

DISCUSSION PAPER SERIES

DP12546

**ON SECOND THOUGHTS, SELECTIVE
MEMORY, AND RESULTING
BEHAVIORAL BIASES**

Philippe Jehiel and Jakub Steiner

INDUSTRIAL ORGANIZATION



ON SECOND THOUGHTS, SELECTIVE MEMORY, AND RESULTING BEHAVIORAL BIASES

Philippe Jehiel and Jakub Steiner

Discussion Paper DP12546
Published 27 December 2017
Submitted 27 December 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Philippe Jehiel and Jakub Steiner

ON SECOND THOUGHTS, SELECTIVE MEMORY, AND RESULTING BEHAVIORAL BIASES

Abstract

A proposed model of information processing generates a prediction about the constrained-optimal stochastic choice that is robust to details of the feasible information structures. A decision-maker processes payoff-relevant information until she reaches her cognitive constraint, at which point she either terminates the decision-making and chooses an action, or restarts the process. By conditioning the probability of termination on the information collected, she controls the correlation between the payoff state and her terminal action. The constrained-optimal choice rule exhibits (i) confirmation bias, (ii) speed-accuracy complementarity, (iii) overweighting of rare events, and (iv) salience effect.

JEL Classification: N/A

Keywords: N/A

Philippe Jehiel - jehiel@pse.ens.fr
PSE and UCL and CEPR

Jakub Steiner - jakub.steiner@cerge-ei.cz
cerge-ei and Edinburgh University and CEPR

Acknowledgements

We thank Mark Dean, Alessandro Pavan, Philip Reny, Colin Stewart, colleagues at the University of Edinburgh and the audiences at Bocconi, Queen Mary, Columbia, and Ecole Polytechnique universities and workshops and conferences in Erice, Alghero, Faro, Gerzensee, and New York. Ludmila Matyskova and Jan Sedek provided excellent research assistance. Deborah Novakova and Laura Strakova have helped with English. Jakub Steiner has received financial support from the Czech Science Foundation grant 16-00703S.

On Second Thoughts, Selective Memory, and Resulting Behavioral Biases*

Philippe Jehiel
PSE and UCL

Jakub Steiner
University of Edinburgh and Cerge-Ei

December 19, 2017

Abstract

A proposed model of information processing generates a prediction about the constrained-optimal stochastic choice that is robust to details of the feasible information structures. A decision-maker processes payoff-relevant information until she reaches her cognitive constraint, at which point she either terminates the decision-making and chooses an action, or restarts the process. By conditioning the probability of termination on the information collected, she controls the correlation between the payoff state and her terminal action. The constrained-optimal choice rule exhibits (i) confirmation bias, (ii) speed-accuracy complementarity, (iii) overweighting of rare events, and (iv) salience effect.

keywords: bounded rationality, cognitive constraints, information processing, stochastic choice, confirmation bias, speed-accuracy complementarity, probability weighting, salience.

JEL codes: D03, D80, D81, D83, D89, D90.

*We thank Mark Dean, Alessandro Pavan, Philip Reny, Colin Stewart, colleagues at the University of Edinburgh and the audiences at Bocconi, Queen Mary, Columbia, and Ecole Polytechnique universities and workshops and conferences in Erice, Alghero, Faro, Gerzensee, and New York. Ludmila Matysková and Jan Šedek provided excellent research assistance. Deborah Nováková and Laura Straková have helped with English. Jakub Steiner has received financial support from the Czech Science Foundation grant 16-00703S.

1 Introduction

Economic agents often acquire information about the state of the economy before making their decisions. The information is typically modelled as a signal that helps the agent refine the distribution of the state and improve the decision-making. Often, signals come over time and agents can absorb only a small number of them. We capture this information-processing friction by assuming that agents can receive as many signals as they wish but remember only one when making their choices. (Extensions to richer but finite observability are discussed later.) An agent in the face of the first signal she observes can either make a choice based on this observation (either in an optimal way or as dictated by some rule of thumb) or rerun the very same stochastic signal extraction process. In the face of the second observed signal she can either make a final decision (after the second observation and according to the same strategy that maps instantaneous observations to choices) or rerun the signal extraction again and so on. Specifically, we allow agents to rerun the decision-making from the start at any time, the agent can randomize over the rerunning/termination decisions, and the restarting probabilities can depend on the currently favored action that the agent would choose if a choice had to be made now. We do not impose time constraints and costs in the basic formulation so that the economic friction comes solely from the limited cognitive capacity of the agent.

We ask ourselves: Should the agent optimally make her choice as soon as she receives the first signal whatever the realization of it is, or can she be better off by rerunning the very same information-acquisition process? Can hesitation—selective repetition of a fixed stochastic decision procedure where the repetition conditions on the procedure’s outcome—be productive?

A general insight is that selective rerunning of the primitive decision procedure is typically optimal. Specifically, we provide a simple necessary condition satisfied by the optimal rerunning strategy. The result is an interim indifference condition imposed on the agent who has concluded her decision-making with a plan to choose a particular action. Given the recommended action, the agent’s posterior expected payoff from implementing this action must be the same as the posterior expected payoff from rerunning the whole decision-making—the whole selective repetitions of the primitive signal extraction—and implementing whichever action the second run of the decision-making will recommend. We refer to this as to the *second-thought-free* condition.

For illustration, consider a binary decision of whether to make an investment of a fixed size. The agent receives payoff 1 if she invests in the good state of the economy, payoff -2 if she invests in the bad state, and receives 0 when she does not invest whatever the state. Both states are a priori equally likely and give rise to a population of good and bad attributes, with the share of the good attributes at 90% in the good state and 10% in the bad state. The agent draws possibly several attributes in sequence but remembers only the last one when making her investment choice. We assume she invests if the last drawn attribute was good (and does not invest upon the last bad attribute). Observe that the decision rule generated by the uniform draw of the representative attribute and immediate termination upon the first signal that comes in does *not* generate the second-thought-free choice rule: An agent who has observed a uniformly drawn good attribute

prefers to rerun the decision process, since the new run will either lead to investing again or will reveal the attribute that conflicts the first observation and will lead to not investing. Since, conditional on two conflicting attributes, not investing is preferred, the agent benefits from having the second thought.

In the sequel, we refer to the probability of terminating the decision-process after receiving a particular signal as a search intensity for the signal; a higher probability of termination at a given information set inflates the likelihood that the agent makes the terminal choice at the set. We show that the failure of the second-thought-free condition with uniform search intensity in the above investment decision example indicates that relative to the uniform search, the agent benefits from increasing the search intensity for the bad attribute.

More generally, the second-thought-free condition follows from the first-order condition imposed on the optimal search intensities. In the above example, a marginal increase in the relative search intensity in favor of the bad attribute implements the second thought: Consider a deviation from the discussed decision strategy that consists of repeating the attribute draw with a small probability whenever the observed attribute is good. This new decision procedure effectively replaces a marginal measure of contingencies in which the agent terminates after observing the good attribute with new draws of the attribute. The indifference to marginal changes in the search intensities at optimum implies the second-thought-free condition.

From an empirical viewpoint, the second-thought-free condition has interesting implications. First, even if the analyst does not know the details of the decision process (including what the agent observes, how this is jointly distributed with the state, and what use the agent makes of the information), the analyst can check whether the second-thought-free condition holds simply based on the joint distribution of chosen actions and states in conjunction with the knowledge of the agent's preferences. Second, when preferences are not known, some properties of those can be inferred from the choice data under the assumption that agents maximize across the search intensities, thereby implying that the second-thought-free condition holds. For example, in the binary state setting with two signal realizations and two actions as in the investment example, we note that the utility can be inferred from the second-thought-free condition up to an affine transformation.

The model provides microfoundations to a range of behavioral stylized facts. The unifying principle of our behavioral insights is the intuition that the agent targets her search towards the type of evidence that would provide her with more informed posteriors under the uniform search. This principle generates *confirmation bias*, since evidence that confirms the agent's prior leads to more informed posterior than does evidence that contradicts the prior. Optimally targeted information search also generates *speed-accuracy complementarity*; that is, accuracy of choice declines with the response time. The effect is generated by the confirmation bias: The agent encountering evidence contradicting her prior is likely to disregard the evidence and to have a second thought. Hence, long response times indicate a surprising state of the world, and the constrained-optimal choice rule commits errors in the surprising state relatively often. *Overweighting rare events* occurs when

the agent’s task is to form a probability belief about an event that is known to be rare, such as a flight accident, by observing a random flight outcome. Since observing a flight accident is far more informative about the probability of future accidents than observing an uneventful flight, the agent optimally biases her search towards eventful flights. In the last behavioral application, we show that distinct states of the world are *salient* in the sense that they attract the agent’s attention. The effect arises because an indistinct perception stimulus that can be generated by several similar states is less informative than a distinct perception stimulus likely generated by a specific distinct state. Hence, the optimal information search targets stimuli indicating distinct states.

It should be stressed that our insights carry over as long as the agent can decide optimally on rerunning/terminating the decision-making as a function of what she is supposed to do. The second-thought-free condition continues to hold when the agent is allowed to use finer termination strategies that condition on details of information collected beyond the recommended action (see Section 6 for elaboration). Since it is tautological that the agent can access information about her own choice, and since this information is not of great complexity, it seems to us plausible that agents would be able to find profitable deviations from the unbiased information search. From this perspective, our approach provides a natural explanation for a number of behavioral biases, as summarized above.

Related literature

When the decision-maker can choose from feasible error distributions, then choice data partially reveal the decision-maker’s objective, since the constrained-optimal choice rule exhibits costly types of error relatively rarely. At various levels of formalization, this error-management idea has occurred repeatedly in biology, psychology and economics; see Johnson et al. (2013) for an interdisciplinary review.¹ A range of economic models make the error-management idea precise by specifying particular sets of feasible error distributions. These models differ greatly in the assumed constraints imposed on decision-making, and hence in the predicted constrained-optimal stochastic choice rules. In what follows we review the variety of the modeling approaches to frictional decision-making.

Sims (2003) and Matějka and McKay (2015) constrain the expected reduction of the entropy of the decision-maker’s belief within the decision process. The entropy-based models generate constrained-optimal choice rules akin to the logit rules used in structural estimation, and hence they have practical applications to identification of utility from choice data, e.g. see the discussion in Steiner et al. (2017). The source of the identification is, however, the assumed entropy-based constraint.

A different information-processing constraint is assumed in drift-diffusion models, e.g. Wald (1947), Arrow et al. (1949), Ratcliff (1978), Hébert and Woodford (2016), and Morris and Strack (2017) that restrict the decision procedure to sequential sampling. By choosing the regions of the stopping beliefs, the decision-maker can trade off the accuracy against the speed of her decision

¹A typical application of the error-management theory appears in Haselton and Buss (2000), who explain men’s overperception of women’s sexual intent by the asymmetry in the fitness cost of the two types of the perceptual error.

procedure.² The class of drift-diffusion models restricts the agents to learning procedures with continuously evolving Bayesian beliefs. In contrast to this assumption, Zhong (2017) argues that learning processes with discontinuous belief evolution are optimal under a broad class of information acquisition costs, and Che and Mierendorff (2016) study one such discontinuous learning in which signals arrive in a Poisson process. Yet another modeling approach to limited cognition conceptualizes decision-making and information processing as finite algorithms with adjustable parameters, e.g. Hellman and Cover (1970), Compte and Postlewaite (2012), and Wilson (2014).

Compared to the above diverse models that fully specify the cognitive friction, our model delivers partial characterization of the constrained-optimal stochastic choice rule without a full specification of the cognitive constraint. Another research program delivering robust predictions about constrained-optimal stochastic behavioral without full specification of the cognitive friction is by Caplin et al. (2017) who provide full behavioral characterization of all posterior-separable information cost models. Yet another alternative to our attempt to deliver predictions robust to the detail of the friction is to estimate the information processing cost from the choice data. The proposed methodologies of information cost identification in Caplin and Dean (2015) and Oliveira et al. (2017) make this approach theoretically feasible when rich data, such as choice data over menus or many decision problems, are available.

This paper belongs to a growing economic literature that explains established behavioral phenomena as the constrained-optimal behavior of decision-makers facing information processing frictions. For instance, Robson (2001), Rayo and Becker (2007), Netzer (2009), and Khaw et al. (2017) provide microfoundations for risk attitudes, Gabaix and Laibson (2017) endogenize discounting, and Wilson (2014), Compte and Postlewaite (2012), and Leung (2017) establish constrained-optimal ignorance of weakly informative news.

A related model of extreme-events oversampling was developed by psychologists Lieder et al. (2014). Their decision-maker has an access to the objective distribution of a payoff state but faces a friction in the computation of the payoff expectations. She forms the expected payoffs for considered actions in a Monte Carlo simulation as a utility average over a finite sample of simulated payoff states. Instead of the uniform sampling from the true belief distribution, she optimally oversamples states in which stakes are high (and corrects for the oversampling by lowering the respective weights), as in the so-called importance-sampling Monte Carlo method, e.g. Glynn and Iglehart (1989). This project and us examine the opposing sides of information processing. While we assume non-representative sampling in the signal collection and allow for frictionless formation of the posterior beliefs, Lieder et al. take available information as given and impose the friction on its interpretation.

Brain imaging technologies have recently advanced to the level that allows to document the process of hesitation in real time for non-human subjects. Studies on rats, Redish (2016), and monkeys, Rich and Wallis (2016), document that animals in some decision problems deliberate in a

²Recently, Hébert and Woodford (2016), and Morris and Strack (2017) have provided the microfoundations of Sim’s entropy-based model and of its generalization in terms of drift-diffusion models.

stochastic series of actions plans, in which their minds move there-and-back between the available options. For instance, researchers are able to link distinct hippocampal activities representing particular parts of a maze, and link in real time those brain patterns to a sequence of deliberation of a rat that has paused in a maze and is choosing how to proceed in the maze. Unlike the neural processes represented by the drift-diffusion model of Ratcliff (1978) that assumes a continuous evolution of a mental state, these deliberation processes, like our model, exhibit discrete transitions between the deliberated actions.³

2 Reduced-form example

Let us illustrate the main insights using an example in which targeting of the information search is modeled in a reduced form and the studied decision problem is binary. We provide an explicit process that implements the targeting in a general decision problem in the next section. An agent chooses a binary action $a \in A = \{0, 1\}$ and receives a positive reward if her action matches a payoff state $\theta \in \Theta = A$; her payoff is $u(a, \theta) = u_\theta > 0$ if $a = \theta$ and it is 0 otherwise. The state θ , unknown to the agent, is drawn from an interior prior distribution $\pi \in \Delta(\Theta)$ with π_θ denoting the prior probability of each θ . The agent learns about the state by observing a random signal x with support $X = \Theta$. The observed signal is drawn from a population of signals with conditional distributions $p(x | \theta)$ for each θ ; $p(x | \theta) > 0$ for all x and θ . To avoid a trivial case, $p(x | \theta)$ differs across θ .

