Rational Inattention Dynamics:

inertia and delay in decision-making

Jakub Steiner[1] Colin Stewart[2] Filip Matějka[3]

[1]Cerge-ei and Edinburgh

[2]Toronto

[3]Cerge-ei

Paris
May 2016

Example

States $\theta_t \in \{\text{good, bad}\}$, positively serially correlated

Actions $a_t \in \{\text{invest, not invest}\}$

DM can

1. relay on the past decision
2. pay a cost to acquire information about the current state

Optimal combination of 1. and 2.:

- endogenous default rule: repeat the last action
- stochastic choice rule with penalties for violating the default rule
- $\Rightarrow$ virtual switching cost

When is the model useful?

Incentives vs. information frictions
Actions can be sluggish because of

- the physical switching costs, or
- reliance on the serial correlations

Extrapolation:
Will actions be serially correlated in an i.i.d. setting?

- material frictions $\Rightarrow$ yes
- information frictions $\Rightarrow$ no

Overview

Dynamic decision-making with costly information acquisition

- general decision processes
- general signal structures
- bayesian, forward-looking DM
- particular cost of information

Intertemporal link

- currently acquired information will be useful in the future

Solution

- dynamic logit
- with a bias towards an endogenous default rule

Literature

RI applications are inherently dynamic:
Sims (2003), Mackowiak & Wiederholt (2009), Tutino (2013), Ravid (2014)

But, existing general solutions are static:
Matějka & McKay (2015), Caplin & Dean (2013)

Our dynamic extension exploits connection to a control problem of Mattsson & Weibull (2002)—predecessor of RI.

Structural estimation: Rust (1987)

Outline

Model

Discrete time $t = 1, 2, \ldots$

States $\theta_t \in \Theta$ (finite) follow a stochastic process

Actions $a_t \in A$ (finite)

Histories $\theta^t = (\theta_1, \ldots, \theta_t)$, $a^t = (a_1, \ldots, a_t)$

Payoffs (excluding information costs): $\sum_t \delta^{(t)} u_t \left(a^t, \theta^t\right)$

(infinite and finite horizons are special cases)

Information acquisition

Signal space X with $|X| \geq |A|$

For each t, DM chooses:

1. signal distributions $f_t(x_t \mid \theta^t, x^{t-1}) \in \Delta(X)$
2. action strategy $s_t : X^t \longrightarrow A$

Let $f = (f_t)_t$, $s = (s_t)_t$

Key assumption: cost of information is proportional to expected reduction in entropy of beliefs

Does DM observe flow payoffs? For free?

Mutual information

Surprisal $-\log p(\theta)$

Entropy—expected surprise (measure of uncertainty):

$$H(\theta) = -\sum_{\theta} p(\theta) \log p(\theta)$$

Mutual information—measure of reduction of uncertainty:

$$I(\theta; X) = H(\theta) - E_X \left[H(\theta \mid X) \right]$$

Captures how much, on average, observation of X reduces uncertainty about θ

Mutual information

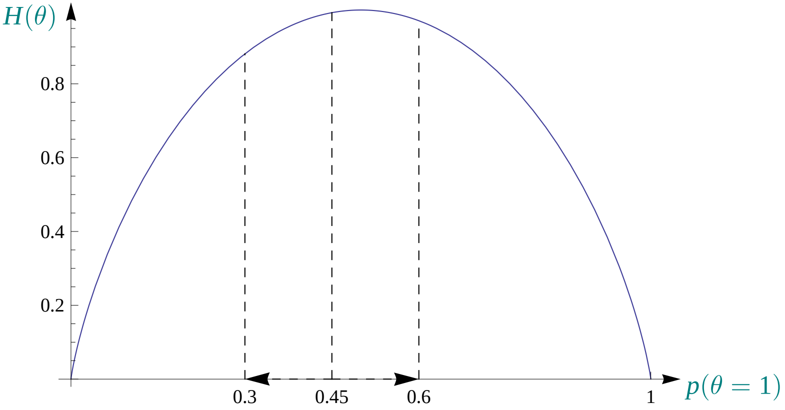

The problem

Definition (Dynamic rational-inattention problem)

$$\max_{f,s} E\left[\sum_t \delta^{(t)}\left(u_t\left(s^t(x^t),\theta^t\right) - I\left(\theta^t; x_t \mid x^{t-1}\right)\right)\right]$$

Why entropy? Tractability.

With free information

DM receives free signal y_t at the end of each period

- $y_t \sim g_t(y_t \mid \theta^t, a^t, y^{t-1}) \in \Delta(Y)$
- e.g. y_t is a noisy signal of the stage payoffs
- $f(x_t \mid \theta^t, x^{t-1}, y^{t-1})$, $a_t = s_t(x^t, y^{t-1})$.
- assumption: $y_t \perp x^t \mid (\theta^t, y^{t-1}, a^t)$.

