Selective Sampling with Information-Storage Constraints

Philippe Jehiel[1] Jakub Steiner[2]

[1]PSE and UCL

[2]CERGE-EI and Edinburgh U

Cambridge
May 2018

sequential sampling with information-aggregation friction

normative result about “reasonable doubt”

robustness to unobservable cognition constraints

match of behavioral stylized facts

Outline

task: $a \in A$, $\theta \in \Theta$, $u(a, \theta)$, prior $\pi \in \Delta(\Theta)$, $|A|, |\Theta| < \infty$

primitive experiments: $p \in \mathcal{P}$, $p(x \mid \theta) > 0$, $|X| < \infty$

termination strategy: $(\beta_x)_{x \in X} \neq (0, \ldots, 0)$

- nature draws θ from π, fixed throughout
- agent draws signal x from $p(x \mid \theta)$
- terminates with probability β_x
- restarts with probability $1 - \beta_x$ and draws x' from p, ...

action strategy: $\sigma : X \longrightarrow A$

- the agent a.s. eventually terminates and
- chooses action σ(last signal)

effective choice rule: $r(a \mid \theta; p, \beta, \sigma)$

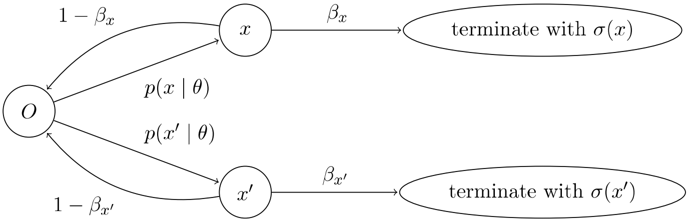
$1 - \beta_x$
x
β_x
terminate with $\sigma(x)$
O
$p(x \mid \theta)$
$p(x' \mid \theta)$
$1 - \beta_{x'}$
x'
$\beta_{x'}$
terminate with $\sigma(x')$

Repeated-cognition problem

$$\max_{\boldsymbol{p},\beta,\sigma} \sum_{\theta\in\Theta, a\in A} \pi_\theta u(a,\theta) r(a \mid \theta; \boldsymbol{p}, \beta, \sigma)$$

Non-trivial memory

the model accommodates agents with non-trivial memory

special case: Wilson (2014) with endogenous terminations

source of the flexibility: general set $\mathcal{P}$

Outline

$r(a \mid \theta)$—a generic stochastic choice rule

a_1, a_2—two conditionally independent action draws from $r(a \mid \theta)$: joint distribution of θ, a_1, and a_2 is

$$\alpha(\theta, a_1, a_2) = \pi_\theta r(a_1 \mid \theta) r(a_2 \mid \theta)$$

Definition

A choice rule $r(a \mid \theta)$ is second-thought-free for a prior π and utility u if for each action a chosen with positive probability:

$$\mathrm{E}_\alpha \left[u\left(a_1, \theta\right) \mid a_1 = a \right] \geq \mathrm{E}_\alpha \left[u\left(a_2, \theta\right) \mid a_1 = a \right].$$

$r(a \mid \theta)$—a generic stochastic choice rule

a_1, a_2—two conditionally independent action draws from $r(a \mid \theta)$: joint distribution of θ, a_1, and a_2 is

$$\alpha(\theta, a_1, a_2) = \pi_\theta r(a_1 \mid \theta) r(a_2 \mid \theta)$$

Definition

A choice rule $r(a \mid \theta)$ is second-thought-free if the agent prefers the choice of each recommended action to rerunning the rule.