We envision an agent who—if she were not restricted by her cognitive limitation—could resolve the payoff-state uncertainty. No external constraints prohibit the agent from drawing arbitrarily many conditionally independent signal realizations, and if she could process sufficiently many such signals, she could learn the realized state with arbitrary precision. The bottleneck prohibiting the agent from reaching complete information is that she can comprehend only one signal draw. Given this constraint, the agent benefits from deviating from a uniform signal draw to a signal draw that partially targets a particular signal realization.

To that end, we distinguish the above $p \in \Delta(X)^{|\Theta|}$, which we refer to as the *primitive* information structure, from an *effective* information structure that differs due to the targeting of the agent’s search. The agent chooses a vector $\beta = (\beta_x)_{x \in X}$ of non-negative search intensities and draws one signal from the effective signal distribution

$$r(x | \theta; \beta) = \frac{\beta_x p(x | \theta)}{\sum_{x' \in X} \beta_{x'} p(x' | \theta)}. \tag{1}$$

That is, the effective probability that the agent observes a signal realization x in state θ is given by the primitive conditional probability $p(x | \theta)$, inflated by the search intensity β_x , and renormalized; $r(\beta) \in \Delta(X)^{|\Theta|}$ is the effective information structure.

³The study of Rich and Wallis (2016) documents that, in contrast with our baseline model but in accord with our extensions in Section 6, information is aggregated across the rounds of deliberation.

Let $B = [0, 1]^{|X|} \setminus \{(0, 0)\}$ be the set of the feasible vectors of the search intensities,⁴ let $\sigma : X \rightarrow A$ be an action strategy that maps each observed signal x to action $\sigma(x)$, and let S be the set of all such strategies. We abuse the notation by letting $r(a | \theta; \beta, \sigma)$ denote the stochastic choice rule generated by the search intensities β and the action strategy σ . That is,

$$r(a | \theta; \beta, \sigma) = \sum_{x: \sigma(x)=a} \frac{\beta_x p(x | \theta)}{\sum_{x' \in X} \beta_{x'} p(x' | \theta)}$$

is the probability that the agent chooses action a in the state θ . The agent selects the search intensities and the action strategy that jointly maximize her expected payoff:

$$\max_{\beta \in B, \sigma \in S} \sum_{\theta \in \Theta, a \in A} \pi_\theta u(a, \theta) r(a | \theta; \beta, \sigma). \quad (2)$$

Below, in Sections 3 and 4, we offer an explicit procedure that implements the directed search in (1). For the purpose of this section, we interpret the agent as consisting of two selves. The first self collects decision-related data. The second self observes a uniform draw from the collected data set and chooses an action. In detail, the first self has access to the signal population with distribution $p(x | \theta)$ for the realized state θ . The first self chooses probabilities β_x , $x \in X$, and reports each signal from the population to the constructed data set with probability β_x . The signal population in the data set is thus $r(x | \theta; \beta)$ given by (1). The second self draws one signal from the data set and chooses an action. The strategies of the two selves form an equilibrium of a team game in which both selves maximize the agent's expected payoff.⁵ Self one optimizes β given self two's choice of action strategy σ . Self two's action choice $\sigma(x)$ is optimal, given the effective posterior belief induced by self one's choice of β .⁶ The two selves may represent two distinct members of a team (e.g. a subordinate and her superior in a firm), or the same person at distinct times. In particular, the model may represent a person with selective memory, where the first self chooses the selectiveness of the memorized data, and the second self chooses action after a recollection of a memorized data point.

Let $r^*(a | \theta) = r(a | \theta; \beta^*, \sigma^*)$ be the choice rule generated by an optimal pair (β^*, σ^*) solving (2). The next result is the first-order condition on r^* .

Proposition 1. *Let a_1 and a_2 be two conditionally independent action draws from the optimal choice rule r^* , with the joint probability distribution of the state and the actions: $\alpha(\theta, a_1, a_2) =$*

⁴The imposed upper bound $\beta_x \leq 1$ is without loss of generality because $r(\beta_0, \beta_1)$ is homogeneous of degree zero, and thus depends only on the ratio β_1/β_0 . The exclusion of $(0, 0)$ guarantees that the denominator in (1) is positive.

⁵Given the aligned interest of the two selves, it may be surprising that self one chooses to trim the data. As recognized in Steiner and Stewart (2016), the conflict of interest between the selves arises because they act at distinct information sets.

⁶We implicitly assume, as in Jehiel (2017), that the agent learns the posterior distributions in repetitions of the decision-making. In contrast to Jehiel (2017), our agent learns the true posteriors; self two does not face any frictions in interpreting the observed signal draw.

$\pi_{\theta} r^*(a_1 | \theta) r^*(a_2 | \theta)$. Then,

$$\mathbb{E}_{\alpha} [u(a_1, \theta) | a_1 = a] = \mathbb{E}_{\alpha} [u(a_2, \theta) | a_1 = a], \quad (3)$$

for all actions a chosen with positive probability under r^* .

The result follows from a more general Proposition 2 in Section 4.

Condition (3) is an indifference condition. The left-hand side is the payoff expectation of the agent who receives and implements the recommendation to choose action a generated by the choice rule r^* and who holds the Bayesian posterior belief about θ . The right-hand side is the payoff expectation of the agent who has also received the recommendation a generated by r^* , who will decline the recommendation, and will generate (and implement) a new action recommendation a_2 in the new run of the optimal decision process r^* . This agent is uncertain both about θ and the new recommendation draw a_2 , and, based on the joint distribution $\alpha(\theta, a_1, a_2)$, forms the posterior belief about the pair (θ, a_2) conditional on the first recommendation $a_1 = a$.

Thus, the agent optimizing her information search behaves *as if* she had an option to rerun the decision process. It is as if, at the end of the decision-making that she is about to conclude with choice a , the agent could discard the information collected so far, run the decision-making for the second time, and implement the new conclusion a_2 . Condition (3) requires the agent to be indifferent between terminating after the first run of the decision-making and having the second thought.

The result follows from the first-order condition. Consider a marginal increase of the search intensity β_x for the signal realization x , such that $\sigma(x) = a$. It is easy to verify that the impact of the increase on the objective in (2) is proportional to $\mathbb{E}_{\alpha}[u(a_1, \theta) - u(a_2, \theta) | a_1 = a]$. The increase in β_x affects the joint distribution of (a, θ) . In particular, the increase in β_x has the same effect as replacement of a marginal measure of contingencies in which the agent concludes the decision process with action $a = \sigma(x)$ with new draws a_2 from the effective choice rule.

For an illustration of how the agent benefits from targeting her information search, consider a setting with uniform prior, uniform rewards $u_0 = u_1 = 1$ and an asymmetric primitive information structure $p(x | \theta)$ that satisfies

$$.9 = \Pr_p(\theta = 1 | x = 1) > \Pr_p(\theta = 0 | x = 0) = .8.$$

Let us fix the action strategy $\sigma(x) = x$.⁷ Consider the untargeted information search with $\beta_0 = \beta_1 = 1$ so that the effective and the primitive information structures coincide. The signal realization $x = 1$ leads to the correct action in 90% of cases, whereas the somewhat less informative signal $x = 0$ leads to the correct choice in only 80%. A marginal decrease of β_0 replaces a marginal measure of the contingencies in which the agent chooses action 0 with new runs of the decision process. To see that such replacement is beneficial, suppose that the agent concludes the decision process with recommendation $a_1 = 0$, and before she implements a_1 , she considers having a second

⁷We show in Section 5 that the action strategy $\sigma(x) = x$ is optimal.

thought—running the decision process again and implementing whichever action a_2 the new run of the process recommends. If $a_2 = a_1$, then the second thought will have been inconsequential. If $a_2 = 1 \neq a_1$, then it is more likely that $\theta = 1$ than $\theta = 0$, and thus the agent benefits from switching from a_1 to a_2 . The marginal deviation from the uniform information search towards targeting the more informative evidence $x = 1$ is thus beneficial. The relative increase in β_1 affects the informativeness of both signal realizations: The precision of $x = 1$ decreases, and the precision of $x = 0$ increases. At optimum, in this example with symmetric payoffs and prior, both signal realizations must be equally informative:

$$\Pr_{r^*}(\theta = 1 \mid x = 1) = \Pr_{r^*}(\theta = 0 \mid x = 0),$$

so that the agent concluding the decision process with either of the two actions is indifferent between implementing the action and having a second thought.

Proposition 1 allows an outside observer who does not know the primitive information structure $p(x \mid \theta)$ to identify the agent’s preferences. Let the outsider observe the joint distribution of the state θ and action a , which is equivalent to observing the optimal choice rule $r^*(a \mid \theta)$ and the prior π_θ . To keep the discussion in the example simple, let the observed prior be uniform and the frequencies of both types of errors in choice equal: $0 < r(0 \mid 1) = r(1 \mid 0) < 1/2$. Then, the outsider can conclude that both types of error are equally costly: $u_1 = u_0 > 0$, which in this binary setting fully identifies the utility function up to affine transformations.⁸ The source of identification is the optimality condition (3). Applied to action $a = 1$ and multiplied by the unconditional probability of action 1, (3) becomes:

$$\pi_1 u_1 r^*(1 \mid 1) = \pi_0 u_0 r^*(0 \mid 0) r^*(1 \mid 0) + \pi_1 u_1 r^*(1 \mid 1) r^*(1 \mid 1),$$

and for the considered choice data, the last condition implies $u_1 = u_0$.

Let us highlight the role of information targeting in the identification of preferences. If the agent were constrained to use the primitive information structure $p(x \mid \theta)$, then the outsider, uninformed about p , could not disentangle the impacts of p and u on the choice data. Nor does the fact that the agent optimizes the effective information structure allow, on its own, the outsider to identify the preferences. If the agent were choosing the information structure $r(x \mid \theta)$ from a set \mathcal{R} , with \mathcal{R} unobserved by the outsider, then the outsider could not distinguish the agent’s preferences from the shadow prices of the constraint $r \in \mathcal{R}$. Below, we argue that the optimality condition (3) holds in a variety of settings that differ by the level of the agent’s sophistication. The robustness of condition (3) to the unobserved details of the cognitive constraints allows an outsider who is ignorant about the agent’s cognitive frontier to partially identify the agent’s preferences from the choice data.

We return to the binary setting of this example and provide its complete analytical solution in Section 5.

⁸The fact that $r(0 \mid 1) = r(1 \mid 0) < 1/2$ shows that $u_1 = u_0$ are positive, and hence the outsider can identify choice $a \neq \theta$ as an error. Otherwise, if $u_1 = u_0$ were negative, the agent would benefit from permuting her action strategy $\sigma(x)$.

3 Model

An agent faces a decision problem under uncertainty. She chooses an action $a \in A$ in a process specified below and receives a payoff $u(a, \theta)$ in the payoff state $\theta \in \Theta$ drawn from an interior prior distribution $\pi \in \Delta(\Theta)$. The state space Θ and the action set A are finite.

The agent can employ any *primitive decision process* $p \in \mathcal{P}$, where each p is a family of conditional distributions $p(a, x | \theta)$, $\theta \in \Theta$. That is, process $p \in \mathcal{P}$ generates an action recommendation $a \in A$ and a *side information* $x \in X$, where X is a finite signal space. We introduce the side information to capture decision processes in which the agent acquires information beyond that encoded in a . The side information will play a role in Section 6. We write $p(a | \theta) = \sum_x p(a, x | \theta)$ for the marginal conditional action distribution. We interpret the set \mathcal{P} as unobservable to analysts and thus are interested in results that do not depend on \mathcal{P} .

Compared to the last, reduced-form section, we model here the information targeting explicitly. We allow the agent to selectively repeat the primitive decision process, and show that the selective repetitions implement the effective choice rules studied in the reduced-form example. To capture the cognition constraint, we assume that each repetition of the decision process involves discarding the information collected so far. We interpret the primitive decision process as a cognition that exhausts the agent's capacity dedicated to the problem being solved. Once the agent hits the constraint at the end of the primitive process, she can continue only after she unclogs her capacity. One particular source of the studied cognition constraint is limited memory that the agent can unclog by amnesia.

The agent is able to condition the termination on the action a recommended by the primitive decision process p and can randomize: She chooses a vector $\beta = (\beta_a)_{a \in A} \in B = [0, 1]^{|A|} \setminus \{(0, \dots, 0)\}$ of termination probabilities β_a for each action a ; we call β a *termination strategy*. The agent runs the primitive process p for the first time, receives action recommendation a_1 with probability $p(a_1 | \theta)$ and terminates with a_1 with probability β_{a_1} . She restarts her decision-making with the complementary probability $1 - \beta_{a_1}$, and receives a new action recommendation a_2 from a new run of the process p with probability $p(a_2 | \theta)$, terminates with probability β_{a_2} or restarts with probability $1 - \beta_{a_2}$, and continues to rerun the primitive process p until she terminates after a random number of repetitions of p ; see Figure 1. To assure that for each $\beta \in B$ and for each state the decision process almost surely eventually terminates, we assume that $p(a | \theta) > 0$ for all a, θ and p , and $\beta \neq (0, \dots, 0)$. One interpretation of β_a is that when the recommendation of a is being made it is immediately forgotten with probability $1 - \beta_a$ thereby triggering a continuation of the thought process. From this perspective, having distinct β_a for different a implements the familiar idea of selective memory according to which some facts (here recommendations of a) are more easily forgotten/remembered than others.

