On Memory Management

preliminary and incomplete

Philippe Jehiel[1] Jakub Steiner[2]

[1]PSE and UCL

[2]Cerge-ei and Edinburgh U

Bocconi
June 2017

people face cognition/information-processing constraints

choice data reflects both (i) preferences, and (ii) the constraints

we know little about the cognitive constraints

behavioral predictions robust to details of the constraints?

(reinterpretation: optimal team/firm decision-making)

Model overview

bounded operational memory: agent comprehends only a limited number of facts

the memory bound $\Rightarrow$ targeted information search

predictions about the stochastic choice rule robust to the details of the cognitive technology

1. reduced-form model of the targeted information search
2. explicit model

Reduced-form example

task: $A = \Theta = \{0, 1\}$, prior π_θ

$u(a, \theta) = u_\theta > 0$ if $a = \theta$, and 0 otherwise

Reduced-form example

task: $A = \Theta = \{0, 1\}$, prior π_θ

$u(a, \theta) = u_\theta > 0$ if $a = \theta$, and 0 otherwise

primitive decision process generates a signal $p(x \mid \theta)$,

$x \in \{0, 1\}$, $p(\theta \mid \theta) > 1/2$, fix strategy $a(x) = x$

Reduced-form example

task: $A = \Theta = \{0, 1\}$, prior π_θ

$u(a, \theta) = u_\theta > 0$ if $a = \theta$, and 0 otherwise

primitive decision process generates a signal $p(x \mid \theta)$,

$x \in \{0, 1\}$, $p(\theta \mid \theta) > 1/2$, fix strategy $a(x) = x$

targeted information search: the agent chooses search intensities $\beta_x > 0$ for each $x \in \{0, 1\}$

effective signal distribution: $r(x \mid \theta; \beta) = \frac{\beta_x p(x|\theta)}{\sum_{x'} \beta_{x'} p(x'|\theta)}$

Reduced-form example

task: $A = \Theta = \{0, 1\}$, prior π_θ

$u(a, \theta) = u_\theta > 0$ if $a = \theta$, and 0 otherwise

primitive decision process generates a signal $p(x \mid \theta)$,

$x \in \{0, 1\}$, $p(\theta \mid \theta) > 1/2$, fix strategy $a(x) = x$

targeted information search: the agent chooses search intensities $\beta_x > 0$ for each $x \in \{0, 1\}$

effective signal distribution: $r(x \mid \theta; \beta) = \frac{\beta_x p(x|\theta)}{\sum_{x'} \beta_{x'} p(x'|\theta)}$

Reduced-form example

task: $A = \Theta = \{0, 1\}$, prior π_θ

$u(a, \theta) = u_\theta > 0$ if $a = \theta$, and 0 otherwise

primitive decision process generates a signal $p(x \mid \theta)$,

$x \in \{0, 1\}$, $p(\theta \mid \theta) > 1/2$, fix strategy $a(x) = x$

targeted information search: the agent chooses search intensities $\beta_x > 0$ for each $x \in \{0, 1\}$

effective choice rule: $r(a \mid \theta; \beta) = \frac{\beta_a p(a|\theta)}{\sum_{a'} \beta_{a'} p(a'|\theta)}$

Reduced-form example

task: $A = \Theta = \{0, 1\}$, prior π_θ

$u(a, \theta) = u_\theta > 0$ if $a = \theta$, and 0 otherwise

primitive decision process generates a signal $p(x \mid \theta)$,

$x \in \{0, 1\}$, $p(\theta \mid \theta) > 1/2$, fix strategy $a(x) = x$

targeted information search: the agent chooses search intensities $\beta_x > 0$ for each $x \in \{0, 1\}$

effective choice rule: $r(a \mid \theta; \beta) = \frac{\beta_a p(a|\theta)}{\sum_{a'} \beta_{a'} p(a'|\theta)}$

the problem:

$$\max_\beta \sum_{\theta,a} \pi_\theta u(a, \theta) r(a \mid \theta; \beta)$$

Solution

optimal effective choice rule

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

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

Proposition

Optimal choice rule $r\left(a \mid \theta; \beta^*\right)$ *satisfies*

$$E_\alpha[u(a_1, \theta) \mid a_1 = a] = E_\alpha[u(a_2, \theta) \mid a_1 = a],$$

for every a *chosen with positive probability.*

Interpretation of the optimality condition

example of a second thought

e.g:

- uniform reward $u_\theta = 1$, uniform prior,
- $.9 = \Pr_p(\theta = 1 \mid x = 1) > \Pr_p(\theta = 0 \mid x = 0) = .8$

consider $\beta = (1, 1)$, so that $r = p$

then $E[u(a_1, \theta) \mid a_1 = 0] < E[u(a_2, \theta) \mid a_1 = 0]$ because:

- if $a_2 = a_1$, then the "second thought" is inconsequential
- if $1 = a_2 \neq a_1 = 0$, then $\theta = 1$ is more likely

$\Rightarrow \beta = (1, 1)$ is not optimal

$\beta_0 \searrow \Rightarrow \nearrow$ informativeness of observing $a = 0$

Second-thought-free condition

Definition

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

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

Second-thought-free condition

Lemma

If a choice rule $r(a \mid \theta)$ is second-thought-free, then for each action a chosen with positive probability:

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

Second-thought-free condition

Lemma

If a choice rule $r(a \mid \theta)$ is second-thought-free, then for each action a chosen with positive probability:

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

proof: Suppose not:

$$E\left[u(a_1, \theta) \mid a_1\right] \geq E\left[u(a_2, \theta) \mid a_1\right]$$

Second-thought-free condition

Lemma

If a choice rule $r(a \mid \theta)$ is second-thought-free, then for each action a chosen with positive probability:

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

proof: Suppose not:

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

Second-thought-free condition

Lemma

If a choice rule $r(a \mid \theta)$ is second-thought-free, then for each action a chosen with positive probability:

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

proof: Suppose not:

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

Solution
reformulation

Proposition

Optimal choice rule is second-thought-free.

this is equivalent to $E_\alpha[u(a_1, \theta) \mid a_1, a_2] = E_\alpha[u(a_2, \theta) \mid a_1, a_2]$

both recommendations are “equally convincing”, regardless of the primitive process $p(a \mid \theta)$

utility identification: e.g. suppose that an outsider observes

- uniform prior, and
- equal shares of both error types, $r(1|0) = r(0|1)$

$\Rightarrow$ losses from both types of error are equal

Literature

error-management (interdisciplinary): Johnson et al. (2013)

rational inattention: Sims (2003), Matejka&McKay (2015)
entropy of beliefs cannot be reduced by more than κ

- identification of preferences from choice data

$$p(a|\theta) = \frac{\exp\left(u(a,\theta) + \log p(a)\right)}{\sum_{a'} \exp\left(u(a',\theta) + \log p(a')\right)}$$

- source of identification: the assumed constraint

unrealistic predictions:

- RI behavior does not reflect perceptual distance (Caplin, Dean, Morris, Woodford,...)
- decision times of the RI agent are uninformative of choice and the state

Literature

microfoundations of cognitive constraints:

Hebert & Woodford (2016), and Morris & Strack build on the drift-diffusion model of Ratcliff (1978)

Zhong (2017): discontinuous belief evolution

bounded memory:

Piccione & Rubinstein (1997), Wilson (2014)

Outline

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

primitive decision processes: $p(a, t, x \mid \theta)$
a—action recommendation
t—time elapsed
x—side information

set of the feasible primitive processes P

repetitions:
termination strategy $(\beta_a)_a \in B = [0, 1]^{|A|} \setminus \{(0, \ldots, 0)\}$

- agent gets action recommendation a from $p(a, t, x \mid \theta)$
- terminates with probability β_a
- restarts with probability $1 - \beta_a$ and draws a', ...

effective choice rule: $r(a \mid \theta; p, \beta)$—probability that the agent terminates with action a

Effective choice rule

technical definition

$$\Pr\left((a_l, t_l, x_l)_{l=1}^{k} | \theta; p, \beta\right) = \beta_{a_k} p(a_k, t_k, x_k | \theta) \prod_{l=1}^{k-1} (1 - \beta_{a_l})\, p\left(a_l, t_l, x_l | \theta\right)$$

$$r(a \mid \theta; p, \beta) = \sum_{k=1}^{\infty} \sum_{(a_l, t_l, x_l)_{l=1}^{k} : a_k = a} \Pr\left((a_l, t_l, x_l)_{l=1}^{k} \mid \theta; p, \beta\right)$$

