

ECONOMIC AGENTS AS IMPERFECT PROBLEM SOLVERS*

COSMIN ILUT AND ROSEN VALCHEV

We develop a novel bounded rationality model of imperfect reasoning as the interaction between automatic (System 1) and analytical (System 2) thinking. In doing so, we formalize the empirical consensus of cognitive psychology using a structural, constrained-optimal economic framework of mental information acquisition about the unknown optimal policy function. A key result is that agents reason less (more) when facing usual (unusual) states of the world, producing state- and history-dependent behavior. Our application is an otherwise standard incomplete-markets model with no a priori behavioral biases. The ergodic distribution of actions and beliefs is characterized by endogenous learning traps, where locally stable state dynamics generate familiar regions of the state space within which behavior appears to follow memory-based heuristics. This results in endogenous behavioral biases that have many empirically desirable properties: the marginal propensity to consume is high even for unconstrained agents, hand-to-mouth status is more frequent and persistent, and there is more wealth inequality than in the standard model. *JEL Codes:* D83, D91, E21, E71, C11.

I. INTRODUCTION

Cognitive psychology describes human reasoning as the result of an interplay between automatic intuitive thought and deliberative analytical thinking. That is, when reasoning through a given issue or problem, intuitive ideas—often driven by associations with past situations which the agent recalls—tend to “automatically” spring to our minds, while deeper insights are usually the outcome of conscious deliberation. Psychologists describe this

*We are deeply grateful to Andrei Shleifer for numerous thoughtful discussions; and to our discussants George-Marios Angeletos, Paolo Bonomolo, Tarek Hassan, and Luigi Iovino; as well as Isaac Baley, Ryan Chahrour, Ian Dew-Becker, Petru Ilut, Greg Kaplan, Spencer Kwon, Guido Lorenzoni, Filip Matějka, David Min, Kristoffer Nimark, Philipp Sadowski, Todd Sarver, Mirko Wiederholt, seminar and conference participants at Bank of Finland, CERGE, CEU, DNB, EEA Congress, Green Line Macro Meetings, NBER EFBEM and EFCE Groups, New Approaches for Modeling Expectations in Economics, Northwestern, Richmond Fed, Society for Economic Dynamics and Computing in Economics and Finance, TMF Workshop, University of Wisconsin, Virtual Macro Seminar, Wake Forest University, and the 13th Workshop on Theoretical and Experimental Macro for their helpful comments. This project is supported by NSF Awards SES-1824367 and SES-1824405.

© The Author(s) 2022. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

The Quarterly Journal of Economics (2022), 1–50. <https://doi.org/10.1093/qje/qjac027>. Advance Access publication on July 1, 2022.

duality in terms of the metaphor of an automatic System 1 type of thinking that relies on intuitive associations, and a System 2 that captures our innate ability to think abstractly and analytically, but is cognitively costly (Stanovich and West 2000; Evans 2008; Kahneman 2011). Moreover, this literature emphasizes that as “cognitive misers,” people economize mental resources and thus the costly analytical System 2 is not always engaged. Rather, it is triggered only when one does not feel confident enough in their intuitive, System 1 solution to the problem at hand. To psychologists, this interaction is in many ways crucial to understanding cognitive biases; however, it is largely absent from economic models.¹

This article formalizes the psychological notion of an interactive dual-reasoning process in an economic framework by modeling reasoning as costly mental information processing about optimal behavior.² Specifically, agents perfectly observe all objective payoff-relevant details of the economic environment, captured by the relevant state variable y_t , but face subjective uncertainty about the policy function $c^*(y)$, which characterizes optimal actions and behavior given the observed state. We model reasoning as the costly process of thinking through the specifics of the (known) situation y_t and drawing conclusions about the unknown optimal action $c^*(y_t)$. Our key methodological insight is to represent the reasoning process as an as if nonparametric learning of the unknown function $c^*(y)$, which allows us to tractably capture the key interactive nature of the reasoning process emphasized by psychologists.

I.A. Dual-Reasoning Model

In each period t , the agent can engage in costly deliberation over the question “What is my best course of action given current circumstances?” This generates a noisy mental signal η_t about the unknown optimal action $c^*(y_t)$. This signal represents a new insight about the agent’s problem, and serves as our as if

1. As put by Evans (2019, 409): “It is now becoming clearer that cognitive biases result largely from false feelings of confidence which perhaps allow the cognitive miser to stick with default intuitions.” In economics, Shleifer (2012, 1084) notes on the dual-process narrative that “turning metaphors into models remains a critical challenge.”

2. More generally, the interest in modeling scarce cognitive resources in economics is long-standing (Simon 1955, 1956). See Sent and Klaes (2005) for a conceptual history of bounded rationality and Conlisk (1996), Rubinstein (1998), Todd and Gigerenzer (2003), and DellaVigna (2009) for reviews.

representation of the analytical System 2 thinking. In addition, agents have a memory database $\{\eta^{t-1}, y^{t-1}\}$ of past deliberations and insights and the situations in which those occurred.³ Agents use their database to produce a nonparametric, beginning-of-period estimate $\mathbb{E}(c^*(y_t)|\eta^{t-1}, y^{t-1})$ of the optimal action. This estimate endogenously puts more weight on past signals derived in circumstances y_{t-k} similar to the current state y_t , thus displaying associative memory. This initial estimate is also our representation of the automatic System 1 thinking—it requires no new deliberation effort and relies on associative memory, as often emphasized by the psychological literature (Tulving 1972; Gennaioli and Shleifer 2010; Kahneman 2011; Kahana 2012).⁴

The agent is “resource-rational” and chooses the precision of any new deliberation signal η_t by trading off the resulting reduction in uncertainty over the optimal action and a cognitive cost proportional to the amount of information carried in the new signal.⁵ Through standard signal extraction results, the beginning-of-period estimate (System 1) and the new signal (System 2) get weighted in the resulting time- t inference of the optimal action $\mathbb{E}(c^*(y_t)|\eta^t, y^t)$. This conditional expectation is the overall outcome of the agent’s (dual) reasoning process and guides the agent’s actions.

A key property of our framework is that due to the nonparametric nature of the learning problem, the subjective uncertainty over the optimal action is state- and history-dependent. In particular, for state realizations y_t close to past situations y_{t-k} which the agent has encountered and deliberated about, the beginning-of-period estimate is already quite precise, as the agent has accumulated a number of useful ideas about how to handle similar situations in the past. As a result, in such familiar situations, agents have a high confidence in the intuition that springs automatically in their minds (System 1) and hence might optimally choose not to incur the cost of active deliberation (System 2) this period.

3. Throughout, a superscript t denotes the history of outcomes up to and including period t .

4. Tulving (2002) and Gershman and Daw (2017) survey the evidence on associative memory as a critical psychological feature. For example, Plonsky, Teodorescu, and Erev (2015) and Bornstein et al. (2017) use experiments and neuroimaging to document how the current decision is influenced more by a past choice that occurred in a circumstance more similar to the current one.

5. The notion that agents are “resource-rational” and exhibit behavior that is the outcome of limited but “appropriate” deliberation, has been advocated as early as Simon (1976).

Specifically, in our model, active deliberation is engaged only when the subjective uncertainty over the beginning-of-period estimate of the optimal action, $\text{Var}(c^*(y_t)|\eta^{t-1}, y^{t-1})$, is high enough. In such sufficiently novel situations agents are not confident in the accuracy of their immediate sense of optimal behavior, prompting them to actively deliberate over their situation and generate a new signal η_t to further inform their thinking. This captures the key psychological evidence that active deliberation is indeed engaged only when there is sufficient “conflict” or uncertainty in the automatic, System 1 type of thinking (Evans 2007b; De Neys 2014; Stanovich 2018).⁶ Otherwise, active deliberation (System 2) is not engaged, and behavior is guided by the agent’s immediate intuition (System 1).

Overall, our representation of the dual-reasoning process effectively follows a default-interventionist scheme (coined by Evans 2007b, and present in many related forms, e.g., Kahneman and Frederick 2002; Stanovich 2011; Bago and De Neys 2017), where the automatic, System 1 type of thinking provides a default estimate of optimal behavior that is further revised by the deliberative System 2 reasoning when engaged.

I.B. Qualitative Implications

Due to this interaction, and the resulting selective engagement of deliberative reasoning, our model can generate persistent behavioral mistakes. In particular, in our framework analytical thinking (System 2) could remain unaware of a mistaken, but subjectively confident intuition (System 1), and thus not jump in to correct it, even if this intuition is not objectively accurate. This phenomenon is actually often intuitively described in psychology. For example, Kahneman (2011, 24) states that since it is only “when System 1 runs into difficulty, it calls on System 2 to support more detailed and specific processing,” then “biases cannot always be avoided, because System 2 may have no clue to the error” (28).

Because our formalization of these intuitive ideas rests on the combination of nonparametric learning and costly cognition,

6. In our framework, a high $\text{Var}(c^*(y_t)|\eta^{t-1}, y^{t-1})$ precisely formalizes the notion that System 1 has many conflicting ideas. See also Thompson, Turner, and Pennycook (2011) and Thompson et al. (2013), who document that the trigger for analytical thought is a lack of subjective “feeling of rightness” in the automatic guess.

our model has the specific implication that active reconsideration of decision rules is only triggered in novel situations, and otherwise behavior is habitual and driven by established intuition. We show that this basic implication connects to evidence that appears across many economic contexts, including managerial actions (Louis and Sutton 1991; Tyre and Orlikowski 1994), health and environmental choices (Verplanken et al. 2008; Wood, Tam, and Witt 2005), labor market (Lee and Mitchell 1991; Holtom et al. 2005; DellaVigna et al. 2017), and financial decisions (Liu 2021).

Furthermore, we show that our dual-reasoning mechanism generates systematic mistakes in behavior, even in the absence of ex ante biases in updating or initial beliefs.⁷ This outcome arises when the economic environment features an endogenous state y_t , and thus the agent's actions affect the state evolution, generating feedback between the states which the agent encounters and her beliefs about optimal behavior. This can lead to learning traps, where established beliefs (System 1) imply a behavior that endogenously keeps the state dynamics stable in a familiar region of the state space. As a result, encountering new situations becomes unlikely, active deliberation (System 2) stops being triggered, and the initial System 1 beliefs are perpetuated. Such stable beliefs are not necessarily close to the unknown optimal policy, however, and this consistent deviation is a novel source of behavioral biases.

I.C. Consumption-Savings Application

The specific observable implications for behavior depend on the economic environment and how actions influence the state evolution. To showcase this, and more broadly the applicability of our framework, we incorporate our reasoning friction into an otherwise standard Aiyagari (1994) setting. There we show how learning traps predict two types of endogenously stable consumption behavior, which jointly and parsimoniously rationalize a number of empirical consumption puzzles.

One type of stable behavior is characterized by an excessively high belief about the level of optimal consumption near the borrowing constraint. Agents with such beliefs are dissaving on

7. To illustrate this cleanly, we abstract from primitive biases in reasoning in the main text. On the one hand, this also allows us to relate to a psychology literature (e.g., Evans and Stanovich 2013; De Neys 2014; Stanovich 2018) arguing that System 1 and 2 are not inherently biased in all cases. On the other hand, we discuss how our framework can accommodate a wide range of ex ante biases without fundamental changes.

average, and thus their stable point in wealth is the constraint itself, as their wealth tends to drift down and the constraint acts as an absorbing barrier. In turn, such agents' wealth is rarely far from the constraint, and this lack of variation in circumstances makes it unlikely that active deliberation is triggered and their consumption behavior reconsidered. To an outside observer, agents in this kind of a learning trap habitually exhibit an excessively high consumption level, and they are also persistently poor and close to the borrowing constraint.

This first type of stable behavior allows our model to match the dual empirical puzzle that (i) a high fraction of households have near-zero net wealth and that (ii) those households also remain in such a hand-to-mouth (HtM) situation persistently (Aguiar, Bils, and Boar 2020). These facts are challenging for standard models where precautionary saving motives make HtM status both rare and temporary. This behavior can also rationalize the undersaving of agents who know they have a relatively high income only temporarily—for example, severance pay and expiring unemployment benefits (Ganong and Noel 2019; Farrell et al. 2020; Gerard and Naritomi 2021). In our model, before the actual income drops, the situation is still familiar to agents, no new active reasoning is triggered, and so they maintain their usual relatively high consumption. This appears as a puzzling lack of preparation for the future, when the resulting low buffer stock leads to a consumption cut.

The second type of stable behavior is characterized by a consumption policy belief $\hat{c}_i(y)$ that intersects the permanent income hypothesis (PIH) policy from below, at a wealth level $\bar{y}_i$ away from the constraint. Thus, the agent's policy estimate is steeper than the PIH policy in the neighborhood of the intersection point $\bar{y}_i$. This implies that for state realizations $y_{i,t} > \bar{y}_i$, the agent dis-saves because a consumption level higher than the PIH generates a negative drift in wealth, and vice-versa for $y_{i,t} < \bar{y}_i$. Hence, this relatively elastic perception of the optimal policy stabilizes state dynamics around $\bar{y}_i$, which in turn can establish a learning trap and perpetuate such beliefs.⁸ A key characteristic of agents in this second type of learning trap is that they appear to be liquid, unconstrained, and yet display a relatively high marginal

8. This result also highlights an important subtlety—in our model there is inertia in the default mode of behavior, that is, unchanged beliefs over the optimal policy function, but not necessarily inertia in actions.

propensity to consume (MPC), which helps rationalize another puzzling empirical fact.⁹

Last, our mechanism cautions against policy makers treating agents' imperfect behavior as a fixed phenomenon that can be easily exploited (à la [Lucas 1976](#) critique). In an illustrative policy experiment, we showcase a simple extension where a policy intervention can increase subjective uncertainty, and thus lead agents to actively deliberate and in turn largely abandon their habitual actions. This significantly lowers the policy's effectiveness.

1.D. Related Modeling Approaches

We share a broad interest of modeling mistakes in decision making with a long tradition of studying the imperfect perception of the relevant state of the world—ranging from the rational-inattention literature inspired by [Sims \(1998, 2003\)](#) where the attention choice is optimal (see [Wiederholt 2010](#); [Gabaix 2019](#); [Mackowiak, Matejka, and Wiederholt forthcoming](#) for surveys), to work inspired by neuroscience evidence on the imperfect perception of stimuli (e.g., [Girshick, Landy, and Simoncelli 2011](#); [Wei and Stocker 2015](#)). That literature assumes the agents know the policy function that maps their (imperfect) beliefs about the state into optimal actions, whereas we model cognitive uncertainty about the optimal action, given beliefs about the state (i.e., uncertainty about the policy function). Moreover, two key results of our application—the large fraction of persistently HtM agents and the high MPC for the unconstrained agents—differentiate the observable implications of our model.¹⁰

In modeling uncertainty over the value-maximizing action, we connect to a growing literature on Bayesian cognitive

9. While liquidity frictions (e.g., [Kaplan and Violante 2014](#)) can produce high MPCs for rich and liquidity-constrained agents, this typically falls short in explaining the evidence for liquid wealthy agents, as documented by [Kuang \(2018\)](#), [Olafsson and Pagel \(2018\)](#), [Fagereng, Holm, and Natvik \(2021\)](#), and [McDowall \(2020\)](#). This shortcoming has also spurred recent interest in behavioral models that can lead to high MPCs even for such rich and liquid agents (e.g., [Lian 2020](#); [Bianchi, Ilut, and Saijo 2021](#); [Boutros 2021](#)).

