes an independent realization of X_i
3. chooses a_t and receives $-\mathbb{E}[(a_t - \omega)^2]$

$\Longrightarrow$ choose source that minimizes posterior variance about ω

All signal realizations are public.

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

$t = 1$

choose X_2

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

$t = 1$

choose X_2

$t = 2$

observe prior realization of X_2

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

$t = 1$

choose X_2

$t = 2$

observe prior realization of X_2

choose X_1

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

$t = 1$

choose X_2

$t = 2$

observe prior realization of X_2

choose X_1

$t = 3$

observe prior realizations of X_1,X_2

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

$t = 1$

choose X_2

$t = 2$

observe prior realization of X_2

choose X_1

$t = 3$

observe prior realizations of X_1,X_2

choose X_8

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

$t = 1$

choose X_2

$t = 2$

observe prior realization of X_2

choose X_1

$t = 3$

observe prior realizations of X_1,X_2

choose X_8

etc.

Summary of Model

Available sources (e.g. different kinds of studies): $X_1, X_2, \ldots, X_N$

etc.

$t = 1$	$t = 2$	$t = 3$
choose X_2	observe prior realization of X_2	observe prior realizations of X_1,X_2
	choose X_1	choose X_8

Q: Which information sources are attended to in the long run?
What is the speed at which society learns ω?

Related Literature

Builds on the **social learning** literature (Banerjee ('92); Bikhchandani, Hirshleifer & Welch ('92))

- but assumption of **public** info shuts down classic learning friction
- **endogenous** info acquisition (Mueller-Frank & Pai ('16); Ali ('17))

Related Literature

Builds on the **social learning** literature (Banerjee ('92); Bikhchandani, Hirshleifer & Welch ('92))

- but assumption of **public** info shuts down classic learning friction
- **endogenous** info acquisition (Mueller-Frank & Pai ('16); Ali ('17))

Not a **multi-armed bandit** problem: separation of action that yields information vs. action that yields payoff

Related Literature

Builds on the **social learning** literature (Banerjee ('92); Bikhchandani, Hirshleifer & Welch ('92))

- but assumption of **public** info shuts down classic learning friction
- **endogenous** info acquisition (Mueller-Frank & Pai ('16); Ali ('17))

Not a **multi-armed bandit** problem: separation of action that yields information vs. action that yields payoff

Informational complementarities: Milgrom & Weber ('82); Myatt & Wallace ('08); Borgers et. al ('13); Chade & Eeckhout ('18)

- literature has primarily focused on **static** settings
- study consequences of these complementarities for **dynamic** learning

Preview of Main Results

- Informational environments fall into two categories:
  1. Guaranteed efficient long-run learning
  2. Potential “learning traps”: slow, less productive paths (depends on prior belief)

Preview of Main Results

- Informational environments fall into two categories:
  1. Guaranteed efficient long-run learning
  2. Potential "learning traps": slow, less productive paths (depends on prior belief)

- These environments are differentiated by the structure of complementarities across available kinds of information.

Example: Learning Trap

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

- Efficient learning: eventually concentrate on $\{X_1, X_2\}$

Example: Learning Trap

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

- Efficient learning: eventually concentrate on $\{X_1, X_2\}$
- But consider a prior where ω and b_1 are independent, variance about b_1 is large.

Example: Learning Trap

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

- Efficient learning: eventually concentrate on $\{X_1, X_2\}$
- But consider a prior where ω and b_1 are independent, variance about b_1 is large.
- X_3 chosen by every agent.
- Socially beneficial to acquire information about b_1, but agents don't invest.

Example: Efficient Learning

$$X_1 = \omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + b_2 + \varepsilon_2$$
$$X_3 = b_2 + \varepsilon_3$$
$$X_4 = 10\omega + b_1 + 2b_2 + \varepsilon_4$$

- Efficient learning: eventually concentrate on $\{X_2, X_3, X_4\}$.
- Starting from any prior, agents eventually end up concentrating all acquisitions on efficient set
- Key difference from the previous example is that all paths for learning ω involve learning b_1 and b_2.