The outcomes of distinct runs of the primitive process are conditionally independent. Thus, the probability that the agent terminates after t repetitions of the process p , with each run $l = 1, \dots, t$

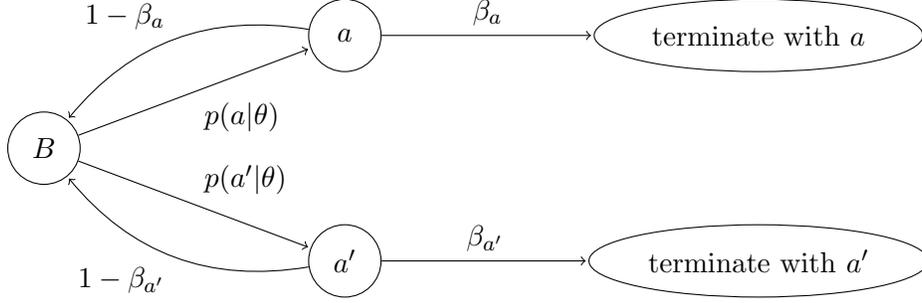


Figure 1: The decision process is a Markov chain evolving on the agent’s states of mind, with transition probabilities that depend on the payoff state θ . The chain begins in the state of mind B and transits to states $a \in A$ with probabilities $p(a | \theta)$. The process returns to B with probability $1 - \beta_a$, or terminates with choice of a with probability β_a .

resulting in recommendation a_l and side information x_l is

$$\rho((a_l, x_l)_{l=1}^t | \theta; p, \beta) = \beta_{a_t} p(a_t, x_t | \theta) \prod_{l=1}^{t-1} (1 - \beta_{a_l}) p(a_l, x_l | \theta).$$

We let

$$r(a, t, x | \theta; p, \beta) = \sum_{(a_l, x_l)_{l=1}^t: a_t=a, x_t=x} \rho((a_l, x_l)_{l=1}^t | \theta; p, \beta) \quad (4)$$

denote the probability that the agent who employs the primitive process p and the termination strategy β terminates with action a in round t with side information x . We call $r(a, t, x | \theta; p, \beta)$ the *effective decision process*. The set of feasible effective processes is $\mathcal{R} = \{r(p, \beta) : p \in \mathcal{P}, \beta \in B\}$. Let $r(a | \theta; p, \beta) = \sum_{t,x} r(a, t, x | \theta; p, \beta)$ denote the probability that the agent terminates with action a in state θ ; we call this the *effective choice rule*. For the baseline model analyzed in the next two sections, the choice rule $r(a | \theta)$ is the sufficient summary statistics of the process $r(a, t, x | \theta)$. The side information x and the response time t will be relevant in extensions in Section 6.

The *repeated-cognition* problem is to choose a feasible decision process r that maximizes the expected payoff:

$$\max_{r \in \mathcal{R}} \sum_{\theta \in \Theta, a \in A} \pi_\theta r(a | \theta) u(a, \theta). \quad (5)$$

The agent solving the problem knows the prior π , payoff function u , and the set \mathcal{P} of the feasible processes p . One natural interpretation of the optimization in (5) is that it is an outcome of selective pressures that favor successful decision procedures via cultural or biological evolution, or via competition of firms differing in their internal decision procedures.

Our focus will be on the optimization in (5) with respect to $\beta = (\beta_a)_a$ while imposing minimal restrictions on the set of the feasible primitive processes. For example, these processes may, as in the illustration from the last section, consist of the acquisition of a signal x and of forming an action

recommendation $\sigma(x)$ conditional on the acquired signal x . In this case, the emerging optimal effective choice rule $r(a | \theta)$ is Bayes-rational in the sense that each action a recommended by r is a posteriori optimal, since otherwise the agent would benefit from permuting $\sigma(x)$. Alternatively, the agent may be less sophisticated and the emerging optimal effective choice rule may fail to exhibit Bayes-rationality. For instance, the primitive decision process $p(a | \theta)$ may be an imperfect control process rather than an information acquisition process. It may be possible to improve upon the optimal effective $r(a | \theta; \beta, p)$ by permuting the action recommendations, but the agent may not be aware of such improvement.

The model imposes seemingly strong restrictions on the agent's termination strategy. On its face, the model restricts the agent to conditioning termination only on the recommended action in the last run of the primitive process, and not on the time elapsed, on the side information, or on information from previous runs of the primitive process. However, in Section 6, we show that the basic model is sufficiently flexible to accommodate agents who can remember more than one signal, and who can partially aggregate information over time. The source of flexibility is the generality of the set of feasible primitive processes. In Section 6 we accommodate the extended settings with sophisticated agents in the baseline model by choosing appropriate definitions of the primitive processes.

4 Main result

The first lemma shows that the selective repetitions of the primitive decision process allow the agent to target her search for information as assumed in the reduced-form formula (1).

Lemma 1. *The probability that the agent who employs a primitive decision process p and a termination strategy $\beta \in B$ terminates with action a in state θ is*

$$r(a | \theta; p, \beta) = \frac{\beta_a p(a | \theta)}{\sum_{a' \in A} \beta_{a'} p(a' | \theta)}. \quad (6)$$

Proof of Lemma 1. The effective choice rule $r(a | \theta; p, \beta)$ satisfies a recursion:

$$r(a | \theta; p, \beta) = \beta_a p(a | \theta) + \sum_{a' \in A} (1 - \beta_{a'}) p(a' | \theta) r(a | \theta; p, \beta),$$

where the first summand is the probability that the agent terminates with a after the first run of the process p , and the second summand is the probability that the agent continues with decision-making after the first run of p and terminates with a later. Solving for r gives (6). \square

The next definition introduces *second-thought-free* choice rules. If the agent's decision process generates such a rule, then she has no incentive to rerun the process regardless of the action with which the process terminates.

Definition 1. Let a_1 and a_2 be two conditionally independent draws of actions from a choice rule $r(a | \theta)$, with the joint probability distribution of the state and the actions: $\alpha(\theta, a_1, a_2) = \pi_\theta r(a_1 | \theta) r(a_2 | \theta)$. The choice rule r is second-thought-free if for each action a chosen with positive probability:

$$\mathbb{E}_\alpha[u(a_1, \theta) | a_1 = a] \geq \mathbb{E}_\alpha[u(a_2, \theta) | a_1 = a], \quad (7)$$

where the expectations are with respect to the random variables θ and a_2 .

Although the definition allows for the strict preference against having a second thought, the next lemma shows that (7) is always met with equality: If a choice rule is second-thought-free, then the terminating agent is indifferent. The lemma is a simple consequence of the law of iterated expectations.

Lemma 2. If the choice rule $r(a | \theta)$ is second-thought-free, then the condition (7) is met with equality for each action a chosen with positive probability:

$$\mathbb{E}_\alpha[u(a_1, \theta) | a_1 = a] = \mathbb{E}_\alpha[u(a_2, \theta) | a_1 = a]. \quad (8)$$

We refer to (8) as the *second-thought-free condition*.

Proof of Lemma 2. Suppose, for contradiction, that (7) holds with strict inequality for some a chosen with positive probability. Then

$$\sum_{a \in A} \mathbb{E}_\alpha[u(a_1, \theta) | a_1 = a] \Pr_\alpha(a_1 = a) > \sum_{a \in A} \mathbb{E}_\alpha[u(a_2, \theta) | a_1 = a] \Pr_\alpha(a_1 = a).$$

The last inequality is equivalent to $\mathbb{E}_\alpha[\mathbb{E}_\alpha[u(a_1, \theta) | a_1]] > \mathbb{E}_\alpha[\mathbb{E}_\alpha[u(a_2, \theta) | a_1]]$, and applying the law of iterated expectation, this simplifies to $\mathbb{E}_\alpha[u(a_1, \theta)] > \mathbb{E}_\alpha[u(a_2, \theta)]$. This establishes the contradiction because a_1 and a_2 are conditionally i.i.d. from $r(a | \theta)$. \square

Our main result follows. It specifies an optimality condition that is independent of the feasible primitive decision processes.

Proposition 2. If a stochastic choice rule solves the repeated-cognition problem (5), then it satisfies the second-thought-free condition (8).

The result follows from the first-order condition with respect to the termination strategy β , and thus it holds even if for some reason the agent fails to optimize across the primitive processes $p \in \mathcal{P}$.

Proof of Proposition 2. Using Lemma 1, the repeated-cognition problem is equivalent to

$$\max_{p \in \mathcal{P}, \beta \in B} \sum_{\theta \in \Theta, a \in A} \pi_\theta \frac{\beta_a p(a | \theta)}{\sum_{a' \in A} \beta_{a'} p(a' | \theta)} u(a, \theta). \quad (9)$$

The constraint $\beta_a \geq 0$ is not binding for a chosen with positive probability. Therefore, the first-order condition of problem (9) with respect to β_a is:

$$\sum_{\theta \in \Theta} \pi_{\theta} \frac{r^*(a | \theta)}{\beta_a} u(a, \theta) - \sum_{\theta \in \Theta, \hat{a} \in A} \pi_{\theta} r^*(\hat{a} | \theta) \frac{r^*(a | \theta)}{\beta_a} u(\hat{a}, \theta) \geq 0, \quad (10)$$

where the inequality can arise if the constraint $\beta_a \leq 1$ is binding. Multiplying this by $\beta_a / \sum_{\theta} \pi_{\theta} r^*(a | \theta)$ gives (7). Lemma 2 implies (8). \square

The intuition for the second-thought-free condition is best grasped when thinking of a marginal decrease of $\beta_a > 0$ by a small ε . This marginal change replaces a small measure of the histories of the decision-making in which the agent has terminated with a by alternative histories in which instead of the termination with a at the given time the agent has continued with the decision-making. The impact of this replacement on the ex ante expected payoff is proportional to $E_{\alpha}[u(a_2, \theta) | a_1 = a] - E_{\alpha}[u(a_1, \theta) | a_1 = a]$, hence yielding the second-thought-free condition at the optimum. Even though the maximization is made at the ex ante stage, the first-order condition leads to an interim condition similar to the one-shot deviation comparison familiar to game theorists.⁹

4.1 Discussion of the second-thought-free condition

The second-thought-free condition is related to Piccione and Rubinstein (1997), who show that the ex ante optimal decision strategy of a forgetful decision-maker can be thought of as a team equilibrium of multiple selves, with each self representing the decision-maker at an information set. Each self takes the strategies of the other selves as given, internalizes the other selves' strategies in her Bayesian inference, and maximizes the decision-maker's payoff. Our agent is forgetful in that she conditions terminations only on the last action recommended, not on the recommendation history. As in Piccione and Rubinstein, the optimal termination strategy is a team equilibrium in the sense of the modified multi-self equilibrium defined there. The self who has received recommendation a makes inferences about the state both from the recommended action a and from the fact that the previous selves (if any) have not terminated. Given her equilibrium posterior belief, each self decides whether to terminate with the current recommendation a or to delegate the decision to the next self. The equilibrium is mixed, with each self indifferent between terminating and passing the decision forward.

The second-thought-free condition is particularly simple to interpret in two specific cases. The first is a setting with binary action space. In this case, both action recommendations must be "equally convincing" at optimum in the following sense:

$$E_{\alpha}[u(a_1, \theta) | a_1 \neq a_2] = E_{\alpha}[u(a_2, \theta) | a_1 \neq a_2]; \quad (11)$$

⁹The small change in β_a leads to additional changes in the distribution of the histories of the decision-making that affect termination after several recommendations of the action a . However, these more complex changes affect the objective by the order of ε^k , $k \geq 2$, and thus can be ignored in the marginal argument.

that is, if the agent were able to observe two conflicting action recommendations generated by the optimal process r^* , then she would be indifferent.¹⁰

The other case with a simple interpretation of the second-thought-free condition arises for arbitrary action and state spaces when the agent has access to almost precise information. Let, for each action a , Θ^a denote the set of states in which a is optimal, and let $a^*(\theta) \in \arg \max_a u(a, \theta)$ stand for an optimal action in state θ .

Corollary 1. *Suppose that the probability that the optimal effective choice rule r^* recommends a suboptimal action is bounded by $\varepsilon > 0$. Further suppose that if a particular action a is optimal at θ , then $\arg \max_a u(a, \theta)$ is a singleton. Then r^* satisfies for the action a :*

$$\sum_{\theta \in \Theta^a, a' \neq a} \pi_\theta r^*(a' | \theta) (u(a, \theta) - u(a', \theta)) = \sum_{\theta \in \Theta \setminus \Theta^a} \pi_\theta r^*(a | \theta) (u(a^*(\theta), \theta) - u(a, \theta)) + o(\varepsilon^2). \quad (12)$$

That is, the expected loss from failing to choose a when a is optimal equals the expected loss from choosing a when a is suboptimal, up to a term of order of ε^2 .

The corollary formalizes the error-management idea in that the agent makes relatively costly errors relatively rarely. A particularly simple result applies in the binary setting. When errors are rare, the loss associated with each error type is inversely related to its frequency.

To see how the assumption that mistakes are rare simplifies the second-thought-free condition (8), recall that the restart of the decision process impacts the agent's payoff only if $a_1 \neq a_2$. If the rule r^* makes optimal recommendations with a high probability and two recommendations are in conflict, then it is very likely that exactly one of them is optimal. This provides a simplification compared to the general case, in which both recommendations may happen to be suboptimal. See Appendix A.1 for the derivation of (12).

Our model is very permissive on the set \mathcal{P} of the feasible primitive decision processes that need not be known to an outside observer. Note that the second-thought-free condition can be checked without any knowledge of this set simply based on the knowledge of the joint distribution of (θ, a) and the agent's preference u . Sometimes the joint distribution of (θ, a) is accessible in the data, but the analyst is ignorant of u . In this case, the second-thought-free condition allows the analyst to partially identify the agent's utility function from the choice data. As in the example, assume that the analyst observes the joint distribution of (θ, a) induced by the prior π_θ and the optimal choice rule $r^*(a | \theta)$. Such data can be generated either by repeated choices of an individual facing i.i.d. payoff states, or by one-shot decisions of a population with idiosyncratic types. The

¹⁰To derive (11), rewrite (8) as below and eliminate the first summand on each side:

$$\begin{aligned} & \mathbb{E}_\alpha[u(a_1, \theta) | a_1 = a_2] \Pr_\alpha(a_2 = a_1) + \mathbb{E}_\alpha[u(a_1, \theta) | a_1 \neq a_2] \Pr_\alpha(a_2 \neq a_1) = \\ & \mathbb{E}_\alpha[u(a_2, \theta) | a_1 = a_2] \Pr_\alpha(a_2 = a_1) + \mathbb{E}_\alpha[u(a_2, \theta) | a_1 \neq a_2] \Pr_\alpha(a_2 \neq a_1). \end{aligned}$$

second-thought-free condition is equivalent to

$$\sum_{\theta \in \Theta} \pi_{\theta} u(a, \theta) r^*(a | \theta) = \sum_{\theta \in \Theta, a' \in A} \pi_{\theta} u(a', \theta) r^*(a' | \theta) r^*(a | \theta),$$

for each action a chosen with positive probability. When the analyst observes N actions, the second-thought-free condition provides $N - 1$ linear, generically independent conditions imposed on the utility function.¹¹ The partial utility identification is based on our formalization of the error-management idea. The constrained-optimal choice rule must avoid costly types of errors. Indeed, if an action a would often lead to a costly loss, then the agent who is about to terminate with a would benefit from a second thought.