10. In such alternative approaches agents typically tend to a unique wealth steady state away from the constraint and consumption underreacts to shocks, as these are imperfectly observed ([Luo 2008](#)). We also note that in our model, learning over $c^*(y)$ and over the optimal savings policy $a^*(y) = y - c^*(y)$ are equivalent, since the total budget y available for the consumption-savings allocation is known. In contrast, such equivalence generally does not hold in models of imperfect information ([Reis 2006](#)).

noise (e.g., Gabaix and Laibson 2017; Enke and Graeber 2019; Woodford 2020; Frydman and Jin 2022). In these approaches, agents update their cognitive prior, typically *ex ante* biased, with an unbiased noisy cognitive signal of the unknown optimal action. In our model, we effectively generate an endogenously biased prior (System 1), through the interaction of associative memory and the endogenous choice of the cognitive signal precision. As such, we further relate and contribute to a broad set of existing work as follows.

First, we relate to alternative modeling approaches focusing on context similarity, like Gilboa and Schmeidler (1995, 2001), Mullainathan (2002), Bordalo, Gennaioli, and Shleifer (2020), and Cerigioni (2021). In our model, a surprise in the encountered situation, that is, a sufficiently different state y_t compared with past states y_{t-k} , leads to both a default System 1 guess that is less anchored by past signals, a property shared with these approaches, and to a novel implication of heightened urgency to engage the analytical but costly reasoning of System 2.

Second, our dual-reasoning interaction predicts endogenous learning traps, differentiating this work from a related literature that appeals to misspecified priors or biased updating to generate persistence in mistaken beliefs (e.g., Berk 1966; Rabin and Schrag 1999; Schwartzstein 2014; Heidhues, Köszegi, and Strack 2018; Gagnon-Bartsch, Rabin, and Schwartzstein 2020). Instead, our model can generate persistent and systematic errors even without any *ex ante* misspecified learning.

Finally, our formalization of System 1 as retrospective and memory-based and of System 2 as prospective and generating new insights also connects to a related duality and interaction in the reinforcement learning literature between model-free and model-based evaluation of the optimal action (e.g., the review by Gershman and Daw 2017; Barberis and Jin 2021 for applications to economic decisions). Complementing such reward-based evaluations, we emphasize noisy subjective learning through human thinking, endogenously triggered by high state- and history-dependent subjective uncertainty (or “conflict”). This conditional trigger implication is furthermore consistent with neuroscience views of cognition initiated by the co-activation of competing responses (e.g., Botvinick, Cohen, and Carter 2004; Pennycook 2017) and the related arbitration between the habitual, cognitively cheap model-free mode and more costly model-based mode (e.g., Lee, Shimojo, and O’Doherty 2014; Kool, Gershman,

and Cushman 2017; and the economic model of Landry, Webb, and Camerer 2021). More broadly, our model contributes to a growing literature emphasizing “resource rationality” in the allocation of costly cognition (e.g., Gershman, Horvitz, and Tenenbaum 2015; Griffiths, Lieder, and Goodman 2015; Shenhav et al. 2017).

II. GENERAL FRAMEWORK

To fix ideas, consider a general recursive economic framework with discrete time indexed by t , where the agent chooses an action c_t and is perfectly aware of all payoff-relevant details of the economic environment. The decision problem can be summarized as

$$(1) \quad V(y_t) = \max_{c_t \in B(y_t)} [u(y_t, c_t) + \beta \mathbb{E}_t V(y_{t+1})],$$

where y_t collects all the relevant state variables, both exogenous and endogenous, and follows the law of motion $y_{t+1} = F(y_t, c_t, \nu_{t+1})$, where ν_{t+1} are innovations to the exogenous states.

The primitives of the economic environment are the per period reward function (e.g., utility) $u(y_t, c_t)$, the budget set $B(y_t)$ that defines the set of currently feasible actions c_t , and the law of motion of the state F . The agent knows these primitives, perfectly observes y_t , and realizes that the recursive optimization problem implicitly defines an optimal policy function $c^*: Y \rightarrow C$, which characterizes the optimal action c_t for any possible state realization y_t .

The standard approach in the literature is to assume outright that the agent knows the optimal policy $c^*(y)$. However, this ignores the fact that solving for $c^*(y)$ is a nontrivial functional optimization problem that can be quite complex and is often not fully tractable even to highly trained economists. Put another way, figuring out optimal behavior takes significant cognitive effort—in terms of comparing the instantaneous value $u(c_t, y_t)$ across all feasible actions in the possibly large budget set $B(y_t)$, as well as the broader question of how any given policy function $c(y)$ affects the present discounted value of future utility $\mathbb{E}_t(V(y_{t+1}))$.¹¹

11. The challenge of comparing a large set of feasible options is familiar from static problems and motivates, for example, the satisficing principle of Simon (1976). Dynamic problems present the additional difficulty of evaluating whole paths of actions, or in recursive terms, an action plan with all possible future con-

In practice, reasoning through all possible options, especially in dynamic settings, is not trivial and thus optimal behavior is not immediately obvious.

To take such real-world difficulties into account, we assume that agents are uncertain about the optimal policy function $c^*(y)$, even though they are perfectly aware of the objective description of their problem in [equation \(1\)](#). In turn, agents can engage costly cognitive resources to reason through their current situation (i.e., problem at hand) and produce (mental) information about the unknown optimal behavior as characterized by $c^*(y)$. Thus, we effectively represent the reasoning process as an as-if problem of learning the optimal policy function $c^*(y)$, given knowledge of the state, as opposed to learning the state y_t itself, as in a large literature on imperfect information (see [Mackowiak, Matejka, and Wiederholt forthcoming](#) for a survey).

II.A. The Dual Reasoning Process as a Learning Problem

We ground the specific modeling ingredients of this learning problem in cognitive psychology evidence, which generally describes reasoning as the outcome of a “dual-system” process driven by the interplay between intuitive thought (typically labeled System 1) and deliberative, analytical thought (typically labeled System 2). We design our framework to tractably, yet flexibly capture three well-documented properties of System 1 and System 2 type of thinking.

First, psychological evidence suggests that two key features of the intuitive System 1 thinking are that it is (i) automatic, in the sense of not requiring active deliberation, or “controlled attention” (e.g., [Stanovich 2011](#); [Evans and Stanovich 2013](#); and (ii) often based on associative memory, meaning that it relies on insights derived in the past, but in a way that depends on the similarity between the current situation and the context in which those past insights were obtained ([Tulving 1972](#); [Kahana 2012](#); [Bornstein et al. 2017](#)). Second, in contrast, the analytical, System 2 type of thinking is the result of active deliberation; it produces novel insights, but is mentally costly ([Evans 2007a, 2019](#); [Kahneman 2011](#); [Evans and Stanovich 2013](#)). Third, the costly analytical reasoning is only occasionally engaged, as triggered by the outcome of the auto-

tingencies. Characterizing the optimal policy $c^*(y)$ in dynamic models is generally a challenging functional equation problem (see [Judd 1998](#); [Bertsekas 2019](#)).

matic intuitive thought, that is, System 1 (Thompson, Turner, and Pennycook 2011; De Neys 2014; Stanovich 2018).

Our key methodological insight is that we can tractably and flexibly capture these features by casting reasoning as an as if Bayesian problem of nonparametric learning of the unknown optimal policy function $c^*(y)$ via costly mental signals. In particular, in each period the agent can engage in costly active deliberation about the unknown optimal action, the outcome of which we model as an informative (but noisy) signal:

$$(2) \quad \eta_t = c^*(y_t) + \sigma_{\eta,t} \varepsilon_t,$$

where $\varepsilon_t \sim N(0, 1)$. This represents the outcome of System 2 type of thinking, which produces new insights about the agent's decision problem through active contemplation.

We remain agnostic about the specific mental methods people use in deliberation and simply represent the end result as an abstract informative signal. Moreover, η_t is the outcome of an internal process through which the agent draws new insights about optimal behavior from already known facts—that is, an outside observer will not measure the arrival of new objective information (similar to a theory literature on deliberation as internal reflection, e.g., Alaoui and Penta 2022). For example, the agent might mentally iterate forward using a limited set of policy functions, constructing associated values $V(y_t; c(y))$ for candidate policies $c(y)$, and pick the one delivering the highest value out of this set; or perhaps the agent draws conclusions about $c^*(y)$ based on mental heuristics that are similarly not exhaustive.¹²

Instead, we are primarily interested in capturing the general (and intuitive) feature, shared by many different deliberation methods, that more cognitive effort increases the precision of the outcome of deliberation. To model this, we let the agent choose the variance of the reasoning signal noise, $\sigma_{\eta,t}^2$, optimally subject to a cognitive cost that we describe later. The resulting signal precision models the chosen deliberation intensity.

12. We abstract from learning by experiencing the current outcomes of actions $u(y_t, c_t)$. Although these outcomes are informative in a static maximization, for a dynamic problem like in equation (1), the relevant object that needs to be learned, $V(y_t; c(y))$, is not directly accessible via experience. It requires computing the present discounted utility of the whole path of future states and actions, not just the instantaneous utility. Thus, we view the mental reasoning process we focus on as particularly applicable to dynamic problems.

In addition to the outcome of this active deliberation, agents keep a memory database of past deliberation outcomes η^{t-1} and the situations y^{t-1} in which those past insights were derived. We represent the agent's automatic intuition about optimal behavior, that is, the System 1 type of thinking, with the beginning-of-period estimate $\mathbb{E}(c^*(y)|\eta^{t-1}, y^{t-1})$. This estimate is immediately available, formed based on associations with past memories and situations and does not require additional conscious effort—in that sense, it is automatic.

Next we describe the setup and key results in detail. To simplify the exposition, we assume that the state y_t and the action c_t are scalars, so that $c^* : \mathbb{R} \rightarrow \mathbb{R}$, but insights are robust to multidimensional state and choice spaces.

1. *Associative Memory.* Without loss of generality, we represent the uncertainty about $c^*(y)$ as uncertainty over the projection coefficients θ_j associated with a complete set of basis functions $\phi_j(y)$, such that:

$$(3) \quad c^*(y) = \sum_{j=1}^{\infty} \theta_j \phi_j(y).$$

To operationalize the notion of associative memory in a way that is both tractable and general, we choose ϕ_j to be the set of Gaussian kernels, that is, $\phi_j(y) = \exp(-\psi(y - r_j)^2)$, where all kernels have the same precision parameter $\psi > 0$ but different means r_j that are uniformly distributed over the real line. This choice of basis is quite general because any continuous function can be represented as a linear combination of this set of Gaussian kernels.

The kernels ϕ_j are known, so the fundamental uncertainty which agents face is over the projection coefficients θ_j that characterize equation (3). Thus, the problem of learning $c^*(y)$ is effectively equivalent to a Bayesian, nonparametric kernel regression with $\{\eta^{t-1}, y^{t-1}\}$ as the data. The nonparametric nature of the learning problem formalizes associative memory, as it endogenously makes information in past signals η_{t-k} local to the specific contexts y_{t-k} in which those insights about $c^*(y)$ were derived, as we show next.

Assuming independent Gaussian priors over θ_j makes the learning problem particularly tractable because it allows us to

directly characterize the prior distribution over the space of functions $c^*(y)$ in closed form, resulting in the following Gaussian process (GP) distribution.¹³

LEMMA 1 (Tractable uncertainty over functions). Assuming Gaussian priors over θ_j , the function $c^*(y)$ as given in equation (3) has a GP distribution, $c^* \sim \mathcal{GP}(\hat{c}_0, \hat{\sigma}_0)$, meaning that for any pair of inputs (y, y') the joint distribution of resulting function values is:

$$\begin{bmatrix} c^*(y) \\ c^*(y') \end{bmatrix} \sim N \left(\begin{bmatrix} \hat{c}_0(y) \\ \hat{c}_0(y') \end{bmatrix}, \begin{bmatrix} \hat{\sigma}_0(y, y) & \hat{\sigma}_0(y, y') \\ \hat{\sigma}_0(y, y') & \hat{\sigma}_0(y', y') \end{bmatrix} \right),$$

where $\hat{c}_0(y) = \sum_{j=1}^{\infty} \mathbb{E}(\theta_j) \phi_j(y)$ depends on the choice of prior means over θ_j , and

$$\hat{\sigma}_0(y, y') = \sigma_c^2 \exp(-\psi(y - y')^2), \quad \text{with } \sigma_c^2 > 0 \quad \text{and } \psi > 0.$$

Although the prior distribution over $c^*(y)$ is fully parameterized by the choice of means and variances of the Gaussian priors over the underlying θ_j , it is equivalent and more convenient to think of the resulting mean and covariance functions $\hat{c}_0(y)$ and $\hat{\sigma}_0(y, y')$ as the parameters of the prior. Naturally, $\hat{c}_0(y)$ characterizes the prior mean of $c^*(y)$ for any state y and could be any continuous function. It is a modeling choice, and any difference between the chosen $\hat{c}_0$ and the true optimal c^* represents an ex ante bias in beliefs, an issue we return to later.

The key associative property of System 1 thinking is captured by $\hat{\sigma}_0(y, y')$, which specifies the agent's priors on how the values of c^* evaluated at any two points y and y' are likely to correlate with one another. Because of the underlying Gaussian kernel structure, $\hat{\sigma}_0(y, y')$ takes on an intuitive exponential form and is characterized by two parameters. First, σ_c^2 controls the prior uncertainty over $c^*(y)$ at any given input y , thus modeling the overall quantity of cognitive uncertainty the agents face. Second, ψ controls the extent to which information about the value of the function

13. Proofs are in [Online Appendix A](#). Lemma 1 captures a well-known result that Bayesian Gaussian kernel regression limits to a GP distribution over c^* (see [Rasmussen and Williams 2006](#)). A GP distribution is the limit of the multivariate Gaussian distribution for infinite-length vectors and is often used in Bayesian nonparametric statistics ([Liu, Principe, and Haykin 2011](#)) and recently in economics for nonparametric learning of the state distribution—for example, [Dew-Becker and Nathanson \(2019\)](#), [Ilut, Valchev, and Vincent \(2020\)](#), and [Bardhi \(2022\)](#).

evaluated at some y is informative about its value at a different point y' . Thus, ψ controls the notion of similarity between contexts. The higher is ψ , the faster the decay in similarity between two different situations y and y' , implying that a deliberation signal centered at $c^*(y)$ (i.e., an insight about optimal behavior given context y) is less informative about the value of $c^*(y')$ in some other situation y' .¹⁴

A powerful feature of our formal framework is that the memory-based evaluation of the optimal policy, that is, the conditional distribution $c^*(y)|\{\eta^{t-1}, y^{t-1}\}$, remains a GP, with conditional moments characterized in closed form by the following recursions.

