The Complexity of Economic Decisions*

Xavier Gabaix Thomas Graeber

August 28, 2025

Abstract

This paper develops a calibrated theory of economic decision complexity. For any decision problem, the framework delivers a measure of its complexity with no free parameters, making it highly portable across settings. We illustrate its use in central economic tasks, including dealing with taxes, forecasting, consumption, intertemporal financial planning, and the choice of financial products with opaque features.

The microfoundation builds on a macro-style production framework that treats the mind as a cognitive economy. Complexity is defined as the inverse productivity of thought: it increases with the number of importance-weighted components and decreases when outcomes hinge on a small set of dominant factors. More complex problems generate larger errors and reduced sensitivity to underlying parameters. We complement the theory with an experimental methodology for measuring subjective perceptions of complexity, and an experiment that confirms its main predictions. The measure provides a portable formula for predicting ex ante and quantifying ex post decision complexity across economic domains, potentially useful both theoretical and empirical applications.

*xgabaix@fas.harvard.edu, tgraeber@hbs.edu. For excellent research assistance we thank Dustin Fichmann, Lilian Hartmann and especially Constantin Schesch, and for helpful comments Pedro Bordalo, Emmanuel Chemla, Alex Imas, Emir Kamenica, David Laibson, Luba Petersen, Indira Puri, Matthew Rabin, Philippe Schlenker, Simeon Schudy, Bennett Smith-Worthington, Dmitry Taubinsky, Richard Thaler, Mike Woodford, Leeat Yariv, and seminar participants at various seminars. This research was approved by Harvard IRB (IRB17-2035). Gabaix gratefully acknowledges financial support from the Sloan Foundation and Ferrante Economics Research Fund.

1 Introduction

This paper provides a closed-form theory of the complexity of standard economic decisions and its implications for behavior. The guiding thread underlying our model is that the insights, thoughts and calculations underlying economic actions and beliefs are *produced* by the mind, akin to a cognitive factory. To represent that, we apply the macroeconomic framework of production functions to the individual decision-making problem and combine it with standard notions of imprecise cognitive processing adapted from the cognitive sciences. As a car firm uses labor, steel and plastic as inputs, the mind, to produce a thought, uses mental labor, and thoughts about subproblems, that in turn rely on more mental labor, and thoughts about sub-subproblems. In a related way, macroeconomics conceptualizes the total efficiency of a process via total factor productivity. Here, complexity – the difficulty of producing a thought – will be a sort of inverse total factor productivity. However, such mental operations tend to be imprecise and are a prone to error, which can lead to systematic and predictable deviations from optimality (Gabaix, 2019; Woodford, 2020). We propose a combination of the concepts of production functions and imprecise cognition, harnessing methods that that the economic toolbox is already well adapted to while accommodating the richness of key insights from the cognitive sciences.

A central goal of our approach is to provide a measure of complexity applicable to any given problem. This measure can be employed in a variety of contexts to assess the complexity of situations, goods, financial products, or consumption decisions—whether static or intertemporal. We anticipate that researchers may adopt our measure and examine its correlation with observable outcomes. In particular, our framework naturally implies that greater complexity is associated with a higher likelihood of errors and, under certain conditions, longer decision times.[1]

We first present the "plug-and-play" version of our measure, followed by a detailed discussion of its microfoundations. We start with a one-dimensional decision a: the agent's problem is to maximize

$$\max_a u\left(a, x_1, \ldots, x_n\right)$$

where the x_i's are disturbances, drawn from mean-zero distributions. Given this, we calculate the following quantities:

- $\frac{\partial a^r}{\partial x_i} := -u_{aa}^{-1} u_{ax_i}$: the normative marginal impact of x_i on a

[1] The condition will be, $\alpha > 1$.

- $y_i := \frac{\partial a^r}{\partial x_i} x_i$: the normative total impact of x_i on a
- $s_i := \frac{y_i^2}{\sum_j y_j^2}$: the "importance share" of variable x_i for decision a — more precisely, the share of action variance coming from x_i. This is

The last input is a "micro" cost c_i to think about variable i. Typically, we set it to $c_i = 1$; but can be endogenized later, or measured experimentally.

Our model yields the following central expression for the macro complexity of a problem as a function of the above primitives:

$$\mathcal{C} = \left(\sum_{i=1}^{N} s_i^{\frac{1}{\alpha}} c_i^{\frac{\alpha-1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}} \tag{1}$$

where α is a curvature parameter that we calibrate later.

So, the macro complexity $\mathcal{C}$ is a sort of non-linear sum of the underlying micro-complexities c_i, weighed by the importance $s_i^{\frac{1}{\alpha}}$ of the various dimensions.

This expression (1) can be used for potentially many decision problems, as we illustrate below. In particular, one of its main raisons d'être is to be useful for empirical work.

We derive the following general propositions about complexity $\mathcal{C}$. First, higher compositionality leads to more complexity: problem complexity increases with the number of dimensions N, but more precisely, as a function of the "effective" number of dimensions, where important dimensions (with high s_i) count more than less important ones. Second, complexity decreases in the relative extremity of individual dimensions: as the importance share of the most important dimensions rises, complexity decreases. Intuitively, the agent can focus their attention more and more on that dominating dimension: in the limit where there is only one very large dimension, only that dimension matters: if $s_1 \simeq 1$, then $\mathcal{C} \simeq c_1$. Third, the central behavioral implications of complexity are larger errors and behavioral attenuation: as complexity rises, people's actions becomes less precise, and less sensitive to variation in problem parameters, relative to the corresponding sensitivity of the rational action. Moreover, complexity increases the subjective feeling of error, cognitive uncertainty (Enke and Graeber, 2023).

Why should economists care about complexity? Even though complexity intuitively matters for key economic outcomes and policymaking, there is a dearth of theories about the complexity *of economic decisions* to guide and discipline empirical work. Existing empirical work (including our own) often invokes complexity as a catch-all concept, a reduced-form notion or a residual explanation, and had to rely on *ad hoc* definitions of complexity for a specific application under consideration. In theory and applied work,

there is a collection of "local" complexity definitions that are not applicable outside of their narrowly defined contexts (see review below).

We attempt to provide a domain-general theory, starting from an elementary definition of complexity. We posit a cognitive production function that takes cognitive labor L as an input to generate an expected level of performance $q \in [0, 1]$ on a decision task by means of a cognitive strategy. The primary objective of this paper is to endogenize this cognitive production function. We define two notions of complexity. *Input-based* complexity is the mental effort required to achieve performance level q. *Output-based* complexity is the performance shortfall – relative to perfect accuracy – of the response reached given mental effort L. Notably, this setup proposes a continuous notion of complexity: complexity is a function, defined for a given level of performance or cognitive labor. Complexity is thus a high-dimensional object, hence potentially quite untractable, so we propose a one-dimensional measure, by imagining a world with a Cobb-Douglas production function of thought. Then, the one-dimensional measure of complexity is a form of inverse total factor productivity. This is the key object on which calculations can be made.

We then proceed to the more concrete task of the paper, deriving complexity from primitives. The basic building blocks of our model are the *components* of a problem: each dimension i captures a state of the world or problem parameter x_i that the utility-maximizing action depends on. We study linear-quadratic setups that cover all smooth problems by Taylor approximation. The agent receives noisy cognitive signals about how each dimension i should change their rational action. We interpret such processing noise not as perceptual noise about problem parameters but rather as integration noise that occurs when people form mental representations of problem quantities in determining how those affect their action. Agents account for such cognitive processing noise in an (as-if) Bayesian fashion by integrating the signal with a suitable prior. This approach to imprecise cognition leverages a large body of work in the cognitive sciences as well as a recent theoretical and empirical literature in economics as reviewed below.

Applications A central objective of our model is to provide an easy-to-use approach to model and test for complexity in concrete economic applications, e.g. empirical or theoretical — not just for behavioral theory. We provide a user's guide that delineates the basic template for applying our model to practical economic settings. We then implement this approach to study, through the lens of our model, the complexity of basic static consumption, of the tax system, and of consumption planning over the life cycle, as well as the complexity of financial products with confusing features.

We revisit the basic theory of consumption, where a consumer chooses a bundle of

goods.[2] We ask: what is the complexity of choosing a good, or N goods?

The conclusions are the following: Even choosing a good is hard, as we typically cannot precisely identify a key parameter, the elasticity of demand. Consider the following question: The price of coffee goes up by 10%, how much would you decrease your demand for coffee? Introspection shows that this is hard to know – even though basic economics essentially always assumes perfect knowledge of this.

Then, we move on to the complexity of choosing two goods, with a "linear good" that absorbs income effects. When utility is separable, the problem is relatively easy: choose one good at a time. As the degree of complementarity (or substitutability) increases, complexity is higher. Suppose that I like eating fish with lemon. Then, my optimal consumption of fish depends on the price of fish, and my consumption of lemon. And vice-versa: my optimal consumption of lemon depends on the price of lemon, and my consumption of fish. So we have a "loop" — a fixed point problem — which adds complexity.

So, one prediction is that people will make more mistakes when goods have higher complementarity of substitutability — this captures that there are many moving parts that interact with each other.[3]

We then study the complexity of intertemporal consumption. Our model captures the intuition that optimally responding to a \$1 change in permanent income is easy, as this is achieved by increasing consumption today by \$1 – the answer exhibits zero "compositionality" in our model. However, reacting optimally to a \$1 change in transitory income is harder, as it requires taking into account the number of consumption periods. Reacting to a change in the interest rate is harder yet, as it requires introspecting about one's elasticity of substitution in addition to the number of periods. Hence, even controlling for objective stakes, the interest rate is more complex to incorporate than income. In practice, that implies that people are more reactive to income shocks than interest rate shocks, compared to a rational benchmark.

Our applications demonstrate how our model delivers testable predictions for classic economic decision problems. Complementing the theoretical approach, we conduct an empirical exercise with two main objectives: first, we provide a portable methodology to experimentally measure subjective perceptions of complexity. Second, we test the basic predictions of our theory in the case of intertemporal consumption choice.

[2]The prior conclusions in Gabaix (2014) and Lian (2021) still hold, but we can now ask about the complexity of such decisions.

[3]This also would give a theory of "menus of bundles" — simplifications where restaurants propose a pre-arranged bundles: this simplifies the decision.

Measuring subjective perceptions of complexity We propose a theory-guided procedure to measure the subjective feeling of complexity in a manner that is widely applicable and portable across decision contexts. We ask subjects to make pairwise complexity comparisons between tasks they have previously worked on. Specifically, we ask them to judge in which past task they found it more difficult to determine an answer that falls within a specific distance to the (suitably defined) optimal response. We estimate complexity scores statistically by maximum likelihood. To obtain cardinal complexity scores from these pairwise judgments, we leverage two calibration tasks that can be freely selected and held constant across experimental contexts (for instance, mentally computing $7 + 53$ and $7 + 53 + 394 + 7893$). By (arbitrarily) normalizing complexity in terms of the calibration tasks, we obtain absolute complexity scores for different decision problems.

We then conduct a simple proof-of-concept experiment on the theory's predictive validity. We find that, when the model predicts that a decision problem is more complex, people indeed (i) perceive it as more complex – this way, validating the model's core feature –, (ii) commit larger errors, (iii) respond in a more attenuated way to variation in parameters, and (iv) spend more time on the problem.

Extensions and limitations We provide the following extensions of our theory. First, we endogenize micro-complexities by modeling further layers in the hierarchy of the cognitive production function and extend the model to multi-dimension actions. Second, while the baseline model is formulated for continuous choice, in Section 5 we provide a way to re-factor all our continuous-action results to apply to the discrete-action case, such as consumer choice between multiple goods or standard choice under risk paradigms. We obtain a generally applicable notion of complexity for discrete choices. When one alternative is clearly dominating and thus more tempting than the others, choosing is easy. When there are several roughly equally tempting alternatives, with many attributes, choosing is more complex. This result squares well with the widely documented empirical finding that response times and error rates increase as stimuli or the value of choice options become more similar (e.g., Li and Camerer, 2022).

We point out several limitations of our approach. Our model is cognitive in nature and thus does not speak to motivational biases, such as explored in the literature on motivated beliefs. We also do not model flawed intuitions in the sense of "what comes to mind," which has been the focus of much classic work (à la Kahneman and Tversky, 1979) as well as more recent advances in behavioral economics (Bordalo et al., 2012). Instead, we take the mental representation of a problem as given and explore the implications of allocating cognitive labor within that representation to different components in a bound-

edly rational fashion. Moreover, while we discipline our theory using two main building blocks – macroeconomic production functions and imprecise cognitive processing – we do not provide an axiomatization of our approach. We propose that the analogy to production functions is useful because behavioral economics sometimes faces the challenge that it relies on different levels of abstractions than most economic theory, such as memories or neuroeconomic processes, making it hard to incorporate these concepts broadly.

Related literature We provide a (necessarily incomplete) overview of connections to previous work here and defer a more extensive discussion of the voluminous existing literature on complexity to Appendix B. First, on the theoretical side, the central distinguishing feature of our model is that it provides a measure of complexity that applies to canonical economic models in a fairly general way. This is largely owed to the model relying on continuous actions and concepts. Different complexity concepts have received significant attention in specific subfields of economics, such as automata in game theory and experimental economics (Rubinstein, 1998; Oprea, 2020; Banovetz and Oprea, 2023; Rubinstein, 1986; Abreu and Rubinstein, 1988; see also the overview in Chatterjee and Sabourian, 2020), notions of computational complexity in computer science (see, e.g., Lloyd, 2001) or the number of states in lottery choice (e.g., Puri, 2024; Bernheim and Sprenger, 2020).

Second, a growing literature abstains from modeling complexity explicitly and focuses on the type of behaviors that are induced by information-processing constraints and complex decision environments. This includes theoretical (e.g., Woodford, 2020) and experimental work (e.g., Frydman and Jin, 2022; Khaw et al., 2021; Enke et al., 2023, 2024; Enke and Graeber, 2023) on noisy (or imprecise) reasoning and heuristic or rule-based decision-making (Nielsen and Rehbeck, 2022; Halevy and Mayraz, 2022; Lacetera et al., 2012a; Tversky and Kahneman, 1974). Empirical work has also identified a pervasive aversion to complexity, e.g., in the case of lotteries (Goodman and Puri, 2022; Fudenberg and Puri, 2023) or multi-attribute choice (Iyengar and Lepper, 2000; Iyengar and Kamenica, 2010). Our model is compatible with this work, especially the behavioral attenuation induced by noisy processing, but attempts to endogenize the origin of such simplification strategies.

Third, the literature on the empirical measurement of complexity spans many disciplines, as reviewed by Oprea (2024a). A recent example in economics, applied to lottery choice complexity, is Enke and Shubatt (2023). Our approach is a belief-based metric that elicits a subjective ranking of tasks, leveraging the idea that comparisons may be cognitively less demanding than absolute assessments (e.g., Stewart et al., 2005) and employing

methods from psychometrics (Bradley and Terry, 1952). The topic of what determines the complexity of a task has a long history in the cognitive sciences (as reviewed, e.g., in Campbell, 1988; Liu and Li, 2012).

Fourth, a growing literature studies field applications of complexity, often leveraging intuitive notions for what makes a choice more or less complex. This line of work primarily includes – but is not limited to – applications to financial decision-making (e.g., Célérier and Vallée, 2017; Carvalho and Silverman, 2023; Colliard and Georg, 2025; Carlin, 2009) and the tax code (e.g., Aghion et al., 2024; Zwick, 2021).