5 Behavioral applications

This section presents four behavioral effects generated by our model. We demonstrate the first three effects—confirmation bias, speed-accuracy complementarity, and overweighting of rare events—in the binary setting from Section 2. To this end, we start this section with the complete solution of the binary setting. The fourth effect—salience of distinct states—will be presented in a generalization of the binary setting.

As in Section 2, $a \in A = \{0, 1\}$, θ is drawn from prior $\pi \in \Delta(\Theta)$, and $\Theta = A$. To avoid a trivial case, we assume that neither action is dominant. Then, without loss of generality, $u(a, \theta) = u_{\theta} > 0$ if $a = \theta$ and $u(a, \theta) = 0$ otherwise.

For expositional reasons, we restrict here the agent’s primitive processes to those that emerge from a use of a single binary information structure. This restriction is only pedagogical. In Section 6 we show that the solution extends to agents with nontrivial memory who can use a rich set of sophisticated decision processes that aggregate information over time.

In each run of the primitive decision process, the agent observes a binary signal x and chooses action $a = \sigma(x)$, $\sigma : X \rightarrow A$. Without loss of generality, we identify the signal, action, and the state spaces: $X = A = \Theta$, and we choose signal labels in such a way that the signal x satisfies the monotone likelihood ratio property: $\tilde{p}(1 | \theta) / \tilde{p}(0 | \theta)$ increases in θ , where $\tilde{p}(x | \theta)$ is the conditional signal distribution (this implicitly rules out non-informative signals). We assume that $\tilde{p}(x | \theta) > 0$ for all x and θ .

We show in Appendix B that the restriction to the binary signal space assumed here is without loss of generality. In the appendix, we consider an arbitrary finite signal space and prove that when the state space is binary, then regardless of the size of the action set, a solution exists in which the agent searches for at most two signal realizations and ignores all other signal values.

¹¹Summation of (8) over all a leads to an identity due to the law of iterated expectations. This reduces the number of the independent conditions by one.

The agent of this section can choose from four primitive processes

$$p(a, x | \theta; \sigma) = \begin{cases} \tilde{p}(x | \theta), & \text{if } \sigma(x) = a, \\ 0 & \text{otherwise,} \end{cases}$$

each one generated by one of the four available action strategies σ . We abuse the notation by letting $r(a | \theta; \beta, \sigma)$ denote the effective choice rule $r(a | \theta; \beta, p(\sigma))$ generated by the termination strategy β and the primitive decision process $p(\sigma)$. The agent chooses β and σ to maximize the expected payoff $U(\beta, \sigma) = \sum_{\theta} \pi_{\theta} u_{\theta} r(\theta | \theta; \beta, \sigma)$. The set of the feasible effective choice rules is compact and thus the solution exists.

Let $\sigma^*(x) = x$ be the identity function, and consider the agent who employs action strategy σ^* . By controlling the termination strategy β , the agent trades off the likelihoods $r(0 | 0; \beta, \sigma^*)$ and $r(1 | 1; \beta, \sigma^*)$ of the correct choice in the states 0 and 1, respectively. The next lemma clarifies the trade-off. Let us introduce a parameter that we dub *perceptual distance*:

$$d = \frac{\tilde{p}(1|1)\tilde{p}(0|0)}{\tilde{p}(0|1)\tilde{p}(1|0)},$$

the larger the perceptual distance is, the more reliably \tilde{p} discriminates between the two states. The monotone likelihood property of \tilde{p} implies $d > 1$. The perceptual distance is preserved under any termination strategy β :

Lemma 3. *The effective choice rule $r(a | \theta; \beta, \sigma^*)$ satisfies, for each positive β :*

$$\frac{r(0|0; \beta, \sigma^*)r(1|1; \beta, \sigma^*)}{r(0|1; \beta, \sigma^*)r(1|0; \beta, \sigma^*)} = d. \quad (13)$$

Equation (13) is a feasibility condition specifying effective choice rules available to the decision maker, and it holds for all, possibly suboptimal termination strategies. Omitted proofs are in Appendix. See Figure 2 for the set of feasible choice rules $r(\beta, \sigma^*)$.

The agent can equate her choice to the observed signal without a loss:¹²

Lemma 4. *The binary repeated-cognition problem has a solution in which σ is the identity function.*

From now on we restrict σ to the identity function and omit it from the notation for $r(a | \theta; \beta, \sigma)$.

We are now ready to solve the binary setting. The optimal effective choice rule $r(a | \theta; \beta^*)$ described by four unknowns is determined by four conditions: It satisfies two normalization conditions, the feasibility condition (13), and the second-thought-free condition (8). Let parameter $R = \frac{\pi_1 u_1}{\pi_0 u_0}$ measure the relative a priori attractiveness of action 1.

Proposition 3. *1. When $R \geq d$ then the agent always chooses action 1,*

¹²In particular, the identity action strategy σ^* allows the agent to generate a deterministic choice rule that always selects action a^* by setting $\beta_a = 0$ for $a \neq a^*$. Vice versa, when the optimal choice rule always selects a^* then there is a trivial multiplicity of solutions since then $\sigma(x) = a^*$ is optimal for any β .

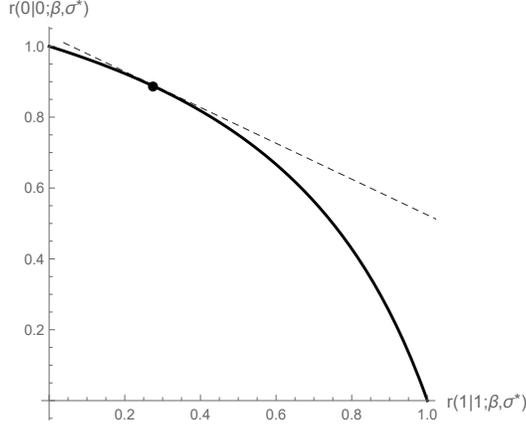


Figure 2: Thick curve: the set of the feasible effective choice rules when $\sigma(x) = x$; dashed line: tangential indifference line; dot: the optimal effective choice rule.

2. when $R \leq 1/d$ then the agent always chooses action 0,

3. when $R \in (1/d, d)$ then the agent chooses both actions with positive probabilities and

$$r^*(1 | 1) = \frac{dR - \sqrt{dR}}{(d-1)R}, \quad r^*(0 | 0) = \frac{d - \sqrt{dR}}{d-1}, \quad (14)$$

$$\frac{\beta_1^*}{\beta_0^*} = \frac{dR - \sqrt{dR} \tilde{p}(0 | 1)}{\sqrt{dR} - R \tilde{p}(1 | 1)}. \quad (15)$$

5.1 Confirmation bias

Psychologists distinguish at least three mechanisms leading to confirmation bias. (i) People search for evidence selectively, targeting the evidence type in accord with their priors, e.g. Nickerson (1998); (ii) they selectively memorize and recall the data supporting their priors, e.g. Oswald and Grosjean (2004); and (iii) they selectively interpret ambiguous evidence. Although the third mechanism is the one most commonly emphasised within economics, e.g. in Rabin and Schrag (1999) and Fryer et al. (2016), we focus on the first two mechanisms and argue that they are constrained-optimal.

Let us introduce $f_\theta = \sum_a \beta_a \tilde{p}(a | \theta)$ and interpret it as the decision rate in state θ ; it is the per-round probability that the agent in state θ terminates the decision-making. Thus, the response time in state θ is geometrically distributed with rate f_θ , and the expected response time is $1/f_\theta$.

Corollary 2. 1. When $R \in (1/d, d)$ then:

- (i) relative search intensity β_1^*/β_0^* for signal 1 increases with R ,
- (ii) for each state θ , the relative likelihood $\frac{r(1|\theta; \beta^*)}{r(0|\theta; \beta^*)}$ of action 1 increases with R ,
- (iii) relative decision rate, f_1/f_0 , increases with R .

2. When action 1 is a priori more attractive, $R \in (1, d)$, and the primitive information structure is symmetric, $\tilde{p}(1 | 1) = \tilde{p}(0 | 0)$, then the agent searches relatively more intensively for signal 1: $\beta_1^* > \beta_0^*$.

To illustrate the benefit of confirmation bias, consider symmetric rewards $u_1 = u_0 = 1$, symmetric primitive signal distribution $\tilde{p}(1 | 1) = \tilde{p}(0 | 0) > 1/2$, and asymmetric prior that favors state 1, $\pi_1 > \pi_0$. Consider in this setting the uniform information search with $\beta_0 = \beta_1 = 1$, and assume that the agent has observed the a priori less likely signal 0. Such an agent prefers to restart the decision-making to terminating with action 0, since if the new conclusion is $a = 1$, then the switch from action 0 to 1 will have been beneficial. This is because, conditional on the two conflicting signals, the a priori more common state 1 is relatively more likely. The agent benefits from a second thought whenever she receives the surprising recommendation, and thus will deviate from the uniform search in favor of the a priori likely signal 1.

Though our results apply when the agent or the selection pressures manage to find the optimal termination strategy, the above intuition suggests that the confirmation bias may appear even when the selection of the termination strategy is imperfect, since a partial deviation from the unbiased strategy in the direction of the confirmation bias is payoff improving.

Let us relate the above confirmation-bias effect to the choice of media outlets. Let each state of the world θ generate the population of attributes x with population shares $\tilde{p}(x | \theta)$. Each media outlet chooses an editorial strategy β and the reader of the outlet with editorial policy β observes one attribute generated by $r(x | \theta; \beta)$. As in Gentzkow and Shapiro (2006), our agent prefers outlets biased towards her prior belief; she prefers editorial policies that over-represent her a priori likely attribute. The source of this bias in Gentzkow and Shapiro is reputation: Their agents evaluate media outlets confirming their prior beliefs as relatively reliable information sources. In our case, all outlets are ex ante identical in that they have access to the same population of attributes and thus reputation does not play a role. Rather, in our case, demand for the prior-confirming outlets is driven by the information aggregation friction.¹³

5.2 Speed-accuracy complementarity

The binary setting generates the speed-accuracy complementarity effect—a stylized fact stating that delayed choices tend to be less accurate than speedy choices; see the psychology studies of Swensson (1972) and Luce (1986). We establish this effect in a setting with the symmetric primitive signal distribution, $\tilde{p}(1|1) = \tilde{p}(0|0) > 1/2$.

For the sake of concreteness and without loss of generality, we let $R > 1$. That is, the agent's

¹³Another microfoundation of confirmation bias in media choice is found in Calvert (1985) and Suen (2004). These models view editorial policy as a map from a rich signal space to a recommendation of action to the reader. The reader's preferred editorial policy is her best-response mapping from the rich signal to the action, given her prior belief. In our case the editorial policy is not a coarsening but rather a non-uniform draw of an attribute. However, the delegation-based intuition of this literature applies to our case. The optimal media outlet chooses the same biased attribute draw as the reader would conduct herself. See Che and Mierendorff (2016) for a recent constrained-optimal media-bias model with information-processing microfoundations.

objective imposes a larger weight $u_1\pi_1$ on the likelihood of the correct choice $r(1|1; \beta)$ in state 1 than the weight $u_0\pi_0$ imposed on $r(0|0; \beta)$.

Corollary 3. 1. *The decision rate is higher in the a priori more important state: $f_1 > f_0$.*

2. *The conditional probability $\Pr_{\pi, r^*}(a = \theta|t)$ of the correct choice generated by the prior π and the optimal decision process $r^*(a, t|\theta)$ decreases with response time t .*

Given its relatively lower weight in the agent’s objective, the constrained-optimal probability of the correct choice is lower in the a priori unimportant state than in the important one. A long response time indicates that the agent has repeatedly arrived at the a priori unattractive action recommendation and has hesitated to terminate. Hence, large t indicates a high likelihood of the a priori unimportant state. The longer the agent has hesitated, the likelier it is that she has encountered the unimportant state, in which she commits an error relatively often.

Fudenberg et al. (2015) have recently generated speed-accuracy complementarity in a sequential sampling model that features uncertainty about a binary payoff state and about the scale of the incentives the agent faces. The agent interprets long spells of inconclusive sequential sampling as an indication suggesting that the incentives are mild, and hence optimally terminates the process with relatively imprecise posteriors. The authors’ proposed mechanism differs from ours. While their agent is “giving up” after long delays, in our model the long delays indicate that the agent has encountered the surprising state for which her decision strategy is not well adapted. Che and Mierendorff (2016) propose yet another explanation for the speed-accuracy complementarity. Their agent learns from inconclusive signals arriving in a Poisson process, and the long delay in the arrival of the signal is counter-evidence to the signal’s value.

For settings with uniform prior, Statement 1 of Corollary 3 implies that response time is longer in the state of the world in which the agent’s incentives are milder. See Konovalov and Krajbich (2016) for a review of the empirical evidence that mild incentives generate long response times, and for an empirical strategy that exploits this effect to identify preferences from the response-time data only.

5.3 Overweighting of rare events

We now consider an agent forming a probability belief. The agent of this subsection has access to rich data that would allow her, in the absence of cognitive constraints, to identify the correct probability distribution by applying the law of large numbers. Our model implies that when the agent is unable to aggregate the data, she optimally overweights rare events, as in Kahneman and Tversky (1979).

The agent’s task is to form a belief about a random variable x with support $\{0, 1\}$. The distribution of x is uncertain and depends on the binary state of the world, $\theta \in \{0, 1\}$, drawn from the uniform prior. The probability of $x = 1$ in the state of the world θ is $\rho_\theta \in (0, 1)$ (and the probability of $x = 0$ is $1 - \rho_\theta$). The agent knows the parameters ρ_0, ρ_1 , and the distribution of θ , but not the realized state θ itself. The agent observes one draw of x and guesses the state

θ —and hence the probability ρ_θ . Her objective is to maximize the ex ante probability of the correct guess (i.e., the agent gets 1 if her guess a is equal to the state and 0 otherwise).¹⁴ She chooses search intensities β_x for $x \in \{0, 1\}$, observes $x = 1$ in state θ with effective probability $r(1 | \theta; \beta) = \beta_1 \rho_\theta / (\beta_1 \rho_\theta + \beta_0 (1 - \rho_\theta))$, and announces the state of the world $\sigma(x)$ where $\sigma : \{0, 1\} \rightarrow \{0, 1\}$ is her action strategy and x is the observed data point. She chooses β and σ optimally given her objective.