LEMMA 2 (Conditional distributions). The conditional distributions are GP distributions, $c^*|\{\eta^{t-1}, y^{t-1}\} \sim \mathcal{GP}(\hat{c}_{t-1}, \hat{\sigma}_{t-1})$, with recursively determined moments

$$(4) \quad \hat{c}_{t-1}(y) = \hat{c}_{t-2}(y) + \alpha_{t-1}(y)(\eta_{t-1} - \hat{c}_{t-2}(y_{t-1})),$$

$$(5) \quad \hat{\sigma}_{t-1}(y, y') = \hat{\sigma}_{t-2}(y, y') - \alpha_{t-1}(y)\hat{\sigma}_{t-2}(y', y_{t-1}),$$

where $\hat{c}_{t-1}(y) \equiv E(c^*(y)|\eta^{t-1}, y^{t-1})$; $\hat{\sigma}_{t-1}(y, y') \equiv \text{Cov}(c^*(y), c^*(y')|\eta^{t-1}, y^{t-1})$;

$$\alpha_{t-1}(y) \equiv \frac{\hat{\sigma}_{t-2}(y, y_{t-1})}{\hat{\sigma}_{t-2}^2(y_{t-1}) + \sigma_{\eta, t-1}^2}$$

is the effective signal-to-noise ratio of η_{t-1} ; and $\hat{\sigma}_{t-1}^2(y) \equiv \text{Var}(c^*(y)|\eta^{t-1}, y^{t-1}) = \hat{\sigma}_{t-1}(y, y)$.

As a result, at any time t , the memory-based estimate of the optimal action at the current state of the world—the output of System 1—is the following Gaussian distribution:

$$c^*(y_t)|\{\eta^{t-1}, y^{t-1}\} \sim N(\hat{c}_{t-1}(y_t), \hat{\sigma}_{t-1}^2(y_t)),$$

where for convenience we have defined separate notation for the conditional variance of the optimal action at a given y as $\hat{\sigma}_{t-1}^2(y) \equiv \text{Var}(c^*(y)|\{\eta^{t-1}, y^{t-1}\}) = \hat{\sigma}_{t-1}(y, y)$.

14. Interestingly, a squared exponential similarity metric similar to our $\hat{\sigma}_0(y, y')$ also emerges in empirical neuroscience results on the way the brain encodes memories in practice, for example, [Gershman, Horvitz, and Tenenbaum \(2015\)](#).

Importantly, the conditional estimate $c^*(y_t) | \{\eta^{t-1}, y^{t-1}\}$ displays associative memory, a key empirical feature of System 1 type of thinking we aim to capture. To intuitively illustrate this associative memory, consider first the update conditional on one signal in memory—that is, the System 1 estimate in the second period, time $t = 2$, which draws on $\{\eta_1, y_1\}$.

COROLLARY 1 (Associative memory). At $t = 2$, the beginning-of-period uncertainty $\hat{\sigma}_1^2(y_2)$ is the following function, which is decreasing in the distance between y_2 and y_1

$$\hat{\sigma}_1^2(y_2) = \sigma_c^2(1 - \alpha_1(y_2)) = \sigma_c^2 \left(1 - \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{\eta,1}^2} \exp(-\psi(y_2 - y_1)^2) \right).$$

The effective signal-to-noise ratio $\alpha_1(y_2) = \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{\eta,1}^2} \exp(-\psi(y_2 - y_1)^2)$ is a decreasing function of the distance between the state the agent faces today, y_2 , and the state at which she had reasoned yesterday, y_1 . Thus, the subjective uncertainty $\hat{\sigma}_1^2(y_2) = \sigma_c^2(1 - \alpha_1(y_2))$ is U-shaped with a minimum at $y_1 = y_2$. That is, the memory-based estimate of $c^*(y)$ is most precise when today's situation is exactly the same as the situation that was already reasoned about yesterday, and uncertainty is increasing in the distance from yesterday's situation.

This U-shape in subjective uncertainty captures the basic idea of associative memory, a feature of System 1 thinking that has been well documented in lab and neuroscience environments (e.g., Plonsky, Teodorescu, and Erev 2015; Bornstein et al. 2017; Bornstein and Norman 2017; Gershman and Daw 2017) and in the field (like chess players in Chase and Simon (1973) and fire fighters in Klein, Calderwood, and Clinton-Cirocco 1986; see Tulving 2002 for a survey).

Updating beliefs with multiple past signals η_{t-k} does not change the basic implication about state- and history-dependent signal weights α_{t-k} and the U-shape of the resulting posterior uncertainty $\hat{\sigma}_{t-1}^2(y_t)$, as we showcase extensively in later sections. However, updating with multiple signals does introduce an additional subtlety: if the contexts at which two past signals have been derived are relatively close to each other, the conditional variance $\hat{\sigma}_{t-1}^2(y)$ achieves its minimum at a state value y in the interval in between the location of these two signals, rather than precisely in the same context as either of the two past signals. Intuitively, for




FIGURE I

Associative Memory and System 1 Uncertainty at Time $t = 3$: $\hat{\sigma}_2^2(y)$

Plotted based on two signals η_1 and η_2 with equal precision, at state realizations y_1 and y_2 .

such in-between situations, the past signals are highly relevant and thus useful.

Let us showcase this with the simple example of the System 1 estimate at time $t = 3$.

LEMMA 3 (A confident System 1). At $t = 3$, memory draws on two separate signals η_1 and η_2 , and there exists a threshold $\tau > 0$ such that if $\|y_1 - y_2\| < \tau$, then for all $y \in (y_1, y_2)$

$$\hat{\sigma}_2^2(y) < \hat{\sigma}_2^2(y_1) = \hat{\sigma}_2^2(y_2).$$

We illustrate [Lemma 3](#) visually in [Figure I](#), where we plot $\hat{\sigma}_2^2(y)$ as a function of possible values of the state y . In Panel A, we show a case where the two signals in memory are derived in situations y_1 and y_2 that are close to each other, so that the agent is particularly confident in her reasoning in intermediate situations $y \in (y_1, y_2)$, where the past signals are highly relevant. In Panel B, we plot a case where in the past the agent has reasoned in disparate situations—in that case, there are two local minima, one at each past context y_1 and y_2 . Both panels also showcase the general feature of associative memory, as captured by the fact that beliefs are most precise in the neighborhoods of past signals.

2. Occasional and Costly System 2. We turn to the choice of active deliberation. We are guided by the fact that analytical thinking is only occasionally engaged, and specifically that

this engagement is triggered by the subjective uncertainty of the automatic System 1 thinking (e.g., [Thompson, Turner, and Pennycook 2011](#)). We model this by introducing a cost-benefit trade-off to analytical reasoning.

We model the cost of deliberation as proportional to the amount of information about $c^*(y_t)$ produced by the deliberation effort. We quantify this information flow as the reduction in entropy achieved by acquiring the new signal η_t (as in information theory, e.g., [Sims 2003](#)):

$$(6) \quad H(c^*(y_t)|\eta^{t-1}) - H(c^*(y_t)|\eta_t, \eta^{t-1}) = \frac{1}{2} \ln \left(\frac{\hat{\sigma}_{t-1}^2(y_t) + \sigma_{\eta,t}^2}{\sigma_{\eta,t}^2} \right),$$

where $H(X)$ denotes the entropy of a random variable X . Thus, we are effectively assuming that more precise signals η_t are the outcome of higher (hence costlier) cognitive effort.

To model the benefit of deliberation, we assume the agent faces a standard tracking problem, and tries to minimize the expected squared deviation of her action, c_t , from the unknown optimal action $c^*(y_t)$. Putting it all together, the agent's problem is

$$(7) \quad \min_{c_t, \sigma_{\eta,t}^2} \mathbb{E}[(c_t - c^*(y_t))^2 | \eta^t, y^t] + \kappa \ln \left(\frac{\hat{\sigma}_{t-1}^2(y_t) + \sigma_{\eta,t}^2}{\sigma_{\eta,t}^2} \right).$$

The parameter κ controls the constant marginal cost of deliberation effort. For example, κ will be higher for individuals with a higher opportunity cost of time invested in reasoning or because their particular deliberation process takes longer to achieve a given improvement in precision. In addition, κ would be higher if the environment is more complex, and it is objectively harder to come up with useful insights on the unknown optimal action.

The solution to [expression \(7\)](#) has two components—a choice for the action c_t and a choice for the intensity of active deliberation, $\sigma_{\eta,t}^2$. Given the quadratic loss function, the agent chooses her action c_t according to the best estimate of the unknown optimal action and sets $c_t = \hat{c}_t(y_t)$.

In turn, the optimal choice of deliberation effort is to choose $\sigma_{\eta,t}^2$ so that

$$\hat{\sigma}_t^2(y_t) = \min [\kappa, \hat{\sigma}_{t-1}^2(y_t)].$$

Thus, the optimal deliberation choice is characterized by a target level of uncertainty, equal to the marginal cost of deliberation κ . This is achieved by setting $\sigma_{\eta,t}^2$ as follows.

PROPOSITION 1 (Trigger of System 2). The optimal reasoning signal noise variance is

$$(8) \quad \sigma_{\eta,t}^{*2} = \begin{cases} \frac{\kappa \hat{\sigma}_{t-1}^2(y_t)}{\hat{\sigma}_{t-1}^2(y_t) - \kappa}, & \text{if } \hat{\sigma}_{t-1}^2(y_t) > \kappa, \\ \infty, & \text{if } \hat{\sigma}_{t-1}^2(y_t) \leq \kappa. \end{cases}$$

This result has two key implications. First, sometimes agents might choose not to engage in active deliberation at all, in which case $\sigma_{\eta,t}^{*2} = \infty$. Second, the choice of whether to actively deliberate depends on the subjective uncertainty of the automatic System 1 thinking.

In particular, if the agent comes into the period with a memory-based estimate of the optimal action that is perceived as sufficiently precise ($\hat{\sigma}_{t-1}^2(y_t) \leq \kappa$), then expending any cognitive resources in additional deliberation is suboptimal. Conversely, if the agent's uncertainty over her immediate intuition is high ($\hat{\sigma}_{t-1}^2(y_t) > \kappa$), she finds it optimal to actively deliberate over the best course of action and obtain a new insight ($\sigma_{\eta,t}^2 < \infty$).

II.B. Cognitive Misers and Dual-Interventionist Behavior

The fact that active deliberation is only engaged when an agent's subjective confidence in her intuition (System 1) is sufficiently low, characterizes our agents as “cognitive misers” in the terminology of [Simon \(1955\)](#) and [Tversky and Kahneman \(1975\)](#)—that is, our agents economize cognitive resources when they feel these resources are not needed. Naturally, the evaluation of whether costly reasoning (System 2) is needed is based on the agent's confidence in her ex ante System 1 estimate.

This is in line with psychological findings that System 1 type of thinking brings to mind not only a candidate action but also a “feeling of confidence or rightness” in that intuition ([Evans 2019](#)). Moreover, [Thompson, Turner, and Pennycook \(2011\)](#) and [Thompson et al. \(2013\)](#) specifically document that a lower “feeling of rightness” in the System 1 guess can trigger System 2 thinking. Related work has shown a similar effect of triggering analytical thinking when there is sufficient “conflict” in the System 1 thinking (e.g., [Evans 2007b](#); [De Neys 2014](#); [Stanovich 2018](#)). These findings are also echoed in the intuitive [Kahneman \(2011, 24\)](#) notion that “when System 1 runs into difficulty, it calls on System 2

to support more detailed and specific processing.”¹⁵ Our model provides a formal economic theory of these results, where a high $\hat{\sigma}_{t-1}^2(y_t)$ articulates and formalizes the psychological notion of a lower subjective confidence, and equivalently higher “conflict,” in the System 1 estimate.

Having established the psychologically appealing features of the model’s optimal deliberation choice, we turn to the implications about the chosen action, c_t , which will allow us to make predictions about observed economic behavior. Recall that the agent acts on her best conditional estimate $\hat{c}_t(y_t)$, which, due to its Bayesian update nature, can be represented as a baseline of the beginning-of-period estimate $\hat{c}_{t-1}(y_t)$ with an adjustment driven by the outcome of any current additional deliberation η_t .

PROPOSITION 2 (Default-interventionist behavior). The agent’s time- t action is

$$(9) \quad c_t = \hat{c}_t(y_t) = \hat{c}_{t-1}(y_t) + \alpha_t^*(y_t)(\eta_t - \hat{c}_{t-1}(y_t)),$$

where the optimal weight put on the new reasoning signal, $\alpha_t^*(y_t)$, depends on the current state y_t and the history $\{y^{t-1}, \sigma_{\eta}^{t-1}\}$ of past signals’ location and precision:

$$(10) \quad \alpha_t^*(y_t) \equiv \frac{\hat{\sigma}_{t-1}^2(y_t)}{\hat{\sigma}_{t-1}^2(y_t) + \sigma_{\eta,t}^{*2}} = \max \left[1 - \frac{\kappa}{\hat{\sigma}_{t-1}^2(y_t)}, 0 \right].$$

Thus, the System 1 type of thinking effectively provides a memory-based, default action $\hat{c}_{t-1}(y_t)$, which is then revised or adjusted by any additional System 2 insights η_t .

In fact, the weight on the System 2 revision, $\alpha_t^*(y_t)$, is state- and history-dependent. Specifically, when the agent’s confidence in her initial intuition is high—that is, when $\hat{\sigma}_{t-1}^2(y_t) \leq \kappa$ —the agent acts entirely based on her automatic, System 1 estimate $\hat{c}_{t-1}(y_t)$ without actively rethinking and revising this action ($\alpha_t^*(y_t) = 0$). By the intuition of [Lemma 3](#), we are likely to observe such default or habitual behavior in situations that are already familiar to the agent—that is, when the current state y_t is within a neighborhood of past situations y_{t-k} that the agent has already reasoned through.

15. Indeed, if costly reasoning was to determine whether costly reasoning is needed is a conceptual paradox, as also noted in the psychological literature (e.g., [Evans 2009](#); [Shleifer 2012](#); [De Neys 2014](#)).

In contrast, additional deliberation and associated revisions in behavior are likely to occur when the state moves into unfamiliar territory, that is, when the agent is confronted with a novel situation y_t , as compared with the types of situations she has deliberated on in the past. In that case, the subjective confidence in the memory-based estimate of the optimal action is low (high $\hat{\sigma}_{t-1}^2(y_t)$, as per its characteristic U-shape), and the agent optimally chooses to engage System 2. This leads to a high $\alpha_t^*(y_t)$ and behavior that is revised by a new insight η_t —the agent does not simply follow the default, beginning-of-period intuition.

Proposition 2 thus formalizes the dual-reasoning interaction of our model as a default-interventionist behavior (as introduced by [Evans 2007b](#)) and present in many related forms, for example, [Kahneman and Frederick \(2002\)](#), [Kahneman \(2011\)](#), [Stanovich \(2011\)](#), and [Bago and De Neys \(2017\)](#).¹⁶ Consistent with those views, in our model the cuing of the default System 1 response may be (i) endorsed if perceived as confident enough ($\alpha_t^*(y_t) = 0$ when $\hat{\sigma}_{t-1}^2(y_t) \leq \kappa$) or (ii) otherwise altered ($\alpha_t^*(y_t) > 0$) by the outcome of System 2 reasoning.

Crucially, the choice of rethinking the default action $\hat{c}_{t-1}(y_t)$ is based on an agent's subjective evaluation of the size of the likely error, as given by $\hat{\sigma}_{t-1}^2(y_t) = \mathbb{E}_{t-1}(c^*(y_t) - \hat{c}_{t-1}(y_t))^2$, and not by the actual error $(c^*(y_t) - \hat{c}_{t-1}(y_t))^2$. As a result, our boundedly rational agents might stick with a default action that may appear like an objectively costly behavioral bias to an outside observer.¹⁷ Thus, our framework formalizes the emerging view in the psychology literature that behavioral mistakes are indeed often due to the interplay in the dual-reasoning system—for example, “cognitive biases result largely from false feelings of confidence which perhaps allow the cognitive miser to stick with default intuitions” ([Evans 2019](#), 409).