Fifth, we relate to the project of behavioral macroeconomics. It is well established by now that inattention has direct implications for macroeconomics, such as fiscal and monetary policy (Gabaix (2020)) and general equilibrium thinking (Angeletos and Lian (2023)). While much prior work models the benefit of attention, modeled as some gains in the traditional utility from more optimal decisions (Gabaix and Laibson (2002), Sims (2003), Gabaix (2014), Woodford (2020)), endogenizing meaningfully the cost of attention (e.g. based on the complexity of the subparts) has proven more challenging. We provide an attempt to endogenize attention, which is a key contribution of our paper.

Several recent papers have been employing the approach in this paper. Fenig and Petersen (2024) studies analytically and experimentally a life-cycle planning model related to the life-cycle we present, with a rich series of variants. Kurtz-David et al. (2024) model the complexity of strategic interactions. Gabaix (2025) derives the optimal simplicity in a number of situations, e.g. the optimal simplicity of a tax system, or a contract.

Outline of this paper. Section 2 provides the basics of the theory. Section 3 outlines some applications, to the basic theory of consumption, the complexity of the tax system, of intertemporal consumption and of planning consumption across the life cycle. Section 4 presents our experimental method, and the results. Section 5 extends the theory to discrete actions, and applies this to the choice between risky lotteries, and between complex financial products. Section 6 extends the model to limited metacognition. Section 7 concludes. The appendices contain proofs and experimental instructions, as well as some extensions, e.g. the complexity of forecasting, of basic arithmetic operations, first and second order complexity aversion, and the interaction between learning and complexity.

2 The Complexity of Decision Problems: Basics

2.1 Cognitive production function

Our model is inspired by the observation that in practice, people approach most problems pragmatically in the precise sense that they do not strive for perfect accuracy but instead aim for a balance of committing *sufficient* resources to get to a *sufficiently* close to optimal response.[4] This idea suggests that people have a somewhat continuous perception of complexity in a given problem, and it calls for a theoretical concept of "pragmatic complexity". Consequently, rather than examining the question of "What's the complexity of getting the problem exactly, or 100%, right?", as much prior work does, we ask "What's the complexity of getting the problem q percent right", for a quality q that can be less than 100%. This motivates the concept of a "complexity curve" as a function of q. This subsection develops this simple line of reasoning and introduces the corresponding analogy of a cognitive production function.

The cognitive production function Take the problem $\max_a u(a)$. Suppose that cognitive effort L leads to decision $\tilde{a}(L)$ (which could be stochastic), hence to utility $U(L) = \mathbb{E}[u(\tilde{a}(L))]$. We normalize

$$Q(L) \coloneqq \frac{U(L) - U(0)}{U^r - U(0)} \tag{2}$$

where $U(\infty) = U^r$ is the utility corresponding to perfect choice, $U(0)$ the utility obtained with 0 thinking (e.g. randomize). So $Q(L) \leq 1$.

The primary objective of this paper is to endogenize the cognitive production function. In our examples below, we employ the following parametrization:

$$Q(L) = \mathcal{M}\left(\frac{L}{C}\right) \coloneqq \min\left(\left(\frac{L}{C}\right)^{1-\alpha}, 1\right) \tag{3}$$

with $\alpha \in [0, 1)$, and $C > 0$ is a complexity parameter, which we will endogenize. So, for $L \in [0, C]$, the function is increasing and concave. It saturates at 1 for $L > C$. This is because the maximal precision is reached for $L = C$. Figure 1 (left panel) illustrates this production function.

We will often want a version with $\alpha > 1$, which will induce "sparsity", the fact that some dimensions are not attended at all (Gabaix (2014)). Then, the production function has the shape, $Q(L) = \max\left(1 - \left(\frac{L}{C} + \Phi\right)^{1-\alpha}, 0\right)$, with $\Phi \geq 0$, see Section A.1.

[4] This observation relates to Simon's idea of "satisficing" (Simon, 1955).

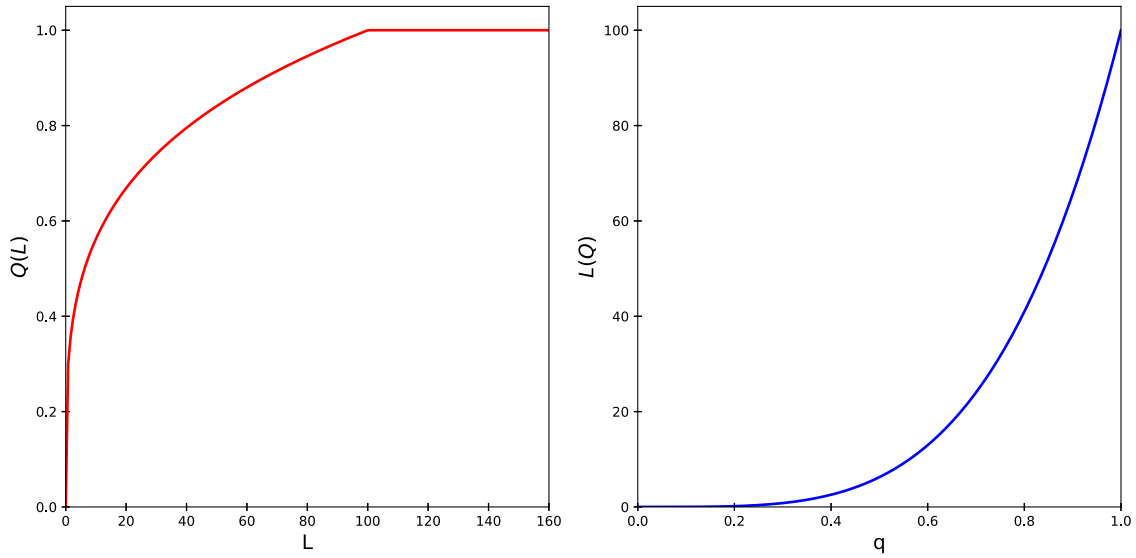


Figure 1: Illustration of a cognitive production function $Q(L)$ and a complexity curve $L(q)$. *Parameters:* $C = 100, \alpha = 3/4$.

The complexity curve We can now define the mental effort required to achieve precision $q \in [0,1]$ – the *complexity curve* as a function of q. This is none other than the cost function, which is here the inverse of the production function: $L(q) = Q^{-1}(q)$. When the production function is concave, the cost function is convex. This is illustrated in Figure 1 (right panel), which uses $C = 100$. To get the problem 100% right the complexity is $L(1) = 100$. But to get the problem 70% right the complexity is $L(0.78) = 25$. In computer science, the "complexity" is typically understood as the complexity to reach a 100% correct answer, $q = 1$. In economic life, people sensibly settle for q less than 1, as we will illustrate. We record those observations in a definition.

Definition 1. *(Complexity)* Call $Q(L)$ the cognitive production function as a function of mental effort L, normalized so that the maximum value is 1. We call *output-based* complexity the performance shortfall $1 - Q(L)$ of the decision reached after mental effort L. We call *input-based* complexity the mental effort $L(q) = Q^{-1}(q)$ needed to reach a performance q.

Those dual notions of complexity are functions, so very high dimensional objects. Empirically, a lower-dimensional parametrization is useful. In our Cobb-Douglas example (3), the main parameter is C. The output-based complexity is (see (3)): $1 - Q(L) = \max\left(1 - \left(\frac{L}{C}\right)^{1-\alpha}, 0\right)$, so that it is weakly increasing in C. The input-based complexity at performance target q is $L(q) = q^{\frac{1}{1-\alpha}} C$, so that it is increasing in C, and, in fact, proportional to it.

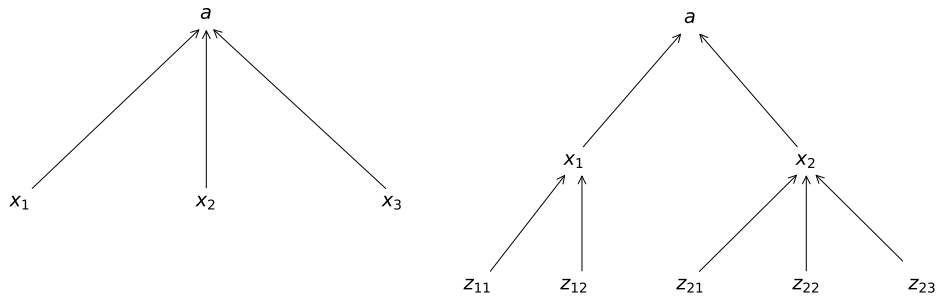


Figure 2: Illustration of two cognition production functions, with one and two layers. *Notes.* In both, the output decision a takes as inputs approximations of x_1, x_2 and x_3. In the two-layer function, each x_i further takes as inputs approximations of the z_{ij}.

The substance of the paper is to actually compute the cognitive production function, this way endogenizing complexity, in particular the parameter C. We now turn to this task.

2.2 Complexity of a problem from the production function of thought

2.2.1 Setup

The task is to maximize over a continuous action a an objective function $u(a, x)$, where u is smooth (three times continuously differentiable) and concave in a. We assume that the x are drawn from a distribution with mean normalized to 0. We take a to be one-dimensional for now, but it is easy to extend to a multidimensional a (Proposition 20). We call default action a^d the optimal action "at the default", i.e. when all x are equal to 0, $a^d = \text{argmax}_a u(a, 0)$. We call a_{x_i} is the partial derivative at a default point. The rational answer is thus (after linearization, so up to second order terms in x_i, so that we assume that the deviations x_i are "small")

$$a^r = \sum_i a_{x_i} x_i = \sum_i y_i, \qquad y_i := a_{x_i} x_i \tag{4}$$

The agent's objective is to cognitively construct those $y_i := a_{x_i} x_i$, which indicates by how much dimension i of the problem should change the rational action. This is schematically illustrated in Figure 2 (left panel); the x_i give rise to the decision a.

We model that people receive noisy signals y_i^s about y_i:[5]

$$y_i^s = m_i y_i + (1 - m_i)\, y_i^d + \sqrt{m_i\,(1 - m_i)}\varepsilon_i \tag{5}$$

where $m_i \in [0, 1]$ is the precision of the signal, y_i^d is a default value, equal to 0 when the mean of x_i is 0, and ε_i a mean-zero noise with variance $\sigma_{\varepsilon_i}^2 = \sigma_{y_i}^2$. If all the shocks are jointly Gaussian, with the prior of y_i equal to y_i^d, we have $\mathbb{E}\left[y_i | y_i^s\right] = y_i^s$.[6] Accordingly, we posit that if the decision maker sees y_i^s, she takes the decision

$$a = \sum_i y_i^s.$$

The next Lemma records the value of information.

Lemma 1. (Value of information: Continuous choice). *The expected utility after acquiring precision m_i is, up to third order terms in σ_{x_i}:*

$$U\,(m) = U\,(0) + \sum_i V_i m_i, \qquad V_i = \frac{1}{2}\left|u''\left(a^d\right)\right|\sigma_{y_i}^2 \tag{6}$$

We imagine that a mental effort L_i allows the decision maker to reach precision $m_i\,(L_i)$ for a cognitive production function that is exogenously given for now. So, the mental allocation problem is

$$\max_{L_1,\ldots L_N} \sum_i V_i m_i\,(L_i) \text{ s.t. } \sum_i L_i \leq L. \tag{7}$$

Let us define the following key quantity:

$$s_i = \frac{V_i}{\sum_j V_j} = \frac{\sigma_{y_i}^2}{\sum_j \sigma_{y_j}^2}. \tag{8}$$

Here, s_i is the relative importance of dimension i: more precisely, it is the share of variance in the action due to dimension i. The cognitive allocation problem is

$$Q\,(L) := \max_{L_1,\ldots L_n} \sum_i s_i m_i\,(L_i) \text{ s.t. } \sum_i L_i \leq L. \tag{9}$$

[5]The noisiness is not central — the fact that signals are imperfect is.

[6]We shall not assume that agents are Bayesian (as traditional information economics) – instead, we will use that benchmark as an inspiration for the model (as e.g in Gabaix (2014), Woodford (2020)). When y_1^s and y_2^s are correlated, a Bayesian agent would use $\mathbb{E}\left[y_1 | y_1^s, y_2^2\right]$ to infer y_1, but instead we model the agent as being a "limited Bayesian", who simply performs $\mathbb{E}\left[y_1 | y_1^s\right]$.

2.2.2 The macro complexity of a problem as arising from its micro parts: formula

We next solve for problem (9), and discuss its economics. We leverage the body of knowledge on how to tractably model production functions – in particular, the aggregation properties of Cobb-Douglas and CES production functions.

Proposition 1. (Macro complexity from micro complexity) *Suppose that the "micro" cognitive production function of component i is*

$$m_i(L_i) = \min\left(\left(\frac{L_i}{c_i}\right)^{1-\alpha}, 1\right) \tag{10}$$

where the complexity of dimension i, c_i, is exogenous for now, and $\alpha \in [0,1)$. If at the optimum of problem (9) *we are in the "interior region" ($m_i\left(L_i^*\right) \in (0,1)$, where L_i^* is the optimal allocation to dimension i), then the macro cognitive production function is*

$$Q(L) = \max\left(\left(\frac{L}{C}\right)^{1-\alpha}, 1\right) \tag{11}$$

with "macro complexity" $C = \mathcal{C}\left((s_i, c_i)_{i=1...N}\right)$, using the "complexity aggregator":

$$\mathcal{C}\left((s_i, c_i)_{i=1...N}\right) := \left(\sum_i s_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}}\right)^{-\frac{\alpha}{1-\alpha}} \tag{12}$$

The key message of Proposition 1 is equation (12), which formulates an endogenous macro complexity C in terms of the exogenous micro-complexities of the components c_i.[7] In terms of Figure 1 (left panel), the "micro cognitive production functions" leading to the micro-components x_i were posited Cobb-Douglas with complexity c_i. Proposition 1 also records that the resulting "macro" cognitive production for decision a is also Cobb-Douglas: simply enough, Cobb-Douglas in, Cobb-Douglas out. To build intuition, we examine a few polar cases.

2.2.3 Simple comparative statics on what makes a problem complex

We next discuss some simple facts to build intuition for the formula, equation (12).

First, take the case where s_1 tends to 1, and the other s_i are close to 0. Then, C tends to $C = c_1$: the effective complexity of the macro problem is that of the one "important"

[7] The appendix records extensions of this proposition. Proposition 17 extends it to the case $\alpha > 1$, which leads to the same complexity aggregator. Proposition 18 derives the allocation of attention component-by-component.

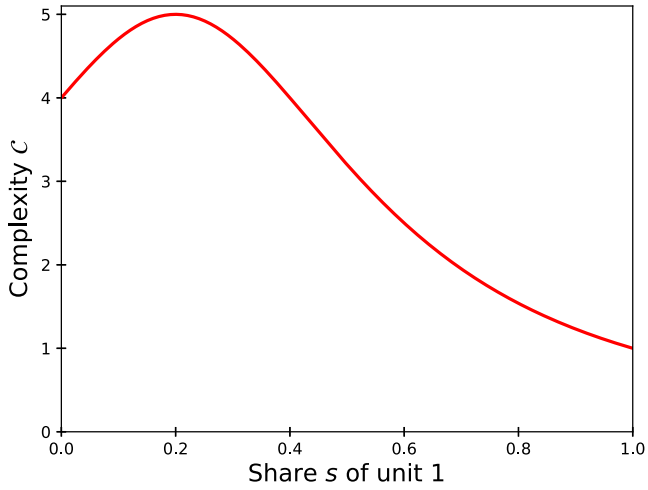


Figure 3: Complexity $\mathcal{C}$ as the share of component 1 varies, in a problem with $N = 5$ components. *Notes*. This graph plots complexity C when component 1 has share s and the other 4 components have share $\frac{1-s}{4}$.

component.