  $\Longrightarrow$ inefficient set such as $\{X_1, X_2, X_3\}$ is not self-reinforcing

Plan for Talk

1. What sampling procedure is socially "best"?
2. How will agents acquire information?
3. What is the extent of welfare loss under learning traps?
4. What can a policymaker do to break learning traps?

What Sampling Procedure Is Socially "Best"?

Definition: Complementary Sets

Definition

$\mathcal{S} \subseteq \{1, \ldots, N\}$ is a **complementary set** if:

- it is possible to recover ω from infinite observations of the sources in $\mathcal{S}$
- every source in $\mathcal{S}$ is necessary for this recovery

$$X_1 = \omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + b_2 + \varepsilon_2$$
$$X_3 = b_2 + \varepsilon_3$$

Characterization of Complementary Sets

Claim

$\mathcal{S}$ is complementary $\iff$ *There are unique, nonzero coefficients $\beta_i^{\mathcal{S}}$ s.t. $(1, 0, \ldots, 0)' = \sum_{i \in \mathcal{S}} \beta_i^{\mathcal{S}} \cdot c_i$*

Characterization of Complementary Sets

Claim

$\mathcal{S}$ is complementary $\iff$ *There are unique, nonzero coefficients $\beta_i^{\mathcal{S}}$ s.t. $(1, 0, \ldots, 0)' = \sum_{i \in \mathcal{S}} \beta_i^{\mathcal{S}} \cdot c_i$*

$$X_1 = \omega + b_1 + \varepsilon_1 \qquad c_1 = (1, 1)$$
$$X_2 = b_1 + \varepsilon_2 \qquad c_2 = (0, 1)$$

unique decomposition $(1, 0) = c_1 - c_2$
set $\{X_1, X_2\}$ is complementary

Characterization of Complementary Sets

Claim

$\mathcal{S}$ is complementary $\iff$ *There are unique, nonzero coefficients* $\beta_i^{\mathcal{S}}$ *s.t.* $(1, 0, \ldots, 0)' = \sum_{i \in \mathcal{S}} \beta_i^{\mathcal{S}} \cdot c_i$

$$X_1 = \omega + \varepsilon_1 \qquad c_1 = (1, 0)$$
$$X_2 = 2\omega + \varepsilon_2 \qquad c_2 = (2, 0)$$

many decompositions $(1, 0) = \beta_1 c_1 + \beta_2 c_2$

set $\{X_1, X_2\}$ is not complementary

Informational Value of Complementary Sets

Claim

Optimal sampling from complementary set $\mathcal{S}$ leads to posterior variance that vanishes like $1/(t \cdot val(\mathcal{S}))$ where

$$val(\mathcal{S}) = \left(\frac{1}{\sum_{i \in \mathcal{S}} |\beta_i^{\mathcal{S}}|} \right)^2$$

- Larger $val(\mathcal{S})$ is better.

Informational Value of Complementary Sets

Claim

Optimal sampling from complementary set $\mathcal{S}$ leads to posterior variance that vanishes like $1/(t \cdot val(\mathcal{S}))$ where

$$val(\mathcal{S}) = \left(\frac{1}{\sum_{i \in \mathcal{S}} |\beta_i^{\mathcal{S}}|} \right)^2$$

- Larger $val(\mathcal{S})$ is better.

Assumption: All complementary sets have distinct values ($\text{val}(\mathcal{S}) \neq \text{val}(\mathcal{S}')$ for all complementary $\mathcal{S}$, $\mathcal{S}'$).

- generically satisfied

Informational Value of Complementary Sets

Claim

Optimal sampling from complementary set $\mathcal{S}$ leads to posterior variance that vanishes like $1/(t \cdot val(\mathcal{S}))$ where

$$val(\mathcal{S}) = \left(\frac{1}{\sum_{i \in \mathcal{S}} |\beta_i^{\mathcal{S}}|} \right)^2$$

- Larger $val(\mathcal{S})$ is better.

Assumption: All complementary sets have distinct values ($\text{val}(\mathcal{S}) \neq \text{val}(\mathcal{S}')$ for all complementary $\mathcal{S}$, $\mathcal{S}'$).

- generically satisfied

Assumption: There is at least one complementary set $\mathcal{S} \subseteq [N]$.