16. In the model of [Cerigioni \(2021\)](#), a related notion of an active reasoning threshold appears. There the agent is assumed to engage in conscious deliberation only when the current context is sufficiently different than the past; otherwise, she follows the same past action. Here, the costly deliberating signal and the nonparametric estimation jointly and endogenously determine the dual-reasoning interaction and persistent mistakes.

17. Our emphasis on learning dynamics and subjective perceptions of payoffs further connects to a recent literature on cognitive adaption, where a noisy neural coding scheme might adapt to a changing environment ([Robson and Whitehead 2019](#); [Aridor, Grechi, and Woodford 2020](#)).

III. KEY QUALITATIVE IMPLICATIONS AND EMPIRICAL EVIDENCE

A key qualitative implication of our model is the switch from habitual behavior, as driven by the default, beginning-of-period estimate of the optimal action $\hat{c}_{t-1}(y_t)$, to active deliberation, and thus adjustment in behavior and thinking, as driven by a new signal η_t in novel situations.

III.A. Switching between Habits of Mind and Active Thinking

In this section we connect to a broad range of empirical evidence on the switch from “habits of mind” to active thinking (Louis and Sutton 1991). To frame the discussion, we first present a simple five-period illustration of the model’s mechanism.

1. *Time 1: A Burst of Active Thinking.* Time $t = 1$ represents a situation or problem for which the agent has no previous insights on how to handle it. This results in a relatively large uncertainty in her immediate, beginning-of-period guess of the optimal action, that is, $\hat{\sigma}_0(y_1) = \sigma_c^2 > \kappa$, and hence the agent decides to actively deliberate (as per Proposition 1). The outcome η_1 of this deliberation then revises the agent’s intuitive guess $\hat{c}_0(y_1)$, and also updates the agent’s confidence in her System 1 estimate going forward, that is, $\hat{\sigma}_1^2(y)$.

2. *Time 2: Less Intense Deliberation.* Entering time $t = 2$, the subjective uncertainty in the beginning-of-period beliefs (System 1) is lower throughout the whole state space, but especially so close to y_1 where the agent has already deliberated. Lemma 3 shows this analytically, and Figure II, Panel A illustrates it by plotting $\hat{\sigma}_1^2(y)$ as a function of potential state realizations y (solid line)—the U-shape embodies the associative memory of System 1.

Because of this characteristic U-shape of $\sigma_1^2(y)$, any change in the state of the world at time $t = 2$ (i.e., if $y_2 \neq y_1$) would trigger a new round of active deliberation, with the intensity of this deliberation and the resulting weight on the new signal $\alpha_2^*(y_2)$ being higher the further y_2 is from last period’s state value y_1 . To illustrate, we plot one specific realization of the time $t = 2$ state y_2 with a vertical dotted line in Figure II, Panel A, and also show the time path of the resulting updating weights $\alpha_t^*(y_t)$ on the new reasoning signals in Panel B.




FIGURE II

Beginning-of-Period Uncertainty ($\hat{\sigma}_{t-1}^2(y)$) and the Weight on the New Signal $\alpha_t^*(y)$

3. *Time 3 and 4: Once Habits Settle In, Further Deliberation Can Endogenously Stop.* The agent enters $t = 3$ having accumulated two distinct insights about optimal behavior at two different state realizations, y_1 and y_2 . Associative memory again results in a U-shaped beginning-of-period subjective uncertainty $\hat{\sigma}_2^2(y)$ (dash-dotted line in Figure II, Panel A).

Importantly, the agent is especially confident in her understanding of the optimal action for state realizations y in between y_1 and y_2 , because in such situations both previous signals η_1 and η_2 are relevant and informative. Thus, perceived uncertainty for such in-between situations is in fact lower than at either y_1 or y_2 (as per Lemma 3), and hence in that region of the state space System 1 perceives uncertainty that is lower than its target, that is, $\hat{\sigma}_2^2(y) < \kappa$.

As a result, for a state realization in this familiar region, like the one labeled y_3 in Figure II, Panel A, the agent feels confident enough in her System 1 guess that she chooses not to expend any further costly mental resources and does not engage System 2 again. Formally, in that part of the state space $\alpha_3^*(y) = 0$ and the agent's actions are driven only by the beginning-of-period guess of the optimal action, that is, $c_3 = \hat{c}_2(y_3)$. Moreover, the memory database underpinning System 1 is also not updated with new signals, since System 2 is disengaged. Hence, there is inertia in the System 1 beliefs, which remain the same next period, as they were this period ($\hat{c}_3(y) = \hat{c}_2(y)$ and $\hat{\sigma}_3^2(y) = \hat{\sigma}_2^2(y)$ for all y).

In turn, if the time $t = 4$ shock is small, so that y_4 remains in the familiar region where $\hat{\sigma}_3^2(y_4) < \kappa$, the agent again does not

actively deliberate. In such a case we say the agent displays habitual behavior, as she acts according to same policy function (or decision rule), as the last few periods—that is, $\hat{c}_2(y) = \hat{c}_3(y) = \hat{c}_4(y)$ for all y .

4. *Time 5: Major Disruptions Trigger Deliberation Again.*

However, if some event alters the context or situation significantly and pushes the agent away from her familiar region of the state space, this habitual type of behavior will change.

For example, if in period $t = 5$ the agent encounters a substantial shock, illustrated by the position of y_5 further to the right in [Figure II](#), Panel A, then the agent is less confident in her System 1, automatic estimate of the optimal action, since the previous deliberations have happened in very different situations. Formally, the variance $\hat{\sigma}_4^2(y_5)$ is high, and the agent chooses to actively deliberate (engage System 2) rather intensely. As a result, she draws a relatively precise new signal η_5 that updates beliefs (and actions) with a substantial weight $\alpha_5^*(y_5) > 0$, as we can see directly in [Figure II](#), Panel B. [Proposition 2](#) also implies that following such active deliberation (as at $t = 1, 2, 5$ in our example) the updated action $\hat{c}_t(y_t)$ is typically closer to the optimal $c^*(y_t)$, than the initial $\hat{c}_{t-1}(y_t)$ was.

5. *Switching Cognitive Gears.* The example illustrates how agents follow default, habitual behavior in familiar situations, but rethink that behavior in novel situations. These fundamental features connect the model with a wide range of field evidence.

For example, in the cognitive managerial literature, [Barr, Stimpert, and Huff \(1992\)](#), [Walsh \(1995\)](#), and [Maitlis and Christianson \(2014\)](#) review how unusual shifts in the external economic environment (“episodic changes” or “jolts”) are critical in changing the “mental models,” “cognitive schemas,” or “routines” of managers, indicating a deliberative rethinking of procedures and actions. This may involve adjusting a hospital’s operation during a doctors’ strike ([Meyer 1982](#)), responding to sudden changes in technology or competition structure ([Bogner and Barr 2000](#)), and generally “unfreezing” the firm’s strategy when dealing with an unfolding crisis (see [Christianson et al. 2009](#) and the review in [Lampel, Shamsie, and Shapira 2009](#)).

In this context, [Tyre and Orlikowski \(1994\)](#) describe new situations as “windows of opportunity” for adjustments. They document that when firms integrate new technologies, there is an

initial large burst of active thinking about the best methods for using such new technologies, like in our illustrative periods $t = 1$ and $t = 2$. They also find that like in our $t = 3$ and $t = 4$, this active thinking slows down relatively quickly, as the situation becomes familiar, and then actions settle in habitual patterns learned in that initial period of active deliberation. Finally, they document that those imperfect but strongly established modes of habitual behavior are only rethought and adjusted in case of a significant shock (e.g., a machine breakdown, corresponding to a shift in y_t as in period $t = 5$ of our example).

Our model's switch from default intuition to active deliberation also connects to a social psychology view of habitual actions in stable contexts as the automatic, low cognitive effort response (Wood, Quinn, and Kashy 2002; Wood et al. 2014). In addition, a range of evidence in health and environmental economics shows that naturally occurring changes in agents' situations indeed precipitate windows of opportunity (Verplanken and Roy 2016) during which habitual modes of behavior are more likely to be reconsidered (Verplanken and Wood 2006). For example, moving to a new location (a significant change in circumstances) is associated with reconsidering the default choice patterns in transportation (Verplanken et al. 2008), energy consumption (Maréchal 2010), and exercise and watching TV (Wood, Tam, and Witt 2005).

Similarly, in the labor literature, the exhaustion of unemployment benefits prompts adjustments in the default job search strategy and consumption behavior (DellaVigna et al. 2017; Ganong and Noel 2019), while large shocks (personal or on the job) often prompt workers to reevaluate their decisions to stay with a particular company. Such unusual shocks are described as "jarring events towards mental deliberations" and are key to understanding the otherwise puzzling behavior of voluntary turnover (see Lee and Mitchell 1994; Holtom et al. 2005). Likewise, in finance, Liu (2021) documents that loan officers faced with surprises in the incoming information on borrowers appear to reconsider the model they use to evaluate borrowers (leading to better loan decisions), akin to our triggering of System 2.¹⁸

Overall, this type of broad evidence on the role of novel situations as key in triggering deliberation and adjusting habitual actions has important public policy implications. In fact, motivated

18. [Online Appendix B](#) discusses action paths in our five-period example that connect to these findings.

by the purely empirical insights above, recent work has advocated for “context-change” policy interventions (Verplanken and Wood 2006; Fernandes, Lynch, and Netemeyer 2014), which not only aim to provide new information (e.g., on health or financial choices) but to do so precisely when agents experience significant changes in their individual circumstances. Our model provides a structural framework for quantifying such policy counterfactuals.

III.B. Behavioral Biases as Stable Habitual Behavior

A further key result is that this characteristic state- and history-dependent engagement of active deliberation can result in learning trap situations, where agents optimally stop reasoning. In such situations, agents will habitually follow a default policy function estimate that is stable and unchanging over time (even if it leads to objectively costly mistakes).

Specifically, learning traps occur when a sequence of past reasoning signals η^{t-1} are (i) concentrated in a particular part of the state space, so that $\hat{\sigma}_{t-1}^2(y) < \kappa$ for an interval of y values within the neighborhood of y^{t-1} ; and (ii) the resulting beginning-of-period estimate of the policy function $\hat{c}_{t-1}(y)$ implies that the law of motion of the state $F(y_t, c_t, v_{t+1})$ is stable and mean-reverting to a value $\bar{y}$ in this familiar region of past experiences y^{t-1} , that is, $\hat{\sigma}_{t-1}^2(\bar{y}) < \kappa$. Formally, such beginning-of-period beliefs define a locally stable steady state around which beliefs do not change as the state fluctuates locally around $\bar{y}$.

PROPOSITION 3 (Habitual behavior). A set of System 1 (beginning-of-period) beliefs $\{\hat{c}(y), \hat{\sigma}^2(y)\}$ and a state value $\bar{y}$ define a locally stable steady state if and only if

- i. The System 1 uncertainty at $\bar{y}$ is below the target threshold level κ : $\hat{\sigma}^2(\bar{y}) < \kappa$.
- ii. The System 1 policy $\hat{c}(y)$ implies stable y_t dynamics around $\bar{y}$: $|\frac{\partial F}{\partial y}(\bar{y}, \hat{c}(\bar{y}), 0)| < 1$.

Hence, for small shocks to y_t around $\bar{y}$ the agent does not actively deliberate, the effective policy function is unchanged $\hat{c}_t(y) = \hat{c}_{t-1}(y) = \hat{c}(y)$, and the state level y_t mean reverts to $\bar{y}$.

At one extreme, consider an environment where the state is purely exogenous so that $\frac{\partial F}{\partial c} = 0$. The second restriction in Proposition 3 is then irrelevant—reasoning eventually stops in the neighborhood of the mean of y ’s distribution regardless of the particular properties of the estimate $\hat{c}_t(y)$ at that time. In such an

environment, the model predicts that all behavior across agents is purely idiosyncratic, because there is nothing special about being in or out of a learning trap and heterogeneity is only driven by the reasoning signal errors $\varepsilon_{i,t}$ that are independent across agents. With no ex ante biases, actions will on average equal $c^*(y)$. This lack of systematic patterns in errors echoes a common argument, discussed for example by [Thaler \(1992\)](#), that a rational model may be a good approximation for aggregate behavior, even if individual agents make mistakes, because errors may wash out on average.

In contrast, our dual-reasoning mechanism provides a caveat to this argument when the state is endogenous, and the agent's actions affect its law of motion—that is, $\frac{\partial F}{\partial c} \neq 0$. As per the second part of [Proposition 3](#), in environments with endogenous states only certain types of policy beliefs $\hat{c}(y)$ can emerge as stable, because only certain policy functions ensure that the endogenous state dynamics are mean reverting. Otherwise, the state is likely to drift away from any given familiar region, which would trigger rethinking and revision in beliefs.

In particular, if the agent comes into the period with beliefs that satisfy the conditions in [Proposition 3](#), then small shocks to y_t keep the beginning-of-period uncertainty below the threshold that triggers active reasoning. Hence, the agent continues to follow the same policy as last period, $\hat{c}_t(y) = \hat{c}_{t-1}(y) = \hat{c}(y)$. In turn, because this policy $\hat{c}(y)$ is such that it keeps y_t in the familiar territory near $\bar{y}$, the agent does not reason next period (for small shocks) either. Overall, agents inside a learning trap exhibit habitual behavior in the sense of following the same decision rule over time, without actively reconsidering it.

Moreover, such habitual behavior is not necessarily close to the optimal $c^*(y)$, but is defined by the fact that it stabilizes the value of the state. Thus, the selective engagement of analytical reasoning effectively leads to persistent behavioral biases even in the absence of ex ante biases in beliefs or updating (i.e., in System 1 or System 2 processes).

We view our System 1/2 interaction mechanism and the resulting endogenous biases as a complement to other psychological features and primitive biases. In particular, in [Online Appendix C](#) we analyze biases in the memory recall and judgment of System 1, such as rehearsal ([Mullainathan 2002](#)), confirmation ([Rabin and Schrag 1999](#)), and recency biases ([Bushong and Gagnon-Bartsch 2016](#)); motivated reasoning of System 2 ([Kunda 1990](#); [Handley and Trippas 2015](#)); and ex ante mean biases in System 1 and/or 2

(Enke et al. forthcoming). There we argue that incorporating these biases in our model is quite tractable, and while such extensions could generate additional interesting behavior, we show they do not affect the dual-systems reasoning interaction we emphasize. Thus, to cleanly illustrate our proposed mechanism and its unique predictions, we opt for a parsimonious benchmark framework that relies only on associative memory and costly deliberation.

IV. CONSUMPTION-SAVINGS MODEL WITH COSTLY DELIBERATION

The specific nature of these endogenous behavioral biases depend on the way the agent's actions affect the law of motion of the state. To showcase this, and more broadly the empirical promise of our framework, in this section we set our mechanism in an otherwise standard Aiyagari (1994) incomplete-markets model. We have chosen this particular application due to its transparency and because consumption-savings decisions are a workhorse mechanism in a number of different macroeconomic settings. At the same time, consumption decisions in the micro data appear to deviate from standard rational models in many important ways.