Let us assume, essentially without loss of generality, that $\rho_0 < \rho_1 < 1 - \rho_0$.¹⁵ Thus, event $x = 1$ is a priori rarer, and the distribution of x is more uncertain—closer to the uniform one—in state $\theta = 1$.

Corollary 4. *The optimal information search discriminates in favor of the rare event $x = 1$: $\beta_1^* > \beta_0^* > 0$, and the agent announces that the probability of $x = 1$ is high if and only if her observed data point is $x = 1$: $s(x) = x$.*

To gain the intuition, consider an agent who is forming her probability belief over a flight accident. The accident probability per flight in the safe state of the world is 10^{-6} , whereas it is 10^{-5} in the dangerous state of the world, and both states are a priori equally likely. The agent, being unable to aggregate data from many past flights, recalls one flight outcome and identifies the state of the world as dangerous if and only if she observes an accident.

Consider first an agent who draws the flight observation uniformly from the population of flights in the realized state of the world. Such an agent benefits from a “second thought” whenever she observes an uneventful flight: Either the second observed flight will be uneventful, in which case the second thought will have been inconsequential, or the redrawn flight will be eventful and the agent will switch her assessment from the safe to the dangerous state. Such a switch is beneficial, since conditional on two contradicting data points, the less certain state of the world—the one with the distribution closer to the uniform one—is more likely. In the above notation, when $\rho_0 < \rho_1 < 1 - \rho_0$ then the signal $x = 0$ —identified here with no accident—is less informative than the signal $x = 1$ —identified with accident. The agent benefits from targeting the more informative flight outcome.

For the assumed accident rates, the optimal ratio of the search intensities is $\beta_1^* / \beta_0^* \approx 316,000$; the agent prefers information sources that heavily over-represent rare accidents. This search-intensities ratio induces effective probabilities of the correct state identification $r(1|1; \beta^*) = r(0|0; \beta^*) \approx 0.76$ in both states of the world. Since the uniform information search identifies the correct state with a probability equal approximately to half, targeting significantly improves the payoff.

Steiner and Stewart (2016) provide an unrelated explanation for distorted probability perception that is based on an asymmetry of losses associated with under- and over-perception of probabilities. The explanation here is the opposite. We assume that the losses from the probability misperceptions are symmetric across the two types of error, and explain the bias by pointing out an asymmetry in the informativeness of different types of evidence.

¹⁴This agent is capable of distinguishing a low or high probability of the outcome $x = 1$ but does not engage in Bayesian updating that would lead to intermediate posterior beliefs.

¹⁵We can always achieve this by relabeling the states θ and the signals x , unless $\rho_0 = \rho_1$ or $\rho_0 = 1 - \rho_1$.

5.4 Saliency

Bordalo, Gennaioli, and Shleifer (2012) interpret saliency as directed attention focus and quote the popular work by Daniel Kahneman (2011):

“Our mind has a useful capability to focus on whatever is odd, different or unusual.”

The quote states a causal relation between the two features of the salient phenomena: These are (i) odd, different or unusual, and because of (i), people benefit from (ii) focusing their attention on such phenomena. Here, we confirm Kahneman’s intuition in a model. Our microfoundation of the saliency effect is related to the insight emerging in psychological research on visual saliency. Itti (2007) conceptualizes the visual saliency effect as attention allocation to a subset of the visual field that is “sufficiently different from its surroundings to be worthy of [one’s] attention”. Similarly, in our model, a payoff state is salient if it stands out sufficiently from similar states to be worthy of the focus of the agent’s information search.

The agent faces a perceptual task that requires her to recognize a random state θ . She is endowed with a primitive perception technology that generates a perception θ' . The primitive perception is informative but noisy: The perception θ' equals the true state θ with a high probability, but mistakes occur sometimes. We view the primitive perception technology as a black-box model of a physiological sensor that generates a noisy first impression of the state. We allow the agent to use the sensor repeatedly, and examine the terminal perception that results from optimal selective repetitions of the primitive perception process.

The agent chooses $a \in A = \Theta$, where $2 < |\Theta| < \infty$, and receives payoff $u(a, \theta) = 1$ if $a = \theta$ and $u(a, \theta) = 0$ if $a \neq \theta$. The prior is uniform. The set of feasible primitive processes is a singleton; the only available primitive process generates perception a in state θ with a probability $p(a | \theta) > 0$. We make two assumptions on p :

Symmetry: $p(\theta' | \theta) = p(\theta | \theta') := s(\theta, \theta')$ for all θ, θ' .

The symmetric function $s(\theta, \theta')$ captures the *similarity* of the two states θ and θ' .

Sufficient precision: $p(\theta | \theta) > p(\theta' | \theta)$ for all $\theta \neq \theta'$.

For two states $\theta_1 \neq \theta_2$, we say that θ_1 is *more distinct* than θ_2 if for each state $\theta_3 \neq \theta_1, \theta_2$, $s(\theta_1, \theta_3) < s(\theta_2, \theta_3)$. Suppose for illustration that the perceptual task involves recognition of a color uniformly drawn from a set {azure, indigo, red}. Say that the two shades of blue are similar, so that the agent confuses them in 10% of cases, $s(\text{azure}, \text{indigo}) = 0.1$, but $s(\theta, \text{red}) = 0.01$ for $\theta \in \{\text{azure}, \text{indigo}\}$. Then, the red state is more distinct than either of the two blue states.

Recall that the decision rate $f_\theta = \sum_{\theta' \in \Theta} \beta_{\theta'} p(\theta' | \theta)$ is the per round probability of termination.

Proposition 4. *Suppose that the optimal termination probabilities β_a are positive for all a . Then:*

1. *The agent’s decision rate in state θ is proportional to $p^{1/2}(\theta | \theta)$: there exists $c > 0$ such that*

$$f_\theta = c \times p^{1/2}(\theta | \theta).$$

2. If state θ_1 is more distinct than state θ_2 , then the optimal effective choice rule satisfies

$$r^*(\theta_1 | \theta_2) > r^*(\theta_2 | \theta_1).$$

Since the primitive perception technology is symmetric by assumption, the asymmetry in favor of the distinct state of the effective choice rule in Statement 2 is driven solely by the optimization of the termination strategy. To gain the intuition for the salience of the distinct states, consider a state θ^* that is similar to many other states. If the agent were using the uniform termination strategy, then the perception θ^* would be relatively uninformative, and hence the agent would benefit from “having a second thought” whenever the primitive process generates the indistinct perception θ^* . The optimal termination strategy involves repeating the primitive process with relatively high probability whenever the agent forms a perception of such an indistinct state, and this shifts the effective choice rule in favor of the distinct states.

Proposition 4 follows from the second-thought-free condition and from the symmetry of the primitive decision process. The proof in Appendix exploits a formal parallel between the problem studied and the balance condition for stationary distributions of Markov chains.

6 Sophisticated agents and other variants

We now examine the robustness of the second-thought-free condition to modifications of the baseline model. The first subsection deals with several variations of the agent’s level of sophistication. In three examples, we vary the coarseness of the agent’s termination strategy, ability to accumulate information across runs of the primitive decision process, and flexibility to discard a chosen subset of accumulated information. We show that all three examples can be accommodated by our baseline model and thus the optimal choice rules solving the examples are second-thought-free. If, additionally, the state and action spaces are binary then the characterisation of the constrained-optimal choice rule from Proposition 3 applies.

The second subsection examines agents who face an opportunity cost of time. A modification of the second-thought-free condition applies: The agent terminating the optimal process is indifferent between terminating and rerunning the decision-making, internalizing the expected opportunity cost of the new run.

6.1 Sophisticated decision processes

The three examples presented below seemingly violate our baseline model in that the agent can condition the termination of her decision-making on information finer than the mere action recommendation, and can aggregate information across rounds. We show that these sophisticated decision processes are accommodated by our baseline model using appropriate sets of the primitive processes. In all three examples, as elsewhere, the agent maximizes payoff $u(a, \theta)$ where a and θ are from finite sets A , Θ , and θ is drawn from a prior π .

Example 1 (side information). Our baseline model restricts the agent from conditioning terminations on the side information x collected in the last run of the primitive process $p(a, x | \theta)$. The agent in this example can condition the termination on all information collected in the last round.

To keep the discussion simple, we assume that the agent has access to a single statistical experiment that generates a signal $x \in X$ with probability $\mu(x | \theta) > 0$ in each state θ . The agent chooses a vector $\gamma = (\gamma_x)_{x \in X}$ of the termination probabilities and the action strategy $\sigma : X \rightarrow A$. She terminates her decision-making after each run of the information acquisition process with the probability γ_x and chooses action $\sigma(x)$ upon the termination. She reruns the experiment μ with the complementary probability $1 - \gamma_x$. The pair (γ, σ) generates the effective choice rule $p(\gamma, \sigma)$ where each $p(a | \theta; \gamma, \sigma)$ is the probability that the decision-making terminates with action a in the state θ . Let $\Gamma_{si} = [0, 1]^{|X|} \setminus \{(0, \dots, 0)\}$ be the set of the feasible termination strategies γ , S_{si} be the set of all mappings $X \rightarrow A$, and $\mathcal{P}_{si} = \{p(\gamma, \sigma) : \gamma \in \Gamma_{si}, \sigma \in S_{si}\}$ be the set of the choice rules feasible to the agent in this example. The agent selects a feasible choice rule to maximize her ex ante expected payoff:

$$\max_{p \in \mathcal{P}_{si}} \sum_{\theta \in \Theta, a \in A} \pi_\theta u(a, \theta) p(a | \theta). \quad (16)$$

This agent is more sophisticated than the one from our baseline model, because the baseline agent can condition her termination probabilities only on the action $\sigma(x)$, and not on x . Yet, our main result applies:¹⁶

Corollary 5. *The choice rule solving Problem (16) is second-thought-free.*

Proof of Corollary 5. We embed Problem (16) into our baseline model from Section 3. That is, we consider an agent who can choose any primitive choice rule from \mathcal{P}_{si} , and a termination strategy $\beta = (\beta_a)_{a \in A} \in B$ that conditions on the recommended action a . A rule $p \in \mathcal{P}_{si}$ and β generate an effective choice rule $r(p, \beta)$ according to (4). Let $\mathcal{R}_{si} = \{r(p, \beta) : p \in \mathcal{P}_{si}, \beta \in B\}$ be the set of all choice rules that can be generated by this embedding. A rule solving

$$\max_{r \in \mathcal{R}_{si}} \sum_{\theta \in \Theta, a \in A} \pi_\theta u(a, \theta) r(a | \theta) \quad (17)$$

is second-thought-free by Proposition 2.

Corollary 5 follows from the fact that Problems (16) and (17) are identical because the sets of feasible rules \mathcal{P}_{si} and \mathcal{R}_{si} coincide. Indeed, for any γ, σ and β , $r(p(\gamma, \sigma), \beta) = p(\gamma', \sigma) \in \mathcal{P}_{si}$, where $\gamma'_x = \gamma_x \beta_{\sigma(x)}$. \square

Example 2 (imperfect information aggregation). We now relax the agent's inability to aggregate information across the runs of the primitive process by endowing her with a finite set of memory states that she can use to represent the histories of the decision process. One interpretation of this

¹⁶In fact, the second-thought-free condition holds in this example in a stronger form than reported in Corollary 5: The agent is indifferent between terminating and a second thought conditional on each signal realization x . When an outsider observes the agent's signals and the signal space is larger than the action space, then this increases the number of conditions identifying the agent's preferences.

example is that the agent keeps track of the number of elapsed rounds, can count up to a finite number, and her count of the rounds is allowed to be stochastic.

The following setting builds on Hellman and Cover (1970) and Wilson (2014). As in the previous example, the agent is endowed with one statistical experiment $\mu(x | \theta)$ and, additionally, with a finite set M of the memory states m . After each run of the experiment μ the agent either terminates or continues with decision-making. If the agent continues, then she transitions from the current memory state to a new memory state and reruns the statistical experiment $\mu(x | \theta)$. That is, the agent selects a (generalization of the) termination strategy: $\gamma : M \times X \rightarrow \Delta(M \cup \{\mathfrak{t}\})$, where $\gamma(m' | m, x)$ is the probability that the agent in the memory state m who has observed signal x in the last run of the experiment μ continues with the decision-making and transitions to the memory state m' , and $\gamma(\mathfrak{t} | m, x)$ is the probability that such an agent terminates. We restrict $\gamma(\mathfrak{t} | m, x)$ to be positive for all m, x to ensure that the decision-making almost surely eventually terminates. The terminating agent chooses action $\sigma(m, x)$ that depends both on the current memory state and on the signal acquired in the last run of μ . The agent starts the decision-making in the memory state m_0 . A pair γ, σ induces a θ -dependent Markov chain over the memory states that eventually terminates with choice $\sigma(m, x)$, where m is the last memory state and x is the last signal received. Let $p(a | \theta; \gamma, \sigma)$ be the effective choice rule, and let \mathcal{P}_{ia} be the set of all stochastic choice rules $p(\gamma, \sigma)$ that this agent can generate. She selects the choice rule from \mathcal{P}_{ia} to maximize her ex ante expected payoff.

Corollary 6. A choice rule solving the problem in this example is second-thought-free.

Proof of Corollary 6. As in the previous example, we embed \mathcal{P}_{ia} into our baseline model. Let $\mathcal{R}_{ia} = \{r(p, \beta) : p \in \mathcal{P}_{ia}, \beta \in B\}$ be the set of all choice rules that can be generated by this embedding, where the effective choice rule $r(p, \beta)$ is generated according to (4). As in the previous example, it suffices to prove that $\mathcal{P}_{ia} = \mathcal{R}_{ia}$. This is indeed so because $r(p(\gamma, \sigma), \beta) = p(\gamma', \sigma)$ where $\gamma'(\mathfrak{t} | m, x) = \gamma(\mathfrak{t} | m, x)\beta_{\sigma(m, x)}$, $\gamma'(m_0 | m, x) = \gamma(m_0 | m, x) + \gamma(\mathfrak{t} | m, x)(1 - \beta_{\sigma(m, x)})$, and $\gamma'(\tilde{m} | m, x) = \gamma(\tilde{m} | m, x)$ for all $\tilde{m} \in M \setminus \{m_0\}$. \square

Wilson (2014) differs from this example mainly in that she assumes exogenous termination probabilities. By adding optimization over the terminations to the model of Wilson, we gained the partial characterization of the optimal choice rule without fully solving the problem: One can conclude that the optimal choice rule is second-thought-free without analyzing the optimal use of the memory states.