Small simplification

Lemma

If a choice rule r is second-thought-free, then for each action a chosen with positive probability:

$$\mathrm{E}[u(a_1, \theta) \mid a_1 = a] \geq \mathrm{E}[u(a_2, \theta) \mid a_1 = a].$$

Small simplification

Lemma

If a choice rule r is second-thought-free, then for each action a chosen with positive probability:

$$\mathrm{E}[u(a_1,\theta) \mid a_1 = a] = \mathrm{E}[u(a_2,\theta) \mid a_1 = a].$$

Small simplification

Lemma

If a choice rule r is second-thought-free, then for each action a chosen with positive probability:

$$\mathrm{E}[u(a_1, \theta) \mid a_1 = a] = \mathrm{E}[u(a_2, \theta) \mid a_1 = a].$$

proof: Suppose not:

$$\mathrm{E}\left[\mathrm{E}\left[u(a_1, \theta) \mid a_1\right]\right] > \mathrm{E}\left[\mathrm{E}\left[u(a_2, \theta) \mid a_1\right]\right]$$

Small simplification

Lemma

If a choice rule r is second-thought-free, then for each action a chosen with positive probability:

$$\mathrm{E}[u(a_1, \theta) \mid a_1 = a] = \mathrm{E}[u(a_2, \theta) \mid a_1 = a].$$

proof: Suppose not:

$$\mathrm{E}\left[u(a_1, \theta)\right] > \mathrm{E}\left[u(a_2, \theta)\right]$$

Outline

Optimality condition

Proposition

If a choice rule solves the repeated-cognition problem, then it is second-thought-free.

optimal decision process:

- repeats experiment p^* according to β^*
- terminates with a in state θ with prob. $r(a \mid \theta; p^*, \beta^*, \sigma^*)$
- indifference between a and new run of the repetitions of p^*

Proof

primitive experiment $p(x \mid \theta)$: prob. of signal x in any run
effective experiment $s(x \mid \theta; p; \beta)$: prob. that terminal signal is x

high termination probability β_x inflates effective likelihood of the choice at the signal x:

lemma: $s(x \mid \theta; p, \beta) = \frac{\beta_x p(x|\theta)}{\sum_{x' \in X} \beta_{x'} p(x'|\theta)}$

follows from the stationarity

Proof

repeated-cognition problem is equivalent to

$$\max_{p,\beta,\sigma} \sum_{\theta\in\Theta, x\in X} \pi_\theta \frac{\beta_x p(x \mid \theta)}{\sum_{x'\in X} \beta_{x'} p\left(x' \mid \theta\right)} u(\sigma(x), \theta)$$

the second-thought-free condition follows from the foc w.r.t. to β

Example

$A = \Theta = \{0, 1\}$, $u(a, \theta) = 1_{a=\theta}$, uniform prior

$\mathcal{P} = \{p\}$, where

$$.9 = \mathsf{Pr}_p(\theta = 1 \mid x = 1) > \mathsf{Pr}_p(\theta = 0 \mid x = 0) = .6.$$

Example

$A = \Theta = \{0, 1\}$, $u(a, \theta) = 1_{a=\theta}$, uniform prior

$\mathcal{P} = \{p\}$, where

$$.9 = \mathsf{Pr}_p(\theta = 1 \mid x = 1) > \mathsf{Pr}_p(\theta = 0 \mid x = 0) = .6.$$

consider $(p, (1, 1), \sigma_I)$ that runs p once and chooses $a = x$
this process is not second-thought-free:

- assume the agent has observed the muddled signal $x = 0$
- the second thought will either
  - deliver the same signal $x = 0$ (inconsequential),
  - or deliver $x = 1$ (beneficial action switch)

$\Rightarrow$ hesitation upon $x = 0$ with some probability is beneficial

Example

optimal rule must satisfy symmetry:

$$r(1 \mid 1; p, \beta^*, \sigma_I) = r(0 \mid 0; p, \beta^*, \sigma_I)$$

$\Rightarrow \beta_0^*/\beta_1^* = 0.15$

the quickest such rule:

- terminates immediately upon $x = 1$
- hesitates with prob. $1 - .15$ upon $x = 0$
- $r(1 \mid 1; p, \beta^*, \sigma_I) = r(0 \mid 0; p, \beta^*, \sigma_I) = 0.79$

  the naive strategy pays $.66$ only

since generic stochastic rules are not second-thought-free, asymmetries in the termination strategy are typical

Example

optimal rule must satisfy symmetry:

$$\frac{\beta_1^* p(1 \mid 1)}{\beta_0^* p(0 \mid 1) + \beta_1^* p(1 \mid 1)} = \frac{\beta_0^* p(0 \mid 0)}{\beta_0^* p(0 \mid 0) + \beta_1^* p(1 \mid 0)}$$