Informational Value of Complementary Sets

Claim

Optimal sampling from complementary set $\mathcal{S}$ leads to posterior variance that vanishes like $1/(t \cdot val(\mathcal{S}))$ where

$$val(\mathcal{S}) = \left(\frac{1}{\sum_{i \in \mathcal{S}} |\beta_i^{\mathcal{S}}|} \right)^2$$

- Larger $val(\mathcal{S})$ is better.

Assumption: All complementary sets have distinct values ($\mathrm{val}(\mathcal{S}) \neq \mathrm{val}(\mathcal{S}')$ for all complementary $\mathcal{S}$, $\mathcal{S}'$).

- generically satisfied

Assumption: There is at least one complementary set $\mathcal{S} \subseteq [N]$.

$\Longrightarrow$ unique best complementary set $\mathcal{S}^*$

Long-Run Frequency Vector

Define $\lambda_i^* = \begin{cases} \dfrac{|\beta_i^{\mathcal{S}^*}|}{\sum_{j \in \mathcal{S}^*} |\beta_j^{\mathcal{S}^*}|} & \forall\ i \in \mathcal{S}^* \\ 0 & \text{otherwise} \end{cases}$

Concentrates all acquisitions on best complementary set $\mathcal{S}^*$.

λ^* is Optimal

For fixed δ, let $d^\delta(t) = (d_1^\delta(t), \ldots, d_N^\delta(t))$ be the vector of signal counts (up to time t) associated with any strategy that maximizes

$$U_\delta := -\mathbb{E}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}\cdot(a_t-\omega)^2\right]$$

Let $\lambda^\delta(t) = d^\delta(t)/t$ be the associated frequencies.

Proposition

There exists $\underline{\delta} < 1$ such that for any $\delta \geq \underline{\delta}$,

$$\lim_{t\to\infty} \lambda^\delta(t) = \lambda^*.$$

Henceforth call λ^* the optimal long-run frequency.

How Will Agents Actually Acquire Information?

Possible Learning Traps

Assumption (Strong Linear Independence)

Suppose every $K \times K$ submatrix of the coefficient matrix C (with rows c_i) has full rank.

Candidate "learning traps" are complementary sets with fewer than K signals (recall that $K =$ number of unknown states):

Proposition

Assume Strong Linear Independence. For every complementary set $\mathcal{S}$ where $|\mathcal{S}| < K$, there is an open set of prior beliefs from which $\mathcal{S}$ is eventually exclusively observed.

Example: Learning Trap

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

Another Example: Disjoint Communities

$$X_1 = \omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$

$$X_3 = 2\omega + b_2 + \varepsilon_3$$
$$X_4 = b_2 + \varepsilon_4$$

$$X_5 = 3\omega + b_3 + \varepsilon_5$$
$$X_6 = b_3 + \varepsilon_6$$

See: Sethi and Yildiz (2019)

Converse: Efficient Information Aggregation

Proposition

Suppose there are no complementary sets with fewer than K sources. Then from any prior belief, long-run frequencies are given by λ^* (and agents eventually exclusively observe $\mathcal{S}^*$).

Example: Efficient Learning

$$X_1 = \omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + b_2 + \varepsilon_2$$
$$X_3 = b_2 + \varepsilon_3$$
$$X_4 = 10\omega + b_1 + 2b_2 + \varepsilon_4$$

all complementary sets are of size $K = 3$

Why K?

Key force is a learning spillover effect:

- $X = \omega + b + \varepsilon$ provides information both about ω and confounding term b. (Sethi and Yildiz, 2016)

Why K?

Key force is a learning spillover effect:

- $X = \omega + b + \varepsilon$ provides information both about ω and confounding term b. (Sethi and Yildiz, 2016)
- Observation of X helps agents interpret all other signals confounded by b.

Why K?

Key force is a learning spillover effect:

- $X = \omega + b + \varepsilon$ provides information both about ω and confounding term b. (Sethi and Yildiz, 2016)
- Observation of X helps agents interpret all other signals confounded by b.
- Observation of K signals $\Rightarrow$ learn all K states, eventually evaluate all sources by prior-independent criterion.

Why K?