The economy consists of a continuum of ex ante identical households, indexed by i , with concave utility over consumption, who inelastically supply their stochastic endowment of labor $s_{i,t}$ at a constant wage w . The income shocks $s_{i,t}$ are i.i.d., drawn from a time-invariant distribution S with a mean of one. The only asset is physical capital that earns a constant rental rate $\tilde{r}$ and depreciates at rate $\delta \in (0, 1)$. The resulting budget constraint agents face is

$$c_{i,t} + a_{i,t} = (1 + r)a_{i,t-1} + ws_{i,t},$$

where $r \equiv \tilde{r} - \delta$, $a_{i,t-1}$ is the amount of capital held at the end of period $t-1$ and $a_{i,t}$ is the current choice of savings. The agents also face the borrowing constraint $a_{i,t} \geq 0$.

The aggregate production function is standard—it takes as inputs the average capital $K = \int a_{i,t} di$ and employment $H = \int s_{i,t} di$, and produces $K^\alpha H^{1-\alpha}$, with $\alpha \in (0, 1)$. The role of this side of the economy is to determine the rental rate and the wage from the firm's usual first-order conditions $\tilde{r} = \alpha K^{\alpha-1} H^{1-\alpha}$ and $w = (1-\alpha)K^\alpha H^{-\alpha}$, respectively. Given the assumed inelastically supplied labor and i.i.d. labor supply shocks $s_{i,t}$, we have $H = 1$.

IV.A. Decision Problem

We can rewrite the budget constraint in terms of cash-on-hand, defined as

$$y_{i,t} \equiv (1+r)a_{i,t-1} + ws_{i,t}.$$

The consumption-savings problem agents want to solve then can be expressed as

$$(11) \quad V(y_{i,t}) = \max_{c_{i,t}} u(c_{i,t}) + \beta \mathbb{E}_t V(y_{i,t+1}),$$

subject to the law of motion for cash-on-hand,

$$(12) \quad y_{i,t+1} = (1+r)(y_{i,t} - c_{i,t}) + ws_{i,t+1},$$

and the borrowing limit $c_{i,t} \leq y_{i,t}$. As is standard, the optimal consumption choice at time t can then be expressed as $c_{i,t}^* = \min(y_{i,t}, c^*(y_{i,t}))$, where the policy $c^*(y_{i,t})$ gives the optimal action taking into account future borrowing constraints, but ignoring today's constraint.

1. *Deliberation Friction.* Applying the general framework developed in Section II, agents perfectly observe their cash-on-hand $y_{i,t}$ and all constraints, but do not know the optimal policy function $c^*(y)$, and estimate it from a history of costly deliberation signals,

$$(13) \quad \eta_{i,t} = c^*(y_{i,t}) + \varepsilon_{i,t}, \quad \varepsilon_{i,t} \sim N(0, \sigma_{\eta,i,t}^2).$$

Agents have the common time-0 prior of $c^* \sim \mathcal{GP}(\widehat{c}_0, \widehat{\sigma}_0)$, which is centered around the truth ($\widehat{c}_0 = c^*$) and has the covariance function $\widehat{\sigma}_0$ as in Lemma 1.

As in Section II, agents face a quadratic tracking problem and thus select the precision of their reasoning signals, $\frac{1}{\sigma_{\eta,i,t}^2}$, so as to achieve the target posterior variance of κ :

$$(14) \quad \widehat{\sigma}_{i,t}^2(y_{i,t}) = \min \left[\kappa, \widehat{\sigma}_{i,t-1}^2(y_{i,t}) \right].$$

By Proposition 2, the conditional expectation of optimal consumption $c^*(y_{i,t})$ is:

$$\widehat{c}_{i,t}(y_{i,t}) = \widehat{c}_{i,t-1}(y_{i,t-1}) + \alpha_{i,t}(y_{i,t})(\eta_{i,t} - \widehat{c}_{i,t-1}(y_{i,t-1})),$$

where the weight on $\eta_{i,t}$ can be expressed as $\alpha_{i,t}(y_{i,t}) = \max[1 - \frac{\kappa}{\widehat{\sigma}_{i,t-1}^2(y_{i,t})}, 0]$. Last, taking into account the borrowing limit is straightforward as the agent perfectly observes $y_{i,t}$, and hence the agent sets actual consumption according to $c_{i,t} = \min(y_{i,t}, \widehat{c}_{i,t}(y_{i,t}))$.

IV.B. Learning Traps

By [Proposition 3](#), the types of beliefs that are likely to emerge as stable, and thus characterize ergodic behavior, are such that the implied consumption choices keep the dynamics of the state y_t stable in a familiar region where no further deliberation takes place.

To understand how this plays out, let us first define the counterfactual policy function $c^{RW}(y)$ which implies that cash-on-hand is a random walk and thus has a zero growth rate on average, that is, $\mathbb{E}_t(y_{i,t+1}) - y_{i,t} = ry_{i,t} - (1+r)c^{RW}(y_{i,t}) + w = 0$ for all $y_{i,t}$. Solving for $c^{RW}(y)$, we see that this is the familiar permanent income hypothesis (PIH) policy,

$$(15) \quad c^{RW}(y) = \frac{r}{1+r}y + \frac{1}{1+r}w.$$

In turn, we can express the wealth growth rate implied by any policy function $\hat{c}_{i,t}(y_t)$ as

$$(16) \quad \mathbb{E}_t(y_{i,t+1}) - y_{i,t} = (1+r)(c^{RW}(y_{i,t}) - \hat{c}_{i,t}(y_{i,t})).$$

This shows that for a given $y_{i,t}$, if an agent follows a consumption rule that lies above $c^{RW}(y_{i,t})$ then her wealth is expected to fall (she is dissaving), whereas if the agent's consumption policy lies below $c^{RW}(y_{i,t})$, then she is saving and assets grow on average. Using this intuitive insight, we can formally prove the following corollary to [Proposition 3](#), which shows that there are only two types of beliefs $\{\hat{c}_i(y), \hat{\sigma}_i^2(y)\}$ that can be sustained as stable steady states (i.e., learning traps) in this consumption-savings application of our framework.

COROLLARY 2 (Learning traps). A set of beliefs $\{\hat{c}_i(y), \hat{\sigma}_i^2(y)\}$ and a level of wealth $\bar{y}_i$ define a locally stable steady state if and only if they satisfy the following two conditions:

- i. Subjective uncertainty at $\bar{y}_i$ is below the threshold level κ : $\hat{\sigma}_i^2(\bar{y}_i) < \kappa$, and
- ii. The policy $\hat{c}_i(y)$ and wealth level $\bar{y}_i$ satisfy one of the following stability conditions:
 - $\bar{y}_i = w$ and $\hat{c}_i(\bar{y}_i) > c^{RW}(\bar{y}_i) = \bar{y}_i$ (i.e., the borrowing limit binds at $\bar{y}_i$)
 - $\bar{y}_i > w$ and $\hat{c}_i(\bar{y}_i) = c^{RW}(\bar{y}_i) < \bar{y}_i$, while at the same time $\frac{\partial \hat{c}_i}{\partial y}(\bar{y}_i) > \frac{\partial c^{RW}}{\partial y}(\bar{y}_i)$.

As per our general discussion in [Section III.B](#), a locally stable steady state emerges when a sequence of past reasoning signals η^{t-1} implies a beginning-of-period consumption policy estimate $\hat{c}_{i,t-1}(y)$ (the System 1 guess) such that the resulting wealth dynamics are stable around some $\bar{y}_i$ in the region of the state space where the signals η^{t-1} were obtained. In that case, by the logic of [Lemma 3](#), System 1's confidence is high in the neighborhood of the locations of the past signals η^{t-1} , and thus would not perceive sufficient uncertainty to engage further System 2 deliberation. Without the arrival of new signals, the System 1 beliefs and its implied behavior that keeps wealth stable in that region are sustained.

Intuitively, there are two types of consumption behavior that ensure the stable wealth dynamics necessary for such learning traps. The first type is described by an excessively high belief about optimal consumption at low levels of assets, that is, $\hat{c}_i(y) > \hat{c}^{RW}(y)$ for low values of y . In this case, near the constraint the agent is dissaving on average, which makes the borrowing limit itself a locally stable point—since wealth is drifting down, the constraint acts as an absorbing barrier. Formally, the stable cash-on-hand point in this case is simply equal to the average wage, $\bar{y}_i = w$. Overall, agents in this kind of a learning trap exhibit a habitually high consumption level and are persistently in a hand-to-mouth (HtM) situation.

A second type of stable behavior is characterized by a stable wealth point $\bar{y}_i > w$ away from the constraint, where the agent's policy function estimate intersects $c^{RW}(y)$ from below—that is, in the neighborhood of $\bar{y}_i$, $\hat{c}_i(y) > c^{RW}(y) \iff y > \bar{y}_i$. In this case, for state realizations $y_{i,t} > \bar{y}_i$, the agent dissaves and wealth drifts down, and vice versa (see [equation \(16\)](#)), so the dynamics of cash-on-hand are mean reverting to $\bar{y}_i$. A key feature of this learning trap is that the habitual policy $\hat{c}_i(y)$ is steeper than the PIH line near the steady-state wealth level $\bar{y}_i$. Hence, agents who find themselves in this type of learning trap would appear to be unconstrained and yet still display a relatively high marginal propensity to consume.

This last result highlights an important subtlety: our model predicts inertia in the policy function, meaning a stable mode of behavior, and not necessarily inertia in the actual actions taken. Intuitively, in the second kind of a learning trap agents are behaving habitually, in the sense of following the same decision rule each period, but that rule is telling them to adjust consumption

levels relatively strongly to income shocks. It is this local sensitivity in consumption that stabilizes wealth and supports the learning trap, as positive income shocks are consumed, while negative ones are offset by lowering consumption.

Finally, the learning traps in [Corollary 2](#) are conceptually robust along two dimensions. First, learning over $c^*(y)$ has the same implications as learning over the optimal savings policy $a^*(y) = y - c^*(y)$ because the perfectly observed budget constraint induces a deterministic translation between $c^*(y)$ and $a^*(y)$. Indeed, an elastic estimate of $c^*(y)$ and a relatively flat estimate of $a^*(y)$ are both equivalent conditions for wealth stability.¹⁹ In contrast, in models where the state $y_{i,t}$ (and implicitly also the budget constraint) is imprecisely perceived, such equivalence may not hold. For example, in [Reis \(2006\)](#) behavior is different depending on whether agents set consumption or savings, since the control variable necessarily underreacts due to imperfect information, but then the residual action must overreact to satisfy the budget constraint.²⁰

Second, our assumption that active deliberation produces signals $\eta_{i,t}$ about the full-information policy $c^*(y)$ implies that our agents are naive about their bounded rationality (in the terminology of [O'Donoghue and Rabin 1999](#)). One can alternatively model sophisticated agents by assuming that $\eta_{i,t}$ are unbiased signals of the policy $\tilde{c}^*(y)$ that solves a version of the dynamic problem in [equation \(11\)](#) (or more generally in [equation \(1\)](#)) that explicitly accounts for the reasoning friction. However, we note that the stability conditions in [Corollary 2](#) are independent from whether the true conditional mean function generating System 2 signals is $c^*(y)$ or $\tilde{c}^*(y)$, and hence the nature of the learning traps remains the same.

V. NUMERICAL ANALYSIS

We analyze the stationary equilibrium of the model using a simulation detailed in [Online Appendix E](#). In calibrating the model, we follow the parametrization in [Aiyagari \(1994\)](#) whenever possible. Households have log utility with a discount factor

19. [Online Appendix D](#) details this equivalence as a counterpart of [Corollary 2](#) when reasoning is over $a^*(y)$.

20. As in [Luo \(2008\)](#), the typical assumption is to let savings be the residual action (see also [Sims 2003](#); [Mackowiak and Wiederholt 2015](#)), leading to this literature's characteristic result of inertia in consumption.

$\beta = 0.96$. The i.i.d. income shock is drawn as $\ln(s_{i,t}) \sim N(-\frac{\sigma_s^2}{2}, \sigma_s^2)$, with $\sigma_s = 0.2$. The capital share and the annual depreciation rate take standard values of $\alpha = 0.36$ and $\delta = 0.08$, respectively.

Next we turn to the parameters governing the reasoning friction. First, to help with the otherwise slow and history-dependent evolution of beliefs in the stationary equilibrium computation, we introduce a form of discounting of past reasoning signals. We opt for the tractable modeling assumption that agents face i.i.d. Poisson death shocks, where they exit with probability θ . Upon death, agents transfer their assets to an offspring household. However, the transfer of knowledge about optimal behavior is imperfect across generations, where for simplicity we assume that the new household starts with no information, only the ex ante priors $\mathcal{GP}(\hat{c}_0, \hat{\sigma}_0)$. We set $\theta = 0.02$, so that the economy is continuously repopulated with agents that essentially restart the learning problem every 50 years, on average.

Second, we set σ_c^2 and ψ equal to what an econometrician would estimate as the best-fit GP distribution for $c^*(y)$, based on the actual distribution of reasoning signals $\eta_{i,t}$ at the stationary equilibrium of the model. This procedure ensures priors are model consistent, in the sense that if the agents were to sample others' beliefs about $c^*(y)$, they will indeed recover a distribution that conforms with their own priors. Overall, this approach defines a fixed point in $\{\sigma_c^2, \psi\}$ that we solve numerically, as detailed in [Online Appendix F](#).

Finally, we calibrate the cost of deliberation κ by matching the average MPC in the top wealth quintile—which is 0.15 in the data (empirical evidence reviewed below). This approach exploits one of the key behavioral implications derived by [Corollary 2](#), which shows that the MPC of unconstrained agents in our model is higher than that implied by the PIH (as is also true in the data). Thus, we effectively calibrate the size of our bounded rationality friction to one of the interesting consumption puzzles it can generate qualitatively. Putting everything together, the calibration of the reasoning parameters yields $\{\sigma_c^2, \psi, \kappa\} = \{0.77, 0.05, 0.48\}$.

V.A. Reasoning and Beliefs in the Stationary Distribution

1. *Learning Traps.* The steady-state stability formalized in [Corollary 2](#) is only guaranteed locally, that is, for small shocks $s_{i,t}$. Nevertheless, we find that learning traps are ubiquitous and quite stable outcomes in the stationary distribution, and thus are




FIGURE III

Stationary Wealth Distribution

In both panels, we also plot the corresponding values for the full-information counterpart of our model.

an important driver of the ergodic behavior. In particular, 71% of agents in the stationary distribution are in a learning trap (defined as in [Corollary 2](#) to mean $\hat{\sigma}_{i,t-1}^2(\bar{y}_i) < \kappa$), and the typical agent first enters a learning trap just one-third into her lifetime. While the extent to which agents remain within a learning trap region is stochastic, we find that the one-step ahead probability of exiting a learning trap is quite small at 0.8%. On the other hand, the one-step-ahead probability of entering a trap conditional on currently being out of one is significantly larger, at 6.9%. The associated transition matrix implies that learning traps are stable long-run outcomes. [Online Appendix G](#) provides additional related moments.

2. Wealth Inequality. In the full-information (FI) version of the model, which has the same sequence of income shocks but no uncertainty over $c^*(y)$, all agents have the same steady-state wealth point, $\bar{y}^*$, given by the unique intersection of $c^*(y)$ and $c^{RW}(y)$. Thus, in the FI economy any wealth heterogeneity is only due to transitory income shocks. In contrast, in our bounded rationality (BR) model, agents also have different long-run or target wealth levels $\bar{y}_i$ due to the dispersion in policy function estimates $\hat{c}_i(y)$. As a result, we find two notable differences in the stationary distribution of assets a_i between the BR and FI models: (i) a large mass of low-wealth agents and (ii) a larger density of rich agents—as visualized in [Figure III](#), Panel A. Overall, we find that our model

features a higher aggregate demand for savings, which results in a lower equilibrium interest rate: $r^{BR} = 3.89\% < r^{FI} = 4.15\%$.