$\Rightarrow \beta_0^* / \beta_1^* = 0.15$

the quickest such rule:

- terminates immediately upon $x = 1$
- hesitates with prob. $1 - .15$ upon $x = 0$
- $r(1 \mid 1; p, \beta^*, \sigma_I) = r(0 \mid 0; p, \beta^*, \sigma_I) = 0.79$

  the naive strategy pays $.66$ only

since generic stochastic rules are not second-thought-free, asymmetries in the termination strategy are typical

Robustness

set $\mathcal{P}$ is unobservable to outside observers

second-thought-free condition makes reference to

- effective choice rule $r^*(a \mid \theta)$
- prior π
- utility function u

no reference to the set $\mathcal{P}$ of the primitive cognition processes

the condition partially

- identifies preferences from joint distribution of θ and a
- predicts stochastic choice for given u

Imperfect recall

team-equilibrium interpretation (Piccione & Rubinstein 1997)

x-self:

- draws inferences about θ from:
  - x drawn from $p(x \mid \theta)$
  - from the fact that previous selves have not terminated
- maximizes the "team payoff"
- termination $\sim$ delegation to the next self

Outline

Binary Setting

$A = \Theta = \{0, 1\}$, $u(a, \theta) = u_\theta 1_{a=\theta} > 0$, $\pi \in \Delta(\Theta)$

$|X| < \infty$, $|\mathcal{P}| < \infty$

lemma: a solution exists in which $\beta_x > 0$ for at most 2 signals

$X = \{0, 1\}$, increasing likelihood ratio

lemma: there exists a solution in which $\sigma(x) = x$

Feasibility condition

for each primitive experiment p define $d_p = \frac{p(0|0)p(1|1)}{p(0|1)p(1|0)}$

the larger d_p is, the better p discriminates between the states

lemma: the effective choice rule preserves d_p:

$$\frac{r(0 \mid 0; p, \beta)r(1 \mid 1; p, \beta)}{r(0 \mid 1; p, \beta)r(1 \mid 0; p, \beta)} = d_p \text{ for every } \beta$$

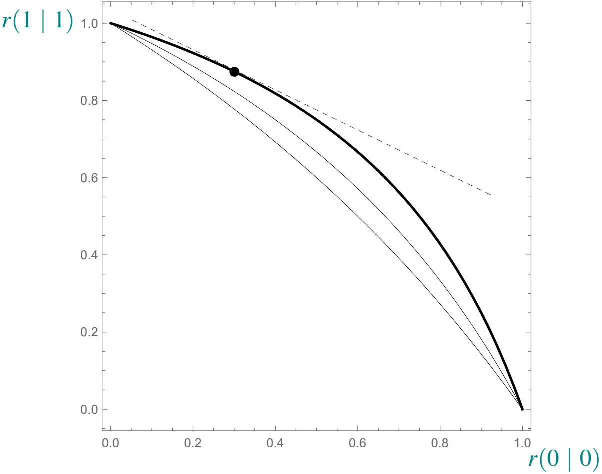

$r(1 \mid 1)$
1.0
0.8
0.6
0.4
0.2
0.0
0.0
0.2
0.4
0.6
0.8
1.0
$r(0 \mid 0)$

Solution

let $\overline{d} = \max_p d_p$, $\overline{p} = \arg\max_p d_p$

let $R = u_1\pi_1/u_0\pi_0$ measure the a priori attractiveness of action 1

Proposition

1. *when $R \geq \overline{d}$, then the agent always chooses 1,*
2. *when $R \leq 1/\overline{d}$, then the agent always chooses 0,*
3. *when $R \in \left(1/\overline{d}, \overline{d}\right)$, then*

$$r^*(1 \mid 1) = \frac{\overline{d}R - \sqrt{\overline{d}R}}{(\overline{d}-1)R}, \; r^*(0 \mid 0) = \frac{\overline{d} - \sqrt{\overline{d}R}}{\overline{d}-1},$$

$$\frac{\beta_1^*}{\beta_0^*} = \frac{\overline{d}R - \sqrt{\overline{d}R}}{\sqrt{\overline{d}R} - R}\frac{\overline{p}(0 \mid 1)}{\overline{p}(1 \mid 1)}.$$