Next, take the case where $c_i = \bar{c}$ for all i, for some $\bar{c} > 0$, which we normalize to 1. When the N components have equal share ($s_i = \frac{1}{N}$), then

$$C = N \tag{13}$$

In the more general case, $C = N^f$ where N^f is the "effective" number of components:

$$C = N^f, \qquad N^f = \left(\sum_{i=1}^{N} s_i^{\frac{1}{\alpha}} \right)^{-\frac{\alpha}{1-\alpha}} \in [1, N]$$

For instance, $N^f = N$ if $s_i = \frac{1}{N}$ for all i, and $N^f = 1$ if $s_1 = 1$, while the other components have 0 share. When $\alpha = 1/2$, $N^f = \left(\sum_{i=1}^{N} s_i^2 \right)^{-1}$, so it is the inverse Herfindhal index.

We illustrate this in Figure 3. It shows the complexity C of the problem when one component has share s, the other $N - 1$ components have share $\frac{1-s}{N-1}$ (here with $N = 5$), and we normalize $\bar{c} = 1$. Start with $s = 0$. Then, there are effectively only 4 components, so $C = 4$. As s increases, C increases, up to the point where $s = \frac{1}{5}$, so that we have 5 equally-sized components, and $C = 5$. In the limit where $s = 1$, there is now just one effective component, and $C = 1$. In between as s increases between $\frac{1}{5}$ and 1, complexity C decreases.

The qualitative phenomena illustrated in Figure 3 are general, and independent of the

curvature parameter α. We state two propositions formalizing them.

Proposition 2. *(*Complexity from compositionality*) Assume c_i is the same for all i. If the effective number of components rises, complexity is higher. More compositionality leads to more complexity.*

Next, consider a thought experiment generalizing Figure 3.

Proposition 3. (Simplicity from extremity) *Assume c_i is the same for all i. Consider a dimension i, change its share ds_i and change other shares proportionally ($ds_j = \left(1_{i=j} - s_j\right) ds_i$). If the share of the largest unit rises, C falls. If the share of the smallest unit rises, C rises.*

Finally, we record the extremal values of C as the shares vary.

Proposition 4. (Complexity measure as α varies) *The complexity aggregator (12), $\mathcal{C}_\alpha = \left(\sum_i s_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}}$, is weakly increasing in $\alpha \in [0, \infty]$, and:*

$$\mathcal{C}_0 = \min_i \frac{c_i}{s_i} \leq \mathcal{C}_1 = \exp\left(\sum_i s_i \ln \frac{c_i}{s_i}\right) \leq \mathcal{C}_\infty = \sum_i c_i.$$

For instance, when $\alpha \to \infty$, we obtain the simple measure $\mathcal{C}_\infty = \sum_i c_i$, the sum of the elementary complexities. When $\alpha = 1$, $\ln \mathcal{C}_1$ is a "complexity-adjusted" entropy. Hence, by varying α, one can recover a number of sensible measures. Still we will keep α constant in a given application.

We next state a Lemma that is useful for calculations.

Lemma 2. (Equivalent formulation with quasi-complexities) *Define the quasi complexities as:*

$$\tilde{c}_i := c_i^{1-\frac{1}{\alpha}}, \qquad \tilde{\mathcal{C}} := \mathcal{C}^{1-\frac{1}{\alpha}} \tag{14}$$

Then, (12) becomes:

$$\tilde{\mathcal{C}} = \sum_i s_i^{\frac{1}{\alpha}} \tilde{c}_i \tag{15}$$

In (15) the macro complexity $\tilde{\mathcal{C}}$ is the weighted sum of the micro complexity $\tilde{c}_i$, with weights $s_i^{\frac{1}{\alpha}}$. The interpretation of quasi-complexity is most intuitive when $\alpha > 1$, as then $\tilde{\mathcal{C}}$ is an increasing function of $\mathcal{C}$. The weights $s_i^{\frac{1}{\alpha}}$ add up to more than 1 when $\alpha > 1$.

2.3 How complexity affects outcomes: actions, errors, deliberation time

Measuring macro-attention and complexity from actions Next, we relate complexity to observables, namely the average actions, the amount of errors, and the deliberation time. Higher complexity attenuates action's responsiveness to incentives

We imagine a researcher gives subjects a series of tasks, $\max_a u(a, x)$, with x drawn from a distribution with mean 0, and x_i having variance $\sigma^2_{x_i}$, with the x_i uncorrelated across i's. For each task x and subject s, the researcher collects $a(s, x)$ and compares it to the rational action $a^r(s, x)$. We suppose that the subjects have a fixed attention L for each task. The next proposition states what the researcher should find, under our model.

Proposition 5. (Higher complexity makes actions less responsive to incentives) *Suppose that one regresses empirical actions a on the rational action a^r:*

$$a - a^d = M\left(a^r - a^d\right) + k + \varepsilon$$

where $M := \frac{cov\left(a-a^d, a^r-a^d\right)}{var\left(a^r-a^d\right)}$ is the composite attention parameter, k is a constant, and ε is extra noise. Then, in the limit of a large sample,

$$M = Q(L),$$

and, as

$$Q = \left(\frac{L}{C}\right)^{1-\alpha},$$

more complexity leads to more attenuation.

Proposition 5 implies that macro attention M can be recovered from choice: it reveals itself as an attenuated response M to incentives. This is consistent with models of imperfect perception, (Gabaix (2014)), field evidence (surveyed in DellaVigna 2009 and Gabaix (2019)), and systematic experimental evidence (Enke et al. (2024)). More importantly, Proposition 5 can allow to (potentially) measure the production function of attention, hence the output-based complexity. In addition, Proposition 5 indicates that we can recover the micro-complexity of the various components, also from choice. Viewed another way, given a theory-predicted complexity C, we can predict the degree of attenuation in choices.[8]

[8]One can also predict attention dimension-by-dimension, something formalized in Appendix A.2. The complexity C is a sufficient statistic for the feeling of the degree of complexity (at least, in the interior region of Proposition 1), but not for the action.

One limitation of this proposition is that the attention L is exogenous. It could be endogenous in practice. If the elasticity of total mental capacity to the task is limited (something that seems likely), then this effect is moderate in size, and does not change drastically the message of Proposition 5—see Proposition 7 for details.

Higher complexity leads to larger errors and higher cognitive uncertainty Next, we study mistakes. We take the root mean square error, $\sigma_a = \mathbb{E}\left[(a - a^r)^2\right]^{1/2}$. Its subjective counterpart (assuming agents are aware of the potential random errors they commit) is related to Enke and Graeber (2023)'s concept of cognitive uncertainty, which in its basic form may be translated in the present context as $\sigma_{CU} = \chi\sigma_a$ for some positive coefficient χ.

Proposition 6. (Higher complexity leads to more errors and more cognitive uncertainty) *More complexity leads to larger errors σ_a and higher cognitive uncertainty $\sigma_{CU} = \chi\sigma_a$ (i.e., subjective feeling of committing errors):*

$$\sigma_a = \sqrt{1 - Q}\sigma_{a^r}, \qquad Q = \left(\frac{L}{C}\right)^{1-\alpha} \tag{16}$$

Hence, we predict that model complexity C leads to larger errors. Moreover, there is a non-trivial functional form, which again could be tested.

Higher complexity leads to higher time spent (in the sparsity-inducing case $\alpha > 1$) The next Proposition 7 shows that effort is decreasing in complexity when $\alpha < 1$, but is increasing in the sparsity-inducing case $\alpha > 1$ (and the constant Φ is low enough). We will assess that the latter is the relevant case.

Proposition 7. (How complexity affects effort and performance) *Assume the same production functions as in Lemma 6. Suppose that one solves the optimum effort, given a shadow benefit $p > 0$ and a shadow cost of effort $w > 0$, $\max_L pQ(L) - wL$, and assume that the solution is interior. The optimum effort is $L^* = C\left(\frac{wC/p}{|\alpha-1|}\right)^{-1/\alpha} - \Phi C$, and the resulting attention is: $Q(L^*) = \left(\frac{wC/p}{|\alpha-1|}\right)^{1-\frac{1}{\alpha}}$ if $\alpha < 1$ and $Q(L^*) = 1 - \left(\frac{wC/p}{|\alpha-1|}\right)^{1-\frac{1}{\alpha}}$ if $\alpha > 1$. As a result, higher complexity C decreases performance $Q(L^*)$, but decreases effort L^* if $\alpha < 1$ and $\Phi \geq 0$, while increasing effort if $\alpha > 1$ and Φ is low enough. Higher benefit of thinking p, and lower cost of effort w increase effort L^* and performance $Q(L^*)$.*

2.4 Recursive complexity: when the complexity of an action depends on the complexity of the other actions

Suppose a multidimensional action, and dimension j, a_j depends on the other actions a_k, as well as some x, as in

$$a_j^r = \sum_i b_i x_i + \sum_{k \neq j} \gamma_{jk} a_k$$

Then, using the formulation with quasi-complexities (14), we obtain:

$$\tilde{C}_{a_j} = \sum_i \left(s_{x_i}^{a_j} \right)^{\frac{1}{\alpha}} \tilde{c}_i + \sum_{k \neq j} \left(s_{a_k}^{a_j} \right)^{\frac{1}{\alpha}} \tilde{C}_{a_k} \tag{17}$$

where the s^{a_j} are are the importance shares of x_i vs a_k in the decision, e.g.

$$s_{x_i}^{a_j} = \frac{b_i^2 \sigma_{x_i}^2}{D}, \qquad s_{a_k}^{a_j} = \frac{\gamma_k^2 \sigma_{a_k}^2}{D}, \qquad D = \sum_i b_i^2 \sigma_{x_i}^2 + \sum_k \gamma_k^2 \sigma_{a_k}^2$$

Hence, the complexity of a_j, $\tilde{C}_{a_j}$, depends on the complexity of action k, $\tilde{C}_{a_k}$. This formulation is tractable, as (17) is a system of linear equations in $\tilde{C}$.[9]

3 Applying the model to some canonical economic problems

We present general guidelines on how to use the model, and then several applications to canonical economic problems.

3.1 Applying the model: a user's guide

In the following we provide some general guidance for applying the model to concrete settings.

1. Formulate the rational, fully specified problem, in the form $\max_a u(a, x)$, where $x = (x_i)_{i=1...N}$ are deviations from a default.

2. Think through a "natural" description of the x_i's. For instance: are they presented in a nominal frame or a real frame; do decision-makers think about the base good

[9] This formulation could be used to think about the complexity of games—we simply interpret a_k as the action of another player. This is the "oneself as another" perspective, to use philosopher Ricœur (1992)'s expression ("soi-même comme un autre"). We are developing this in another paper.

price x_1 and the "shrouded attribute" price x_2 separately, or do they have direct access to the sum $(x_1 + x_2)$? This step generally requires some application-specific assumptions (in much the same way a researcher specifies e.g. an agent's utility function).[10]

3. Then, apply the model: calculate the importance shares of the dimensions, s_i as in (8), and obtain the complexity measure C from (12).

   (a) In the most basic case, use a Laplacian ignorance prior and posit that all micro-parts have the same complexity $c_i = \bar{c}$. One can normalize $\bar{c} = 1$.

   (b) However, one can go one step further, and apply the model recursively to find the complexity c_i of subproblem i as a function of its components, as in Section D.1. In the end, the complexity of the sub-sub-components may be arbitrary, also at $\bar{c}$.[11]

With the model-based complexity prediction at hand, one can then relate it to observables (e.g. error rate by consumers, markups by firms), and use the theoretical guidance about the theoretical correlates of complexity, e.g. errors and attenuation (Section 2.3), and the subjective feeling of complexity (measured via the procedure of Section 4.1).

A "zero free parameters" measure We discuss the value of α in Section D.4. In the end, we recommend taking $\alpha = 0.5$.

This way, the measure is has *zero free parameters*. That makes the analysis easy to port across settings.

There is still some leeway in the termination of the "natural" description, as in Step 2 above. Still, in principle it could be elicited, and in practice, in many cases the natural frame is reasonably clear.

The measure we propose is just a "baseline" measure. In any given case, one can feel that some features are "easier" than others, so that using a uniform micro cost $c_i = \bar{c}$ is crude. We consider that those are useful refinement of our baseline, but that it is still very useful to have a unified baseline across models.

We now show a series of progressively richer examples to illustrate concrete use cases of the model.

[10] Using the identifiability results of Section 2.3, one could in principle elicit the agents' frame.

[11] Conceivably, one could imagine that the complexity c_i of basic components will be estimated systematically, and that researchers can use estimates from prior results, much in the way that economists now use "standard" parameters for loss aversion, adjustment costs in investment, or the elasticity of substitution between goods.

3.2 The complexity of the tax system

The tax system is clearly very complex – but how to quantify this? We now show how the model offers a simple and novel way to address this question.

Suppose that there are N taxes (local, federal, social security, unemployment, child benefits, food support...), each with its tax rate $T_i(e)$ given earnings e, so that total tax is $T(e) = \sum_i T_i(e)$ and the total marginal tax rate is $T'(e) = \sum_i T_i'(e)$.[12] The importance share for tax i is

$$s_i^{T'}(e) = \frac{T_i'^2(e)}{\sum_j T_j'^2(e)}$$

If the micro complexity of each is just $c_i = \bar{c}$, the complexity of the marginal tax rate is just the effective number of components (from (12)):[13]

$$\mathcal{C}^{T'(e)} = \bar{c}\left(\sum_i \left(s_i^{T'}(e)\right)^{\frac{1}{\alpha}}\right)^{-\frac{\alpha}{1-\alpha}}. \tag{18}$$

For instance, take Figure 4, which plots the marginal tax rate in California (for a single parent of two). The total marginal tax rate is widely variable.

The reason is that it is the outcome of many taxes (positive and negative): income tax, federal and state, child care credit, a variety of acronym schemes with names like TANF and SNAP, all of which vary with income e, and all of which "should" be taken into account by a rational household. Formula (18) allows to calculate the complexity of the tax system across those income categories; and then, in future work, to compare how well people do at optimizing, as a function of the complexity, and determinants such as demographics, numeracy, experience, etc. Down the line, one would also want to optimize the tax system taking into account complexity costs, something that (Gabaix, 2025) explores.

Suppose next that we have richer data — for instance, we know that it takes τ_i hours to perfectly determine the amount of tax i.[14] Then, we could set $c_i = \tau_i$, and the complexity

[12] The agent's problem $u(a,x) = a - T(a,x) - \frac{a^2}{2\psi}$, where a is the labor supply, $a - T(a,x)$ is disposable income, i.e. consumption, $\frac{a^2}{2\psi}$ is the disutility of effort, and the wage is normalized to 1. So the optimal rational action is: $a = \psi(1 - T_a(a,x))$. We have $T_a(a,x) = \sum x_i$, $x_i = T_i'(a)$, where x_i is the marginal tax rate. We assume for simplicity that marginal tax rates are constant across labor levels a. This could easily be extended, at the cost of more notations.

[13] The complexity of seeing the average tax rate is calculated in the same way, replacing $T_i'(e)$ by $\frac{T_i(e)}{e}$.

[14] A rough proxy might also be the number of words in the tax code for tax i, or some concave transform of it.