Example 3 (partial forgetting). The agent in this example comprehends up to $N > 1$ signals sequentially drawn from the statistical experiment $\mu(x | \theta)$. She can discard some of the accumulated data at any interim stage of the decision process. Unlike in the baseline model, she is not restricted to discarding all her information.

Let $H = \bigcup_{l=0}^N X^l$ be the set of the signal histories h of length $|h| \leq N$. The agent at a history h can (i) terminate her decision-making; (ii) discard some of the information accumulated;¹⁷ or,

¹⁷The formalization below specifies that the agent can discard an arbitrary number of the most recent signals. The

if $|h| < N$, (iii) acquire a new signal. (i) An agent terminating at h chooses action $\sigma(h)$. (ii) An agent who discards some information transitions to a truncation h' of her current history h . (iii) An agent who acquires a new signal transitions to a history hx , where x is the signal drawn from $\mu(x | \theta)$.

More precisely, the decision-making is governed by a pair of mappings $\gamma : H \times \Theta \rightarrow \Delta(H \cup \{\mathfrak{t}\})$ and $\sigma : H \rightarrow A$, where $\gamma(h' | h, \theta)$ stands for the probability that the agent at history h in state θ continues decision-making and transitions to h' , and $\gamma(\mathfrak{t} | h, \theta)$ is the probability of termination at history h in state θ . The mapping γ is constrained to satisfy 1. $\gamma(h' | h, \theta)$ is independent of θ if h' is a truncation of h , 2. $\gamma(\mathfrak{t} | h, \theta)$ is independent of θ , 3. $\frac{\gamma(hx | h, \theta)}{\gamma(hx' | h, \theta)} = \frac{\mu(x | \theta)}{\mu(x' | \theta)}$, 4. $\gamma(h' | h, \theta) = 0$ unless h' is a truncation of h , or $h' = hx$ for some $x \in X$ and $|hx| \leq N$. Constraints 1 and 2 require the agent to condition the decision to discard information or to terminate only on her current history independently of the state. Constraint 3 allows the agent to expand her information set only by running the experiment $\mu(x | \theta)$. Constraint 4 restricts each step of information acquisition to one draw from $\mu(x | \theta)$ or to a partial discarding of the accumulated information. Let $p(a | \theta; \gamma, \sigma)$ be the probability that the agent who employs (γ, σ) terminates with action a in the state θ . The agent chooses γ and σ to maximize her ex ante expected payoff.

Corollary 7. A choice rule solving the problem in this example is second-thought-free.

Proof of Corollary 7. Let \mathcal{P}_{pf} be the set of all choice rules feasible in this example. As in the previous examples, we show that embedding \mathcal{P}_{pf} into our baseline model does not expand the set of feasible rules. Indeed, $r(p(\gamma, \sigma), \beta) = p(\gamma', \sigma)$, where $\gamma'(\mathfrak{t} | h, \theta) = \gamma(\mathfrak{t} | h, \theta)\beta_{\sigma(h)}$, $\gamma'(\emptyset | h, \theta) = \gamma(\emptyset | h, \theta) + \gamma(\mathfrak{t} | h, \theta)(1 - \beta_{\sigma(h)})$, and $\gamma'(\tilde{h} | h, \theta) = \gamma(\tilde{h} | h, \theta)$ for $\tilde{h} \in H \setminus \{\emptyset\}$. \square

Solving for the constrained optimal rules in the three examples requires, beyond the second-thought-free condition, an analysis of the set of the feasible rules. When the state and action spaces are binary then, despite that sophisticated agents can construct many choice rules, their behavior complies with Proposition 3 derived for non-sophisticated agent with access to only a single primitive signal.

Corollary 8. *For each Example 1-3, there exists d such that for all incentives $u_0, u_1 > 0$ and all priors π the characterization of the constrained optimal choice rule from Proposition 3 applies with d and $R = \frac{\pi_1 u_1}{\pi_0 u_0}$.*

For a primitive choice rule $p(a | \theta)$, let $d_p = \frac{p(1|1)p(0|0)}{p(1|0)p(0|1)}$ be the associated perceptual distance. The proof is based on the observation that a primitive rule with higher perceptual distance dominates that with a lower distance because it leads to set of feasible effective rules that are more precise in both payoff states. The agent thus chooses the primitive rule with the maximal feasible perceptual distance, and Proposition 3 applies with respect to that primitive rule.

Proof of Corollary 8. Consider two primitive rules p', p such that $d_{p'} > d_p$. Then, the effective rule $r(\beta, p)$ cannot be a solution of the repeated cognition problem for any β , any $u_0, u_1 > 0$ and any result continues to hold if forgetting is not chronologically restricted.

π . To see this, note that for any β , there exists β' such that $r(1 | 1; p', \beta') = r(1 | 1; p, \beta)$. Since, by Lemma 3,

$$\frac{r(1 | 1; p', \beta')r(0 | 0; p', \beta')}{r(1 | 0; p', \beta')r(0 | 1; p', \beta')} = d_{p'} > d_p = \frac{r(1 | 1; p, \beta)r(0 | 0; p, \beta)}{r(1 | 0; p, \beta)r(0 | 1; p, \beta)},$$

and thus $r(0 | 0; p', \beta') > r(0 | 0; p, \beta)$. Therefore, the rule $r(p', \beta')$ yields a higher payoff than $r(p, \beta)$. Let \mathcal{P}_e be the set of the rules feasible in Example $e \in \{si, iia, pf\}$. Then, the constrained optimal rule in each example equals the solution of the setting with a single primitive choice rule that maximizes d_p in P_e . \square

In particular, all the applications derived in the binary setting—the confirmation bias, the speed-accuracy complementarity, and the overweighting of rare events—extend to sophisticated agents.

6.2 Opportunity cost

Our baseline model abstracts from the opportunity cost of time in that the agent is only concerned with how the repetitions of the decision process affect the correlation of her choice with the state. When a linear time cost is added to the agent’s objective, a natural modification of the second-thought-free condition continues to hold.

As in the baseline model, the agent in this subsection chooses a primitive process $p \in \mathcal{P}$ and a termination strategy $\beta \in B$, which generate the effective process $r(a, t | \theta; p, \beta)$ according to (4). She maximizes the expected gross payoff less the opportunity cost:

$$\max_{p \in \mathcal{P}, \beta \in B} \sum_{\theta \in \Theta, a \in A, t \in \mathbb{N}} \pi_\theta r(a, t | \theta) (u(a, \theta) - ct), \quad (18)$$

where $c > 0$ is the unit time cost.

Recall that the expected response time in state θ is $1/f_\theta$ where $f_\theta = \sum_{a \in A} \beta_a p(a | \theta)$ is the decision rate. The optimal rule r^* solving Problem (18) satisfies the modified second-thought-free condition:

$$\mathbb{E}_\alpha [u(a_1, \theta) | a_1 = a] \geq \mathbb{E}_\alpha \left[u(a_2, \theta) - \frac{c}{f_\theta} | a_1 = a \right], \quad (19)$$

for all actions a chosen with positive probability under r^* . The expectation is, as before, with respect to the joint distribution $\alpha(\theta, a_1, a_2) = \pi_\theta r^*(a_1 | \theta) r^*(a_2 | \theta)$. If $\beta_a \in (0, 1)$, then (19) holds with equality. The result differs from the second-thought-free condition in that the agent internalizes the additional opportunity cost of the second thought, and in that, when the agent always terminates upon receiving recommendation a —when $\beta_a = 1$ —then she may strictly prefer termination. Condition (19) follows directly from the first-order condition of Problem (18) with respect to β_a .

Let us, in the binary setting with $|A| = |\Theta| = 2$, illustrate how (19) can be used to identify preferences from choice data in the presence of the opportunity cost. Compared to the baseline

setting without the opportunity cost, the outside observer faces two difficulties: (i) She needs to find out whether the modified second-thought-free condition (19) holds with equality, for which she needs to identify from the data whether $\beta_a < 1$; and (ii) to identify the preferences, the outsider needs to identify the decision rates f_θ . Both these difficulties can be addressed if the outsider observes the response times jointly with the choices.

Let $m(a, t, \theta)$ denote the joint distribution of the terminal action, response time, and state. Assume that the outsider knows the agent's opportunity cost c and observes the marginal distribution $m(a, t) = \sum_\theta m(a, t, \theta)$ of the terminal action a and of the response time t . The outsider need not observe the prior state distribution π_θ nor the correlation between θ and (a, t) . Since

$$m(a, t) = \sum_\theta \pi_\theta r^*(a | \theta) f_\theta (1 - f_\theta)^{t-1},$$

the outsider can identify the choice rule $r^*(a | \theta)$, the decision rate f_θ and the prior π_θ in both states from the observation of $m(a, t)$.¹⁸ As $r^*(a | \theta) = \beta_a p(a | \theta) / f_\theta$, the outsider has also identified ratios $p(a | \theta) / p(a | \theta')$ and thus she has also identified the primitive choice rule $p(a | \theta)$, and therefore also the termination probabilities β_a . If one of the identified termination probabilities is interior, then (19) holds with equality, and thus the outsider has obtained an identifying condition on the agent's utility function that differs from the second-thought-free condition only by factor $E[c/f_\theta | a_1 = a]$ known to the outsider.

7 Summary and discussion

Agents, who cannot comprehend all facts available to them, benefit from selective attention. We show that agents can implement a targeted information search in a process that resembles the natural phenomenon of hesitation. Like a hesitating person, the agent can, conditional on the action contemplated, decide whether she implements the action or whether she will have a second thought, and run the cognition process once more. Such hesitation can be productive, despite consisting of repetitions of the same stochastic cognition process. By conditioning the probability of the repetition on the conclusion of the process, the agent controls the correlation of her terminal choice and the payoff state. The optimal decision process arising in our model exhibits natural hesitation patterns: The agent will have second thoughts—that is, she will repeat her cognition—whenever the expected payoff for the currently favored choice is inferior to the expected payoff for continuing decision-making. At optimum, the agent terminating the decision-making must be indifferent between terminating with the currently contemplated action, and repeating the process. In a sense, the condition formalizes the concept of a reasonable doubt. Abstracting from many considerations such as information aggregation, a jury deciding a trial under the common law should be, if using the optimal decision procedure, indifferent between declaring a verdict and

¹⁸Let $\underline{f} = \min_\theta f_\theta$ and $\underline{\theta} = \arg \min_\theta f_\theta$. The outsider can identify \underline{f} , $\pi_{\underline{\theta}}$, and $r^*(a | \underline{\theta})$ from the tail behavior of $m(a, t)$ because $m(a, t) \approx \pi_{\underline{\theta}} r^*(a | \underline{\theta}) \underline{f}_{\underline{\theta}} (1 - \underline{f}_{\underline{\theta}})^{t-1}$ for large t . This allows for the identification of $r^*(a | \theta)$, f_θ and of π_θ in the other state $\theta \neq \underline{\theta}$.

announcing a hung jury and initiate retrial.

Let us conclude by reviewing the limitations of our main result. The central assumption—the ability of the agent to freely repeat her decision process—may fail for several reasons. One reason is that the agent may only have access to a limited data set that constrains her to a finite number of repetitions of the primitive decision process, making the optimal termination strategy non-stationary. Another complication arises if the outcomes of distinct runs of the same cognition process are not conditionally independent as assumed in our model; this may arise if some cognition errors are systematic and are likely to emerge in distinct repetitions of the cognition. We conjecture that the second-thought-free condition holds in such a case, with the agent internalizing the correlations between the cognition runs.

A Omitted Proofs

A.1 Proofs for Section 4

Proof of Corollary 1. Let $\hat{u}(a, \theta)$ stand for $\pi_\theta u(a, \theta)$. Recall that $\Theta^a = \{\theta \in \Theta : a \in \arg \max_a u(a, \theta)\}$, and let $\Theta_a = \Theta \setminus \Theta^a$, and $A_\theta^* = \arg \max_a u(a, \theta)$.

The second-thought-free condition (8) implies:

$$\sum_{\theta \in \Theta} \hat{u}(a, \theta) r(a|\theta) = \sum_{\theta \in \Theta, a' \in A} \hat{u}(a', \theta) r(a'|\theta) r(a|\theta).$$

Rewrite this as

$$\sum_{\theta \in \Theta^a} \hat{u}(a, \theta) r(a|\theta) + \sum_{\theta \in \Theta_a} \hat{u}(a, \theta) r(a|\theta) = \sum_{\theta \in \Theta_a, a' \in A_\theta^*} \hat{u}(a', \theta) r(a'|\theta) r(a|\theta) + \sum_{\theta \in \Theta^a, a' \in A} \hat{u}(a', \theta) r(a'|\theta) r(a|\theta) + o(\varepsilon^2), \quad (20)$$

where the term $o(\varepsilon^2)$ stands for all terms in which both a' and a are suboptimal in state θ . Rewrite the right-hand side as

$$o(\varepsilon^2) + \sum_{\theta \in \Theta_a, a' \in A_\theta^*} \hat{u}(a', \theta) r(a'|\theta) r(a|\theta) + \sum_{\theta \in \Theta^a, a' \neq a} \hat{u}(a', \theta) r(a'|\theta) r(a|\theta) + \sum_{\theta \in \Theta^a} \hat{u}(a, \theta) r^2(a|\theta),$$

and subtract the last term from both sides of (20) to get:

$$\begin{aligned} \sum_{\theta \in \Theta^a} \hat{u}(a, \theta) r(a|\theta) (1 - r(a|\theta)) + \sum_{\theta \in \Theta_a} \hat{u}(a, \theta) r(a|\theta) = \\ \sum_{\theta \in \Theta_a, a' \in A_\theta^*} \hat{u}(a', \theta) r(a'|\theta) r(a|\theta) + \sum_{\theta \in \Theta^a, a' \neq a} \hat{u}(a', \theta) r(a'|\theta) r(a|\theta) + o(\varepsilon^2). \end{aligned}$$

Rearrange the left-hand side as:

$$\sum_{\theta \in \Theta^a, a' \neq a} \hat{u}(a, \theta) r(a|\theta) r(a'|\theta) + \sum_{\theta \in \Theta_a} \hat{u}(a, \theta) r(a|\theta).$$

Next, we use that $r(a|\theta)r(a'|\theta) = r(a'|\theta) + o(\varepsilon^2)$ when a is the unique optimal action at θ and a' is not optimal at θ , and similarly, for a suboptimal at θ , $\sum_{a' \in A_\theta^*} r(a'|\theta)r(a|\theta) = r(a|\theta) + o(\varepsilon^2)$. We get:

$$\sum_{\theta \in \Theta^a, a' \neq a} \hat{u}(a, \theta)r(a'|\theta) + \sum_{\theta \in \Theta_a} \hat{u}(a, \theta)r(a|\theta) = \sum_{\theta \in \Theta_a} \hat{u}(a^*(\theta), \theta)r(a|\theta) + \sum_{\theta \in \Theta^a, a' \neq a} \hat{u}(a', \theta)r(a'|\theta) + o(\varepsilon^2).$$

Rearrangement of this gives (12). □

A.2 Proofs for Section 5

Proof of Lemma 3. For any positive β ,

$$\frac{r(1|1; \beta, \sigma^*)r(0|0; \beta, \sigma^*)}{r(0|1; \beta, \sigma^*)r(1|0; \beta, \sigma^*)} = \frac{\frac{\beta_1 \tilde{p}(1|1)}{\sum_a \beta_a \tilde{p}(a|1)} \frac{\beta_0 \tilde{p}(0|0)}{\sum_a \beta_a \tilde{p}(a|0)}}{\frac{\beta_0 \tilde{p}(0|1)}{\sum_a \beta_a \tilde{p}(a|1)} \frac{\beta_1 \tilde{p}(1|0)}{\sum_a \beta_a \tilde{p}(a|0)}}} = \frac{\tilde{p}(1|1)\tilde{p}(0|0)}{\tilde{p}(0|1)\tilde{p}(1|0)} = d.$$

□

Proof of Lemma 4. The statement is trivial when the solution involves only one action chosen with probability 1. Accordingly, assume that the problem has only a solution in which both actions are chosen with positive probabilities. Then, Lemma 3 implies that the effective choice rule $r(a | \theta; \beta, \sigma^*)$ inherits the monotone likelihood ratio property from \tilde{p} .

