Optimal and Myopic Information Acquisition

Annie Liang[1] Xiaosheng Mu[2] Vasilis Syrgkanis[3]

[1]UPenn

[2]Harvard [3]MSR

Introduction

- Classic problem:
  - DM repeatedly acquires information and takes actions.
  - Payoff depends on the actions taken, as well as on an unknown payoff-relevant state.
- We consider additional features:
  - acquisition of information from flexibly correlated sources
  - limited attention: fixed number of observations each period.
- Simple strategy for information acquisition: act as if each period were the last. (myopic)
- Main result: in a canonical setting—jointly normal signals—this is sometimes the best you can do.

Preview of Results

We show that the myopic rule is optimal:

1. if signal observations are acquired in sufficiently large blocks each period.
2. and for all block sizes:
   - in “separable” environments.
   - eventually in generic environments.

These results hold across all payoff functions (and in particular, independently of discounting).

Preview of Results

We show that the myopic rule is optimal:

1. if signal observations are acquired in sufficiently large blocks each period.
2. and for all block sizes:
   - in "separable" environments.
   - eventually in generic environments.

These results hold across all payoff functions (and in particular, independently of discounting).

Implications:

- Exactly characterization of dynamically optimal solution.
- Robustness of myopic rule to uncertainty about payoff function and timing of decision.

Model

- K States: $\underbrace{(\theta_1, \ldots, \theta_K)}_{\text{fixed over time}} \sim \mathcal{N}(0, V)$.
- $t = 1, 2, \ldots$.
- N Signals: $X_i^t = \sum_{k=1}^{K} c_{ik}\theta_k + \varepsilon_i^t, \quad \underbrace{\varepsilon_i^t \sim \mathcal{N}\left(0, \sigma_i^2\right)}_{\text{i.i.d. over time}}$.

  c_{ik} and σ_i^2 are known
- Each period t, the DM
  - samples B signals
  - chooses an action $a_t \in A_t$
- Payoff is arbitrary function $U(a_1, a_2, \ldots; \theta_1)$.

Special Cases

Exogenous Final Date:

$$U(a_1, a_2, \ldots; \theta_1) = u_T(a_T, \theta_1)$$

where T is random exogenously determined final time period.

Endogenous Stopping with Per-Period Costs.

- Each a_t specifies both the decision of whether to stop, and also the action to be taken if stopped.
- Payoff as above, but T is endogenously chosen.

Restrictions on Environment

Note assumptions in model:

- One-dimensional payoff-relevant state (some linear combination of $\theta_1, \dots, \theta_K$)
- No feedback from actions

Additionally impose:

Assumption (Non-Redundant Signals)

Infinite observations of each signal are necessary and sufficient to fully learn the payoff-relevant unknown state.

Strategy

A strategy consists of:

- information acquisition strategy strategy: signal choices in each period given history of signal choices and realizations,
- decision strategy: action choice after each history

(Without loss, consider only pure strategies.)

Hence focus on information acquisition strategy.

Myopic Information Acquisition

Definition

An information acquisition strategy is myopic, if at every next period, it prescribes choosing the B signals that (combined with the history of observations) lead to the lowest posterior variance about the payoff-relevant state.

- Blackwell dominates any other multi-set of B signals (Hansen and Torgersen, 1974)$\rightarrow$ best for all payoff criteria.
- Optimal if the current period is the last chance for information acquisition.

Myopic Rule is Optimal Given Large Batch Sizes

Theorem (Immediate Optimality under Many Observations)

Fix any prior and signal structure, and suppose B is sufficiently large. Then the DM has an optimal strategy that acquires information myopically.

Myopic Rule is Optimal in Separable Environments

Definition

The informational environment is separable if there exist convex functions $g_1, \ldots, g_K$ and a strictly increasing function F such that

$$Var(q_1, \ldots, q_K) = F(g_1(q_1) + \cdots + g_K(q_K))$$

where $q_i = \#$ of times signal i observed.