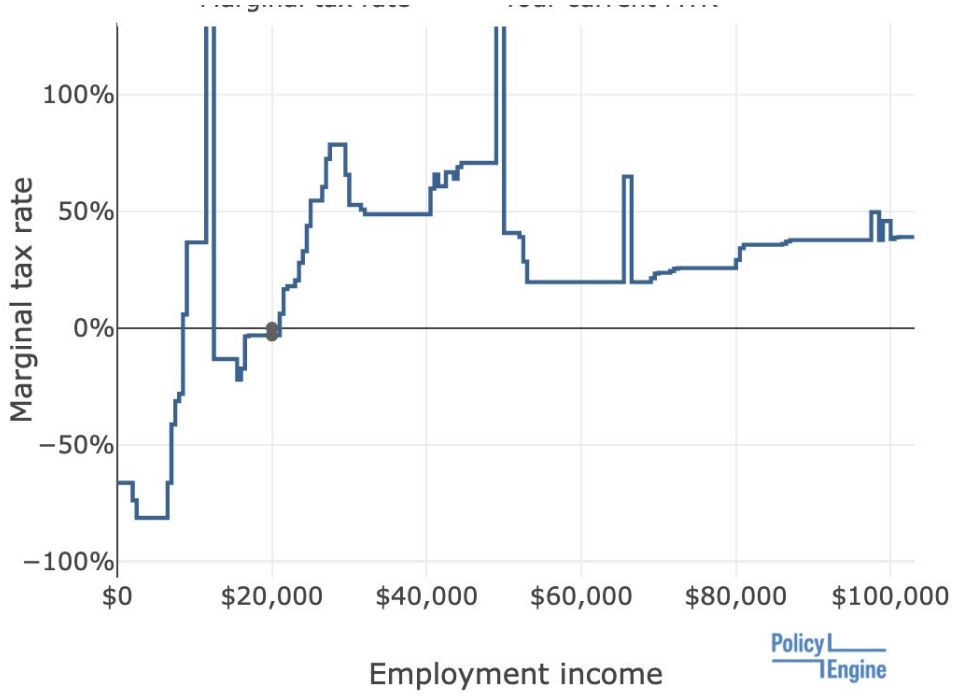


Figure 4: Marginal tax rate in California. This is for a single parent of two, as a function of income, in 2023. *Source*: PolicyEngine.org

of the tax system emerges as:

$$\mathcal{C}^{T'(e)} = \left(\sum_i \left(s_i^{T'}(e) \right)^{\frac{1}{\alpha}} \tau_i^{1-\frac{1}{\alpha}} \right)^{-\frac{\alpha}{1-\alpha}} \tag{19}$$

It is expressed as an "effective number of hours". The fact that it is less than $\sum_i \tau_i$ reflects that agents don't need to perfectly think through all the minutiae linked to tax i.

The conclusion is that the model gives a new measure of complexity of the tax system, (18), offering a way to do simple empirics on it.

3.3 Complexity of forecasting

We want to capture that some forecasts are "more complex" than others. Suppose that we're interested in variable $y_t = \sum_i x_{it}$, composed of several components x_{it}, each of which follows an AR(1): $x_{it} = e^{-\phi_i} x_{i,t-1} + \varepsilon_{it}$.[15] We want to forecast value of y_t in h

[15] We explore a more general version in Section D.8

periods:

$$f_{t,h}^{r} \coloneqq \mathbb{E}_t \left[y_{t+h} \right] = \sum_i e^{-\phi_i h} x_{it} \tag{20}$$

We write the "frame" of the forecast as:

$$f_{t,h}^{s} = a_{-1} y_t + a_0 \left(e^{-\phi^d h} - 1 \right) y_t + \sum_i a_i \left(e^{-\phi_i h} - e^{-\phi^d h} \right) x_{it} \tag{21}$$

where the rational weights are to have all $a_j = 1$: in that case, $f_{t,h}^{s} = f_{t,h}^{r}$ is the rational forecast. This equation captures the following. One predictor is simply y_t. Another is a more refined predictor, $e^{-\phi^d h} y_t$, which replaces the process by a simple AR(1). Finally, each component i can be refined, using its correct autocorrelation: the correction is then $\sum_i \left(e^{-\phi_i h} - e^{-\phi^d h} \right) x_{it}$.

Hence, shares are $s_{-1} = k \mathbb{E} \left[y_t^2 \right]$,$s_0 = k \left(e^{-\phi^d h} - 1 \right)^2 \mathbb{E} \left[y_t^2 \right]$, $s_i = k \left(e^{-\phi_i h} - e^{-\phi^d h} \right)^2 \sigma_{x_i}^2$, where k ensure that the shares add up to 1. For simplicity, we model that the micro costs c_i, are all 1, except for the simplest predictor that gives the current value y_t, whose complexity is 0, $c_{-1} = 0$. Hence, the complexity of forecasting is are usual aggregator, $\mathcal{C}(s, c)$.

Proposition 8. (Complexity of forecasting). *The complexity of forecasting is 0 at very short horizons (as agents just use the current values as their forecast), 1 at very high horizons (as only one component matters), and is between 0 and $N + 1$ at intermediate horizons.*

As is plains processes composed of "features" with very different time-scales are more complex to forecast. They're of course trivial to "nowcast" (the predictor for y_t tomorrow is just today's y_t, unless the autocorrelation is nil). They're reasonably simple to forecast at very large horizons (just one component matters, the most persistent one). They're most complex at intermediate horizons, where different component, with different degree of medium-size persistence, varies.

3.4 The complexity of choosing different goods

We now study the complexity of a basic building block in economics: consumer theory. This is useful, because that problem is central to most models, and other problems (e.g. producer theory) are entirely similar, so that the arguments and effects can be transposed. To simplify the discussion and avoid income effects (for now), we center on quasi-linear utility, $U(c, c_0) = v(c) + c_0$, i.e. where there is a special good 0 with marginal utility and price of 1, and the consumption of the other goods are represented by vector c. So, the

rational problem is simply

$$\max_{c \in \mathbb{R}^n} \quad v(c) - p \cdot c \tag{22}$$

We denote proportional, i.e. log, changes using hats (e.g. $\hat{c} = \frac{dc}{c}$).

3.4.1 Do we know our own demand elasticity? Choosing the quantity one good

We start with the case of one good. If a good's price changes by $\hat{p}$ percent, consumption should change by $\hat{c}$ percent:

$$\hat{c} = -\psi \hat{p} \tag{23}$$

where $\psi = -\frac{u'(c)}{cu''(c)}$ is the elasticity of substitution.[16]

If the agent simply needs to pay attention to the price, then the complexity of choosing consumption is just that of looking up the price, $C^{\hat{c}} = c^{\hat{p}}$.

The newer case is that the agent may not know ψ. Traditional economics assumes that people know their preferences, including the elasticity of demand ψ. Readers can introspect about their elasticity of substitution for mineral water, or coffee, or fish. Intuitively, those quantities seem hard to know.

We will model that via some effortful introspection, people can get a better estimate of their rational elasticity of demand. Then, the complexity is $\mathcal{C}\left(s_{\hat{p}}, s_{\psi}, c_p, c_{\psi}\right)$, where $s_{\hat{p}} = \frac{V_{\hat{p}}}{V_{\hat{p}} + V_{\psi}}$, $s_{\psi} = 1 - s_{\hat{p}}$ are the importance shares, and the values of information are

$$V_{\hat{p}} = \left(\bar{\psi}^2 + \sigma_{\psi}^2\right)\sigma_{\hat{p}}^2, \qquad V_{\psi} = \sigma_{\psi}^2 \sigma_{\hat{p}}^2. \tag{24}$$

When the difficulty is knowing one's elasticity, the complexity is higher when people are more uncertain about their tastes. Indeed, more generally, the difficulty of a choice is not to simply to know the elements of the choice (such as the price), but to simulate the hedonic consequences of it.

It may be worth pondering how people introspect about their rational demand elasticity. One technique is the "limit case" method: for instance, if we go without food for days, life is miserable. Through that thought experiment (or real world experiment), we conclude that $\lim_{c \to 0} u(c) = -\infty$, so the macro-elasticity of the demand for aggregate food is less than 1, $\psi < 1$.[17] But not eating one specific food for days (e.g. strawberries) is fine — it lowers utility by just a bit. So the micro-elasticity of one given food item is

[16] Indeed, $u'(c) = p$, so that $u''(c)\, dc = dp$, implying $\frac{dc}{c} = -\psi \frac{dp}{p}$ with $\psi = -\frac{u'(c)}{cu''(c)}$.

[17] We have in mind a utility such as $u(c) = \frac{c^{1-\frac{1}{\psi}}}{1-\frac{1}{\psi}}$.

greater than 1, $\psi^{\perp} > 1$.[18] Via this sort of thought experiment, people may get a sense of their elasticities of demand. Modeling that process explicitly would be interesting in future research (see Gabaix and Laibson (2022) and Imas et al. (2022) for some progress on this).

3.4.2 Complexity of multi-good consumption

With N goods, again expressing proportional changes using hats:

$$\hat{c}_i = -\psi_i \hat{p}_i - \sum_{j \neq i} \psi_{ij} \hat{c}_j, \tag{25}$$

where we call $\psi_{ij} = \frac{u_{ij} c_j}{u_{ii} c_i}$ the cross-elasticity of consumption.[19] Now, we study how having many goods affects complexity, so that we assume that the elasticities ψ_i and ψ_{ij} are known to the agent. Then, the complexity of choosing i is:

$$\tilde{C}_{c_i} = \left(s_{\psi_i}^{c_i}\right)^{\beta} \tilde{c}_{\psi_i} + \left(s_{p_i}^{c_i}\right)^{\beta} \tilde{c}_{p_i} + \sum_{j \neq i} \left(s_{c_j}^{c_i}\right)^{\beta} \tilde{C}_{c_j} \tag{26}$$

and it is recursive, as in Section 2.4.

Example: two goods To gain intuition, let us examine in some detail the case of two goods. Calling $x_1 = -\psi_1 \Delta p_1$ and $a_1 = -\Delta c_1$, (25) becomes:

$$a_1 = x_1 + \gamma_1 a_2, \qquad a_2 = x_2 + \gamma_2 a_1 \tag{27}$$

where $\gamma_1 = \psi_{12}$, $\gamma_2 = \psi_{21}$ express the degree of complementarity, and

$$a_1^r = \frac{x_1 + \gamma_1 x_2}{1 - \gamma_1 \gamma_2}, \qquad a_2^r = \frac{x_2 + \gamma_2 x_2}{1 - \gamma_1 \gamma_2} \tag{28}$$

---

[18] We have in mind something like nested CES here, e.g. $u(c_1, \ldots, c_n) = \frac{C^{1-\frac{1}{\psi}}}{1-\frac{1}{\psi}}$ with $C = \left(\sum_i c_i^{1-\frac{1}{\psi^{\perp}}}\right)^{\frac{\psi^{\perp}}{\psi^{\perp}-1}}$, with a macro-elasticity $\psi < 1$ and a micro-elasticity $\psi^{\perp} > 1$.

[19] Here is the derivation. We decompose $u'' = u''^{D} + u''^{ND}$, where u''^{DD} contains the diagonal terms, and u''^{ND} the non-diagonal terms (i.e., $u''^{D}_{ij} = u''_{ij} 1_{i=j}$, $u''^{ND}_{ij} = u''_{ij} 1_{i \neq j}$). Then, $u''^{D} dc = dp$ gives $u''^{D} dc = dp - u''^{ND} dc$ and

$$dc = \left(u''^{DD}\right)^{-1} dp - \left(u''^{DD}\right)^{-1} u''^{ND} dc$$

We assume that x_1 and x_2 are uncorrelated, with same variance σ_x^2. So, $\sigma_{a_1}^2 = \frac{1+\gamma_1^2}{(1-\gamma_1\gamma_2)^2}\sigma_x^2$ and in (27) the share of variance in a_1 coming from x_1 vs a_2 are $s_{x_1}^{a_1}$ and $s_{a_2}^{a_1}$ with:

$$s_{x_1}^{a_1} = \frac{1}{1 + \frac{\gamma_1^2(1+\gamma_1^2)}{(1-\gamma_1\gamma_2)^2}}, \qquad s_{a_2}^{a_1} = 1 - s_{x_1}^{a_1} \tag{29}$$

The expressions for $s_{x_2}^{a_2}$ and $s_{a_1}^{a_2}$ are similar, swapping the 1 and 2.

We call $\tilde{C} := C^{1-\beta}$, with $\beta := \frac{1}{\alpha}$. We normalize the complexity of x_1 and x_2 to 1. So:

$$\tilde{C}_{a_1} = \left(s_{x_1}^{a_1}\right)^{\beta} \tilde{C}_{x_1} + \left(s_{a_2}^{a_1}\right)^{\beta} \tilde{C}_{a_2}, \qquad \tilde{C}_{a_2} = \left(s_{x_2}^{a_2}\right)^{\beta} \tilde{C}_{x_2} + \left(s_{a_1}^{a_2}\right)^{\beta} \tilde{C}_{a_1}$$

Hence, we have a simple linear system in $\tilde{C}^{a_1}$, $\tilde{C}^{a_2}$, so we can solve it. In what follows, we take the simplified case where $\gamma_1 = \gamma_2 = \gamma$.

Proposition 9. (Complexity of choosing between two goods as a function of their complementarity) *The complexity of choosing good 1, when the complementarity between goods 1 and 2 is $\gamma_1 = \gamma_2 = \gamma$, is:*

$$C_a = \left(\frac{s_x^{\beta} + s_a^{\beta} s_x^{\beta}}{1 - s_a^{2\beta}}\right)^{\frac{1}{1-\beta}} C_x, \qquad s_x = \frac{1}{1 + \frac{\gamma^2(1+\gamma^2)}{(1-\gamma^2)^2}}, \qquad s_a = 1 - s_x \tag{30}$$

The complexity is increasing in $|\gamma|$. It is lowest where there is no complementarity ($\gamma = 0$), so that $C_a = C_x$. When $\gamma \to 1$, the complexity becomes unboundedly large, $C_a \to \infty$.

Figure 5 shows the result. We see that complexity is minimal at $\gamma = 0$, and increases in $|\gamma|$, i.e. as the goods become more complements or substitutes.

3.5 The complexity of life: Complexity of consumption planning over the life cycle

We have completed our tour of the complexity of static consumption. We now extend our analysis to address the polar opposite: What's the complexity of consumption planning across the life cycle?

Our framework offers a partial answer, with the complexity of planning consumption across one's life span illustrated in Figure 6. Suppose that the agent is born at time 0, dies at time $T-1$, deterministically, and has utility $\sum_{t=0}^{T-1} \beta^t u\left(a_t - h_t\right)$, where a_t is consumption, h_t is an i.i.d. random taste or needs shock. The agent has earnings $e_t = \bar{e} + \hat{e}_t$,

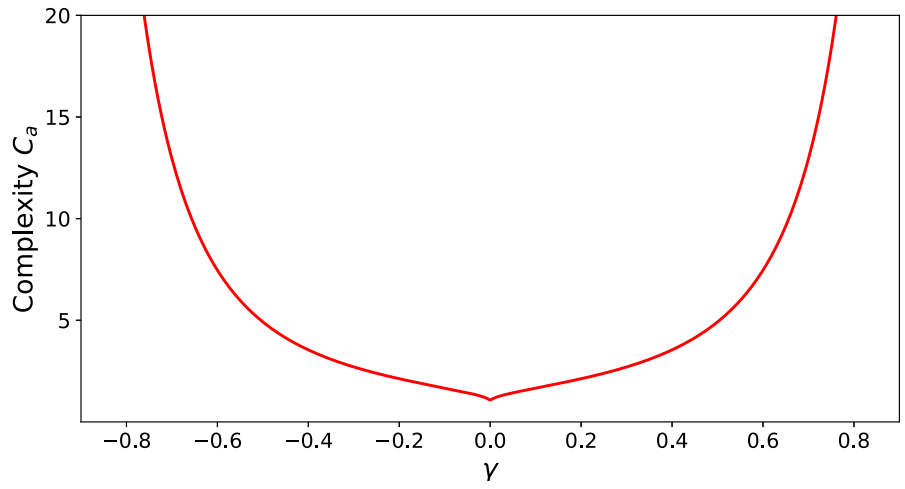


Figure 5: Complexity of choosing between two goods, as a function of their complementarity γ. *Notes*. This graph plots C_a from Proposition 9 as a function of γ, with $\beta = 0.2$ and $C_x = 1$.

where $\bar{e}$ is mean earnings, and $\hat{e}_t$ is a predictable deviation (known at time 0), but requires mental effort to contemplate. So, the rational problem is:

$$\max_{(a_t)} \sum_{t=t_0}^{T-1} \beta^t u\left(a_t - h_t\right) \text{ s.t. } w_{t+1} = R\left(w_t - a_t + e_t\right)$$

What's the complexity of solving this problem, for an agent at age t_0? We assume small shocks, or a quadratic utility function to remove precautionary saving, and $\beta = R = 1$ to remove mechanical "time horizon" effects. We say complexity of all parameters is the same, 1, but complexity of cash is 0.[20] The complexity of life-cycle consumption planning as a function of age is illustrated in Figure 6.