Key force is a learning spillover effect:

- $X = \omega + b + \varepsilon$ provides information both about ω and confounding term b. (Sethi and Yildiz, 2016)
- Observation of X helps agents interpret all other signals confounded by b.
- Observation of K signals $\Rightarrow$ learn all K states, eventually evaluate all sources by prior-independent criterion.
- Observation of $k < K$ signals allows for persistent uncertainty about some confounding terms, and hence some sources.

Technical Intuition: Part 1/2

Each period, increase count vector $(q_1(t), \ldots, q_N(t))$ in coordinate that maximally reduces variance $V(q_1, \ldots, q_N)$.

Technical Intuition: Part 1/2

Each period, increase count vector $(q_1(t), \ldots, q_N(t))$ in coordinate that maximally reduces variance $V(q_1, \ldots, q_N)$.

Consider "asymptotic normalized variance" $V^*(\lambda) = \lim_{t\to\infty} t \cdot V(\lambda t)$.

- Strictly convex in λ, unique minimum is (optimal) λ^*.
- Efficient info aggregation obtains if $\lambda(t) \to \lambda^*$.

Technical Intuition: Part 1/2

Each period, increase count vector $(q_1(t), \ldots, q_N(t))$ in coordinate that maximally reduces variance $V(q_1, \ldots, q_N)$.

Consider "asymptotic normalized variance" $V^*(\lambda) = \lim_{t \to \infty} t \cdot V(\lambda t)$.

- Strictly convex in λ, unique minimum is (optimal) λ^*.
- Efficient info aggregation obtains if $\lambda(t) \to \lambda^*$.

Acquisitions follow "pseudo"-gradient descent: finite set of feasible directions.

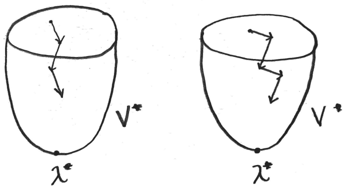

Technical Intuition: Part 2/2

- Restricted movement is without loss if V^* is differentiable
  - directional derivatives can be rewritten as convex combination of partial derivatives along basis vectors

Technical Intuition: Part 2/2

- Restricted movement is without loss if V^* is differentiable
  - directional derivatives can be rewritten as convex combination of partial derivatives along basis vectors
- V^* is not differentiable everywhere; descent can become stuck.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

e.g. consider frequency vector $(0, 0, 1)$.

Technical Intuition: Part 2/2

- Restricted movement is without loss if V^* is differentiable
  - directional derivatives can be rewritten as convex combination of partial derivatives along basis vectors
- V^* is not differentiable everywhere; descent can become stuck.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

e.g. consider frequency vector $(0, 0, 1)$.

- Sufficient condition for differentiability: λ has at least K nonzero coordinates.

Technical Intuition: Part 2/2

- Restricted movement is without loss if V^* is differentiable
  - directional derivatives can be rewritten as convex combination of partial derivatives along basis vectors
- V^* is not differentiable everywhere; descent can become stuck.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

e.g. consider frequency vector $(0, 0, 1)$.

- Sufficient condition for differentiability: λ has at least K nonzero coordinates.
- When all complementary sets are at least size K, the sufficient condition is met on path.

Generalizing these Results

Definition

Say that a set $\mathcal{S}$ is **strongly complementary** if:

- $\mathcal{S}$ is complementary
- $\text{val}(\mathcal{S}) > \text{val}(\mathcal{S}')$ for all sets $\mathcal{S}'$ where $|\mathcal{S} - \mathcal{S}'| = |\mathcal{S}' - \mathcal{S}| = 1$

Generalizing these Results

Definition

Say that a set $\mathcal{S}$ is **strongly complementary** if:

- $\mathcal{S}$ is complementary
- $\text{val}(\mathcal{S}) > \text{val}(\mathcal{S}')$ for all sets $\mathcal{S}'$ where $|\mathcal{S} - \mathcal{S}'| = |\mathcal{S}' - \mathcal{S}| = 1$

Example:

$$X_1 = \omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = 2b_1 + \varepsilon_3$$

$\{X_1, X_2\}$ is not strongly complementary, since

$$\text{val}(\{X_1, X_3\}) > \text{val}(\{X_1, X_2\}).$$

Characterization of Long-Run Outcomes

Theorem

Agents eventually exclusively observe $\mathcal{S}$ (from some set of priors)	$\iff$	*$\mathcal{S}$ is a strongly complementary set*

Characterization of Long-Run Outcomes

Theorem

Agents eventually exclusively observe $\mathcal{S}$ (from some set of priors) $\iff$ *$\mathcal{S}$ is a strongly complementary set*

- Note: best set $\mathcal{S}^*$ is always strongly complementary.