Definition (Dynamic rational-inattention problem)

$$\max_{f,s} E\left[\sum_t \delta^{(t)}\left(u_t\left(s^t(x^t, y^{t-1}), \theta^t\right) - I\left(\theta^t; x_t \mid x^{t-1}, y^{t-1}\right)\right)\right]$$

Example 1: "sunk cost" behavior

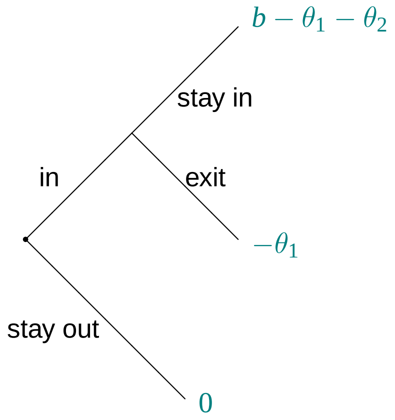

Example 2: tracking problem

$\theta_t \in \{0,1\}$ follows a Markov chain

$a_t \in \{0,1\}$

$$u_t(a_t, \theta_t) = \begin{cases} u_\theta & \text{if } a_t = \theta_t \\ 0 & \text{if } a_t \neq \theta_t \end{cases}$$

Infinite horizon, exponential discounting

We compare settings without and with free information:
free signals can hamper choice precision (but improve welfare)

Example 3: response times

$\theta \in \{0, 1\}$, fixed over time

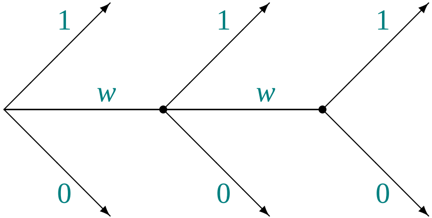

$$u_t(\theta, a_t) = \begin{cases} 0 & \text{if } a_t = w \\ 1 - ct & \text{if } a_t = \theta \\ -ct & \text{if } a_t \neq \theta \end{cases}$$

Finite horizon, no discounting

Simplifying the problem

Lemma

There exists a solution with a one-to-one mapping between signal realizations and actions

Idea: any information that is not used now can be delayed

Ravid (2014) calls such strategies recommendation strategies

Stochastic choice rule: $p_t(a_t \mid \theta^t, a^{t-1})$

Definition (Reduced problem)

$$\max_p E\left[\sum_t \delta^{(t)}\left(u_t(a^t, \theta^t) - I(\theta^t; a_t \mid a^{t-1})\right)\right]$$

Outline

Dynamic logit

Matějka & McKay (2015):
Static RI problem $\longrightarrow$ static logit

Standard dynamic programming would lead to stage payoffs

$$u_t\left(a^t, \theta^t\right) + \delta E\left[V_{t+1}\left(\text{belief}_t\right)\right]$$

This paper:
Recursive characterization with stage payoffs

$$u_t\left(a^t, \theta^t\right) + \delta E\left[V_{t+1}\left(a^t, \theta^{t+1}\right)\right]$$

Definition (Rust, 1987)

The dynamic logit rule under payoffs $u_t(a^t, \theta^t)$ is

$$p_t(a_t \mid \theta^t, a^{t-1}) = \frac{e^{u_t(a^t,\theta^t)+\delta E[V_{t+1}(a^t,\theta^{t+1})|\theta^t]}}{\sum_{\hat{a}_t} e^{u_t(\hat{a}^t,\theta^t)+\delta E[V_{t+1}(\hat{a}^t,\theta^{t+1})|\theta^t]}}$$

where

$$V_t\left(a^{t-1}, \theta^t\right) = \log\left(\sum_{\hat{a}_t} e^{u_t(\hat{a}^t,\theta^t)+\delta_t E\left[V_{t+1}\left(\hat{a}^t,\theta^{t+1}\right)|\theta^t\right]}\right)$$

Default rule

Generalization of the "conservative heuristic" from the example

Definition

Default rule q: system of conditional probabilities $q_t(a_t \mid a^{t-1})$

We call $q_t(a_t \mid a^{t-1})$ a predisposition to choose a_t at a^{t-1}

Solution

Proposition

The optimal choice rule $p_t(a_t \mid a^{t-1}, \theta^t)$ *in the dynamic RI problem with payoffs* u_t *is the dynamic logit rule under payoffs*

$$\tilde{u}_t(a^t, \theta^t) = u_t(a^t, \theta^t) + \log q_t(a_t \mid a^{t-1}),$$

where the default rule q *satisfies*

$$q_t\left(a_t \mid a^{t-1}\right) = p_t\left(a_t \mid a^{t-1}\right).$$

$\log q_t(a_t \mid a^{t-1})$ term: actions that are rarely chosen tend to be associated with high information cost