Theorem

Suppose the informational environment is separable. Then for every $B \in \mathbb{N}^+$, the DM has an optimal strategy that acquires information myopically.

Examples of Separable Environments

1. multiple biases — you care about x. first signal tells you $x + b_1 + b_2$, three other signals inform about each individual b_i.

Examples of Separable Environments

1. multiple biases — you care about x. first signal tells you $x + b_1 + b_2$, three other signals inform about each individual b_i.

2. hierarchy of biases — you care about x. signals are

$$\begin{aligned} x + b_1 \quad\quad & + \epsilon_1 \\ b_1 + b_2 & + \epsilon_2 \\ b_2 & + \epsilon_3 \end{aligned}$$

Examples of Separable Environments

1. multiple biases — you care about x. first signal tells you $x + b_1 + b_2$, three other signals inform about each individual b_i.

2. hierarchy of biases — you care about x. signals are

$$\begin{aligned} x + b_1 \qquad\quad &+ \epsilon_1 \\ b_1 + b_2 &+ \epsilon_2 \\ b_2 &+ \epsilon_3 \end{aligned}$$

3. symmetric signals — you care about $\theta_1 + \theta_2 + \theta_3$. signals are about $\theta_1 + \theta_2$, $\theta_1 + \theta_3$, $\theta_2 + \theta_3$ respectively.

Eventual Optimality

Theorem (Eventual Optimality)

For generic coefficient matrices C, there exists a time $T^* \in \mathbb{N}$ s.t. for every batch size B, the DM has an optimal strategy that acquires information myopically after T^* periods.

- At all late periods, optimal rule proceeds myopically.
- In paper, complementary result: myopic acquisition eventually leads to optimal signal path.

Summary of Results

Three results regarding optimality of the myopic information acquisition rule:

- Thm 1: Myopic information acquisition is optimal from period 1 if B is sufficiently large.
- Thm 2: For class of separable environments, myopic information acquisition is optimal from period 1 given any B.
- Thm 3: For every B, generically the optimal rule is eventually myopic.

Intuition for Results

One-shot version of problem: optimally allocate t observations across signals.

t-optimal "division vector":

$$n(t) = (n_1(t), \ldots, n_K(t)) \in \underset{q_i \in \mathbb{Z}_+, \sum_{i=1}^{K} q_i = t}{\operatorname{argmin}} \operatorname{Var}(q_1, \ldots, q_K)$$

$n(t)$ Does Not Always Evolve Sequentially

Example 1

$$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0, 1)$$

payoff-relevant state: θ_1

$$\begin{aligned} X_1 &= \theta_1 - \theta_2 + \epsilon_1 \\ X_2 &= \theta_2 - \theta_3 + \epsilon_2 \\ X_3 &= \theta_3 + \epsilon_3 \end{aligned}$$

complementarities across signals

$n(t)$ Does Not Always Evolve Sequentially

Example 1

$$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0, 1)$$

payoff-relevant state: θ_1

$n(4)$	
3	$X_1 = \theta_1 - \theta_2 + \epsilon_1$
1	$X_2 = \theta_2 - \theta_3 + \epsilon_2$
0	$X_3 = \theta_3 + \epsilon_3$

complementarities across signals

$n(t)$ Does Not Always Evolve Sequentially

Example 1

$$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0,1)$$

payoff-relevant state: θ_1

$n(4)$	$n(5)$	
3	4	$X_1 = \theta_1 - \theta_2 + \epsilon_1$
1	1	$X_2 = \theta_2 - \theta_3 + \epsilon_2$
0	0	$X_3 = \theta_3 + \epsilon_3$

complementarities across signals

$n(t)$ Does Not Always Evolve Sequentially

Example 1

$$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0,1)$$

payoff-relevant state: θ_1

$n(4)$	$n(5)$	$n(6)$		
3	4	3	X_1	$= \theta_1 - \theta_2 + \epsilon_1$
1	1	2	X_2	$= \theta_2 - \theta_3 + \epsilon_2$
0	0	1	X_3	$= \theta_3 + \epsilon_3$

complementarities across signals

Sequentiality of $n(t)$ Produces Desired Result

But suppose $(n(t))_{t\geq 1}$ could be achieved by sequential sampling.