Proposition 10. (Complexity of the life-cycle problem, as a function of age) *Call* $e_{1,t} = \left(1 - \frac{1}{T-t}\right) h_t$, $e_{it} = \frac{\hat{e}_{t+i-2}}{T-t}$ *for* $i = 2, \ldots, T-t+1$. *Form, for* $i \geq 1$, $s_{it} = \frac{\mathbb{E}\left[e_{it}^2\right]}{\sum_{j=1}^{T-t+1} \mathbb{E}\left[e_{jt}^2\right]}$. *Then, the complexity of the life-cycle at age* t *is* $\mathcal{C}\left(s_t, 1\right)$, *see Figure 6.*

The life cycle problem is easier (i) when you're very old, because few periods are left to think about, life has reached a stark simplicity; (ii) when you're very young, because even though there are lots of predictable events $\hat{e}_t$ in the future, they don't matter much, as they are smoothed over many periods, so the normative impact $\frac{\partial a_t}{\partial e_s} = \frac{1}{T-t}$ is low. The hard spot is middle age.

[20] This just for expositional clarity. This way, we do not have to average on the realization of cash.

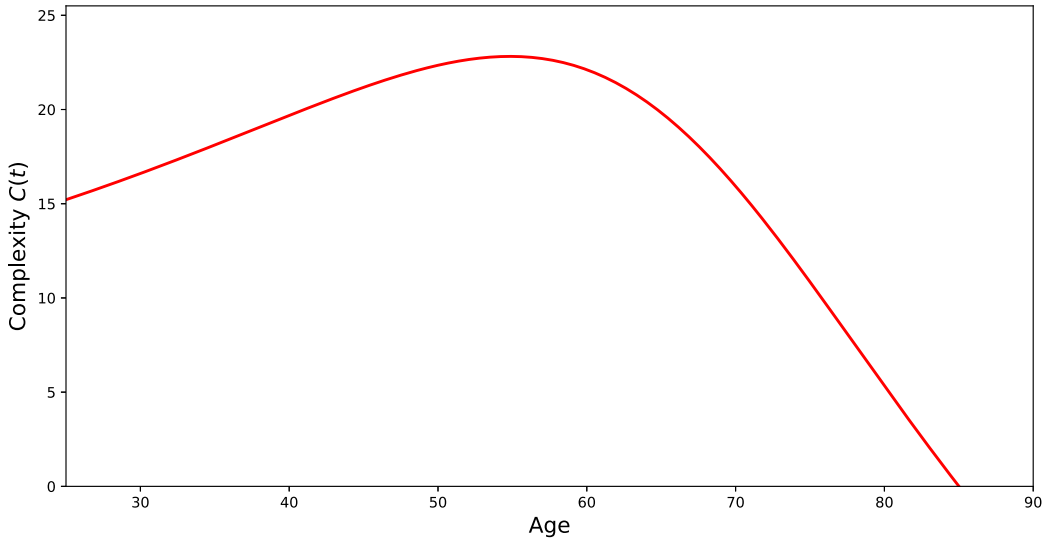


Figure 6: Complexity of life, as a function of age. *Notes*. This graph plots the complexity of the life-cycle problem at age t. Calibration parameters are in Section C.

The model could be enriched, with big decisions early in life (e.g., choice of a profession), leading to the fact that "consumption choice" still keeps the shape derived here; but with the complexity of "human capital choice" being much larger early in life. Down the road of research, one could imagine correlating those predictions with survey measures of happiness, stress etc. Our point is that the model allows a simple angle to attack those rich issues.

3.6 Complexity of intertemporal consumption when interest rate and income vary

We finally introduce a last enrichment to our study of consumption, with a non-trivial interest rate to take into account. It will guide our experimental investigation.

Setup We now present an application, which is intertemporal consumption choice, where the interest rate varies. Agents live for $T+1$ periods $t = 0 \dots T$, have earnings e_t in period t, and can borrow and save at interest rate r. Their flow utility function is $u(c) = \frac{c^{1-\frac{1}{\psi}}}{1-\frac{1}{\psi}}$ and there is no discounting. So the rational problem is:

$$\max_{c_1,c_2} \sum_{t=0}^{T} u(c_t) \text{ s.t. } \sum_{t=0}^{T} \frac{c_t - e_t}{(1+R)^t} = 0. \tag{31}$$

Optimal consumption at time 0 is:

$$c_0^r = \frac{\sum_{t=0}^{T} \frac{e_t}{(1+R)^t}}{\sum_{t=0}^{T} (1+r)^{(\psi-1)t}} \tag{32}$$

To see it more clearly, we call $\hat{e}_t = e_t - e_0$, and assume r and the $\hat{e}_t$ to be small. Then, the Taylor expansion of c_0 is (with residual error $O\left(\|\hat{e}\| \, |r| + r^2\right)$)

$$c_0^r = e_0 + \sum_{t=1}^{T} b^e \hat{e}_t + b^r r, \qquad b^e = \frac{1}{1+T}, \qquad b^r = -\psi \frac{T}{2} e_0 \tag{33}$$

Intuition Before presenting the formal analysis, let us inspect our intuition about the effect of different parameters and pit it against the model predictions. First, (33) confirms that incorporating a change in permanent income e_0 is easy (keeping the $\hat{e}_t$ constant): one just consumes it one for one.

However, incorporating a change in $\hat{e}_t$ is harder: one needs to divide it by the total number of periods, which is evidently more challenging.

Incorporating the interest rate is even harder, because one needs to take into account three quantities: the number of periods in the future, T, and the intertemporal elasticity of substitution (IES), ψ. The latter is, we submit, particularly hard to assess.

Formal analysis Let us think about $b_e = \frac{1}{1+T}$ (see (33)). It's composed of just one element, T, so by Proposition 1 its complexity is just the complexity of that element. We call it $\bar{c}$. This gives

$$C^e = \bar{c}$$

The complexity of $b_r = -\psi \frac{T}{2} y_1$ is more recondite. It is composed of two non-obvious elements, T and ψ. We here assume for simplicity that their complexity is the same, $\bar{c}$. This a form of ignorance prior. We suspect that the elasticity is harder to introspect about, so that more generally $c_\psi \geq c_T$. Given the multiplicative functional form, $b_T^r T = b_\psi^r \psi$, so that the shares are equal to $\frac{1}{2}$. So $C_r = \mathcal{C}\left(\frac{1}{2}, \frac{1}{2}, \bar{c}, \bar{c}\right) = 2\bar{c}$

$$C^r = 2\bar{c}$$

Then, for the complexity of the whole problem, take

$$C = \mathcal{C}\left(s_r, s_{\hat{e}}, C_r, C_e\right)$$

where s_r is the share of the interest rate in the consumption problem:

$$s_r = \frac{(b^r r)^2}{(b^r r)^2 + (b^e \hat{e})^2}$$

and $s_e = 1 - s_r$. When $\hat{e} = r = 0$, s_r is undefined, but the complexity is zero. Hence (by Lemma 4) total complexity is between 0 and $C_r + C_e = 3\bar{c}$.

All in all, we see how the model makes precise the intuition above: the interest rate is more complex (actually, twice as complex) as income. It gives the complexity of the intertemporal consumption problem as a function of the parameters (how much interest rates matter compared to future income deviations), which will allow for experimental investigation below.

3.7 Complexity of choosing between different financial products

Consider an agent who chooses between K financial products. What's the complexity of this decision problem? We use the empirical examples of Célérier and Vallée (2017), who study "equity protected" financial products. An example is "At maturity, if the underlying index registers a level equal to or higher than 70% of its strike level, the product offers a capital return of 104% of the initial investment; otherwise, the product offers a capital return of 70% the initial investment."

Formally, if r is the return on the index, the return of the product k is $R^k(r, x_k)$, where $x_k = (x_{kd})_{d=1...D_k}$ is a set of features or dimensions of the problem that parametrize that payoff. Célérier and Vallée (2017) use as their main complexity measure the number of features ($\#x^k = D_k$).[21] How would they do it if they used our measure? Applying our model, the consumer's utility buying product k is, omitting background risk and the like:

$$v^k\left(x^k\right) = \mathbb{E}\left[U\left(R^k(r, x_k)\right)\right] \tag{34}$$

This fits into our general framework (38). Hence, we apply the general formula, including

$$v^k_{x_{kd}} = \mathbb{E}\left[U'\left(R^k(r, x_k)\right) R^k_{x_{kd}}(r, x_k)\right]$$

This shows how the "important" features will matter more, where "important" means a large impact on the real return R^k, especially when marginal utility is high, i.e. for low returns.

In the case where $\alpha \to \infty$ and the "micro" complexities c_{kd} are all the same (say, equal

[21] They further use the number of scenarios and the number of characters used to describe the product.

to 1), then our measure is the same as Célérier and Vallée (2017)'s: the number of features. However, in the general case, the measures differ. By calibrating the complexity of different types of features one might obtain a structural model of complexity in this context of financial products. Of course, the same could be done, in principle, for other types of goods, e.g., the selection of health care products (Abaluck and Adams-Prassl, 2021a).

3.8 Complexity of mortgage refinancing

Knowing when to refinance a mortgage seems very tricky, as it involves fixed costs ("points" is the mortgage jargon), the option value of doing it later, forecasting interest rates etc. To quantify this decision, one can apply our formula. Call x the vector of determinants, and v the value of refinancing (whose normative value is complicated, though in some cases admits a closed form, Agarwal et al. (2013)). Then, we calculate the shares as $s_i = k\left(v_{x_i} x_i\right)^2$, and apply our formula.

4 Experimental Evidence on Complexity

Turning to the experimental side of complexity, we have two goals. First, in Section 4.1 we propose a novel measurement of complexity that is motivated by and tightly linked to our model. Second, in Section 4.2 we harness this measurement tool to conduct a simple experiment that acts as a proof of concept for the main behavioral implications of the model.

4.1 A portable methodology for measuring subjective complexity

We propose a new experimental methodology for measuring perceived complexity, with three key objectives. First, the measure should be disciplined by the theory. Second, it should capture fine-grained variation in people's subjective perceptions of complexity across variants of a task. Third, it should be simple to implement and portable across decision environments.

To meet these criteria, our methodology leverages simple pairwise comparisons between decision problems that subjects have previously completed. Consider a participant who first works on an intertemporal consumption problem with an interest rate of zero and then completes a problem that has a positive interest rate but is otherwise identical. On a subsequent screen, we again show both decision screens, presented side-by-side. The participant is then asked in which of the two problems they would find it "more dif-

ficult" to come up with an answer that falls within a specified window of, e.g., $3 to either side of their subjectively optimal, i.e., expected-payoff-maximizing answer.

A few remarks on this basic version of a complexity judgment are in order. First, previous work relies on non-comparative assessments of individual tasks on absolute scales. For instance, Enke and Graeber (2023) use judgments on a probability scale and Oprea (2020) has subjects choose a monetary amount. A host of evidence from the cognitive sciences suggests that people find comparative assessments easier than absolute assessments (e.g., Stewart et al., 2005) and that absolute judgments are, in fact, often unduly influenced by comparisons, as documented in the vast literature on cognitive comparison effects, such as anchoring (Kahneman et al., 1982). We therefore rely on a pairwise comparison technique that we propose is cognitively less challenging than an absolute assessment.[22]

Second, the measurement implements the notion of *pragmatic complexity* introduced above that reflects the practically relevant notion of getting a task approximately – but not necessarily exactly – right. Specifically, rather than asking subjects to assess the complexities of getting a task exactly right, which is often near-impossible in continuous problems and irrelevant for many practical purposes, we ask them to evaluate how difficult it is to get close to the payoff-maximizing answer. The measure is therefore defined up to a suitable definition of proximity to optimality in a given problem. In effect, this is equivalent to picking a specific point in the complexity curve developed above. The measurement, however, is meaningful independent of which exact window the experimenter chooses. When assessing task variants within a given problem class such as the intertemporal consumption problem, the width of the optimality window is kept constant. When making comparisons across different problem classes, we suggest choosing window sizes in a way that they are large enough so that obtaining that level of accuracy is still meaningful (achievable in principle) to subjects, yet not too large so that most subjects believe they can achieve this performance with reasonable effort. A portable way of choosing the window size across problem classes is to pick it so that some fraction (e.g., 50%) of actual responses are expected to fall within this level of accuracy.

Third, the basic measure is unincentivized, although we also provide and test an incentivized variant. As recent experimental work on this topic suggests, unincentivized measures often do equally well or even outperform incentivized measures (see, e.g., the discussion in Enke and Graeber, 2023). At the same time, unincentivized measures circumvent the need to explain a complicated payment structure to subjects (potentially

[22]We validate this by comparing the performance of our measure against that of cognitive uncertainty as suggested by Enke and Graeber (2023).

interfering with "behavioral incentive compatibility"; Danz et al. (2022)), and they avoid potential confounds from incentive dilution or interactions with the main experimental task of interest. We also provide an incentivized variant, in which subjects are asked to forecast other subjects' subjective feeling of complexity and receive a higher payoff for a more accurate forecast.

Fourth, our measure is easy and almost costless to implement for researchers, as it requires virtually no additional instructions.

Fifth, despite its plain binary nature in its elementary form, our methodology can be used to obtain numerical complexity scores that are readily comparable across decision classes and experimental environments. To do so, we implement the following protocol. After each new decision problem, we ask subjects to make a number of pairwise complexity assessments between this problem and previous decision problems the subject has faced. We pick the iterative comparison pairs efficiently using the *Quicksort* algorithm.[23] This allows us to construct a subject-level complexity ranking of the decisions tasks from pairwise comparisons alone. To convert these rankings to an absolute score, we add two canonical "anchor tasks", which merely serve to normalize the scale in a simple, domain-independent fashion. Empirical researchers can freely choose anchor tasks in a way that is domain-general or adequate for their specific setting of interest. Here, we employ two simple mental algebra tasks: subjects were asked to compare the complexity of a task to that of computing 7+53 as well as to computing 7+53+394+7893 in their heads.

We obtain numerical complexity scores from these pairwise comparison data by applying a well-established method from psychometrics, popularized by Bradley and Terry (1952). Specifically, we assume that subject i facing task τ perceives complexity as:

$$c_{i\tau} = c_\tau + \varepsilon_{i\tau} + b_i \tag{35}$$

where c_τ is the average complexity of task τ, b_i is a subject-specific complexity, and $\varepsilon_{i\tau}$ is noise. [24] Hence, the probability that task τ is judged more complex than τ', i.e. $c_{i\tau} > c_{i\tau'}$, is:

$$p_{\tau,\tau'} = \mathbb{P}\left(c_\tau + \varepsilon_{i\tau} > c_{\tau'} + \varepsilon_{i\tau'}\right) \tag{36}$$

We assume that the noise term $\varepsilon_{i\tau}$ is i.i.d. extreme value-distributed, so that the model reduces to Logit on pairs and can be estimated using conventional software. To obtain interpretable units, we rescale estimated parameters so that the simple addition task has

[23] On average, this entails approximately $\log_2(n!)$ comparisons for n tasks. In our experiment, subject perform 13.5 comparisons on average for 7 tasks.