Characterization of Long-Run Outcomes

Theorem

Agents eventually exclusively observe $\mathcal{S}$ (from some set of priors)	$\iff$	*$\mathcal{S}$ is a strongly complementary set*

- Note: best set $\mathcal{S}^*$ is always strongly complementary.
- Unique strongly complementary set $\Rightarrow$ all priors lead to same (best) long-run outcome.

Characterization of Long-Run Outcomes

Theorem

Agents eventually exclusively observe $\mathcal{S}$ (from some set of priors)	$\Longleftrightarrow$	*$\mathcal{S}$ is a strongly complementary set*

- Note: best set $\mathcal{S}^*$ is always strongly complementary.
- Unique strongly complementary set $\Rightarrow$ all priors lead to same (best) long-run outcome.
- Multiplicity $\Rightarrow$ different priors can lead to different outcomes.

Comments

In general cannot order the sizes of learning traps and efficient sets.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

Efficient set $\{X_1, X_2\}$ is larger than learning trap $\{X_3\}$.

Comments

In general cannot order the sizes of learning traps and efficient sets.

$$
\begin{aligned}
X_1 &= \omega + b_1 + \varepsilon_1 \\
X_2 &= b_1 + b_2 + \varepsilon_2 \\
X_3 &= b_2 + \varepsilon_3
\end{aligned}
$$

$$
\begin{aligned}
X_4 &= \omega + b_3 + \varepsilon_4 \\
X_5 &= b_3 + \varepsilon_5
\end{aligned}
$$

Efficient set $\{X_4, X_5\}$ is smaller than learning trap $\{X_1, X_2, X_3\}$.

Comments

Adding a source can worsen speed of learning.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$

Comments

Adding a source can worsen speed of learning.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

Comments

Degrading a source can improve speed of learning.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

Comments

Degrading a source can improve speed of learning.

$$
\begin{aligned}
X_1 &= 5\omega + b_1 + \varepsilon_1 \\
X_2 &= b_1 + \varepsilon_2 \\
X_3 &= \omega + b_1 + \varepsilon_3
\end{aligned}
$$

Comments

Degrading a source can improve speed of learning.

$$X_1 = 5\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + b_1 + \varepsilon_3$$

In general, adjustments to the signal structure can change the efficient set, so welfare comparisons are not straightforward.

Welfare Loss Under Learning Traps

Criteria

- speed of information aggregation achieved by community (relative to efficiency)
- payoffs achieved by community (relative to social planner)

Information Aggregation

Extension of previous learning traps example:

$$X_1 = \zeta\omega + b_1 + \varepsilon_1$$
$$X_2 = b_1 + \varepsilon_2$$
$$X_3 = \omega + \varepsilon_3$$

$$\frac{\mathrm{val}(\{X_1, X_2\})}{\mathrm{val}(\{X_3\})} = \zeta^2/4 \text{ can be made arbitrarily large.}$$

Information aggregation can be arbitrarily inefficient.

Payoff Inefficiency Under Learning Traps

$$\underbrace{U_{SP}^{\delta} := -\mathbb{E}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}(a_{SP}^{t}-\omega)^2\right]}_{\delta\text{-discounted average payoffs under social planner rule}}$$

$$\underbrace{U_{M}^{\delta} := -\mathbb{E}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}(a_{M}^{t}-\omega)^2\right]}_{\delta\text{-discounted average payoffs given (myopic) social acquisitions}}$$

Payoff Inefficiency Under Learning Traps

Compare (social planner payoff) U_{SP}^{δ} and (myopic payoff) U_{M}^{δ}:

1. payoff ratio $\lim_{\delta \to 1} U_M^{\delta} / U_{SP}^{\delta}$
2. payoff difference $\lim_{\delta \to 1} (U_{SP}^{\delta} - U_M^{\delta})$

Payoff Inefficiency Under Learning Traps

Compare (social planner payoff) U_{SP}^{δ} and (myopic payoff) U_M^{δ}:

1. payoff ratio $\lim_{\delta \to 1} U_M^{\delta} / U_{SP}^{\delta}$ can be arbitrarily large
2. payoff difference $\lim_{\delta \to 1} (U_{SP}^{\delta} - U_M^{\delta})$

Payoff Inefficiency Under Learning Traps

Compare (social planner payoff) U_{SP}^{δ} and (myopic payoff) U_{M}^{δ}:

1. payoff ratio $\lim_{\delta \to 1} U_M^{\delta} / U_{SP}^{\delta}$ can be arbitrarily large
2. payoff difference $\lim_{\delta \to 1} (U_{SP}^{\delta} - U_M^{\delta})$ always zero(!)

Payoff Inefficiency Under Learning Traps

Compare (social planner payoff) U_{SP}^{δ} and (myopic payoff) U_{M}^{δ}:

1. payoff ratio $\lim_{\delta \to 1} U_M^{\delta} / U_{SP}^{\delta}$ can be arbitrarily large
2. payoff difference $\lim_{\delta \to 1} (U_{SP}^{\delta} - U_M^{\delta})$ always zero(!)
    - depends heavily on assumption of persistent state

Payoff Inefficiency Under Learning Traps

Compare (social planner payoff) U_{SP}^{δ} and (myopic payoff) U_{M}^{δ}:

1. payoff ratio $\lim_{\delta \to 1} U_{M}^{\delta} / U_{SP}^{\delta}$ can be arbitrarily large
2. payoff difference $\lim_{\delta \to 1} (U_{SP}^{\delta} - U_{M}^{\delta})$ always zero(!)
   - depends heavily on assumption of persistent state
   - next: relax this assumption, allow for (slowly evolving) states

Extension to Evolving States

- Suppose instead the state vector $\theta^t = (\omega^t, b_1^t, \ldots, b_{K-1}^t)$ evolves according to:

$$\theta^1 \sim \mathcal{N}(0, \Sigma^0)$$
$$\theta^{t+1} = \sqrt{\alpha} \cdot \theta^t + \sqrt{1-\alpha} \cdot \mathcal{N}(0, M).$$

- In each period, available signals are
$$X_i^t = c_i^1 \omega^t + c_i^2 b_1^t + \cdots + c_i^K b_{K-1}^t + \mathcal{N}(0, 1).$$

- Each agent t chooses one signal to (myopically) optimize prediction of ω^t.

Extension of Learning Traps (Informal)

Theorem

Fix any strongly complementary set S. There is an M and Σ_0, such that for all α sufficiently large:

1. *social acquisitions eventually concentrate on the set $\mathcal{S}$*

Extension of Learning Traps (Informal)

Theorem

Fix any strongly complementary set S. There is an M and Σ_0, such that for all α sufficiently large:

1. *social acquisitions eventually concentrate on the set $\mathcal{S}$*
2. *and yield a discounted average payoff of (approximately)*

$$-\sqrt{(1-\alpha)\left(\frac{M_{11}}{\mathsf{val}(\mathcal{S})}\right)}$$

while it is feasible to achieve (approximately)

$$-\sqrt{(1-\alpha)\left(\frac{M_{11}}{\mathsf{val}(\mathcal{S}^*)}\right)}$$

by sampling from the best set $\mathcal{S}^*$.

Extension of Learning Traps (Informal)

Theorem

Fix any strongly complementary set S. There is an M and Σ_0, such that for all α sufficiently large:

1. *social acquisitions eventually concentrate on the set $\mathcal{S}$*
2. *and yield a discounted average payoff of (approximately)*

$$-\sqrt{(1-\alpha)\left(\frac{M_{11}}{\mathrm{val}(\mathcal{S})}\right)}$$

while it is feasible to achieve (approximately)

$$-\sqrt{(1-\alpha)\left(\frac{M_{11}}{\mathrm{val}(\mathcal{S}^*)}\right)}$$

by sampling from the best set $\mathcal{S}^*$.