Then,

- myopic information acquisition will produce this sampling rule.
- Lemma: the sampling rule is best for all payoff criteria (intuitively: no conflict across periods)

$\Longrightarrow$ myopic rule is optimal

Example in Which $n(t)$ is Sequential

Example 2

$$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0,1)$$

payoff-relevant state: $\theta_1 + \theta_2 + \theta_3$

$$\begin{aligned} X_1 &= \theta_1 + \epsilon_1 \\ X_2 &= \theta_2 + \epsilon_2 \\ X_3 &= \theta_3 + \epsilon_3 \end{aligned}$$

"independent signal structure"
independent prior

Example in Which $n(t)$ is Sequential

Example 2

$$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0, 1)$$

payoff-relevant state: $\theta_1 + \theta_2 + \theta_3$

$n(1)$	$n(2)$	$n(3)$		
1	1	1	X_1	$= \theta_1 + \epsilon_1$
0	1	1	X_2	$= \theta_2 + \epsilon_2$
0	0	1	X_3	$= \theta_3 + \epsilon_3$

"independent signal structure"
independent prior

Relationship Between Two Examples

Let's rewrite Example 1 to look more like Example 2.

Example 1	Example 1, Rewritten
$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0,1)$	
payoff-relevant: θ_1	
$X_1 = \theta_1 - \theta_2 + \epsilon_1$	
$X_2 = \theta_2 - \theta_3 + \epsilon_2$	
$X_3 = \theta_3 + \epsilon_3$	

Relationship Between Two Examples

Let's rewrite Example 1 to look more like Example 2.

Example 1	Example 1, Rewritten
$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0, 1)$	
payoff-relevant: θ_1	
$X_1 = \theta_1 - \theta_2 + \epsilon_1$	$X_1 = \tilde{\theta}_1 + \epsilon_1$
$X_2 = \theta_2 - \theta_3 + \epsilon_2$	
$X_3 = \theta_3 + \epsilon_3$	

Relationship Between Two Examples

Let's rewrite Example 1 to look more like Example 2.

Example 1	Example 1, Rewritten
$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0,1)$	
payoff-relevant: θ_1	
$X_1 = \theta_1 - \theta_2 + \epsilon_1$	$X_1 = \tilde{\theta}_1 + \epsilon_1$
$X_2 = \theta_2 - \theta_3 + \epsilon_2$	$X_2 = \tilde{\theta}_2 + \epsilon_2$
$X_3 = \theta_3 + \epsilon_3$	

Relationship Between Two Examples

Let's rewrite Example 1 to look more like Example 2.

Example 1	Example 1, Rewritten
$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0,1)$	
payoff-relevant: θ_1	
$X_1 = \theta_1 - \theta_2 + \epsilon_1$	$X_1 = \tilde{\theta}_1 + \epsilon_1$
$X_2 = \theta_2 - \theta_3 + \epsilon_2$	$X_2 = \tilde{\theta}_2 + \epsilon_2$
$X_3 = \theta_3 + \epsilon_3$	$X_3 = \tilde{\theta}_3 + \epsilon_3$

Relationship Between Two Examples

Let's rewrite Example 1 to look more like Example 2.