Memory-management problem

the agent chooses the primitive process p and the termination probabilities β to maximize her expected payoff:

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

Outline

Effective choice rule characterization

selective repetitions allow for targeted information search

Lemma

The agent terminates with action a in state θ with probability

$$r(a \mid \theta; \beta) = \frac{\beta_a p(a \mid \theta)}{\sum_{a' \in A} \beta_{a'} p(a' \mid \theta)}$$

proof:

$$r(a \mid \theta) = \beta_a p(a \mid \theta) + \left(\sum_{a' \in A} p(a' \mid \theta) \left(1 - \beta_{a'}\right) \right) r(a \mid \theta),$$

and rearrange

The main result

Proposition

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

proof: the first-order condition

Team-equilibrium reinterpretation

Piccione & Rubinstein (1997)

each self:

- takes strategies of the other selves as given
- draws inferences about θ from:
  - a drawn from $p(a \mid \theta)$
  - from the fact that previous selves have not terminated
- maximizes the "team payoff"
- is indifferent btwn terminating and passing to the next self

Error management

consider almost precise information:
$r(a \mid \theta) = o(\varepsilon)$ for all a that is not optimal at θ

Proposition (loss equalization)

Optimal choice rule r satisfies for each a

$$\sum_{\theta \in \Theta^a, a' \neq a} \pi_\theta r(a'|\theta)\big(u(a,\theta) - u(a',\theta)\big) = \sum_{\theta : \in \Theta_a} \pi_\theta r(a|\theta)\big(u(a^*_\theta), \theta) - u(a,\theta)\big),$$

up to $o(\varepsilon^2)$.

for each a, the losses from the two types of the error are equal

Error management

consider almost precise information:
$r(a \mid \theta) = o(\varepsilon)$ for all a that is not optimal at θ

Proposition (loss equalization)

the loss from failing to choose a when a is optimal

$$=$$

the loss from choosing a when a is suboptimal

for each a, the losses from the two types of the error are equal

Preference identification

an outsider observes:

- choice rule $r(a \mid \theta)$
- prior π_θ

conditions linear in u:

$$E[u(a_1, \theta) \mid a_1] = E[u(a_2, \theta) \mid a_1] \text{ for all } a_1 \in A$$

$|A| - 1$ independent conditions

$|\Theta| = 2$: full utility identification up to affine transformations

$|\Theta| > 2$: partial identification

Preference identification

an outsider observes:

- choice rule $r(a \mid \theta)$
- prior π_θ

conditions linear in u:

$$\sum_\theta \pi_\theta u(a_1,\theta) r(a_1 \mid \theta) = \sum_{\theta,a_2} \pi_\theta u(a_2,\theta) r(a_1 \mid \theta) r(a_2 \mid \theta) \text{ for all } a_1 \in A$$

$|A| - 1$ independent conditions

$|\Theta| = 2$: full utility identification up to affine transformations

$|\Theta| > 2$: partial identification

Outline

Generality

The model accommodates agents who

- condition on side information
- condition on time elapsed
- remember some information from previous runs
- choose how much to forget
- ...

Side information

$\mu(x \mid \theta)$—the primitive process

agent conditions terminations on x

action recommendations $a = \sigma(x)$

new model:

$(\gamma_x)_x$—termination probabilities

effective choice rule $p(a \mid \theta; \gamma, \sigma) = \sum_{x:\sigma(x)=a} \frac{\gamma_x \mu(x|\theta)}{\sum_{x'} \gamma_{x'} \mu(x'|\theta)}$

P—set of all feasible choice rules in the new model

embed the new model into the old model:

agent chooses any $p \in P$ and $(\beta_a)_a$

this results in effective choice rule $r(a \mid \theta; p, \beta)$

the set of all feasible $r(a \mid \theta; p, \beta)$ coincides with P:
$r(p(\gamma, \sigma), \beta) = p(\gamma', \sigma)$ where $\gamma'_x = \gamma_x \beta_{\sigma(x)}$