Assume for contradiction that the only solution involves action strategy $\sigma^{**}(x) = 1 - x$. Then the payoff difference between the optimal choice rule and the choice rule that always selects $a = 1$ must be positive:

$$\pi_0 u_0 r(1 | 0; \beta, \sigma^*) + \pi_1 u_1 r(0 | 1; \beta, \sigma^*) - \pi_1 u_1 > 0,$$

where we have used $r(a | \theta; \beta, \sigma^*) = r(1 - a | \theta; \beta, \sigma^{**})$. Similarly, the payoff difference between the optimal rule and the rule that always selects $a = 0$ is positive:

$$\pi_0 u_0 r(1 | 0; \beta, \sigma^*) + \pi_1 u_1 r(0 | 1; \beta, \sigma^*) - \pi_0 u_0 > 0.$$

The last two inequalities imply

$$\frac{r(0 | 1; \beta, \sigma^*)}{r(0 | 0; \beta, \sigma^*)} > \frac{\pi_0 u_0}{\pi_1 u_1} > \frac{r(1 | 1; \beta, \sigma^*)}{r(1 | 0; \beta, \sigma^*)},$$

which violates the monotone likelihood ratio property of $r(a | \theta; \beta, \sigma^*)$, as needed for the contradiction. □

Proof of Proposition 3. The agent's objective is linear with respect to the choice rule $r(\beta)$. Thus, the optimal choice rule is the point of tangency of the set of the feasible effective rules and of the indifference line; see Figure 2. The slope $\frac{dr(0|0;\beta)}{dr(1|1;\beta)}$ is decreasing in $r(1|1; \beta)$ and attaining value $-1/d$

for $r(1|1; \beta) = 0$, and value $-d$ for $r(1|1; \beta) = 1$. Thus, when $R < 1/d$ or $R > d$, the problem has a corner solution as specified in Statements 1 and 2 of the proposition.

When $R \in (1/d, d)$, the optimal choice rule r^* satisfies the feasibility condition (13), the second-thought-free condition (8) (applied to action $a = 1$):

$$\pi_1 u_1 r^*(1 | 1) = \pi_0 u_0 r^*(0 | 0) r^*(1 | 0) + \pi_1 u_1 r^*(1 | 1) r^*(1 | 1),$$

and two normalization conditions $\sum_a r^*(a | \theta) = 1$, for $\theta \in \{0, 1\}$. These four conditions jointly imply the explicit solution for the optimal choice rule in (14). The expression (15) for β_1^*/β_0^* follows from (14) and the condition $\frac{r(1|\theta; \beta)}{r(0|\theta; \beta)} = \frac{\beta_1 \tilde{p}(1|\theta)}{\beta_0 \tilde{p}(0|\theta)}$. \square

Proof of Corollary 2. Comparative statics results (i) and (ii) follow from the signs of the derivatives of the explicit solutions (14) and (15). For the relative response rate, the monotone likelihood property of \tilde{p} implies that

$$\frac{f_1}{f_0} = \frac{\beta_0^* \tilde{p}(0 | 1) + \beta_1^* \tilde{p}(1 | 1)}{\beta_0^* \tilde{p}(0 | 0) + \beta_1^* \tilde{p}(1 | 0)}$$

increases in β_1^*/β_0^* .

For the last statement of the corollary, since β_1^*/β_0^* increases in R , it suffices to show that $\beta_1^*/\beta_0^* = 1$ when $R = 1$ and the primitive information structure is symmetric. Indeed, when $R = 1$, then by (15),

$$\frac{\beta_1^*}{\beta_0^*} = \sqrt{d} \frac{\tilde{p}(0 | 1)}{\tilde{p}(1 | 1)} = \sqrt{\frac{\tilde{p}(0 | 0) \tilde{p}(0 | 1)}{\tilde{p}(1 | 1) \tilde{p}(1 | 0)}} = 1,$$

where the last equality follows from the symmetry of \tilde{p} . \square

Proof of Corollary 3. $\beta_1^* > \beta_0^*$ by Corollary (2), since $\pi_1 u_1 > \pi_0 u_0$. Recall that $f_\theta = \beta_1^* \tilde{p}(1 | \theta) + \beta_0^* \tilde{p}(0 | \theta)$ denote the probability of termination per each round in state θ , and that the response time t in the state θ is geometrically distributed with the decision rate f_θ : $\Pr(t | \theta) = f_\theta (1 - f_\theta)^t$. Since $\tilde{p}(1|1) = \tilde{p}(0|0) > 1/2$ and $\beta_1^* > \beta_0^*$, the decision rate is higher in state 1 than in state 0: $f_1 > f_0$. Thus, the likelihood ratio $\Pr(t | \theta = 1)/\Pr(t | \theta = 0)$ decreases with t , and hence $\Pr(\theta = 1 | t)$ decreases in t . The fact that $\beta_1^* > \beta_0^*$, and the symmetry of \tilde{p} implies that the probability of the correct choice is larger in state 1 than in state 0:

$$r(1|1; \beta^*) = \frac{\beta_1^* \tilde{p}(1|1)}{\beta_0^* \tilde{p}(0|1) + \beta_1^* \tilde{p}(1|1)} > \frac{\beta_0^* \tilde{p}(0|0)}{\beta_0^* \tilde{p}(0|0) + \beta_1^* \tilde{p}(1|0)} = r(0|0; \beta^*).$$

Since $\Pr(a = \theta | t) = \Pr(\theta = 1 | t) r(1|1; \beta^*) + \Pr(\theta = 0 | t) r(0|0; \beta^*)$, the result 2 obtains. \square

Proof of Corollary 4. The belief formation problem studied is a special case of our binary setting with the primitive information structure $\tilde{p}(x | \theta) = \rho_\theta$ if $x = 1$, $\tilde{p}(x | \theta) = 1 - \rho_\theta$ if $x = 0$ and with equally a priori attractive actions, $R = 1$. Since $\rho_0 < \rho_1$, the setting satisfies the monotone likelihood property, and thus by Lemma 4, there exists a solution with $\sigma(x) = x$. Since $R = 1 \in (1/d, d)$, Proposition 3 implies that the agent's behavior is stochastic, both β_0^* and β_1^* are

positive, and the ratio of the search intensities β_1^*/β_0^* satisfies (15). Since $R = 1$, (15) simplifies to

$$\frac{\beta_1^*}{\beta_0^*} = d^{1/2} \frac{\tilde{p}(0|1)}{\tilde{p}(1|1)} = \left(\frac{\tilde{p}(0|1)\tilde{p}(0|0)}{\tilde{p}(1|1)\tilde{p}(1|0)} \right)^{1/2} = \left(\frac{(1-\rho_1)(1-\rho_0)}{\rho_1\rho_0} \right)^{1/2}.$$

The inequality $\beta_1^* > \beta_0^*$ follows from $\rho_0 < \rho_1 < 1 - \rho_0$. \square

The next result is an auxiliary lemma used in the proof of Proposition 4.

Lemma 5. *Suppose that termination probabilities β_θ are positive for all θ . Then, the optimal effective choice rule r^* satisfies for any pair of states θ, θ' :*

$$r^*(\theta | \theta)r^*(\theta' | \theta) = r^*(\theta | \theta')r^*(\theta' | \theta'). \quad (21)$$

Condition (21) is a strong version of the second-thought-free condition; it is equivalent to the indifference condition

$$E_\alpha[u(a, \theta) | a_1 = a, a_2 = a'] = E_\alpha[u(a', \theta) | a_1 = a, a_2 = a'] \text{ for all } a, a',$$

which requires indifference for any given pair of the action recommendations, whereas the second-thought-free condition (8) only requires indifference under the uncertainty regarding the second recommendation.

Proof of Lemma 5. The optimal effective choice rule satisfies the second-thought-free condition (8), equivalent to:

$$r^*(\theta | \theta) = \sum_{\theta' \in \Theta} r^*(\theta | \theta')r^*(\theta' | \theta') \text{ for all } \theta \in \Theta,$$

which after two algebraic steps gives:

$$r^*(\theta | \theta)(1 - r^*(\theta | \theta)) = \sum_{\theta' \neq \theta} r^*(\theta | \theta')r^*(\theta' | \theta') \text{ for all } \theta \in \Theta,$$

$$\sum_{\theta' \neq \theta} r^*(\theta | \theta)r^*(\theta' | \theta) = \sum_{\theta' \neq \theta} r^*(\theta' | \theta')r^*(\theta | \theta') \text{ for all } \theta \in \Theta.$$

The last system of equations is formally equivalent to the system of balance conditions for a Markov chain. To see this, consider an ergodic Markov chain with transition probabilities from θ to θ' equal to $r^*(\theta' | \theta)$. The balance condition for the stationary distribution $\mu(\theta)$ of this chain is

$$\sum_{\theta' \neq \theta} \mu(\theta)r^*(\theta' | \theta) = \sum_{\theta' \neq \theta} \mu(\theta')r^*(\theta | \theta') \text{ for all } \theta \in \Theta,$$

and thus $r^*(\theta | \theta)$ is proportional to the ergodic probability $\mu(\theta)$ of the state θ for the chain with transition probabilities $r^*(\theta' | \theta)$.

Recall that if a Markov chain with transition probabilities $m(\theta' | \theta)$ is reversible, then its stationary distribution $\mu(\theta)$ satisfies detailed balance conditions

$$\mu(\theta)m(\theta' | \theta) = \mu(\theta')m(\theta | \theta') \text{ for all } \theta \neq \theta'.$$

Thus, it suffices to prove that the probabilities $r^*(\theta' | \theta)$ constitute a reversible Markov chain.

Recall that a Markov chain $m(\theta' | \theta)$ is reversible if and only if it satisfies the Kolmogorov criterion, which requires for all sequences of states $\theta_1, \theta_2, \dots, \theta_n$,

$$\frac{m(\theta_2 | \theta_1)m(\theta_3 | \theta_2) \dots m(\theta_n | \theta_{n-1})m(\theta_1 | \theta_n)}{m(\theta_n | \theta_1)m(\theta_{n-1} | \theta_n) \dots m(\theta_2 | \theta_3)m(\theta_1 | \theta_2)} = 1. \quad (22)$$

The Markov chain with transition probabilities $p(\theta' | \theta)$ given by the primitive decision process satisfies the Kolmogorov criterion (22) because p is symmetric by assumption. Finally, for any positive termination strategy β , the effective choice rule $r(\theta' | \theta; \beta)$ satisfies the Kolmogorov criterion too. This is because $r(\theta' | \theta; \beta) = \frac{\beta_{\theta'} p(\theta' | \theta)}{\sum_{\tilde{\theta}} \beta_{\tilde{\theta}} p(\tilde{\theta} | \theta)}$, and when the expressions for $r(\theta' | \theta; \beta)$ are substituted into (22) then the terms $\beta_{\theta'}$ and the denominators cancel out, and hence

$$\frac{r(\theta_2 | \theta_1; \beta)r(\theta_3 | \theta_2; \beta) \dots r(\theta_1 | \theta_n; \beta)}{r(\theta_n | \theta_1; \beta)r(\theta_{n-1} | \theta_n; \beta) \dots r(\theta_1 | \theta_2; \beta)} = \frac{p(\theta_2 | \theta_1)p(\theta_3 | \theta_2) \dots p(\theta_1 | \theta_n)}{p(\theta_n | \theta_1)p(\theta_{n-1} | \theta_n) \dots p(\theta_1 | \theta_2)} = 1,$$

as needed. □

Proof of Proposition 4. Statement 1.: Lemma 5 implies for all θ, θ' :

$$r^*(\theta | \theta)r^*(\theta' | \theta) = r^*(\theta | \theta')r^*(\theta' | \theta').$$

Substituting (6) from Lemma 1 gives

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

Using symmetry of p we get

$$\frac{\sum_{\tilde{\theta}}\beta_{\tilde{\theta}}p(\tilde{\theta}|\theta')}{\sum_{\tilde{\theta}}\beta_{\tilde{\theta}}p(\tilde{\theta}|\theta)} = \frac{p^{1/2}(\theta'|\theta')}{p^{1/2}(\theta|\theta)},$$

as needed.

Statement 2.:

$$\frac{r^*(\theta_1|\theta_2)}{r^*(\theta_2|\theta_1)} = \frac{\beta_{\theta_1}p(\theta_1|\theta_2)}{\sum_{\tilde{\theta}}\beta_{\tilde{\theta}}p(\tilde{\theta}|\theta_2)} = \frac{\beta_{\theta_1}p(\theta_1|\theta_2)}{p^{1/2}(\theta_2|\theta_2)} = \frac{\beta_{\theta_1}p^{1/2}(\theta_1|\theta_1)}{\beta_{\theta_2}p^{1/2}(\theta_2|\theta_2)},$$

where we have used Lemma 1 in the first step, Statement 1 in the second step, and symmetry of p in the last step. Define $\hat{\beta}_{\theta} := \beta_{\theta}p^{1/2}(\theta | \theta)$. We need to prove that if θ_1 is more distinct than θ_2 ,

then $\hat{\beta}_{\theta_1} > \hat{\beta}_{\theta_2}$.