Moreover, the Gini coefficient of the wealth distribution in our model is 0.58, significantly higher than its FI counterpart (at 0.39) and closer to the U.S. data (at 0.77, as reported by [Krueger, Mitman, and Perri 2016](#) for the PSID in 2006). A key reason for the increased wealth inequality is our model's characteristic dispersion in stable wealth points $\bar{y}_i$, which itself has a Gini coefficient of 0.49—and is also plotted in [Figure III](#), Panel B.²¹

V.B. Joint Distribution of Wealth and Consumption

[Corollary 2](#) implies that the typical consumption behavior should be systematically different across agents in the left and the right tail of the wealth distribution.

To evaluate this implication quantitatively, we aim to characterize the typical shape of the individual policy function estimates $\hat{c}_i(y)$ in the stationary distribution of the calibrated model. Given the significant heterogeneity in stable wealth levels $\bar{y}_i$, however, agents are not equally likely to visit all points of the state space y , but mostly stay in their individual-specific stable regions of the state space (e.g., see [Figure III](#) and [Online Appendix Table G.1](#)). Thus, to understand the typical behavior of agents, we cannot simply average policy functions state by state across all agents of the economy and compute $\int \hat{c}_i(y) di$ but need to account for the different positions of the stable wealth levels $\bar{y}_i$.

To illustrate the issue, in [Figure IV](#) we plot the long-run beliefs $\hat{c}_i(y)$ of three different agents at our benchmark calibration, labeled as $i \in \{c, u_1, u_2\}$, that have settled in three different learning trap regions of the state space. In Panel A, the thick solid line plots the ergodic estimate $\hat{c}_c(y)$ of an agent that has converged to a stable wealth level $\bar{y}_c$ against the borrowing constraint, while the dashed and the dot-plus lines plot the policy estimates $\hat{c}_{u_1}(y)$ and $\hat{c}_{u_2}(y)$ of two different agents that have established long-run wealth levels away from the borrowing limit (and are typically unconstrained). The filled-in circles denote the individual-specific

21. This heterogeneity in $\bar{y}_i$ reflects the persistent differences in the perceptions of the optimal consumption policy across agents in our model, and is consistent with empirical work that emphasizes the role of persistent unobserved heterogeneity ([Bernheim, Skinner, and Weinberg 2001](#); [Ameriks, Caplin, and Leahy 2003](#); [Hendricks 2007](#)).




FIGURE IV

Conditional Beliefs Evolution

Long-run conditional beliefs for three example agents, based on simulating 100 periods for each agent (without replacement) at benchmark calibration. The filled-in circles denote $\hat{c}_i(\bar{y}_i)$ (Panel A) and $\hat{\sigma}_i^2(\bar{y}_i)$ (Panel B).

locations of the stable wealth levels $\bar{y}_i$ and corresponding consumption choices there $\hat{c}_i(\bar{y}_i)$ for these three example agents.

Similarly, in Panel B, the solid, dashed, and dot-plus lines plot the corresponding ergodic uncertainty functions $\hat{\sigma}_i^2(y)$ of each agent. We denote the locations of the stable wealth levels $\bar{y}_i$ with filled-in circles, and can see that the characteristic U-shape of the uncertainty functions define large learning trap regions around each agent's specific stable wealth point $\bar{y}_i$, as the respective intervals where each $\hat{\sigma}_i^2(y) \leq \kappa$.

Two systematic differences emerge across the policy estimates in Panel A. First, the ergodic policy $\hat{c}_c(y)$ reflects the first type of stable behavior predicted by Corollary 2, characterized by an overly high level of consumption and being constrained in the long run—that is, $\hat{c}_c(\bar{y}_c) > \bar{y}_c > c^*(\bar{y}_c)$. On the other hand, the other two agents are unconstrained in the long run, with their ergodic wealth points instead defined by an upward crossing of their policy estimates and the PIH line $c^{RW}(y)$. If we simply average policy functions across all three agents, these differences will tend to wash away, and the resulting typical policy will not properly capture the fact that there are some high-consumption, often-constrained agents. To account for this difference, when computing typical policy shapes, we average policy functions within groups of agents stratified by their ergodic wealth levels.




FIGURE V

Typical Shape of Individual Policy Functions: $\int_i \hat{c}_i(y_i + x)di$

Second, now contrasting just the two example unconstrained agents with one another (dashed and dot-plus lines), notice that the shapes of their policy functions in the neighborhood of their individual $\bar{y}_i$ are quite similar, both being roughly linear and crossing the PIH from below. As such, these agents have similar consumption-savings behavior—for example, similar savings rates and MPCs—around the typical states each of them is likely to encounter.

To capture this commonality, we control for the individual fixed effects in typical cash-on-hand values different agents face, by averaging across policy estimates while holding constant the deviation from each agent's ergodic cash-on-hand value:

$$(17) \quad \hat{c}_{group}(x) \equiv \int_{i \in group} \hat{c}_i(y_i + x)di.$$

Here y_i is the wealth of agent i at the end of our simulation, and *group* is defined by wealth bins. This gives us the typical shape of an agent's consumption policy around the agent's wealth level in the stationary equilibrium and can speak to the average behavior of agents around the types of situations (i.e., cash on hand) they typically face.

In Figure V, Panel A we plot the policy shape thus extracted for agents who are classified as HtM (representative of the left tail of the wealth distribution), defined as agents who have less than two months of income in terms of assets, as per the usual way HtM status is defined in the data (e.g., Aguiar, Bils, and Boar 2020). In Panel B, we focus on agents in the top 20% of assets. The x -axis

TABLE I
KEY QUANTITATIVE RESULTS

Moments	Data (1)	Benchmark (2)	Full info (3)	MPC = 0.2 (4)
Panel A: Constrained agents				
Fraction of HtM	23%	22%	0.7%	21%
β_1 univariate reg.	0.2%	0.4%	1.5%	0.3%
β_2 multivariate reg.	2.3%	2.9%	1.5%	2.9%
γ multivariate reg.	-3.7%	-3.6%	-1.3%	-3.5%
Δ_{t+1}^c	-10.7%	-9.8%	-0.01%	-11.3%
Panel B: MPC				
Average MPC	0.2–0.6	0.30	0.05	0.34
Average MPC Q5(assets)	0.1–0.2	0.15	0.04	0.20
Average MPC non-HtM		0.16	0.05	0.21
Average MPC HtM		0.83	0.27	0.86

Notes. We report empirical moments in column (1) and from the model's stationary distribution in column (2). We set $\sigma_\epsilon^2 = 0$ in column (3) to obtain the full-information counterpart to our model. The data moments in Panel A are documented by Aguiar, Bils, and Boar (2020) Table 3, and Ganong and Noel (2019), Online Appendix Table 8. In Panel B the range of credible estimates of the aggregate MPC is from Carroll et al. (2017), and the MPC estimates of the rich and liquid agents cover a typical range documented by Gelman (2021) and McDowall (2020).

is in terms of deviations from the ergodic values of cash-on-hand (x in equation (17)), so zero represents the steady-state wealth; we plot a range of $\pm$ five standard deviations of income shocks (that is, $\sigma_l w = 0.24$).

1. *Large and Persistent Mass of HtM Agents.* First, we find that the policy $\hat{c}_{HtM}(y)$ differs from $c^*(y)$ mainly in its level. In particular,

$$(18) \quad \hat{c}_{HtM}(x) > c_{HtM}^{RW}(x) > c_{HtM}^*(x),$$

so these low-wealth agents tend to overestimate optimal consumption and on average decumulate assets (recall $\hat{c}_{HtM}(x) > c_{HtM}^{RW}(x)$ implies a negative saving rate). Thus, the behavior of these agents reflects the first type of local stability in Corollary 2 and is characterized by high average consumption which endogenously keeps assets low continuously.

i. *Low Target Wealth Level versus Transitory Low Income.* Due to these characteristic features of $\hat{c}_{HtM}(x)$, our model matches the empirical result that HtM status is both prevalent and persistent at the individual level, two challenging facts for standard models. To show this, Table I, Panel A reports that in the

stationary distribution of our model 22% of agents are HtM, lining up well with the 23% found in the data. In contrast, in the full-information model the fraction of HtM agents is just 0.7%. Second, to evaluate the persistence of HtM status, we connect to the empirical analysis of Aguiar, Bils, and Boar (2020). They regress the future two-year consumption growth of individual agents on a dummy variable that measures if the agent is currently HtM (i.e., $a_{i,t} < \frac{w}{6}$):

$$(19) \quad \ln c_{i,t+2} - \ln c_{i,t} = \beta_0 + \beta_1 \mathbb{1}(HtM_{i,t}) + \varepsilon_{i,t+2}.$$

Aguiar, Bils, and Boar (2020) estimate $\beta_1 = 0.2\%$ (Table 3 of their paper), which, being close to zero, runs counter to the standard FI model's implication that agents near the constraint are expected to save, and thus experience high future consumption growth (see Figure V, Panel A— $c^*(y)$ is below $c^{RW}(y)$ near the constraint). Consistent with this intuition, in the FI version of the model we find $\beta_1 = 1.5\%$. In contrast, we match the data (generating $\beta_1 = 0.4\%$) because in our model HtM status is primarily due to a low steady-state wealth level $\bar{y}_i$. In other words, the bulk of the HtM agents in our model are not actively saving their way out of the constraint—they are content to stay there, as that is what their stable estimate of optimal consumption implies. In turn, they do not experience high future consumption growth like in the FI model.

In fact, the hypothesis that HtM agents might indeed have different, lower steady-state (or target) wealth leads Aguiar, Bils, and Boar (2020) to augment the regression with the propensity to be HtM, defined as the fraction of periods that an agent is observed in a HtM situation

$$(20) \quad \ln c_{i,t+2} - \ln c_{i,t} = \beta_0 + \beta_2 \mathbb{1}(HtM_{i,t}) + \gamma \overline{\mathbb{1}(HtM_{i,t})}_i + \varepsilon_{i,t+2}.$$

Aguiar, Bils, and Boar (2020) find a strongly negative γ (at -3.7%), indicating that agents with frequent HtM status indeed have a systematically lower saving rate, and thus a low target wealth level. Moreover, after controlling for this propensity to be HtM, they find a significantly positive β_2 , equal to 2.3% , implying that temporarily poorer agents do save more.

This is what happens in our model, too. Many agents do not grow out of the constraint specifically because of their low stable wealth levels $\bar{y}_i$ and our model matches the data well, delivering $\beta_2 = 2.9\%$ and $\gamma = -3.6\%$. On the other hand, the FI model is sharply at odds with the data, as there agents save aggressively and HtM status is temporary.

ii. *The Puzzling Lack of Saving by Constrained Agents.*

The same characteristic features of our model also help it speak to the (related) challenging empirical fact that agents near the constraint undersave, even when they know that their established level of income is about to expire—for example, severance pay or expiring unemployment benefits (Ganong and Noel 2019; Farrell et al. 2020; Gerard and Naritomi 2021). Put differently, at the otherwise predictable moment the extra income is no longer available, consumption drops sharply, reflecting an apparent lack of preparation in terms of savings.

Our model can rationalize this observation because the majority of agents near the constraint are there due to their habitual high consumption behavior, as driven by a confident System 1. Until their income actually falls, such agents remain in a familiar financial context, no new reasoning is triggered, and they maintain their habitual, high-consumption behavior. However, to an outside observer, this might appear as a puzzling lack of preparation.

To quantify this force, we design an experiment meant to capture a predictable 40% drop in income to match the evidence on income drop following UI benefit exhaustion of Ganong and Noel (2019). Starting at the stationary distribution, we compute the average percentage change in consumption, as income falls by 40%: $\Delta_{t+1}^c = \int \frac{c_{i,t+1} - c_{i,t}}{c_{i,t}} di$. We perform the same experiment in the FI version and compute the implied change in consumption when agents follow the optimal consumption policy: $\Delta_{t+1}^{c^*} = \int \frac{c_{i,t+1}^* - c_{i,t}^*}{c_{i,t}^*} di$. To calibrate the experiment, we use the empirical moments Ganong and Noel (2019) provide in their Online Appendix Table 8, leading us to focus on the response of agents in the bottom half of the wealth distribution. We detail the mechanics of this exercise in Online Appendix H.

Our key result is that on average, the boundedly rational agents experience a substantial drop in consumption between t and $t + 1$ of $\Delta_{t+1}^c = -9.8\%$. This result is puzzling from the perspective of the full-information model, since those agents save most of the temporary income and successfully smooth consumption, resulting in $\Delta_{t+1}^{c^*} = -0.01\%$. The main driver of the result in our model is the behavior of the (large) fraction of HtM agents. As discussed earlier, until income actually changes, they are confident in their System 1 high consumption behavior and do not save a buffer stock of assets. When their labor income actually falls at

$t + 1$, they end up having to cut consumption significantly. As we report in Table I, this large drop in consumption is quite close to the average 10.7% drop in spending documented by Ganong and Noel (2019) at the exhaustion of unemployment insurance benefits.

2. High MPCs of Unconstrained Wealthy Agents. In standard models, as an agent's assets grow, her MPC converges to that implied by the PIH, that is, $\frac{r}{1+r} \approx 0.04$. In contrast to this result, in the data highly liquid households exhibit MPCs that are significantly larger than that (even though, naturally, they remain lower than those of poor, HtM agents). For example, Gelman (2021) and McDowall (2020) use two different types of detailed micro data to estimate the MPCs across the distribution of liquid assets. They find that even the households in the top 20% of liquid wealth have MPCs in the range of 0.1 to 0.2, which is significantly higher than $\frac{r}{1+r} \approx 0.04$.²²

Our model is consistent with this puzzle thanks to the second type of learning trap characterized in Corollary 2—away from the borrowing constraint, a stable steady state in individual beliefs must feature an effective policy function $\hat{c}_i(y)$, which is steeper than the PIH policy. In turn, this implies that MPCs are higher than $\frac{r}{1+r}$, the slope of the PIH line.

To visualize this high MPC result, in Figure V, Panel B plots the typical policy function shape of agents in the top quintile of wealth. It resembles the illustrative unconstrained agents plotted in Figure IV in two intuitive ways. First, local to the (high) ergodic wealth levels of these agents, their estimated policies lay mostly below the optimal policy ($\hat{c}_{Q5}(x) < c^*(x)$). In other words, the rich agents in the stationary distribution of the model are endogenously selected as those that underestimate the level of optimal consumption, have a high savings rate, and thus a high stable wealth level $\bar{y}_i > \bar{y}^*$. Second, in the neighborhood of their steady-state wealth levels, the typical effective policy of these agents $\hat{c}_{Q5}(x)$ is indeed steeper than both the PIH policy and $c^*(y)$.

As our benchmark calibration in Table I, column (2) shows, our model is able to match the puzzlingly high MPCs of the top 20% wealthiest agents in the data (we chose to target 0.15, the

22. It is important to differentiate between liquid and illiquid assets (Kaplan and Violante 2014), so we focus on the evidence of high MPCs at high liquid wealth (see also Lewis, Melcangi, and Pilossoph 2019; Fagereng, Holm, and Natvik 2021).

midpoint of the empirical estimates). While the model is calibrated to this moment, its ability to match it perfectly is still a notable quantitative success. Because this targets behavior in the right tail of the wealth distribution, the previous discussion of the model's ability to match well a number of challenging empirical regularities for agents in the left tail of the wealth distribution effectively constitutes evidence of external validity.