Time elapsed

$\mu(x \mid \theta)$—the primitive process, each run lasts 1 round
agent conditions terminations on (x, t)
action recommendations $a = \sigma(x, t)$

new model:
$(\gamma_{x,t})_{x,t}$—termination probabilities
effective choice rule $p(a \mid \theta; \gamma, \sigma)$
P—set of all feasible choice rules in the new model

embed the new model into the old model:
agent chooses any $p \in P$ and $(\beta_a)_a$
this results in effective choice rule $r(a \mid \theta; p, \beta)$

the embedding expands the set of the feasible rules:
it allows the agent to “restart” the clock

Information from previous runs

agent carries on some information from the previous runs

new model:

primitive process $\mu(a \mid \theta)$

set M of mental states m

the agent starts at m_0

at the end of each run of the primitive process, the agent can terminate or transition into mental state m and continue mixing, conditioning on (a, m)

embed the new model into the old model:

embedding does not expand the set of feasible choice rules: the restart is equivalent to the transition into m_0

Opportunity cost

linear opportunity cost of time $c \times t$

the memory management problem becomes

$$\max_{\beta,p} E_{r(\beta,p)}[u(a,\theta) - ct]$$

the second-thought-free condition:

$$E[u(a_1,\theta) \mid a_1] = E[u(a_2,\theta) - ct \mid a_1] \text{ if } \beta_{a_1} \in (0,1)$$

the problem simplifies if each run of p lasts one period

then $t \mid \theta$ is geometrically distributed with rate $\sum_a \beta_a p(a \mid \theta)$

Opportunity cost

linear opportunity cost of time $c \times t$

the memory management problem becomes

$$\max_{\beta,p} E_{r(\beta,p)}[u(a,\theta) - ct]$$

the second-thought-free condition:

$$E[u(a_1,\theta) \mid a_1] \geq E[u(a_2,\theta) - ct \mid a_1] \text{ if } \beta_{a_1} = 1$$

the problem simplifies if each run of p lasts one period

then $t \mid \theta$ is geometrically distributed with rate $\sum_a \beta_a p(a \mid \theta)$

Opportunity cost

linear opportunity cost of time $c \times t$

the memory management problem becomes

$$\max_{\beta,p} E_{r(\beta,p)}[u(a,\theta) - ct]$$

the second-thought-free condition:

$$E[u(a_1,\theta) \mid a_1] = E\left[u(a_2,\theta) - \frac{c}{\sum_a \beta_a p(a \mid \theta)} \mid a_1\right] \text{ if } \beta_{a_1} \in (0,1)$$

the problem simplifies if each run of p lasts one period

then $t \mid \theta$ is geometrically distributed with rate $\sum_a \beta_a p(a \mid \theta)$

Outline

Binary setting

$A = \Theta = \{0, 1\}$, arbitrary prior

$u(a, \theta) = u_\theta > 0$ if $a = \theta$ and 0 otherwise

$p(a \mid \theta)$, increasing likelihood ratio: $\frac{p(1|0)}{p(0|0)} < \frac{p(1|1)}{p(0|1)}$

each run of p lasts one round

comparative statics of β_1/β_0?

Three interpretations of β_1/β_0

1. relative search intensities
2. relative termination probabilities
3. relative decision speeds:
   - $t|\theta$ is geometrically distributed with rate $E[\beta_a \mid \theta]$
   - $E[t \mid \theta] = \frac{1}{E[\beta_a|\theta]}$
   - $\frac{E[t|1]}{E[t|0]}$ decreases in β_1/β_0

Comparative statics

analytical solution

the solution exhibits confirmation bias: $\frac{\beta_1}{\beta_0} \nearrow$ in $\frac{\pi_1 u_1}{\pi_0 u_0}$

an increase of the prior probability of θ increases:

- search intensity for the evidence favoring θ
- decision speed in the state θ

Comparative statics

analytical solution

the solution exhibits confirmation bias: $\frac{\beta_1}{\beta_0} \nearrow$ in $\frac{\pi_1 u_1}{\pi_0 u_0}$

an increase of the incentive to choose θ increases:

- search intensity for the evidence favoring θ
- decision speed in the state θ

Solving the binary setting

feasibility: $\frac{r(1|1)r(0|0)}{r(0|1)r(1|0)} = \frac{\frac{\beta_1 p(1|1)}{\sum_a \beta_a p(a|1)} \frac{\beta_0 p(0|0)}{\sum_a \beta_a p(a|0)}}{\frac{\beta_0 p(0|1)}{\sum_a \beta_a p(a|1)} \frac{\beta_1 p(1|0)}{\sum_a \beta_a p(a|0)}}$

optimality: $E_\theta[u(a_1, \theta) \mid a_1] = E_{a_2,\theta}[u(a_2, \theta) \mid a_1]$

two quadratic conditions & two normalization conditions for r

$\Rightarrow$ analytical solution

Solving the binary setting

feasibility: $\frac{r(1|1)r(0|0)}{r(0|1)r(1|0)} = \frac{p(1|1)p(0|0)}{p(0|1)p(1|0)}$

optimality: $E_\theta[u(a_1, \theta) \mid a_1] = E_{a_2,\theta}[u(a_2, \theta) \mid a_1]$

two quadratic conditions & two normalization conditions for r

$\Rightarrow$ analytical solution

Solving the binary setting

feasibility: $\frac{r(1|1)r(0|0)}{r(0|1)r(1|0)} = \frac{p(1|1)p(0|0)}{p(0|1)p(1|0)}$

optimality: $\sum_\theta \pi_\theta u(a_1, \theta) r(a_1|\theta) = \sum_{\theta, a_2} \pi_\theta u(a_2, \theta) r(a_2|\theta) r(a_1|\theta)$

two quadratic conditions & two normalization conditions for r

$\Rightarrow$ analytical solution

Intuition for the confirmation bias

assume:

- symmetric incentives $u_1 = u_0$
- symmetric primitive process $p(1 \mid 1) = p(0 \mid 0) > 1/2$
- asymmetric prior $\pi_1 > \pi_0$

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

recommendation $a = 0$ is less convincing than $a = 1$

$\Rightarrow$ agent benefits from the "second thought" when $a = 0$

$\Rightarrow$ optimal "bias" for $a = 1$

Outline

Motivation

Woodford (2016) discusses an experiment: a subject is rewarded for a correct identification of an acoustic tone θ

evidence: distr. of the guess θ' is bell-shaped around θ

but

entropy-based model: distr. of θ' is constant for all $\theta' \neq \theta$

Perceptual tasks

$\Theta = A$, $|\Theta| > 2$

uniform prior π_θ

$u(\theta', \theta) = 1$ if $\theta' = \theta$ and 0 otherwise

$p(\theta' \mid \theta)$—probability of observing θ' in state θ under the primitive decision process

define perceptual "distance":

$$d(\theta, \theta') = \log \frac{p(\theta|\theta)p(\theta'|\theta')}{p(\theta|\theta')p(\theta'|\theta)}$$

Conservation of the perceptual distance

Proposition

Perceptual distance is preserved for any β:

$$\log \frac{r(\theta|\theta;\beta)r(\theta'|\theta';\beta)}{r(\theta|\theta';\beta)r(\theta'|\theta;\beta)} = d(\theta,\theta').$$

Conservation of the perceptual distance

Proposition

Perceptual distance is preserved for any β:

$$\log \frac{r(\theta|\theta;\beta)r(\theta'|\theta';\beta)}{r(\theta|\theta';\beta)r(\theta'|\theta;\beta)} = d(\theta,\theta').$$

proof:

$$\frac{r(\theta|\theta;\beta)r(\theta'|\theta';\beta)}{r(\theta|\theta';\beta)r(\theta'|\theta;\beta)}$$

Conservation of the perceptual distance

Proposition

Perceptual distance is preserved for any β:

$$\log \frac{r(\theta|\theta;\beta)r(\theta'|\theta';\beta)}{r(\theta|\theta';\beta)r(\theta'|\theta;\beta)} = d(\theta,\theta').$$

proof:

$$\frac{\frac{\beta_\theta p(\theta|\theta)}{\sum_{\tilde\theta}\beta_{\tilde\theta}p(\tilde\theta|\theta)}\frac{\beta_{\theta'}p(\theta'|\theta')}{\sum_{\tilde\theta}\beta_{\tilde\theta}p(\tilde\theta|\theta')}}{\frac{\beta_\theta p(\theta|\theta')}{\sum_{\tilde\theta}\beta_{\tilde\theta}p(\tilde\theta|\theta')}\frac{\beta_{\theta'}p(\theta'|\theta)}{\sum_{\tilde\theta}\beta_{\tilde\theta}p(\tilde\theta|\theta)}}$$