Example 1	Example 1, Rewritten
$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0, 1)$	$\tilde{\theta}_1, \tilde{\theta}_2, \tilde{\theta}_3 \sim \mathcal{N}(\mu, V)$
payoff-relevant: θ_1	$\tilde{\theta}_1 + \tilde{\theta}_2 + \tilde{\theta}_3$
$X_1 = \theta_1 - \theta_2 + \epsilon_1$	$X_1 = \tilde{\theta}_1 + \epsilon_1$
$X_2 = \theta_2 - \theta_3 + \epsilon_2$	$X_2 = \tilde{\theta}_2 + \epsilon_2$
$X_3 = \theta_3 + \epsilon_3$	$X_3 = \tilde{\theta}_3 + \epsilon_3$

Signals "De-Correlate"

Example 1, Rewritten	Example 2
$\tilde{\theta}_1, \tilde{\theta}_2, \tilde{\theta}_3 \sim \mathcal{N}(\mu, V)$	$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0, 1)$
$\tilde{\theta}_1 + \tilde{\theta}_2 + \tilde{\theta}_3$	$\theta_1 + \theta_2 + \theta_3$
$X_1 = \tilde{\theta}_1 + \epsilon_1$	$X_1 = \theta_1 + \epsilon_1$
$X_2 = \tilde{\theta}_2 + \epsilon_2$	$X_2 = \theta_2 + \epsilon_2$
$X_3 = \tilde{\theta}_3 + \epsilon_3$	$X_3 = \theta_3 + \epsilon_3$
correlated prior	independent prior

Signals "De-Correlate"

Example 1, Rewritten	Example 2
$\tilde{\theta}_1, \tilde{\theta}_2, \tilde{\theta}_3 \sim \mathcal{N}(\mu, V)$	$\theta_1, \theta_2, \theta_3 \sim \mathcal{N}(0, 1)$
$\tilde{\theta}_1 + \tilde{\theta}_2 + \tilde{\theta}_3$	$\theta_1 + \theta_2 + \theta_3$
$X_1 = \tilde{\theta}_1 + \epsilon_1$	$X_1 = \theta_1 + \epsilon_1$
$X_2 = \tilde{\theta}_2 + \epsilon_2$	$X_2 = \theta_2 + \epsilon_2$
$X_3 = \tilde{\theta}_3 + \epsilon_3$	$X_3 = \theta_3 + \epsilon_3$
correlated prior	independent prior

As observations of signals accumulate, beliefs over $\tilde{\theta}_1, \tilde{\theta}_2, \tilde{\theta}_3$ tend to independence, returning Example 2.

Intuition for Results

- De-correlation of signals follows from a Bayesian version of the Central Limit Theorem.
  - not special to normality!
- At late periods we have a setting much like Example 2, and $n(t)$ evolves approximately sequentially.
- Two different conditions allow us to strengthen this to (eventual) *exact* optimality of myopic information acquisition.
  - Larger batch sizes: Demonstrate that for sufficiently large B, division vectors $n(Bt)$ are attainable using a sequential rule.
  - Quantify "typicality" of failures of sequentiality. Show that at late periods t, $n(t)$ generically evolves sequentially.

Accuracy vs. Correlation

Plausible intuition: since agents learn all states,
myopic strategy will be optimal.

Accuracy vs. Correlation

Plausible intuition: since agents learn all states,
myopic strategy will be optimal.

$\longrightarrow$ confused, depends on what we mean by "learn".

- As DM learns, beliefs simultaneously become more precise and less correlated, and these two effects are confounded in our main results.
- We show that the block size B needed in Theorem 1 depends on how many observations are required for $\tilde{\theta}_1, \ldots, \tilde{\theta}_K$ to "de-correlate".
- De-correlation is quicker when prior is:
  - less accurate
  - less correlated
- These also lead to the myopic rule becoming optimal sooner.

Summary

- We consider optimal dynamic information acquisition with normal signals that are flexibly correlated.
- Complementarity/substitution could generate intertemporal tradeoff.
- But we provide conditions under which these complementarities eventually vanish $\Longrightarrow$ myopic strategy becomes optimal.
- When signals acquired in large batches, optimality holds immediately.
- Optimality extends to endogenous sampling intensity and to a class of multi-agent games (see paper).

Thank You!