Estimation

Structural estimation of preferences:

- econometrician observes sequence $(a_t, \theta_t)_t$
- estimates utilities u_t using the dynamic logit model

If our model is correct, the econometrician identifies:

$$\tilde{u}_t(a^t, \theta^t) = u_t\left(a^t, \theta^t\right) + \log q_t\left(a_t \mid a^{t-1}\right)$$

The blue term may resemble switching costs

Makes no difference for fit within the same environment, but matters for extrapolation

Extrapolation

Rust (1987):

- Bus-engine replacement
- Estimated replacement cost: engine price $+$ other costs
- Extrapolation exercise: demand for engines as a function of price

Our model:

- The log-predisposition term is part of the estimated “cost”
- Predispositions change depending on the price
- If the price goes up:
  - for a fixed belief, DM replaces less often (in Rust)
  - DM checks less often whether replacement is needed (not in Rust)
  - $\Rightarrow$ Rust underestimates price elasticity

With free information

Proposition

The optimal choice rule $p_t(a_t \mid a^{t-1}, \theta^t, y^{t-1})$ *in the dynamic RI problem with payoffs* u_t *is the dynamic logit rule under payoffs*

$$\tilde{u}_t\left(a^t, \theta^t, y^{t-1}\right) = u_t\left(a^t, \theta^t\right) + \log q_t\left(a_t \mid a^{t-1}, y^{t-1}\right),$$

where the default rule q *satisfies*

$$q_t\left(a_t \mid a^{t-1}, y^{t-1}\right) = p_t\left(a_t \mid a^{t-1}, y^{t-1}\right).$$

Outline

Definition (Control problem)

$$\max_p E\left[\sum_t \delta^{(t)}\left(u_t(a^t,\theta^t) - \log\frac{p_t(a_t \mid \theta^t, a^{t-1})}{q_t(a_t \mid a^{t-1})}\right)\right],$$

for an exogenous default rule q.

Control problem is simpler than RI problem
(states are observable)

Mattsson and Weibull (2002) study the static version

Reformulation

Lemma

The optimal dynamic RI choice rule p solves

$$\max_{q,p} E\left[\sum_t \delta^{(t)}\left(u_t(a^t,\theta^t) - \log\frac{p_t(a_t \mid \theta^t, a^{t-1})}{q_t(a_t \mid a^{t-1})}\right)\right]$$

(together with some q)

Control problem

deviation from q
control cost

RI problem

information acquisition
information cost

Sketch of proof

The lemma relies on two properties of entropy:

1. Symmetry of mutual information:

$$I(X;Y) = I(Y;X)$$

2. Properness of the logarithmic scoring rule:

$$\begin{aligned} H(X) &\equiv E_p[-\log p(x)] \\ &= \min_{q(x)} E_p[-\log q(x)] \end{aligned}$$

Sketch of proof

$$
\begin{aligned}
&E\left[\sum_t \delta^{(t)}\left(u(a^t,\theta^t) - I(\theta^t; a_t \mid a^{t-1})\right)\right] \\
&= E\left[\sum_t \delta^{(t)}\left(u(a^t,\theta^t) - I(a_t; \theta^t \mid a^{t-1})\right)\right] \\
&= E\left[\sum_t \delta^{(t)}\left(u_t(a^t,\theta^t) - \log \frac{p_t(a_t \mid \theta^t, a^{t-1})}{p_t(a_t \mid a^{t-1})}\right)\right]
\end{aligned}
$$

By property 2, p maximizes the above iff it solves

$$
\max_{p,q} E\left[\sum_t \delta^{(t)}\left(u_t(a^t,\theta^t) - \log \frac{p_t(a_t \mid \theta^t, a^{t-1})}{q_t(a_t \mid a^{t-1})}\right)\right]
$$

Consequences

Simple dynamic programming

- when optimizing p at each a^{t-1} given q, we do not have to consider how changes in posterior beliefs affect continuation values since states are observable in the control problem
- when optimizing q we do not have to consider posteriors because these are determined by p

Outline

Setting

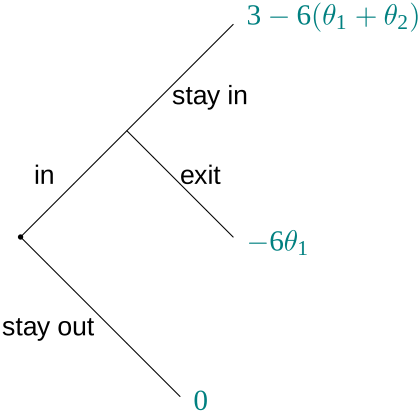

$\theta_t \in \{0, 1\}$

θ_1 uniform, $\Pr(\theta_2 = \theta_1) = \gamma$

Correlated costs ($\gamma = .97$)