Conservation of the perceptual distance

Proposition

Perceptual distance is preserved for any β*:*

$$\log \frac{r(\theta|\theta;\beta)r(\theta'|\theta';\beta)}{r(\theta|\theta';\beta)r(\theta'|\theta;\beta)} = d(\theta,\theta').$$

proof:

$$\frac{p(\theta|\theta)p(\theta'|\theta')}{p(\theta|\theta')p(\theta'|\theta)}$$

Further assumptions and definitions

assumptions:

symmetry: $p(\theta_2 \mid \theta_1) = p(\theta_1 \mid \theta_2)$

sufficient precision:

- $p(\theta_2 \mid \theta_1) \leq p(\theta_1 \mid \theta_1)$
- $\Rightarrow d(\theta_1, \theta_2) > 0$ for all $\theta_1 \neq \theta_2$

definitions:

state θ_1 is more distinct than θ_2 if for all $\theta_3 \neq \theta_1, \theta_2$, $d(\theta_1, \theta_3) < d(\theta_2, \theta_3)$.

decision rate in state θ is $\lambda_\theta := \sum_{\theta'} \beta_{\theta'} p(\theta' \mid \theta)$

Perception results

Proposition

1. *the agent reacts faster in distinct states:*

$$\lambda_\theta \propto p^{1/2}(\theta \mid \theta)$$

2. *(similarity curse) if θ_1 is more distinct than θ_2, then the optimal perception discriminates in favor of θ_1:*

$$r(\theta_1 \mid \theta_2) > r(\theta_2 \mid \theta_1)$$

Perception results

Proposition

1. *the agent reacts faster in distinct states:*

$$\sum_{\theta'} \beta_{\theta'} p(\theta' \mid \theta) \propto p^{1/2}(\theta \mid \theta)$$

2. *(similarity curse) if* θ_1 *is more distinct than* θ_2*, then the optimal perception discriminates in favor of* θ_1*:*

$$r(\theta_1 \mid \theta_2) > r(\theta_2 \mid \theta_1)$$

Similarity curse

intuition

assume uniform termination probabilities

consider a state θ^* that is very similar to many other states

perception θ^* is quite a noisy signal of the true state

agent who draws perception θ^* wants to restart

information search is biased against indistinct states

Proof

second-thought-free condition:

$$E_\theta[u(\theta_1, \theta) \mid \theta_1] = E_{\theta,\theta_2}[u(\theta_2, \theta) \mid \theta_1]$$

Proof

second-thought-free condition:

$$\sum_{\theta' \neq \theta} r(\theta \mid \theta) r(\theta' \mid \theta) = \sum_{\theta' \neq \theta} r(\theta' \mid \theta') r(\theta \mid \theta') \text{ for all } \theta$$

formally, this is a balance condition for a Markov chain $r(\theta' \mid \theta)$, where $r(\theta \mid \theta)$ stands for the ergodic probability of the state θ

Proof

claim: chain $r(\theta' \mid \theta)$ is reversible.

recall that a chain m is reversible iff it satisfies Kolmogorov's criterion for any cycle of states:

$$\frac{m_{j_1,j_2} m_{j_2,j_3} \ldots m_{j_{n-1},j_n} m_{j_n},j_1}{m_{j_1,j_n} m_{j_n,j_{n-1}} \ldots m_{j_3,j_2} m_{j_2,j_1}} = 1.$$

"chain" $p(\theta' \mid \theta)$ is reversible by the symmetry

Kolmogorov's ratios are preserved for any β

hence $r(\theta' \mid \theta)$ is reversible.

reversible Markov chains satisfy detailed balance conditions

$$\sum_{\theta' \neq \theta} r(\theta \mid \theta) r(\theta' \mid \theta) = \sum_{\theta' \neq \theta} r(\theta' \mid \theta') r(\theta \mid \theta') \text{ for all } \theta$$

Proof

claim: chain $r(\theta' \mid \theta)$ is reversible.