[24] One could imagine various refinements, along the lines of $c_{i\tau} = a_i c_\tau + \varepsilon_{i\tau} + b_i$, with a variable a_i, or equivalently a subject-specific noise, but we do not pursue that here.

a complexity of $c_{\text{easy}} = 1$ and the harder addition has $c_{\text{hard}} = 10$. Standard errors are adjusted accordingly.

4.2 Experiment on intertemporal consumption

4.2.1 Design

We implement a naturalistic version of the intertemporal consumption task outlined in Section 3.6. Participants are informed that their choices in this experiment will affect the amounts of two time-dated food delivery vouchers by UberEats, each of which is only valid during a specific future time window: the *Early Period* is the week starting after the day of the survey, and the *Late Period* is the week starting six months following the survey.[25] To ensure concave preferences via real satiation, respondents are further told at the beginning of the study that the vouchers are only valid for a specific type of cuisine, Mexican restaurants.

Participants receive two separate voucher budgets, one for each period, and can then re-allocate these budgets between the two periods by saving or borrowing at a certain interest rate. Moreover, even though the time difference between *Early* and *Late Periods* is fixed across all tasks, the interest rate compounds a specific number of times that varies across problems. On the decision screen, participants are asked to enter their desired consumption in the *Early Period*, as shown in Figure F.3. The complexity comparison screen is designed to look like two decision screens side by side, as shown in Figure F.4.

We call e_1 the *Early Period* budget, e_2 the *Late Period* budget, r the interest rate and T is the number of rounds of compounding. Subjects are reminded that while the interest rate compounds T times, the *Early* and *Late* periods are, in fact, always six months apart. To estimate utility curvature ψ, we complement our design with six simple price list choices, see Appendix F.2.

Participants completed five different consumption problems. The selection of task parameters was fully randomized, with $e_1, e_2 \in \{\$150, \$200, \$250\}$, $r \in \{0\%, 8\%, 16\%\}$ and $T \in \{1, 2, 3\}$, leaving us with $3^4 = 81$ task configurations in total. For analyses at the task level, we only include tasks that have been completed by at least three respondents.

The computation of our measure of complexity via the Bradley-Terry procedure is detailed in Appendix F.3. To provide an example, our approach yields the following complexities for tasks with initial endowments (e_1, e_2, r, T):

[25] The food delivery voucher paradigm is adapted from Enke et al. (2023). Fenig and Petersen (2024) presents an alternative rich design with many periods, and also contains a useful survey of the literature in experimental macro.

$$C(200, 200, 0\%, 1) = 0.00 \qquad C(200, 250, 8\%, 1) = 2.98$$
$$C(150, 200, 0\%, 1) = 1.00 \qquad C(200, 250, 8\%, 3) = 3.83$$

Hence, among those four configurations, the problem with identical *Early* and *Late Period* budget as well as no interest rate is the simplest. The problem with different budgets, a non-zero interest rate and multiple compoundings is the most complex.

We conducted online experiments using the survey platform Prolific. The experiment was approved by Harvard IRB (IRB17-2035). Our experimental instructions are reproduced in Appendix F.1. We screened out participants who failed a simple attention check or incorrectly answered any one of a series of comprehension questions about the task instructions twice. After applying these exclusion restrictions we end up with a sample of 194 subjects (out of 207 who started the experiment). Subjects were paid $3.5 for completing the study in its entirety. Respondents spent a median time of 20 minutes on the survey.

4.2.2 Results

The objective of this empirical exercise is to bring our theoretical measure of complexity to the data by analyzing how it relates to (i) experimentally measured complexity as well as (i) behavioral implications as predicted by the model. Our analysis proceeds in three steps. First, we examine how well our model of complexity predicts respondents subjectively perceived complexity across tasks. Second, we test to what extent model-based complexity of a task is related to the error rate of a task as well as response times. Third, we test whether model-based complexity predicts attenuated responses to variation in task parameters. All analyses are performed on task-level averages or estimates, involving only tasks which have been completed by at least three respondents.

Model complexity predicts subjective perceptions of complexity Figure 7 shows average complexity scores for each task configuration. Visual inspection provides the first piece of evidence in line with model predictions: tasks in which the exchange rate can be neglected ($r = 0$) and where the budgets in the two periods coincide ($e_1 = e_2$) are indeed perceived as least complex, by a margin. Conversely, tasks with a non-zero exchange rate and large differences in the budgets for the two periods tend to be judged as being most complex.

Figure 8 makes the relationship between model complexity and estimated complexity explicit, plotting average subjective complexity against model complexity by task config-

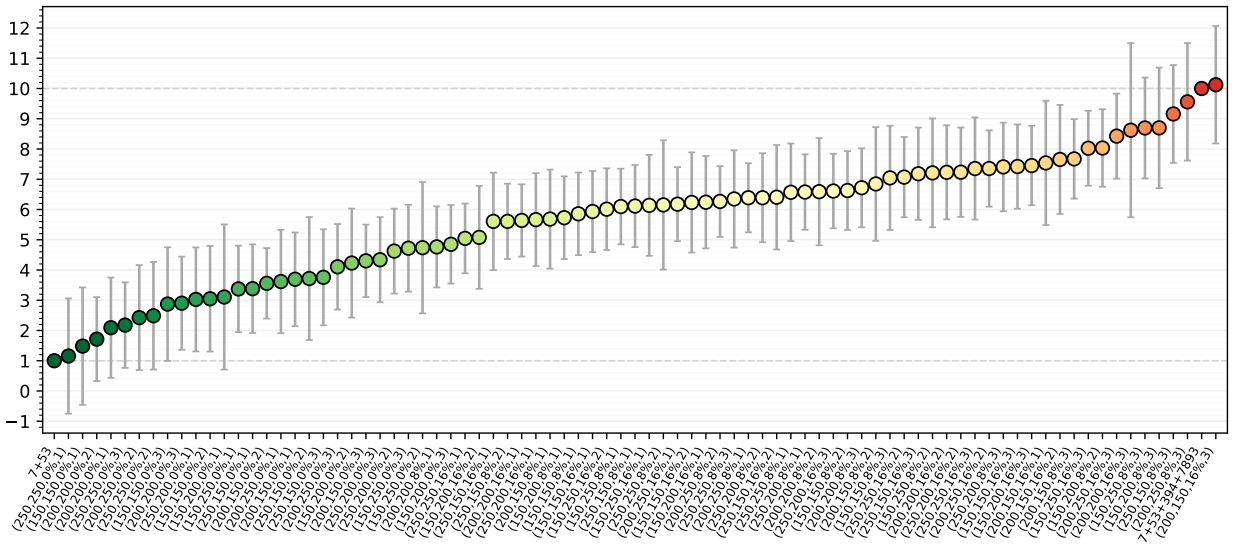


Figure 7: Subjective complexity across tasks. *Notes*. This graph shows respondents' subjective complexities of the consumption tasks, as estimated by the model detailed in Section 3.6. Tasks are labelled as (e_1, e_2, r, T), e.g. in $(150, 250, 8\%, 3)$ subjects receive an *Early Period* and *Late Period* budgets of $e_1 = 150$ and $e_2 = 250$, the interest rate is $r = 8\%$ and it compounds $T = 3$ times. Bars show 95% confidence intervals. *Interpretation.* In line with the predictions of the model, tasks with an interest rate of 0 and with a small difference between e_1 and e_2 are judged easiest, while those with a large difference in budgets $e_2 - e_1$, a high interest rate r and multiple compoundings T are judged most complex.

uration. Recall that the most basic prediction of the model is that, when the model predicts that a task is complex, subjects should also report that they feel that it is complex. Reassuringly, Figure 8 confirms this, showing a pronounced positive relationship. Column 1 of Table 1 shows corresponding regression results. This simple regression achieves an R^2 of 79%, with a positive and significant coefficient on model complexity. Our experimental results thus validate the fundamental thrust of our model on perceptions of complexity: when the model deems a task complex, subjects broadly concur.

Model complexity predicts errors and responses times Based on individual calibrations of utility curvature, we estimate participants' optimal responses in each task and use that to compute the average absolute error implied by their decisions. Column 2 of Table 1 shows that model complexity has a significant positive effect on average absolute errors in a task.

We also find that model complexity predicts thinking time, as proxied by the time spent on the decision screen. Interestingly, model complexity has slightly greater explanatory power than respondents' subjective complexities across both measures. Our findings

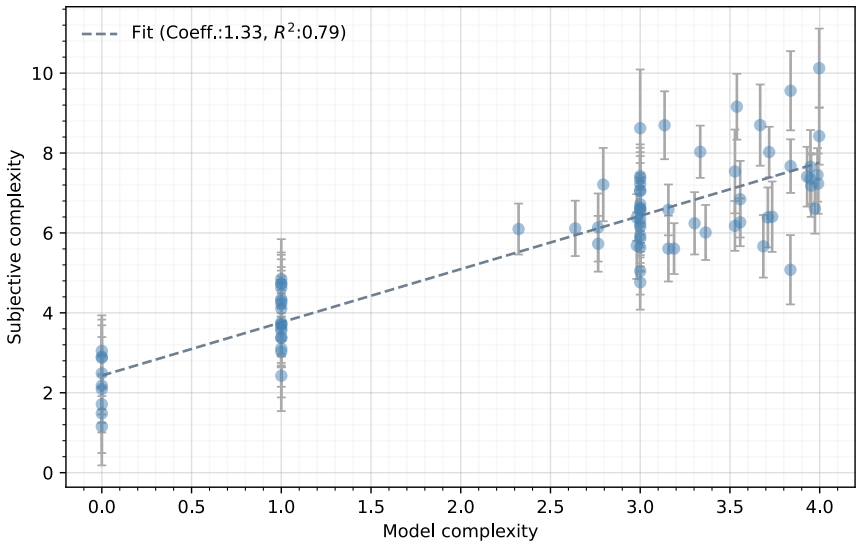


Figure 8: Model complexity predicts subjective complexity. *Notes*. This graph plots the estimated subjective complexities of consumption tasks (see Section 4.1) against the complexity predicted by our model (see Section F.3). Bars show 95% confidence intervals. *Interpretation.* There is a positive correlation between complexity predicted by the model and complexity as subjectively reported by participants in the experiment. Moreover, the model explains 79% of the variance in subjective complexity by task.

thus confirm the insight from our model that tasks that are more complex through the lens of our model lead cognitively constrained agents to commit larger errors on average.

Model complexity predicts attenuation in responses Table 2 presents results on the relationship between model complexity and behavioral dampening of reactions to problem parameters. Column 1 shows a regression of the average first-period allocation in a task on the rational consumption, computed using the globally calibrated utility parameter. Column 2 adds an interaction between the rational consumption and model complexity, finding a significantly negative coefficient: complexity makes agents underreact to problem parameters. Column 3 replaces the continuous measure of model complexity by a dummy equal to 1 if model complexity is higher than the median, again documenting a significantly negative coefficient: its magnitude is now more interpretable, and shows that the size of this effect is large. Note that the increase in explanatory power, R^2, from adding either of the interaction terms is also substantial, between 35% and 46%. Our experimental results thus confirm the key behavioral predictions of our model that more complex problems lead agents to underreact to variation in the rational response. Appendix F contains a more detailed discussion of the experimental design and findings.

Table 1: Model-predicted complexity increases subjective feeling of complexity, errors, and decision time

	Subjective complexity	Absolute Error		Decision time	
Model complexity	1.333***	9.979***		1.893***	
	(0.075)	(1.679)		(0.590)	
Subjective complexity			5.053***		1.056**
			(1.275)		(0.458)
Constant	2.429***	41.178***	36.811***	19.371***	17.989***
	(0.169)	(4.105)	(7.174)	(1.490)	(2.587)
Observations	81	81	81	81	81
R^2	0.787	0.275	0.159	0.087	0.061

Notes. This table displays different regressions on model complexity, computed as in Section F.3, or subjective complexity, computed as in Section 4.1. Absolute error is the average absolute error within a task, computed for each respondent as the difference between their stated answer and the rational answer using their elicited ψ_i (see Section F.2). Decision time is the average time in seconds spent on the decision screen. All regressions are done at the task level, averaging over participants who answered that task. *Interpretation.* Higher model complexity leads to greater subjective complexity, larger absolute errors and longer decision times, as predicted by the theory (see Section 2.3).

Table 2: Higher model complexity leads to more attenuation in responses to incentives

	First Period Allocation		
c_1^r	0.973***	0.982***	0.950***
	(0.135)	(0.119)	(0.122)
Mod. comp. $\times$ c_1^r		-0.070***	
		(0.018)	
$\mathbb{1}_{\text{Mod. comp.}>\text{Median}} \times c_1^r$			-0.156***
			(0.047)
Constant	-27.542	-5.751	-8.934
	(24.561)	(24.188)	(24.571)
Observations	81	81	81
R^2	0.357	0.456	0.450

Notes. Column 1 regresses the average *Early Period* allocation for each task on the rational allocation, computed using the calibrated ψ. Column 2 adds an interaction term with model complexity. Column 3 adds an interaction term with a dummy variable equal to 1 if a task's model complexity is greater than the median. *Interpretation.* The coefficient the interaction terms is significantly negative, showing that greater model complexity leads to attenuated reaction to the rational allocation, as predicted by the theory (see Section 2.3).

5 Complexity of Discrete Choice

In many situations, we face a choice between two or several actions (e.g. buy or rent an apartment), rather than a continuous choice (e.g. how large of an apartment to buy). Hence, we now extend the model to discrete choice.

5.1 Transposing our behavioral theory of continuous choice into a behavioral theory of discrete choice

We use a theoretical device to transform a behavioral theory of continuous choice (e.g. the one we proposed) into a behavioral theory of discrete choice.

Position of the Problem Consider a theory for a behavioral treatment of smooth problems, where the action a is continuous – it lives in a "nice" subset A (e.g., connected, open) of $\mathbb{R}^d$

$$\text{Choice of continuous action } a: \max_{a:a\in A\subseteq\mathbb{R}^d} u(a,x). \tag{37}$$

Next, let us consider a discrete problem, where the action lives in a discrete subset K:

$$\text{Choice of discrete action } k: \max_{k:k\in K\subseteq\mathbb{N}} v(k,x), \tag{38}$$

We still maintain that for each k, $v(k,x)$ is twice continuously differentiable in x. How do we transpose our theory of a our continuous action problem (37) to a discrete choice problem (38)?

Proposed solution We define as $a(k)$ the probability of playing action k, so $\sum_a a(k) = 1$, $a \in \Delta(K)$. We then define the utility function, concave in a, as

$$u(a,x) := \sum_k a(k)\, v(k,x) - \beta D\left(a, a^d\right), \tag{39}$$

where $D\left(a, a^d\right)$ is a divergence measure that is convex in a, $a^d \in \Delta(K)$ is a default action probability, and $\beta > 0$. We shall take the Kullback-Leibler divergence, or relative entropy, between a and a^d:

$$D\left(a, a^d\right) = 1 + \sum_k a(k)\left(\ln\frac{a(k)}{a^d(k)} - 1\right) \tag{40}$$

We could take other divergences $D\left(a, a^d\right)$.[26] We discuss below the specification of a^d and β.

We propose to use the machinery developed for $u\left(a, x\right)$, e.g. the allocation of attention to x, and the feeling of complexity, for the discrete problem $v\left(k, x\right)$.[27]

Definition 2. *(Complexity: going from continuous choice to discrete choice).* The complexity of the discrete problem (38) is the complexity of the associated smooth problem (37), as seen in Section 2, with the associated utility function defined in (39).

For instance, suppose that we ask "what is the mental effort L_i associated with variable x_i, and its impact on utility?" in the discrete problem. We apply the continuous-choice framework of Section 2 to the utility function $u\left(a, x\right)$, and obtain an attention to x_i, and a gain from attention; this applies to the discrete problem (38).