By Statement 1, $(\hat{\beta}_\theta)_\theta$ satisfy the system of linear equations:

$$\sum_{\theta'} D_{\theta'\theta} \hat{\beta}_{\theta'} = 1 \text{ for all } \theta,$$

where $D_{\theta'\theta} = \frac{p(\theta'|\theta)}{p^{1/2}(\theta'|\theta')p^{1/2}(\theta|\theta)}$. Let us show that if θ_1 is more distinct than θ_2 , then $D_{\theta_3\theta_1} < D_{\theta_3\theta_2}$ for all $\theta_3 \neq \theta_1, \theta_2$. This follows from $p(\theta_3|\theta_1) < p(\theta_3|\theta_2)$ and from symmetry of p :

$$p(\theta_1|\theta_1) = 1 - p(\theta_2|\theta_1) - \sum_{\theta_3 \neq \theta_1, \theta_2} p(\theta_3|\theta_1) > 1 - p(\theta_1|\theta_2) - \sum_{\theta_3 \neq \theta_1, \theta_2} p(\theta_3|\theta_2) = p(\theta_2|\theta_2),$$

and therefore,

$$D_{\theta_3\theta_1} = \frac{p(\theta_1|\theta_3)}{p^{1/2}(\theta_1|\theta_1)p^{1/2}(\theta_3|\theta_3)} < \frac{p(\theta_2|\theta_3)}{p^{1/2}(\theta_2|\theta_2)p^{1/2}(\theta_3|\theta_3)} = D_{\theta_3\theta_2}.$$

Thus,

$$D_{\theta_1\theta_1} \hat{\beta}_{\theta_1} + D_{\theta_2\theta_1} \hat{\beta}_{\theta_2} = 1 - \sum_{\theta_3 \neq \theta_1, \theta_2} D_{\theta_3\theta_1} \hat{\beta}_{\theta_3} > 1 - \sum_{\theta_3 \neq \theta_1, \theta_2} D_{\theta_3\theta_2} \hat{\beta}_{\theta_3} = D_{\theta_2\theta_2} \hat{\beta}_{\theta_2} + D_{\theta_1\theta_2} \hat{\beta}_{\theta_1}.$$

Using that $D_{\theta\theta} = 1$ and $D_{\theta\theta'} = D_{\theta'\theta}$, we have

$$\hat{\beta}_{\theta_1} + D_{\theta_2\theta_1} \hat{\beta}_{\theta_2} > \hat{\beta}_{\theta_2} + D_{\theta_2\theta_1} \hat{\beta}_{\theta_1}.$$

The assumption of sufficient precision of p and symmetry of p imply that $D_{\theta_2\theta_1} < 1$, and thus $\hat{\beta}_{\theta_1} > \hat{\beta}_{\theta_2}$, as needed. \square

B Binary state and large signal space

We demonstrate here that the restriction to the binary signal space in Section 5 is without loss of generality. $|X| > 2$, $|\Theta| = 2$, the action set A is of arbitrary finite size, and $\tilde{p}(x | \theta)$ denotes the primitive conditional signal distributions. As in the first example of Section 6.1, the agent can condition the termination on the received signal x : The agent chooses non-negative search intensities $\beta = (\beta_x)_{x \in X}$, observes a signal draw from the effective signal structure

$$r(x | \theta; \beta) = \frac{\beta_x \tilde{p}(x | \theta)}{\sum_{x' \in X} \beta_{x'} \tilde{p}(x' | \theta)},$$

and chooses action $a = \sigma(x)$, where $\sigma : X \rightarrow A$ is the action strategy. The agent chooses β and σ to maximize her ex ante expected payoff under the prior π_θ .

Proposition 5. *The problem of this section has a solution in which β_x is positive for at most two signals $x \in X$.*

Proof of Proposition 5. Assume that there exists a solution with β_x positive for $n > 2$ signals $x \in X$. We show that then there exists a solution with $n - 1$ positive signals. The proposition follows from the induction on n .

Let us prove the induction step. Let β be an optimal termination strategy, and let X' be the set of signals with positive β_x , and write shortly $r(x | \theta)$ for the effective signal distributions induced by β . Let $r(x) = \sum_{\theta} \pi_{\theta} r(x | \theta)$ be the unconditional effective probability of x . For $x \in X'$ let $q_x \in \Delta(\Theta)$ be the posterior belief upon receiving x ; $q_x(\theta) = \pi_{\theta} r(x | \theta) / r(x)$.

Since $|X'| > 2$ and the state space Θ is binary, there exists a signal $x^* \in X'$ such that q_{x^*} is in the convex hull of the posteriors q_x , $x \in X' \setminus \{x^*\}$. Let μ_x be the coefficients that decompose q_{x^*} into q_x , $x \in X' \setminus \{x^*\}$. That is, $\mu \in \Delta(X' \setminus \{x^*\})$ such that $q_{x^*} = \sum_{x \in X' \setminus \{x^*\}} \mu_x q_x$.

We will construct an alternative feasible effective information structure $\tilde{r}(x | \theta)$ with unconditional probabilities of x denoted by $\tilde{r}(x)$ and the posteriors $\tilde{r}(\theta | x)$ denoted by \tilde{q}_x such that:

$$\tilde{r}(x) = \begin{cases} r(x) + r(x^*)\mu_x & \text{if } x \in X' \setminus \{x^*\}, \\ 0 & \text{otherwise,} \end{cases} \quad (23)$$

and

$$\tilde{q}_x = q_x \text{ for all } x \in X' \setminus \{x^*\}. \quad (24)$$

Since $\tilde{r}(x | \theta)$ is more informative than $r(x | \theta)$ (in the sense of the Blackwell comparison), there exists a solution with this alternative feasible effective information structure $\tilde{r}(x | \theta)$, as needed for the induction step.

It remains to construct \tilde{r} . Note that if an effective information structure $r(x | \theta; \beta)$ is induced by some β , then for any vector of non-negative $\tilde{\beta}_x$, the information structure

$$\tilde{r}(x | \theta) = \frac{\tilde{\beta}_x r(x | \theta; \beta)}{\sum_{x'} \tilde{\beta}_{x'} r(x' | \theta; \beta)}$$

is also feasible.

We claim that if

$$\tilde{\beta}_x = \begin{cases} 1 + \frac{r(x^*)\mu_x}{r(x)} & \text{if } x \in X' \setminus \{x^*\}, \\ 0 & \text{otherwise,} \end{cases}$$

then the resulting \tilde{r} satisfies the properties (23) and (24). Check:

$$\begin{aligned}
\tilde{r}(x \mid \theta) &= \frac{\tilde{\beta}_x r(x \mid \theta)}{\sum_{x' \in X' \setminus \{x^*\}} \tilde{\beta}_{x'} r(x' \mid \theta)} \\
&= \frac{\tilde{\beta}_x r(x \mid \theta)}{\sum_{x' \in X' \setminus \{x^*\}} r(x' \mid \theta) + \sum_{x' \in X' \setminus \{x^*\}} \frac{r(x^*) \mu_{x'}}{r(x')} r(x' \mid \theta)} \\
&= \frac{\tilde{\beta}_x r(x \mid \theta)}{\sum_{x' \in X' \setminus \{x^*\}} r(x' \mid \theta) + \sum_{x' \in X' \setminus \{x^*\}} \frac{r(x^*) \mu_{x'}}{\pi_\theta} q(\theta \mid x')} \\
&= \frac{\tilde{\beta}_x r(x \mid \theta)}{\sum_{x' \in X' \setminus \{x^*\}} r(x' \mid \theta) + \frac{r(x^*)}{\pi_\theta} q(\theta \mid x^*)} \\
&= \frac{\tilde{\beta}_x r(x \mid \theta)}{\sum_{x' \in X' \setminus \{x^*\}} r(x' \mid \theta) + r(x^* \mid \theta)} \\
&= \tilde{\beta}_x r(x \mid \theta).
\end{aligned}$$

The property (23) holds since for all $x \in X' \setminus \{x^*\}$:

$$\tilde{r}(x) = \tilde{\beta}_x r(x) = \left(1 + \frac{r(x^*) \mu_x}{r(x)}\right) r(x) = r(x) + r(x^*) \mu_x.$$

To establish the property (24), check that for all $x \in X' \setminus \{x^*\}$:

$$\tilde{q}(\theta \mid x) = \frac{\pi_\theta \tilde{r}(x \mid \theta)}{\sum_{\theta' \in \Theta} \tilde{r}(x \mid \theta')} = \frac{\pi_\theta \tilde{\beta}_x r(x \mid \theta)}{\sum_{\theta' \in \Theta} \tilde{\beta}_x r(x \mid \theta')} = \frac{\pi_\theta r(x \mid \theta)}{\sum_{\theta' \in \Theta} r(x \mid \theta')} = q(\theta \mid x).$$

□

References

- Arrow, K. J., D. Blackwell, and M. A. Girshick (1949). Bayes and minimax solutions of sequential decision problems. *Econometrica* 17(3), 213–244.
- Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics* 127(3), 1243–1285.
- Calvert, R. L. (1985). The value of biased information: A rational choice model of political advice. *The Journal of Politics* 47(2), 530–555.
- Caplin, A. and M. Dean (2015). Revealed preference, rational inattention, and costly information acquisition. *The American Economic Review* 105(7), 2183–2203.
- Caplin, A., M. Dean, and J. Leahy (2017). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. Technical report, National Bureau of Economic Research.
- Che, Y.-K. and K. Mierendorff (2016). Optimal sequential decision with limited attention. *unpublished, Columbia University*.

- Compte, O. and A. Postlewaite (2012). Belief formation. PIER Working Paper 12-027, University of Pennsylvania.
- Fryer, R. G., P. Harms, and M. O. Jackson (2016). Updating beliefs when evidence is open to interpretation: Implications for bias and polarization.
- Fudenberg, D., P. Strack, and T. Strzalecki (2015). Stochastic choice and optimal sequential sampling. *unpublished, Harvard University*.
- Gabaix, X. and D. Laibson (2017). Myopia and discounting. Technical report, National Bureau of Economic Research.
- Gentzkow, M. and J. M. Shapiro (2006). Media bias and reputation. *Journal of Political Economy* 114(2), 280–316.
- Glynn, P. W. and D. L. Iglehart (1989). Importance sampling for stochastic simulations. *Management Science* 35(11), 1367–1392.
- Haselton, M. G. and D. M. Buss (2000). Error management theory: A new perspective on biases in cross-sex mind reading. *Journal of Personality and Social Psychology* 78(1), 81–91.
- Hébert, B. and M. Woodford (2016). Rational inattention with sequential information sampling. *unpublished, Stanford University*.
- Hellman, M. E. and T. M. Cover (1970). Learning with finite memory. *The Annals of Mathematical Statistics*, 765–782.
- Itti, L. (2007). Visual salience. *Scholarpedia* 2(9), 3327.
- Jehiel, P. (2017). Investment strategy and selection bias: An equilibrium perspective on overoptimism. *unpublished, Paris School of Economics*.
- Johnson, D. D., D. T. Blumstein, J. H. Fowler, and M. G. Haselton (2013). The evolution of error: Error management, cognitive constraints, and adaptive decision-making biases. *Trends in Ecology & Evolution* 28(8), 474–481.
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.
- Khaw, M. W., Z. Li, and M. Woodford (2017). Risk aversion as a perceptual bias. Technical report, National Bureau of Economic Research.
- Konovalov, A. and I. Krajbich (2016). Revealed indifference: Using response times to infer preferences. Technical report.
- Leung, T. K. (2017). Limited cognitive ability and selective information processing. *unpublished, Toulouse School of Economics*.
- Lieder, F., M. Hsu, and T. L. Griffiths (2014). The high availability of extreme events serves resource-rational decision-making. In *Proceedings of the Cognitive Science Society*, Volume 36.

- Luce, R. D. (1986). *Response times: Their role in inferring elementary mental organization*. Number 8. Oxford University Press on Demand.
- Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *The American Economic Review* 105(1), 272–298.
- Morris, S. and P. Strack (2017). The wald problem and the equivalence of sequential sampling and static information costs. *unpublished, Princeton University*.
- Netzer, N. (2009). Evolution of time preferences and attitudes toward risk. *The American Economic Review* 99(3), 937–955.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology* 2(2), 175.
- Oliveira, H., T. Denti, M. Mihm, and K. Ozbek (2017). Rationally inattentive preferences and hidden information costs. *Theoretical Economics* 12(2), 621–654.
- Oswald, M. E. and S. Grosjean (2004). Confirmation bias. *Cognitive illusions: A handbook on fallacies and biases in thinking, judgement and memory* 79.
- Piccione, M. and A. Rubinstein (1997). On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20(1), 3–24.
- Rabin, M. and J. L. Schrag (1999). First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics* 114(1), 37–82.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review* 85(2), 59.
- Rayo, L. and G. S. Becker (2007). Evolutionary efficiency and happiness. *Journal of Political Economy* 115(2), 302–337.
- Redish, A. D. (2016). Vicarious trial and error. *Nature reviews. Neuroscience* 17(3), 147.
- Rich, E. L. and J. D. Wallis (2016). Decoding subjective decisions from orbitofrontal cortex. *Nature neuroscience* 19(7), 973.
- Robson, A. J. (2001). The biological basis of economic behavior. *Journal of Economic Literature* 39(1), 11–33.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Steiner, J. and C. Stewart (2016). Perceiving prospects properly. *The American Economic Review* 106(7), 1601–31.
- Steiner, J., C. Stewart, and F. Matějka (2017). Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica* 85(2), 521–553.
- Suen, W. (2004). The self-perpetuation of biased beliefs. *The Economic Journal* 114(495), 377–396.

- Swensson, R. G. (1972). The elusive tradeoff: Speed vs accuracy in visual discrimination tasks. *Perception & Psychophysics* 12(1), 16–32.
- Wald, A. (1947). *Sequential analysis*. John Wiley & Sons.
- Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica* 82(6), 2257–2294.
- Zhong, W. (2017). Optimal dynamic information acquisition. *unpublished, Columbia University*.