Moreover, [Table I](#), column (4) considers an alternative calibration that targets a higher MPC of 0.2 for the top 20% wealthiest agents, a value at the top end of the empirical estimates. There are two key takeaways. First, the model can indeed match this higher target, confirming that our mechanism can reliably deliver even larger MPCs. Second, this calibration does not jeopardize the (untargeted) behavioral implications for poorer agents, still delivering a large fraction of HtM agents (now 21% versus 22% in the benchmark), with the HtM status remaining similarly persistent.

V.C. Further Experiments

Here we briefly discuss a few additional experiments that highlight the endogenous nature of mistakes predicted by our model, with further details provided in the [Online Appendix](#).

1. *Higher Cost of Mistakes Near the Constraint.* One extension we consider is to allow for higher incentives to reason near the constraint, since the utility cost of mistakes is likely higher there due to the steepness in marginal utility. In [Online Appendix I](#) we detail this extension, calibrating the augmented model to the same target of an average 0.15 MPC for the agents in the top quintile of assets. Importantly, this extension still delivers a large fraction of HtM agents (18%) and the fact that HtM status is highly persistent.²³

2. *Control Cost.* To showcase the endogenous nature of mistakes in our model, we also consider another popular model of imperfect decision making, where the mistakes arise from a control-cost of implementing the known optimal action

23. In turn, evidence by [Mullainathan and Shafir \(2013\)](#) and [Mani et al. \(2013\)](#) suggests that due to anxiety and lack of free time, the cognitive cost of poor households may actually be higher than otherwise. Such a modeling feature would counteract the higher incentives of these agents to reason.

(Mattsson and Weibull 2002).²⁴ In this class of models, actions are probabilistic and follow a distribution $q(c_{i,t} | c^*(y_{i,t}))$ that the agent chooses optimally subject to an entropy cost. A key observable difference is that in our model, the dual-reasoning interaction leads to persistent biases. In contrast, in the control-cost version errors wash out on average and thus there are no learning traps and no systematic deviations from $c^*(y)$, such as the persistent HtM status and high MPC for the rich obtained by our model (Online Appendix J provides a more detailed discussion).

3. *Policy Effectiveness.* Our model's aggregate MPC of 0.30 (Table I), is in line with the empirical estimates, which range from 0.2 to 0.6 as per the review in Carroll et al. (2017), in contrast with the 0.05 value in the FI version. Decomposing the aggregate MPC in our model, we find that 70% of it is due to the HtM agents and 30% from the relatively high MPCs of rich agents. The effect of the HtM is partly due to their larger mass, but also partly due to the fact that our HtM agents have significantly higher MPCs than HtM agents under FI (0.83 versus 0.27). This is because in our model, poor agents largely do not try to save and grow out of the constraint. In that sense they are true HtMs in their consumption choices and freely spend any extra income, whereas poor FI agents are likely to save it.

Last, the endogenous nature of mistakes in our model relates to the Lucas (1976) critique. First, policy interventions that lead to large changes in the agents' state $y_{i,t}$ can jolt them to update their effective policy function. In addition, whether or not a policy maker is taking actions at all can itself be a separate state variable that may trigger reasoning.

For example, consider an augmented version of the model with the possibility of a one-time, unexpected stimulus payment τ_t . Due to its unexpected nature, the FI optimal policy rule of this augmented model remains $c^*(y_{i,t})$, as a function of cash on hand.

The boundedly rational agents, however, may not know ex ante that τ_t does not change the nature of the optimal policy function (which is unknown), and allow for the fact that optimal behavior might be different in a world where the government has decided the situation is dramatic enough to warrant fiscal stimulus. We model this by augmenting the state vector with an additional

24. See Flynn and Sastry (2020), (2021) for recent related analysis and applications to business cycle models.

dimension, modeled as an indicator function of whether there is stimulus, so that the effective state of the world is perceived as $\omega_{i,t} = [y_{i,t}, \mathbb{1}(\tau_t > 0)]$, and optimal behavior as an unknown $c^*(\omega_{i,t})$. [Online Appendix K](#) details the implementation.

In this augmented model, the act of receiving stimulus payment shifts the second dimension of the state variable. Thus, agents perceive an increase in uncertainty over the optimal action, even if τ_t and the resulting shift in $y_{i,t}$ are relatively small. This prompts agents to engage in costly active deliberation, rethinking their actions so that rather than continuing with the established automatic (System 1) behavior, beliefs shift toward $c^*(y)$ (and thus lower MPCs). Implementing this exercise numerically, we find that the aggregate MPC (and thus the effect of the stimulus) falls from 0.3 to 0.22. Virtually all of the remaining 0.22 aggregate MPC comes from the fact that while people reconsider behavior (which shifts their policy estimate toward $c^*(y)$ and lower MPCs), the model still inherits a wealth distribution featuring a large fraction of HtM agents. Those agents, even after actively rethinking behavior, remain in a steep part of their consumption function and display sizable MPCs. Overall, this exercise showcases that policy makers cannot always rely on agents mechanically sticking to their System 1 behavior, which in turn affects conclusions about policy effectiveness.

VI. CONCLUSION

We conclude by briefly discussing some further research directions that can build on the theoretical and applied contributions of this article. First, the model offers a concrete framework to guide novel measurement and identification of a wide set of realistic cognitive frictions, based on both experimental and field data. In particular, the model can be used to study and recover uncertainty about the optimal action, even conditional on given beliefs about the state,²⁵ and biases that arise endogenously, purely due to the dual-reasoning interaction, in addition to any static, primitive biases in System 1 and 2. Second, while the framework proposed here was developed to model costly reasoning about individual

25. This type of uncertainty has recently garnered attention in macroeconomics. In particular, [D'Acunto et al. \(forthcoming\)](#) find field evidence that is consistent with our model on the role of cognitive frictions in how agents map macroeconomic expectations to their implied optimal actions.

optimal policy functions, we envision extensions of this framework that also study agents' reasoning about how shocks are mapped into aggregate endogenous outcomes.²⁶

More broadly, our formalization of bounded rationality bridges the narrative and evidence of a broad psychology literature with the economics tools of information acquisition subject to costly cognitive resources. In addition, our application shows cases that the framework is tractable, providing tight and rich predictions on economic behavior driven by persistent cognitive biases that are not necessarily policy invariant. Altogether, we believe that the framework proposed here provides a promising step toward integrating psychologically founded bounded rationality within otherwise standard, quantitative economic models.

DUKE UNIVERSITY AND NATIONAL BUREAU OF ECONOMIC RESEARCH,
UNITED STATES
BOSTON COLLEGE, UNITED STATES

SUPPLEMENTARY MATERIAL

Supplementary material is available at the *Quarterly Journal of Economics* online.

DATA AVAILABILITY

Code replicating the tables and figures in this article can be found in [Ilut and Valchev \(2022\)](#) in the Harvard Dataverse, <https://doi.org/10.7910/DVN/PZMPBD>.

REFERENCES

- Aguiar, Mark, Mark Bilal, and Corina Boar, "Who Are the Hand-to-Mouth?," NBER Working Paper no. 26643, 2020, <https://doi.org/10.3386/w26643>.
- Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 109 (1994), 659–684.
- Alaoui, Larbi, and Antonio Penta, "Cost-Benefit Analysis in Reasoning," *Journal of Political Economy*, 130 (2022), 881–925.
- Ameriks, John, Andrew Caplin, and John Leahy, "Wealth Accumulation and the Propensity to Plan," *Quarterly Journal of Economics*, 118 (2003), 1007–1047.

26. Recent empirical interest in such extensions is illustrated by [Andre et al. \(forthcoming\)](#) who aim to identify agents' subjective perceptions over equilibrium propagation mechanisms. Consistent with our theory, that paper finds significant heterogeneity in those beliefs and a crucial role for associative memory in their formation.

- Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart, "Subjective Models of the Macroeconomy: Evidence from Experts and a Representative Sample," *Review of Economic Studies* (forthcoming).
- Aridor, Guy, Francesco Grechi, and Michael Woodford, "Adaptive Efficient Coding: A Variational Auto-Encoder Approach," bioRxiv (2020), <https://doi.org/10.1101/2020.05.29.124453>.
- Bago, Bence, and Wim De Neys, "Fast Logic?: Examining the Time Course Assumption of Dual Process Theory," *Cognition*, 158 (2017), 90–109.
- Barberis, Nicholas, and Lawrence Jin, "Model-Free and Model-Based Learning as Joint Drivers of Investor Behavior," Mimeo, 2021.
- Bardhi, Arjada, "Optimal Discovery and Influence through Selective Sampling," Mimeo, Duke University, 2022.
- Barr, Pamela S., John L. Stimpert, and Anne S. Huff, "Cognitive Change, Strategic Action, and Organizational Renewal," *Strategic Management Journal*, 13 (1992), 15–36.
- Berk, Robert H., "Limiting Behavior of Posterior Distributions When the Model Is Incorrect," *Annals of Mathematical Statistics*, 37, (1966), 51–58.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg, "What Accounts for the Variation in Retirement Wealth among US Households?," *American Economic Review*, 91 (2001), 832–857.
- Bertsekas, Dimitri P., *Reinforcement Learning and Optimal Control* (Belmont, CA: Athena Scientific, 2019).
- Bianchi, Francesco, Cosmin L. Ilut, and Hikaru Saijo, "Diagnostic Business Cycles," NBER Working Paper no. 28604, 2021, <https://doi.org/10.3386/w28604>.
- Bogner, William C., and Pamela S. Barr, "Making Sense in Hypercompetitive Environments: A Cognitive Explanation for the Persistence of High Velocity Competition," *Organization Science*, 11 (2000), 212–226.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Memory, Attention, and Choice," *Quarterly Journal of Economics*, 135 (2020), 1399–1442.
- Bornstein, Aaron M., Mel W. Khaw, Daphna Shohamy, and Nathaniel D. Daw, "Reminders of Past Choices Bias Decisions for Reward in Humans," *Nature Communications*, 8 (2017), 1–9.
- Bornstein, Aaron M., and Kenneth A. Norman, "Reinstated Episodic Context Guides Sampling-Based Decisions for Reward," *Nature Neuroscience*, 20 (2017), 997–1003.
- Botvinick, Matthew M., Jonathan D. Cohen, and Cameron S. Carter, "Conflict Monitoring and Anterior Cingulate Cortex: An Update," *Trends in Cognitive Sciences*, 8 (2004), 539–546.
- Boutros, Michael, "Windfall Income Shocks with Finite Planning Horizons," Mimeo, Duke University, 2021.
- Bushong, Benjamin, and Tristan Gagnon-Bartsch, "Learning with Misattribution of Reference Dependence," Mimeo, Harvard University, 2016.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White, "The Distribution of Wealth and the Marginal Propensity to Consume," *Quantitative Economics*, 8 (2017), 977–1020.
- Cerigioni, Francesco, "Dual Decision Processes: Retrieving Preferences When Some Choices Are Automatic," *Journal of Political Economy*, 129 (2021), 1667–1704.
- Chase, William G., and Herbert A. Simon, "Perception in Chess," *Cognitive Psychology*, 4 (1973), 55–81.
- Christianson, Marlys K., Maria T. Farkas, Kathleen M. Sutcliffe, and Karl E. Weick, "Learning through Rare Events: Significant Interruptions at the Baltimore & Ohio Railroad Museum," *Organization Science*, 20 (2009), 846–860.
- Conlisk, John, "Why Bounded Rationality?," *Journal of Economic Literature*, 34 (1996), 669–700.
- D'Acunto, Francesco, Daniel Hoang, Maritta Paloviita, and Michael Weber, "IQ, Expectations, and Choice," *Review of Economic Studies* (forthcoming).
- De Neys, Wim, "Conflict Detection, Dual Processes, and Logical Intuitions: Some Clarifications," *Thinking & Reasoning*, 20 (2014), 169–187.

- DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, 47 (2009), 315–372.
- DellaVigna, Stefano, Attila Lindner, Balázs Reizer, and Johannes F. Schmieder, "Reference-Dependent Job Search: Evidence from Hungary," *Quarterly Journal of Economics*, 132 (2017), 1969–2018.
- Dew-Becker, Ian, and Charles G. Nathanson, "Directed Attention and Nonparametric Learning," *Journal of Economic Theory*, 181 (2019), 461–496.
- Enke, Benjamin, Uri Gneezy, Brian Hall, David Martin, Vadim Nelidov, Theo Offerman, and Jeroen van de Ven, "Cognitive Biases: Mistakes or Missing Stakes?," *Review of Economics and Statistics* (forthcoming).
- Enke, Benjamin, and Thomas Graeber, "Cognitive Uncertainty," NBER Working Paper no. 26518, 2019, <https://doi.org/10.3386/w26518>.
- Evans, Jonathan St. B.T. *Hypothetical Thinking: Dual Processes in Reasoning and Judgement* (Washington, DC: Psychology Press, 2007a).
- , "On the Resolution of Conflict in Dual Process Theories of Reasoning," *Thinking & Reasoning*, 13 (2007b), 321–339.
- , "Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition," *Annual Review of Psychology*, 59 (2008), 255–278.
- , "How Many Dual-Process Theories Do We Need? One, Two, or Many?," in *In Two Minds: Dual Processes and Beyond* Jonathan St. B. T. Evans and Keith Frankish, eds. (Oxford: Oxford University Press, 2009), 33–54.
- , "Reflections on Reflection: The Nature and Function of Type 2 Processes in Dual-Process Theories of Reasoning," *Thinking & Reasoning*, 25 (2019), 383–415.
- Evans, Jonathan St. B. T., and Keith E. Stanovich, "Dual-Process Theories of Higher Cognition: Advancing the Debate," *Perspectives on Psychological Science*, 8 (2013), 223–241.
- Fagereng, Andreas, Martin Blomhoff Holm, and Gisle James James Natvik, "MPC Heterogeneity and Household Balance Sheets," *AEJ: Macroeconomics*, 13 (2021), 1–54.
- Farrell, Diana, Peter Ganong, Fiona Greig, Max Liebeskind, Pascal Noel, and Joseph Vavra, "Consumption Effects of Unemployment Insurance during the COVID-19 Pandemic," BFI Working Paper, 2020, <http://dx.doi.org/10.2139/ssrn.3654274>.
- Fernandes, Daniel, John G. Lynch Jr., and Richard G. Netemeyer, "Financial Literacy, Financial Education, and Downstream Financial Behaviors," *Management Science*, 60 (2014), 1861–1883.
- Flynn, Joel P., and Karthik Sastry, "Attention Cycles," MIT Working Paper, 2020, <http://dx.doi.org/10.2139/ssrn.3592107>.
- , "Strategic Mistakes," MIT Working Paper, 2021.
- Frydman, Cary, and Lawrence J. Jin, "Efficient Coding and Risky Choice," *Quarterly Journal of Economics*, 137 (2022), 161–213.
- Gabaix, Xavier, "Behavioral Inattention," in "Handbook of Behavioral Economics: Applications and Foundations 1," vol. 2, Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds. (Amsterdam: Elsevier, 2019), 261–343.
- Gabaix, Xavier, and David Laibson, "Myopia and Discounting," NBER Working Paper no. 23254, 2017, <https://doi.org/10.3386/w23254>.
- Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein, "Channeled Attention and Stable Errors," Mimeo, Harvard University, 2020.
- Ganong, Peter, and Pascal Noel, "Consumer Spending during Unemployment: Positive and Normative Implications," *American Economic Review*, 109 (2019), 2383–2424.
- Gelman, Michael, "What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics," *Journal of Monetary Economics*, 117 (2021), 521–542.
- Gennaioli, Nicola, and Andrei Shleifer, "What Comes to Mind," *Quarterly Journal of Economics*, 125 (2010), 1399–1433.