This approach is conceptually simple. To be useful, it needs to lead to relatively simple concrete calculations. We verify this in the next section, and proceed to illustrations in Section 5.3.

5.2 Complexity of discrete choice model: Concrete measures

We now detail the machinery of discrete choice. The reader eager to see concrete results may wish to go straight to the concrete illustrations of Section 5.3. For notational simplicity, we replace $a\left(k\right)$ by a^k and $v\left(k, x\right)$ by $v^k\left(x\right)$. We start with a well-known fact.

Lemma 3. *The solution of* $\max_a u\left(a, x\right)$ *where* u *is in (39) is* $a = a^F\left(a^d, v\left(x\right), \beta\right)$*, where the probability of playing action* k *is:*

$$a^{F,k}\left(a^d, v\left(x\right), \beta\right) = \frac{a^{k,d} e^{\frac{1}{\beta} v^k}}{\sum_{k'} a^{k',d} e^{\frac{1}{\beta} v^{k'}}} \tag{41}$$

When $\beta \to 0$, the model converges to frictionless choice, i.e. the agent chooses $\text{argmax}_k\, v^k$.

We next derive the associated value of information in discrete-choice, which is analogue of Lemma 1, which covered the value of information in continuous-choice.

[26] For instance, we could take $D\left(a, a^d\right) = \sum_k \frac{\left(a(k) - a^d(k)\right)^2}{2a^d(k)}$, which yields the same divergence up to third order terms, when $a - a^d$ is small.

[27] The use of regularization is standard in the machine learning literature (Goodfellow et al., 2016), and also has been used in economics (e.g. Matějka and McKay, 2015) to study choice, but not yet for the feeling of complexity.

Lemma 4. (Value of information: Discrete choice) *The expected utility after acquiring precision m_i is, up to third order terms in σ_{x_i}:*

$$U(m) = U\left(0^+\right) + \sum_i V_i m_i, \qquad U\left(0^+\right) = \sum_k a^{d,k}\mu^k \tag{42}$$

where $V_i = \frac{V_i^0}{2\bar{\beta}}$ is the value of information in dimension i, with:

$$V_i^0(a) = \sum_k a^k \left(v_{x_i}^k - \bar{v}_{x_i}\right)^2 \sigma_{x_i}^2 \tag{43}$$

If information in dimension i only affects an option $k(i)$ (i.e., if $v_{x_i}^{k'} = 0$ for $k' \neq k$), V_i^0 reduces to:

$$V_i^0(a) = a^{k(i)}\left(1 - a^{k(i)}\right)\left(v_{x_i}^{k(i)}\right)^2 \sigma_{x_i}^2 \tag{44}$$

This value is more informative if it affects an option with high $a^k\left(1-a^k\right)$, which is maximal at $a^k = \frac{1}{2}$. This is because there is no point in furthering the search if an option has very low prior probability a^k of being used (as very likely it will not be chosen), or if an option has a very high probability of being chosen (as very likely it will be chosen anyway).

Specification of the noise parameter β and the default action a^d We need to specify β. It should be proportional the scale of the problem, and intuitively to the uncertainty. So, we we define the following function β^F, that returns a value of β given the parameters of the problem:

$$\beta^F\left(a^d\right) := \left(\sum_i V_i^0\left(a^d\right)\right)^{1/2} \bar{\beta}, \tag{45}$$

where function V_i^0 is in (43) and $\bar{\beta}$ is a unitless parameter. The intuition is the following. The expected value of full information prefactor is $\left(\frac{1}{2}\sum_i V_i^0\right)^{1/2}$, so has the interpretation of the average amount of change in relative valuations that can be learned if it was costless. If $\bar{\beta} = 0$, the agent's choice has zero trembling. Typically, we recommend the parametrization $\bar{\beta} = 1$. We shall make the following assumption.

Assumption 1. (Setting the default action a^d and the noise parameter β in discrete choice) *We set a^d and β in the following manner. First, we define $a^{d,0}(k) := \frac{1}{K}$ the "ignorance prior" on the actions, assigning equal probability to all actions; and compute an initial value of β, $\beta^0 :=$*

$\beta^F\left(a^{d,0}\right)$, *using (45). Then, we obtain the default action probability,* $a^d = a^F\left(a^{d,0}, v\left(x^d\right), \beta^0\right)$, *where* a^F *is the function in (41), and its associated* $\beta = \beta^F\left(a^d\right)$.

This action responds to incentives, as visible in $v\left(x^d\right)$. This procedure captures, we submit, some of the thought process of the agent. $a^{d,0}$ is the "complete ignorance priori", which allows to still calculate some sense of the dispersion of valuations, and the typical size β^0. This then allows to put more weight on the "more tempting" options, as in a^d, and a more refined β.[28]

Complexity of discrete choice We can finally obtain the complexity of discrete choice.

Proposition 11. (Complexity of discrete choice, input-based) *As in Proposition 1, suppose that the micro cognitive production function* $m_i\left(L_i\right)$ *is in (10), and that and the optimum we have an interior solution* ($m_i\left(L_i^*\right) \in (0,1)$). *Then the complexity is as in Proposition 1, with complexity aggregator* $C = \mathcal{C}\left(s_i, c_i\right)$ *where* $s_i = \frac{V_i^0}{\sum_j V_j^0}$ *and* V_i^0 *is defined (43).*

We next move to the output-based complexity.

Proposition 12. (Complexity of discrete choice, output-based) *The output-based complexity is, with* $\bar{\mu} = \frac{1}{K}\sum_{k=1}^K \mu_k$,

$$1 - Q\left(L\right) = \frac{\sum_i V_i^0\left(1 - m_i\left(L\right)\right)}{\sum_i V_i^0 + 2\beta\left(\sum_k a^{d,k}\mu^k - \bar{\mu}\right)}. \tag{46}$$

Finally, we record the allocation of attention to each dimension of the problem.

Proposition 13. (Discrete choice: allocation of attention across dimensions) *Given a shadow cost of cognitive effort* w, *the attention* m_i *to dimension* i *is* $m_i = \mathcal{A}\left(\frac{V_i}{wc_i}\right)$, *where* $\mathcal{A}\left(v\right) = \operatorname{argmax}_m vm - Q^{-1}\left(m\right)$ *is the optimal attention given stakes* v, *and* V_i *is the value of information* i, *given in (43).*

5.3 Complexity of discrete choice: examples

We present some applications, deferring others to the online appendix (see Section D.10 for lotteries).

[28] One could imagine richer variants. For instance, if some actions are in "red" or otherwise visually salient (as in Li and Camerer (2022)), they will capture a large initial attention $a^{d,0}$. This would change our modeling of $a^{d,0}$,via a "bottom-up process", but not the modeling of a^d given $a^{d,0}$.

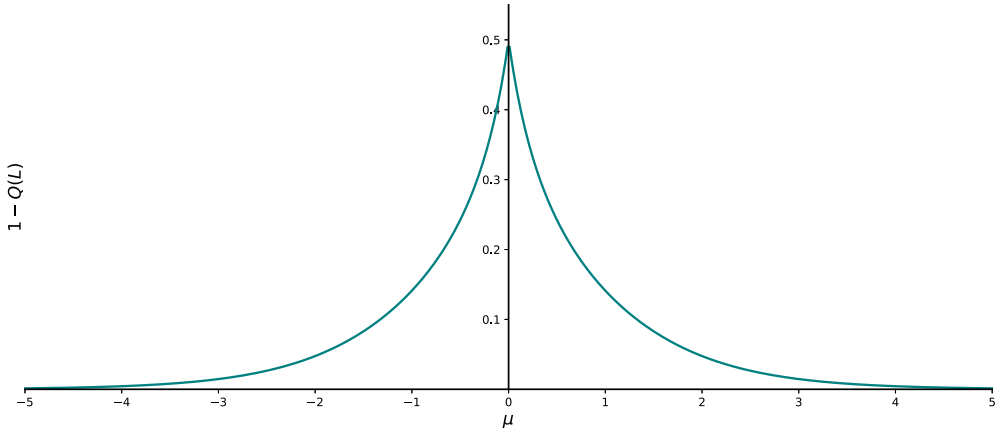


Figure 9: Complexity of discrete choice. *Notes*. The task is to choose between good 1 (with value $\mu + x$) and good 2 (with value 0). This plot shows the output-based complexity of a discrete choice problem $1 - Q(L)$ as a function of μ as in (47). *Interpretation*. When the first impression μ of the value of good 1 is from 0, the problem is easy: for instance, if μ is very positive (resp. very negative), it is "clear" that one ought to take good 1 (resp. good 2). However, when μ is close to 0, the choice is "complex", as more exploration is warranted. *Parameters*: $\sigma_x = 1$, $m(L) = \frac{1}{2}$ and $a = \frac{e^{\beta\mu}}{1+e^{\beta\mu}}$.

5.3.1 Complexity of choosing between two goods: a simple one and a complex one

Suppose that there are two goods (so, $K = 2$). Good 2 has a known value of 0, while good 1 has a value v^1 that requires thinking. We split $v^1 = \mu + v$, where μ is the default value (the one obtained with a very superficial glance at the problem), and $v = \sum_i b_i x_i$ captures the non-trivial dimensions to think about, normalized so that the x_i have mean 0. Defining as above $y_i \coloneqq b_i x_i$, so that $v = \sum_{i=1}^N y_i$ incorporates the impact of the N dimensions.

We apply the model. Applying (45) we have $\beta = \sqrt{\sum_{i=1}^N \sigma_{y_i}^2}\bar{\beta}$, and then the initial probability of choosing the first good, obtained after simply glancing at μ, is $a = a_1 = \frac{e^{\beta\mu}}{1+e^{\beta\mu}}$. Then, $V = \frac{1}{\beta} a(1-a)\sigma_y^2$, and (46) gives the output-based complexity:

$$1 - Q(L) = \frac{V(\iota) - V(m)}{\sum_k a^k \mu^k + V(\iota) - \bar{\mu}} = \frac{a(1-a)\sigma_x^2(1 - m(L))}{a(1-a)\sigma_x^2 + \frac{1}{2}|\mu|} \tag{47}$$

We see that the complexity is lower when $|\mu|$ is higher. This is plausible: if $\mu \gg \sigma_x$ is very high, the decision is a "no brainer": then a is close to 1, as clearly, the agent should choose good 1 over good 2 (which, we recall, has a value of 0). Likewise, if μ is very negative, good 1 is clearly inferior and the agent should choose good 2 (a is close to 0).

5.3.2 Complexity of choosing between several goods

With K uncorrelated goods and D dimensions Suppose that we have K goods. Each good k has D dimensions x_{kd}, and its value is

$$v_k = \mu_k + \sum_{d=1}^{D} b_d x_{kd}$$

with $var\left(b_d x_{kd}\right) = \sigma_{kd}^2$ the same across goods and all x_{kd} uncorrelated across dimensions and goods. So there are KD dimensions x_{kd} to potentially examine.[29] Let us take the case where "looking up dimension d for good k" has a cost c_{kd}.

Proposition 14. (Complexity of choosing among K goods). *The prior probability of using good k is $a^{d,k} \propto e^{\mu_k/\beta^0}$ with $\beta^0 = \left(\frac{1}{K}\left(1-\frac{1}{K}\right)\sum_{k,d}\sigma_{kd}^2\right)^{1/2}\bar{\beta}$. The value of information for good k's dimension d is $V_{kd}^0 = a_k^d\left(1-a_k^d\right)\sigma_{kd}^2$, and the importance is $s_{k\delta} = \frac{V_{k\delta}^0}{\sum_{k',\delta'}V_{k'\delta'}^0}$. The complexity of choosing among the K goods is then*

$$C = \mathcal{C}\left(s_{kd}, c_{kd}\right) \tag{48}$$

By Lemma 13, the attention to dimension d of good k as $m_{kd} = \mathcal{A}\left(\frac{V_{kd}^0}{2\beta w c_{kd}}\right)$.[30]

When the goods are symmetrical, in the sense that $\sigma_{kd} = \sigma_d$ and $c_{kd} = c_d$, one can write $s_{kd} = s_k s_d$, where $s_k = \frac{a_k^d\left(1-a_k^d\right)}{\sum_{k'} a_{k'}^d\left(1-a_{k'}^d\right)}$ is the relative importance of good k, and $s_d = \frac{\sigma_d^2}{\sum_{d'}\sigma_{d'}^2}$ is the relative importance of dimension d. The complexity of the problem is:

$$C = K^f C^f, \qquad K^f = \left(\sum_k s_k^{\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}}, \qquad C^f = \left(\sum_d s_d^{\frac{1}{\alpha}} c_d^{1-\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}} \tag{49}$$

where $K^f \in [1, K]$ is effective number of goods, and C^f is the effective complexity of each good.

In conclusion, we also provide a theory of the complexity of discrete choice, and of the allocation of attention. Its predictions seem sensible – more goods and more compositionality increase the complexity. But having an "extremely tempting option" does decrease the complexity, as K_f becomes 1.

[29] Formally, the dimensions are indexed by $i = (k,d) \in I = \{1,\ldots,K\} \times \{1,\ldots,D\}$.

[30] Hence, we recover the allocation of attention to the dimension d in K symmetric goods, as in Gabaix (2019), eq. (38), but can treat much more general cases, with asymmetric good.

6 Imperfect Metacognition: Solving the Infinite Regress Problem

So far we have modeled imperfect cognition (i.e., imperfect choice of action a) with perfect metacognition (perfect choice of the mental effort L_i). But it is clear intuitively that we all exhibit imperfect metacognition (see Bronchetti et al. (2023) for measurement). We next show how the model handles imperfect metacognition, after a simple enrichment. Put simply, we re-use the model "at the meta level", i.e. at the level of choosing cognitive actions ("inside, in the mind") rather than at the level of the physical action ("outside, in the world"). This has the advantage of solving the infinite regress problem—at some point, the benefit of thinking about thinking etc stops.We next show how this works.

6.1 The metacognitive problem

The "level 0" problem was to choose outside action a: $\max_a u(a, x)$, which gives a value of information $p^a = \frac{1}{2}|u_{aa}| \sum_i |a_{x_i} x_i|^2$, $p^{x_i} = \frac{1}{2}|u_{aa}| \, |a_{x_i} x_i|^2$. The "level 1" problem is "meta": choose inner mental effort L_i

$$\max_{L_i} V(L_i, p^{x_i}), \qquad V(L_i, p^{x_i}) := p^{x_i} Q(L_i; c_i) - wL_i \tag{50}$$

It has the same form as the level 0 problem, with decision variable L_i and hard-to-think about variable p^{x_i}, c_i, so that:

$$a^{Meta} = L_i, \qquad x^{Meta} = (p^{x_i}, c_i)$$

Hence, we can apply the model to this meta choice, where the agent imperfectly perceives the benefits and costs of attention, which are the components of x^{Meta}.

We next examine two questions: will this optimization process stop, and, does this extension yield testable predictions?

6.2 How the infinite regress stops

By Lemma 5 in Section A.1, the optimal attention is 0 when the gain from attention p^{x_i} is small enough. The next proposition shows that at each level of metacognition, the value of the metacognitive step decreases geometrically. This gives a concrete solution to the "infinite regress" problem of Simon (1955). Reasoning at the meta-meta-level (iterating the number of metas) brings exponentially smaller benefits, so agents will stop quickly.