Comparative statics

if $R \in \left(1/\overline{d}, \overline{d}\right)$, then:

1. relative bias β_1^*/β_0^* for signal 1 increases with R
2. $\frac{r^*(1|\theta)}{r^*(0|\theta)}$ increases with R
3. relative decision rate f_1/f_0 increases with R
4. posteriors $\mathrm{Pr}_{r^*}(\theta = 1 \mid a)$ increase with π_1

Outline

Confirmation bias

symmetric primitive experiment: $p(1 \mid 1) = p(0 \mid 0)$

symmetric incentives: $u_1 = u_0$

asymmetric prior: $\pi_1 > \pi_0$

Corollary

the termination strategy favors the a priori likely signal: $\beta_1^* > \beta_0^*$.

intuition: set $\beta_1 = \beta_0$ so that $r = p$

consider the surprising recommendation. repetition leads to

- the same recommendation (inconsequential), or to
- action switch (beneficial)

Speed-accuracy complementarity

the setting from the previous slide

Corollary

$Pr_{\pi,r^*}(a = \theta \mid t)$ decreases with the response time t.

intuition:
delay indicates that the agent has repeatedly received the surprising recommendation

$\Rightarrow$ delay indicates the surprising state $\theta = 0$

correct choice is less likely in the surprising state

Experimental data

Ratcliff and McKoon (2008)

binary choice, visual recognition tasks, varied priors

1. surprising responses are delayed and less accurate
2. posterior distributions monotone w.r. to the prior

- Fundenberg et al (2015) get 1. if uncertainty about stakes
- Wald (1945) and RI models get 2. wrong

Outline

Wilson (2014) with endogenous termination

primitive statistical experiment $\mu(x \mid \theta)$

set M of mental states m, the process starts at m_0

at the end of each run of μ, the agent can terminate or transition into a mental state m and continue

mixing conditional on (x, m)

action choice $a = \sigma(x, m)$

let $\mathcal{P}$ be the set of all constructible rules in this model

Solution

let $\mathcal{R}(\mathcal{P})$ be the set of rules constructed as follow:

- the agent chooses $p \in \mathcal{P}$
- repeats $p(a \mid \theta)$ according to the termination strategy $(\beta_a)_a$
- chooses $\sigma(a)$ once she terminates with a

lemma: this embedding creates no new rules: $\mathcal{R}(\mathcal{P}) = \mathcal{P}$

$\Rightarrow$ the optimal rule $p^* \in \mathcal{P}$ is second-thought-free

$\Rightarrow$ in the binary setting, $p^* \in \mathcal{P}$ satisfies our analytical solution

Exponential discounting

confirmation bias with impatient agent

$\pi_1/\pi_0 = 10$, $p(1 \mid 1) = p(0 \mid 0) = .9$

the unsurprising signal leads to immediate termination: $\beta_1^* = 1$

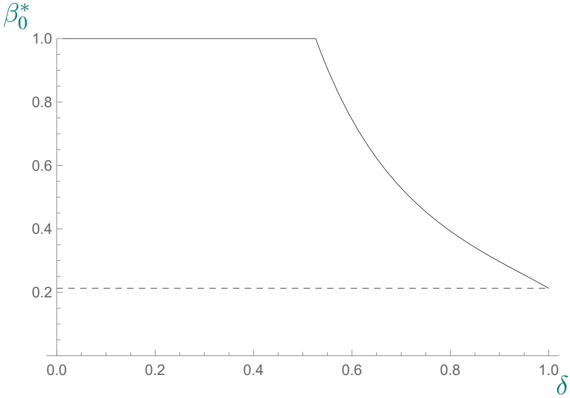

Summary

a model of the reasonable doubt

productive hesitation:

selective repetitions of a decision process affect correlation between the choice and state

the optimal effective decision process is second-thought-free: the condition doesn't refer to the unobservable P

confirmation bias, speed-accuracy complementarity, ...