- Gerard, François, and Joana Naritomi, "Job Displacement Insurance and (the Lack of) Consumption-Smoothing," *American Economic Review*, 111 (2021), 899–942.
- Gershman, Samuel J., and Nathaniel D. Daw, "Reinforcement Learning and Episodic Memory in Humans and Animals: An Integrative Framework," *Annual Review of Psychology*, 68 (2017), 101–128.
- Gershman, Samuel J., Eric J. Horvitz, and Joshua B. Tenenbaum, "Computational Rationality: A Converging Paradigm for Intelligence in Brains, Minds, and Machines," *Science*, 349 (2015), 273–278.
- Gilboa, Itzhak, and David Schmeidler, "Case-Based Decision Theory," *Quarterly Journal of Economics*, 110 (1995), 605–639.
- , *A Theory of Case-Based Decisions* (Cambridge: Cambridge University Press, 2001).
- Girshick, Ahna R., Michael S. Landy, and Eero P. Simoncelli, "Cardinal Rules: Visual Orientation Perception Reflects Knowledge of Environmental Statistics," *Nature Neuroscience*, 14 (2011), 926–932.
- Griffiths, Thomas L., Falk Lieder, and Noah D. Goodman, "Rational Use of Cognitive Resources: Levels of Analysis between the Computational and the Algorithmic," *Topics in Cognitive Science*, 7 (2015), 217–229.
- Handley, Simon J., and Dries Trippas, "Dual Processes and the Interplay between Knowledge and Structure: A New Parallel Processing Model," *Psychology of Learning and Motivation*, 62 (2015), 33–58.
- Heidhues, Paul, Botond Köszegi, and Philipp Strack, "Unrealistic Expectations and Misguided Learning," *Econometrica*, 86 (2018), 1159–1214.
- Hendricks, Lutz, "Retirement Wealth and Lifetime Earnings," *International Economic Review*, 48 (2007), 421–456.
- Holtom, Brooks C., Terence R. Mitchell, Thomas W. Lee, and Edward J. Inderrieden, "Shocks as Causes of Turnover: What They Are and How Organizations Can Manage Them," *Human Resource Management*, 44 (2005), 337–352.
- Ilut, Cosmin, and Rosen Valchev, "Replication Data for: 'Economic Agents as Imperfect Problem Solvers'," (2022), Harvard Dataverse, <https://doi.org/10.7910/DVN/PZMPBD>.
- Ilut, Cosmin, Rosen Valchev, and Nicolas Vincent, "Paralyzed by Fear: Rigid and Discrete Pricing under Demand Uncertainty," *Econometrica*, 88 (2020), 1899–1938.
- Judd, Kenneth L., *Numerical Methods in Economics* (Cambridge, MA: MIT Press, 1998).
- Kahana, Michael Jacob, *Foundations of Human Memory* (Oxford: Oxford University Press, 2012).
- Kahneman, Daniel, *Thinking, Fast and Slow* (New York: Macmillan, 2011).
- Kahneman, Daniel, and Shane Frederick, "Representativeness Revisited: Attribute Substitution in Intuitive Judgment," in *Heuristics and Biases: The Psychology of Intuitive Judgment*, Thomas Gilovich, Dale W. Griffin, and Daniel Kahneman, eds. (Cambridge: Cambridge University Press, 2002), 49–81.
- Kaplan, Greg, and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82 (2014), 1199–1239.
- Klein, Gary A., Roberta Calderwood, and Anne Clinton-Cirocco, "Rapid Decision Making on the Fire Ground," in *Proceedings of the Human Factors Society Annual Meeting* (Los Angeles, CA: Sage Publications, 1986), 576–580.
- Kool, Wouter, Samuel J. Gershman, and Fiery A. Cushman, "Cost-Benefit Arbitration between Multiple Reinforcement-Learning Systems," *Psychological Science*, 28 (2017), 1321–1333.
- Krueger, Dirk, Kurt Mitman, and Fabrizio Perri, "Macroeconomics and Household Heterogeneity," in *Handbook of Macroeconomics*, vol. 2, John B. Taylor and Harald Uhlig, eds. (Amsterdam: Elsevier, 2016), 843–921.
- Kueng, Lorenz, "Excess Sensitivity of High-Income Consumers," *Quarterly Journal of Economics*, 133 (2018), 1693–1751.
- Kunda, Ziva, "The Case for Motivated Reasoning," *Psychological Bulletin*, 108 (1990), 480.

- Lampel, Joseph, Jamal Shamsie, and Zur Shapira, "Experiencing the Improbable: Rare Events and Organizational Learning," *Organization Science*, 20 (2009), 835–845.
- Landry, Peter, Ryan Webb, and Colin Camerer, "A Neural Autopilot Theory of Habit," Mimeo, 2021.
- Lee, Sang Wan, Shinsuke Shimojo, and John P. O'Doherty, "Neural Computations Underlying Arbitration between Model-Based and Model-Free Learning," *Neuron*, 81 (2014), 687–699.
- Lee, Thomas W., and Terence R. Mitchell, "The Unfolding Effects of Organizational Commitment and Anticipated Job Satisfaction on Voluntary Employee Turnover," *Motivation and Emotion*, 15 (1991), 99–121.
- , "An Alternative Approach: The Unfolding Model of Voluntary Employee Turnover," *Academy of Management Review*, 19 (1994), 51–89.
- Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph, "Latent Heterogeneity in the Marginal Propensity to Consume," FRB of New York Staff Report no. 902, 2019, <http://dx.doi.org/10.2139/ssrn.3489434>.
- Lian, Chen, "Mistakes in Future Consumption, High MPCs Now," MIT Working Paper, 2020, <http://dx.doi.org/10.3386/w29517>.
- Liu, Miao, "Assessing Human Information Processing in Lending Decisions: A Machine Learning Approach," Mimeo, The University of Chicago, 2021, <http://dx.doi.org/10.2139/ssrn.3483699>.
- Liu, Weifeng, Jose C. Principe, and Simon Haykin, *Kernel Adaptive Filtering: A Comprehensive Introduction* (New York: John Wiley & Sons, 2011).
- Louis, Meryl Reis, and Robert I. Sutton, "Switching Cognitive Gears: From Habits of Mind to Active Thinking," *Human Relations*, 44 (1991), 55–76.
- Lucas, Robert E., "Econometric Policy Evaluation: A Critique," *Carnegie-Rochester Conference Series on Public Policy*, 1 (1976), 19–46.
- Luo, Yulei, "Consumption Dynamics under Information Processing Constraints," *Review of Economic Dynamics*, 11 (2008), 366–385.
- Mackowiak, Bartosz, Filip Matejka, and Mirko Wiederholt, "Rational Inattention: A Review," *Journal of Economic Literature* (forthcoming).
- Maćkowiak, Bartosz, and Mirko Wiederholt, "Business Cycle Dynamics under Rational Inattention," *Review of Economic Studies*, 82 (2015), 1502–1532.
- Maitlis, Sally, and Marlys Christianson, "Sensemaking in Organizations: Taking Stock and Moving Forward," *Academy of Management Annals*, 8 (2014), 57–125.
- Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao, "Poverty Impedes Cognitive Function," *Science*, 341 (2013), 976–980.
- Maréchal, Kevin, "Not Irrational But Habitual: The Importance of 'Behavioural Lock-in' in Energy Consumption," *Ecological Economics*, 69 (2010), 1104–1114.
- Mattsson, Lars-Göran, and Jörgen W. Weibull, "Probabilistic Choice and Procedurally Bounded Rationality," *Games and Economic Behavior*, 41 (2002), 61–78.
- McDowall, Robert A., "Consumption Behavior across the Distribution of Liquid Assets," NYU Working Paper, 2020.
- Meyer, Alan D., "Adapting to Environmental Jolts," *Administrative Science Quarterly*, 27 (1982), 515–537.
- Mullainathan, Sendhil, "A Memory-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 117 (2002), 735–774.
- Mullainathan, Sendhil, and Eldar Shafir, *Scarcity: Why Having Too Little Means so Much* (New York: Macmillan, 2013).
- O'Donoghue, Ted, and Matthew Rabin, "Doing It Now or Later," *American Economic Review*, 89 (1999), 103–124.
- Olafsson, Arna, and Michaela Pagel, "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *Review of Financial Studies*, 31 (2018), 4398–4446.

- Pennycook, Gordon, "A Perspective on the Theoretical Foundation of Dual Process Models," *Dual Process Theory*, 2 (2017), 34.
- Plonsky, Ori, Kinneret Teodorescu, and Ido Erev, "Reliance on Small Samples, the Wavy Recency Effect, and Similarity-Based Learning," *Psychological Review*, 122 (2015), 621.
- Rabin, Matthew, and Joel L. Schrag, "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics*, 114 (1999), 37–82.
- Rasmussen, Carl Edward, and Christopher K. I. Williams, *Gaussian Processes for Machine Learning*, vol. 1 (Cambridge, MA: MIT Press, 2006).
- Reis, Ricardo, "Inattentive Consumers," *Journal of Monetary Economics*, 53 (2006), 1761–1800.
- Robson, Arthur J., and Lorne A. Whitehead, "Adaptive Cardinal Utility," Simon Fraser University Working Paper, 2019.
- Rubinstein, Ariel. *Modeling Bounded Rationality* (Cambridge, MA: MIT press, 1998).
- Schwartzstein, Joshua, "Selective Attention and Learning," *Journal of the European Economic Association*, 12 (2014), 1423–1452.
- Sent, Esther-Mirjam, and Matthias Klaes, "A Conceptual History of the Emergence of Bounded Rationality," *History of Political Economy*, 37 (2005), 27–59.
- Shenhav, Amitai, Sebastian Musslick, Falk Lieder, Wouter Kool, Thomas L. Griffiths, Jonathan D. Cohen, and Matthew M. Botvinick, "Toward a Rational and Mechanistic Account of Mental Effort," *Annual Review of Neuroscience*, 40 (2017), 99–124.
- Shleifer, Andrei, "Psychologists at the Gate: A Review of Daniel Kahneman's Thinking, Fast and Slow," *Journal of Economic Literature*, 50 (2012), 1080–1091.
- Simon, Herbert A., "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69 (1955), 99–118.
- , "Rational Choice and the Structure of the Environment," *Psychological Review*, 63 (1956), 129.
- , "From Substantive to Procedural Rationality," in *25 Years of Economic Theory*, T. J. Kastelein, S. K. Kuipers, W. A. Nijenhuis, and G. R. Wagenaar, eds. (Leiden: Martinus Nijhoff, 1976), 65–86.
- Sims, Christopher A., "Stickiness," in *Carnegie-Rochester Conference Series on Public Policy*, 49 (1998), 317–356.
- , "Implications of Rational Inattention," *Journal of Monetary Economics*, 50 (2003), 665–690.
- Stanovich, Keith, *Rationality and the Reflective Mind* (Oxford: Oxford University Press, 2011).
- , "Miserliness in Human Cognition: The Interaction of Detection, Override and Mindware," *Thinking & Reasoning*, 24 (2018), 423–444.
- Stanovich, Keith E., and Richard F. West, "Individual Differences in Reasoning: Implications for the Rationality Debate?," *Behavioral and Brain Sciences*, 23 (2000), 645–665.
- Thaler, Richard H., *The Winner's Curse: Paradoxes and Anomalies of Economic Life* (New York: Simon and Schuster, 1992).
- Thompson, Valerie A., Jamie A. Prowse Turner, and Gordon Pennycook, "Intuition, Reason, and Metacognition," *Cognitive Psychology*, 63 (2011), 107–140.
- Thompson, Valerie A., Jamie A. Prowse Turner, Gordon Pennycook, Linden J. Ball, Hannah Brack, Yael Ophir, and Rakefet Ackerman, "The Role of Answer Fluency and Perceptual Fluency as Metacognitive Cues for Initiating Analytic Thinking," *Cognition*, 128 (2013), 237–251.
- Todd, Peter M., and Gerd Gigerenzer, "Bounding Rationality to the World," *Journal of Economic Psychology*, 24 (2003), 143–165.
- Tulving, Endel, "Episodic and Semantic Memory," in *Organization of Memory*, Endel Tulving and W. Donaldson, eds. (New York: Academic, 1972), 381–403.
- , "Episodic Memory: From Mind to Brain," *Annual Review of Psychology*, 53 (2002), 1–25.

- Tversky, Amos, and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," in *Utility, Probability, and Human Decision Making*, Dirk Wednt and Charles Vlek, eds. (Dordrecht: Reidel, 1975), 141–162.
- Tyre, Marcie J., and Wanda J. Orlikowski, "Windows of Opportunity: Temporal Patterns of Technological Adaptation in Organizations," *Organization Science*, 5 (1994), 98–118.
- Verplanken, Bas, and Deborah Roy, "Empowering Interventions to Promote Sustainable Lifestyles: Testing the Habit Discontinuity Hypothesis in a Field Experiment," *Journal of Environmental Psychology*, 45 (2016), 127–134.
- Verplanken, Bas, Ian Walker, Adrian Davis, and Michaela Jurasek, "Context Change and Travel Mode Choice: Combining the Habit Discontinuity and Self-Activation Hypotheses," *Journal of Environmental Psychology*, 28 (2008), 121–127.
- Verplanken, Bas, and Wendy Wood, "Interventions to Break and Create Consumer Habits," *Journal of Public Policy & Marketing*, 25 (2006), 90–103.
- Walsh, James P., "Managerial and Organizational Cognition: Notes from a Trip Down Memory Lane," *Organization Science*, 6 (1995), 280–321.
- Wei, Xue-Xin, and Alan A. Stocker, "A Bayesian Observer Model Constrained by Efficient Coding Can Explain 'Anti-Bayesian' Percepts," *Nature Neuroscience*, 18 (2015), 1509.
- Wiederholt, Mirko, "Rational Inattention," in *The New Palgrave Dictionary of Economics*, vol. 4, Steven N. Durlauf and Lawrence E. Blume, eds. (London: Palgrave Macmillan, 2010).
- Wood, Wendy, Jennifer S. Labrecque, Pei-Ying Lin, and Dennis Runger, "Habits in Dual Process Models," In *Dual Process Theories of the Social Mind*, Jeffrey W. Sherman, Bertram Gawronski, and Yaacov Trope, eds. (New York: Guilford Press, 2014), 371–85.
- Wood, Wendy, Jeffrey M. Quinn, and Deborah A. Kashy, "Habits in Everyday Life: Thought, Emotion, and Action," *Journal of Personality and Social Psychology*, 83 (2002), 1281.
- Wood, Wendy, Leona Tam, and Melissa Guerrero Witt, "Changing Circumstances, Disrupting Habits," *Journal of Personality and Social Psychology*, 88 (2005), 918.
- Woodford, Michael, "Modeling Imprecision in Perception, Valuation, and Choice," *Annual Review of Economics*, 12 (2020), 579–601.