$\longrightarrow$ inefficient average payoffs away from perfect persistence.

Information Interventions For Breaking Learning Traps

Information Interventions

- Have observed the possibility of (large) welfare inefficiencies under learning traps.
- What kind of policy interventions might "break" learning traps?

More Precise Signals

Each agent receives B observations of the signal they choose.

Corollary

Suppose that for $B = 1$, there is a set of priors given which signals in $\mathcal{S}$ are (exclusively) viewed in the long run. Then, for every $B \in \mathbb{Z}_+$, there is a set of priors given which $\mathcal{S}$ is exclusively viewed in the long run.

More Precise Signals

Each agent receives B observations of the signal they choose.

Corollary

Suppose that for $B = 1$, there is a set of priors given which signals in $\mathcal{S}$ are (exclusively) viewed in the long run. Then, for every $B \in \mathbb{Z}_+$, there is a set of priors given which $\mathcal{S}$ is exclusively viewed in the long run.

In practice: funding for additional lab subjects per experiment

More Precise Signals

Each agent receives B observations of the signal they choose.

Corollary

Suppose that for $B = 1$, there is a set of priors given which signals in $\mathcal{S}$ are (exclusively) viewed in the long run. Then, for every $B \in \mathbb{Z}_+$, there is a set of priors given which $\mathcal{S}$ is exclusively viewed in the long run.

In practice: funding for additional lab subjects per experiment

- But sets of priors are different.
- Can be that at the given prior, increasing B breaks the learning trap.
- See examples for both cases in paper.

More Signal Observations

Each agent can allocate B observations across the available signals.

Proposition

There is a B such that given acquisition of B signals every period, efficient information aggregation obtains.

In practice: evaluation of body of work, lab structure

More Signal Observations

Each agent can allocate B observations across the available signals.

Proposition

There is a B such that given acquisition of B signals every period, efficient information aggregation obtains.

In practice: evaluation of body of work, lab structure

- But precise number of signals B needed depends on the specific environment.
- E.g. simple rules like "allow each DM to acquire K signals" are not sufficient.

Free Information

- Policy-maker chooses M signals of form

$$p_1\omega + p_2 b_1 + \cdots + p_K b_{K-1} + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0,1)$$

  where $\|p\|_2 \leq \gamma$.

- Free information revealed to all agents at $t = 0$.

Proposition

Let $k \leq K$ be the size of the optimal set $\mathcal{S}^*$.

There exist $k-1$ signals of finite precision such that with these free signals provided at $t = 0$, efficient information aggregation obtains.

Free Information

- Policy-maker chooses M signals of form
$$p_1\omega + p_2 b_1 + \cdots + p_K b_{K-1} + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0,1)$$
where $\|p\|_2 \leq \gamma$.
- Free information revealed to all agents at $t = 0$.

Proposition

Let $k \leq K$ be the size of the optimal set $\mathcal{S}^*$.

There exist $k - 1$ signals of finite precision such that with these free signals provided at $t = 0$, efficient information aggregation obtains.

In practice: funding agencies can encourage investment in learning about unknowns not directly of social interest (e.g. methodological work).

Extensions

- General payoff functions $u_t(a_t, \omega)$
  - All results extend, but interpretation of optimal benchmark more limited.
- Incomplete learning (can't recover ω from available sources)
  - Define new payoff-relevant state $\widetilde{\omega}$ as the "learnable component" of ω, results extend with some subtleties
- Multiple payoff-relevant states
  - All results extend, need to adapt idea of "complementary."

Conclusion

Other interpretations and directions:

- Multiple priors
  - So far: common prior. Consider instead multiple sequences of DMs with different priors.
    - e.g. researchers in Israel, US, China...
  - Results explain what features of the signal structure determine whether these groups will end up observing the same signals:
    - Unique strongly complementary set $\Rightarrow$ same long-run observations.
    - Otherwise, persistent differences in what sources are listened to across groups.
- Choice between actions with complementarities.
  - e.g. managers take actions that have externalities for future managers, each manager maximizes profit during his tenure.
  - But need analogue to "variance"...

Thank You