recall that a chain m is reversible iff it satisfies Kolmogorov's criterion for any cycle of states:

$$\frac{m_{j_1,j_2}m_{j_2,j_3}\dots m_{j_{n-1},j_n}m_{j_n},j_1}{m_{j_1,j_n}m_{j_n,j_{n-1}}\dots m_{j_3,j_2}m_{j_2,j_1}} = 1.$$

"chain" $p(\theta' \mid \theta)$ is reversible by the symmetry

Kolmogorov's ratios are preserved for any β

hence $r(\theta' \mid \theta)$ is reversible.

reversible Markov chains satisfy detailed balance conditions

$$r(\theta \mid \theta)r(\theta' \mid \theta) = r(\theta \mid \theta')r(\theta' \mid \theta') \text{ for all } \theta' \neq \theta$$

Proof

claim: chain $r(\theta' \mid \theta)$ is reversible.

recall that a chain m is reversible iff it satisfies Kolmogorov's criterion for any cycle of states:

$$\frac{m_{j_1,j_2}m_{j_2,j_3}\ldots m_{j_{n-1},j_n}m_{j_n},j_1}{m_{j_1,j_n}m_{j_n,j_{n-1}}\ldots m_{j_3,j_2}m_{j_2,j_1}} = 1.$$

"chain" $p(\theta' \mid \theta)$ is reversible by the symmetry

Kolmogorov's ratios are preserved for any β

hence $r(\theta' \mid \theta)$ is reversible.

reversible Markov chains satisfy detailed balance conditions

$$\frac{\beta_\theta \beta_{\theta'} p_{\theta\theta} p_{\theta'\theta}}{\left(\sum_{\tilde{\theta}} \beta_{\tilde{\theta}} p_{\tilde{\theta}\theta}\right)^2} = \frac{\beta_\theta \beta_{\theta'} p_{\theta\theta'} p_{\theta'\theta'}}{\left(\sum_{\tilde{\theta}} \beta_{\tilde{\theta}} p_{\tilde{\theta}\theta'}\right)^2}$$

Proof

claim: chain $r(\theta' \mid \theta)$ is reversible.

recall that a chain m is reversible iff it satisfies Kolmogorov's criterion for any cycle of states:

$$\frac{m_{j_1,j_2} m_{j_2,j_3} \ldots m_{j_{n-1},j_n} m_{j_n},j_1}{m_{j_1,j_n} m_{j_n,j_{n-1}} \ldots m_{j_3,j_2} m_{j_2,j_1}} = 1.$$

"chain" $p(\theta' \mid \theta)$ is reversible by the symmetry

Kolmogorov's ratios are preserved for any β

hence $r(\theta' \mid \theta)$ is reversible.

reversible Markov chains satisfy detailed balance conditions

$$\frac{\sum_{\tilde{\theta}} \beta_{\tilde{\theta}} p_{\tilde{\theta}\theta'}}{\sum_{\tilde{\theta}} \beta_{\tilde{\theta}} p_{\tilde{\theta}\theta}} = \frac{p_{\theta'\theta'}^{1/2}}{p_{\theta\theta}^{1/2}}$$

Proof

claim: chain $r(\theta' \mid \theta)$ is reversible.

recall that a chain m is reversible iff it satisfies Kolmogorov's criterion for any cycle of states:

$$\frac{m_{j_1,j_2} m_{j_2,j_3} \dots m_{j_{n-1},j_n} m_{j_n},j_1}{m_{j_1,j_n} m_{j_n,j_{n-1}} \dots m_{j_3,j_2} m_{j_2,j_1}} = 1.$$

"chain" $p(\theta' \mid \theta)$ is reversible by the symmetry

Kolmogorov's ratios are preserved for any β

hence $r(\theta' \mid \theta)$ is reversible.

reversible Markov chains satisfy detailed balance conditions

$$\frac{\lambda'_\theta}{\lambda_\theta} = \frac{p^{1/2}_{\theta'\theta'}}{p^{1/2}_{\theta\theta}}$$

Conclusion

we provide a model of memory management

micro-foundation of the targeted search for information

robust behavioral predictions

preference identification

predictions on decision times

behavioral phenomena:

- confirmation bias
- bias towards distinct states