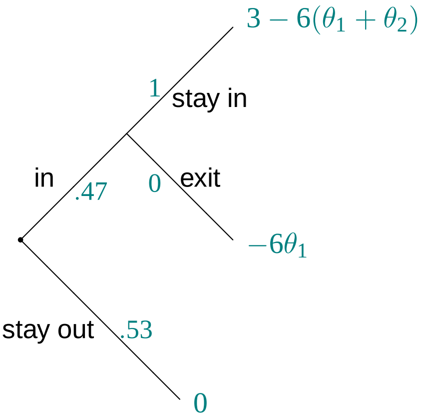

Uncorrelated costs ($\gamma = 1/2$)

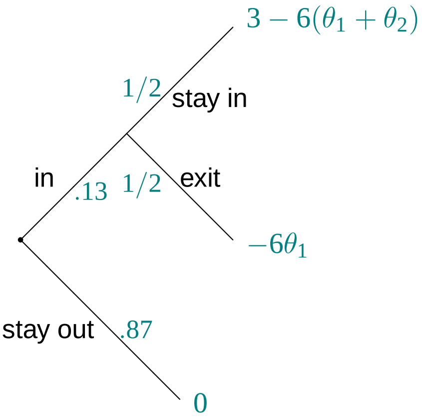

Outline

Setting

$\theta_t \in \{0, 1\}$

Markov chain with transition probabilities $\gamma(\theta, \theta')$

$a_t \in \{0, 1\}$

$$u_t(a_t, \theta_t) = \begin{cases} u_\theta & \text{if } a_t = \theta_t \\ 0 & \text{if } a_t \neq \theta_t \end{cases}$$

Exponential discounting

Benchmark

Exogenous signals: conditionally independent x_t

$\Rightarrow$ Belief at t depends on the whole signal history x^t, which in turn depends on the whole history θ^t

Benchmark

Exogenous signals: conditionally independent x_t

$\Rightarrow$ Belief at t depends on the whole signal history x^t, which in turn depends on the whole history θ^t

Dynamic RI:
Stationary solution: Markov chain with states (θ_t, a_t)

Choice rule $p_t(a_t \mid \theta_t, a_{t-1})$ depends only on current state and last action

Optimal signals are not conditionally independent

Comparative statics

States have positive persistence if $\gamma(0,0) + \gamma(1,1) > 1$

Proposition

Suppose states have positive persistence and the solution is eventually interior. Then

1. *actions have positive persistence, and moreover, the choice rule satisfies* $\hat{p}(a_t \mid \theta_t, a_{t-1}) > \hat{p}(a_t \mid \theta_t, a'_{t-1})$ *whenever* $a_t = a_{t-1} \neq a'_{t-1}$*; and*
2. *the posterior probability* $\pi^p(\theta_t = a \mid a_t = a')$ *is nonincreasing in the payoff* u_a *for all* $a, a' \in \{0, 1\}$.

Free information

What if the agent also receives costless signals?

In a static RI problem

- no change of behavior
- marginal condition on signal precision unaffected

Free information

What if the agent also receives costless signals?

In a static RI problem

- no change of behavior
- marginal condition on signal precision unaffected

In a dynamic problem

- costless signal can more than crowd out information acquisition $\Rightarrow$ behavior becomes "worse"
- receiving free signals in the future reduces incentive to acquire information today
- agent's payoff nonetheless increases

Outline

Setting

$t = 1, \ldots, T$

$\theta \in \{0, 1\}$ uniform, fixed over time

$a_t \in \{0, 1, w\}$

$$u_t(a^t, \theta) = \begin{cases} 1 & \text{if } a^t = (w^{t-1}, \theta) \\ 0 & \text{if } a^t = (w^{t-1}, 1-\theta) \\ -c & \text{if } a^t = w^t \\ 0 & \text{otherwise} \end{cases}$$

No discounting

Information constraint: $I\left(\theta; x_t \mid x^{t-1}\right) \leq \kappa$ for all x^{t-1}

- κ not too large

Solution

Can replace the constraints with a weaker constraint corresponding to capacity being storable

Dual to this relaxed problem fits in our framework

- information cost in each period is weighted by the shadow price of the constraint for that period

Solution

Can replace the constraints with a weaker constraint corresponding to capacity being storable

Dual to this relaxed problem fits in our framework

- information cost in each period is weighted by the shadow price of the constraint for that period

Solution:

- terminal decisions are made gradually, with constant hazard rate in periods $1, \ldots, T-1$
- accuracy of decisions is constant over time
- speed of decision-making is non-monotone in capacity: fast when capacity is high or low

Conclusion

General dynamic rational inattention problem

Solution combines a default rule and adjustment

Behavior resembles dynamic logit, but the model suggests caution when extrapolating across problems

Reduction to a collection of static RI problems