Proposition 15. (Benefits of metacognition, meta-metacognition etc.) *Take the sparsity-inducing cognitive production function with $\alpha > 1$. The value of the meta stage falls geometrically*

$$p^{x_i,Meta} = \gamma p^{x_i}, \qquad \gamma = \left(1 - \frac{1}{\alpha}\right) \frac{1 - m^d}{2}$$

For instance, when $\alpha = 2$, $m^d = \frac{1}{2}$,then $\gamma = \frac{1}{8}$. As a result, the price of the $k-th$ round of meta-reasoning is

$$p^{x_i,Meta} = \gamma^k p^{x_i}$$

So the process of meta-cognitive iteration will stop at the earliest iteration k such that $\gamma^k p^{x_i} < wcA(\alpha)$, where $A(\alpha)$, given in (53), depends solely on α.

6.3 How agents with limited metacognition differ: Empirical predictions

We now derive some consequences of limited metacognition. Consider people who have very limited metacognition. Then, instead of true importance share s_i, they see $s_i^d = \frac{1}{N}$. Instead of true complexity c_i of each subpart, they see $c^d = \bar{c}$. Then, perceived complexity is

$$C^d = \mathcal{C}\left(s_i^d, c^d\right) = N\bar{c} = \text{number of non-zero elements}$$

So, we can predict that the less sophisticated the subjects, the more (i) they rely on "number of non-zero elements" for their feeling of complexity; (ii) their allocation of effort is more uniform across inputs and hence less sensitive to incentives; (iii) they're more influenced by visual cues, e.g. salient information in red or at the center. The latter point is due to the fact that when not directed by stakes, people react more to the "default" attention directed by visual cues, as in Li and Camerer (2022).

We formalize this as follows. Instead of the determinants $s_i \propto V_i$, people perceive V_i with dampening $M \in [0,1]$, i.e. $V_i^s = V_i^{\mathcal{M}^s}\left(V^d\right)^{1-\mathcal{M}^s}$, so their perception is $s_i(M) \propto s_i^{\mathcal{M}^s}$. Likewise, they perceive the complexities with dampening, so $c_i^{\mathcal{M}^c}\left(c^d\right)^{1-\mathcal{M}^c}$. Hence, the perceived complexity is given as follows.

Proposition 16. (Complexity measure with imperfect metacognition) *Call $\mathcal{M}^s, \mathcal{M}^c \in [0,1]$ the degree of precision in metacognition regarding shares vs micro complexities, so that $\mathcal{M}^s = 1$ (respectively $\mathcal{M}^s = 0$) means perfect (respectively fully imperfect) metacognition with respect to*

shares (and similar for $\mathcal{M}^c$. Then, the complexity measure with imperfect metacognition becomes:

$$C = \mathcal{C}\left(s_i\left(\mathcal{M}\right), c_i\left(\mathcal{M}\right)\right) = \left(\sum_i s_i^{\frac{\mathcal{M}^s}{\alpha}} c_i^{\mathcal{M}^c\left(1-\frac{1}{\alpha}\right)}\right)^{\frac{\alpha}{\alpha-1}} \left(\sum_j s_j^{\mathcal{M}^s}\right)^{\frac{1}{\alpha-1}} \left(c^d\right)^{1-\mathcal{M}^c} \tag{51}$$

Indeed, when metacognition is perfect ($\mathcal{M}^s = \mathcal{M}^c = 1$), measure (51) becomes our original measure (12). Bronchetti et al. (2023) shows that the allocation of attention is directionally sensible (people's attention increases in the stakes) but not quite optimal (their attention doesn't react to the stakes as much as it ought to). This corresponds to $\mathcal{M}^s, \mathcal{M}^c < 1$ in our model. More generally, one could imagine that researchers would routinely measure both the quality of cognition (via m) and that of metacognition (via $\mathcal{M}^s, \mathcal{M}^c$). One could also endogenize the metacognitive attentions $\mathcal{M}^s$, $\mathcal{M}^c$, applying again the model, at the metacognitive level.

7 Conclusion

We propose a simple, tractable theory of complexity for basic economic decision problems. The theory allows us to formalize intuitions about subjective perceptions of complexity. It provides tools to predict and measure complexity – e.g., the complexity of evaluating a cup of coffee, goods with complementarity, the tax system, lotteries, or the complexity of planning intertemporal consumption. Guided by the model, we then propose a simple and portable experimental methodology to measure perceptions of complexity. Our experimental analysis confirms the core tenets of the theory in an intertemporal consumption task. We hope that our measure will be useful to predict and measure complexity, and study its impact on economic decisions.

References

Abaluck, J. and A. Adams-Prassl (2021a). What do consumers consider before they choose? identification from asymmetric demand responses. *Quarterly Journal of Economics 136*(3), 1611–1663.

Abeler, J. and S. Jäger (2015). Complex tax incentives. *American Economic Journal: Economic Policy 7*(3), 1–28.

Abreu, D. and A. Rubinstein (1988). The structure of nash equilibrium in repeated games with finite automata. *Econometrica*, 1259–1281.

Agarwal, S., J. C. Driscoll, and D. I. Laibson (2013). Optimal mortgage refinancing: a closed-form solution. *Journal of Money, Credit and Banking 45*(4), 591–622.

Aghion, P., U. Akcigit, M. Gravoueille, M. Lequien, and S. Stantcheva (2024). Tax simplicity or simplicity of evasion? evidence from self-employment taxes in france. Working paper, National Bureau of Economic Research.

Alaoui, L. and A. Penta (2022). Cost-benefit analysis in reasoning. *Journal of Political Economy 130*(4), 881–925.

Angeletos, G.-M. and C. Lian (2023). Dampening general equilibrium: incomplete information and bounded rationality. In *Handbook of economic expectations*, pp. 613–645. Elsevier.

Aragones, E., I. Gilboa, A. Postlewaite, and D. Schmeidler (2005). Fact-free learning. *American Economic Review 95*(5), 1355–1368.

Banovetz, J. and R. Oprea (2023). Complexity and procedural choice. *American Economic Journal: Microeconomics 15*(2), 384–413.

Bernheim, B. D. and C. Sprenger (2020). On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting. *Econometrica 88*(4), 1363–1409.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012). Salience theory of choice under risk. *Quarterly Journal of Economics 127*(3), 1243–1285.

Bordalo, P., N. Gennaioli, and A. Shleifer (2013). Salience and consumer choice. *Journal of Political Economy 121*(5), 803–843.

Bossaerts, P. and C. Murawski (2017). Computational complexity and human decision-making. *Trends in cognitive sciences 21*(12), 917–929.

Bossaerts, P., N. Yadav, and C. Murawski (2019). Uncertainty and computational complexity. *Philosophical Transactions of the Royal Society B 374*(1766), 20180138.

Bradley, R. A. and M. E. Terry (1952). Rank analysis of incomplete block designs: The method of paired comparisons. *Biometrika 39*(3/4), 324–345.

Bronchetti, E. T., J. B. Kessler, E. B. Magenheim, D. Taubinsky, and E. Zwick (2023). Is attention produced optimally? theory and evidence from experiments with bandwidth enhancements. *Econometrica 91*(2), 669–707.

Byrne, R. M. and P. N. Johnson-Laird (1989). Spatial reasoning. *Journal of Memory and Language 28*(5), 564–575.

Campbell, D. J. (1988). Task complexity: A review and analysis. *Academy of Management Review 13*(1), 40–52.

Caplin, A., D. Csaba, J. Leahy, and O. Nov (2020). Rational inattention, competitive supply, and psychometrics. *The Quarterly Journal of Economics 135*(3), 1681–1724.

Carlin, B. I. (2009). Strategic price complexity in retail financial markets. *Journal of Financial Economics 91*(3), 278–287.

Carvalho, L. and D. Silverman (2023). Complexity and sophistication. *Journal of Political Economy: Micro*, Forthcoming.

Célérier, C. and B. Vallée (2017). Catering to investors through security design: Headline rate and complexity. *Quarterly Journal of Economics 132*(3), 1469–1508.

Chatterjee, K. and H. Sabourian (2020). Game theory and strategic complexity. *Complex Social and Behavioral Systems: Game Theory and Agent-Based Models*, 639–658.

Colliard, J.-E. and C.-P. Georg (2025). Measuring regulatory complexity. Working paper, HEC Paris Research Paper Series.

Crawford, V. P. and N. Iriberri (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica 75*(6), 1721–1770.

Danz, D., L. Vesterlund, and A. J. Wilson (2022). Belief elicitation and behavioral incentive compatibility. *American Economic Review 112*(9), 2851–2883.

Dean, M. and N. L. Neligh (2023). Experimental tests of rational inattention. *Journal of Political Economy*.

Dehaene, S. (2003). The neural basis of the Weber–Fechner law: a logarithmic mental number line. *Trends in Cognitive Sciences 7*(4), 145–147.

Dehaene, S. (2011). *The Number Sense: How the Mind Creates Mathematics*. Oxford University Press.

DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature 47*(2), 315–372.

Enke, B. and T. Graeber (2023). Cognitive Uncertainty. *Quarterly Journal of Economics 138*(4), 2021–2067.

Enke, B., T. Graeber, and R. Oprea (2023). Complexity and time. Working paper, National Bureau of Economic Research.

Enke, B., T. Graeber, R. Oprea, and J. Yang (2024). Behavioral attenuation. Working paper, National Bureau of Economic Research.

Enke, B. and C. Shubatt (2023). Quantifying lottery choice complexity. Technical report, National Bureau of Economic Research.

Farhi, E. and X. Gabaix (2020). Optimal Taxation with Behavioral Agents. *American Economic Review 110*(1), 298–336.

Fenig, G. and L. Petersen (2024). Dynamic optimization meets budgeting: Unraveling financial complexities. Working paper, National Bureau of Economic Research.

Frydman, C. and L. J. Jin (2022). Efficient coding and risky choice. *Quarterly Journal of Economics 137*(1), 161–213.

Fudenberg, D. and I. Puri (2023). Evaluating and extending theories of choice under risk. Working paper, MIT Economics.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics 129*(4), 1661–1710.

Gabaix, X. (2019). Behavioral inattention. *Handbook of Behavioral Economics 2*, 261–344.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review 110*(8), 2271–2327.

Gabaix, X. (2025). A theory of complexity aversion. *Available at SSRN 5185671*.

Gabaix, X. and D. Laibson (2002). The 6D bias and the equity-premium puzzle. *NBER Macroeconomics Annual 16*, 257–312.

Gabaix, X. and D. Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics 121*(2), 505–540.

Gabaix, X. and D. Laibson (2022). Myopia and discounting. Working paper, National Bureau of Economic Research.

Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006, September). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review 96*(4), 1043–1068.

Gershman, S. J., E. J. Horvitz, and J. B. Tenenbaum (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science 349*(6245), 273–278.

Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep learning*. MIT press.

Goodman, A. and I. Puri (2022). Bulls and binaries: Price anomalies and behavioral biases. Working paper.

Griffiths, T. L., F. Lieder, and N. D. Goodman (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science 7*(2), 217–229.

Halevy, Y. and G. Mayraz (2022). Identifying rule-based rationality. *Review of Economics and Statistics*, 1–44.

Heidhues, P. and B. Kőszegi (2018b). Behavioral industrial organization. *Handbook of Behavioral Economics: Applications and Foundations 1*, 517–612.

Huck, S. and G. Weizsäcker (1999). Risk, complexity, and deviations from expected-value maximization: Results of a lottery choice experiment. *Journal of Economic Psychology 20*(6), 699–715.

Imas, A., M. A. Kuhn, and V. Mironova (2022). Waiting to choose: The role of deliberation in intertemporal choice. *American Economic Journal: Microeconomics 14*(3), 414–440.

Iyengar, S. S. and E. Kamenica (2010). Choice proliferation, simplicity seeking, and asset allocation. *Journal of Public Economics 94*(7-8), 530–539.

Iyengar, S. S. and M. R. Lepper (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology 79*(6), 995.

Jin, G. Z., M. Luca, and D. Martin (2021a). Complex disclosure. *Management Science*.

Kahneman, D., P. Slovic, and A. Tversky (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.

Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica 47*(2), 263–291.

Kendall, C. and R. Oprea (2024). On the complexity of forming mental models. *Quantitative Economics 15*(1), 175–211.

Khaw, M. W., Z. Li, and M. Woodford (2021). Cognitive imprecision and small-stakes risk aversion. *The Review of Economic Studies 88*(4), 1979–2013.

Kurtz-David, V., A. Brandenburger, and P. Glimcher (2024). The limits of social cognition: Production functions and reasoning in strategic interactions.

Lacetera, N., D. G. Pope, and J. R. Sydnor (2012a). Heuristic thinking and limited attention in the car market. *The American Economic Review 102*(5), 2206–2236.

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review 107*(11), 3257–87.

Li, X. and C. F. Camerer (2022). Predictable effects of visual salience in experimental decisions and games. *Quarterly Journal of Economics 137*(3), 1849–1900.

Lian, C. (2021). A theory of narrow thinking. *Review of Economic Studies 88*(5), 2344–2374.

Lipman, B. L. (1995). Information processing and bounded rationality: A survey. *Canadian Journal of Economics*, 42–67.

Liu, P. and Z. Li (2012). Task complexity: A review and conceptualization framework. *International Journal of Industrial Ergonomics 42*(6), 553–568.

Lloyd, S. (2001). Measures of complexity: A nonexhaustive list. *IEEE Control Systems Magazine 21*(4), 7–8.

Manzini, P. and M. Mariotti (2007). Sequentially rationalizable choice. *American Economic Review 97*(5), 1824–1839.

Martínez-Marquina, A., M. Niederle, and E. Vespa (2019). Failures in contingent reasoning: The role of uncertainty. *American Economic Review 109*(10), 3437–3474.

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–298.

Molavi, P., A. Tahbaz-Salehi, and A. Vedolin (2024, 07). Model Complexity, Expectations, and Asset Prices. *The Review of Economic Studies 91*(4), rdad073.

Murawski, C. and P. Bossaerts (2016). How humans solve complex problems: The case of the knapsack problem. *Scientific Reports 6*(1), 1–10.

Nielsen, K. and J. Rehbeck (2022). When choices are mistakes. *American Economic Review 112*(7), 2237–68.

Oprea, R. (2020). What makes a rule complex? *American Economic Review 110*(12), 3913–51.

Oprea, R. (2024a). Complexity and its measurement. Working paper, Mimeo.

Oprea, R. (2024b). Decisions under risk are decisions under complexity. *American Economic Review 114*(12), 3789–3811.

Ortoleva, P. (2013). The price of flexibility: Towards a theory of thinking aversion. *Journal of Economic Theory 148*(3), 903–934.

Puri, I. (2024). Simplicity and risk. *Journal of Finance, Forthcoming*.

Ricœur, P. (1992). *Oneself as another*. University of Chicago Press.

Rubinstein, A. (1986). Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory 39*(1), 83–96.

Rubinstein, A. (1998). *Modeling Bounded Rationality*. MIT press.

Salant, Y. and J. L. Spenkuch (2022). Complexity and satisficing: Theory with evidence from chess. Working Paper 30002, National Bureau of Economic Research.

Schipper, B. (2015). Awareness. In H. van Ditmarsch, J. Halpern, W. van der Hoek, and B. Kooi (Eds.), *Handbook of Epistemic Logic*, Chapter 3, pp. 77–146. College Publications.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics 69*(1), 99–118.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690.

Sonsino, D., U. Benzion, and G. Mador (2002). The complexity effects on choice with uncertainty–experimental evidence. *Economic Journal 112*(482), 936–965.

Stewart, N., G. D. Brown, and N. Chater (2005). Absolute identification by relative judgment. *Psychological review 112*(4), 881.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *Science 185*, 1124–30.

Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics 12*, 579–601.

You, H. and X.-j. Zhang (2009). Financial reporting complexity and investor underreaction to 10-k information. *Review of Accounting studies 14*(4), 559–586.

Zhang, H. and L. T. Maloney (2012). Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in neuroscience 6*, 1.

Zwick, E. (2021). The costs of corporate tax complexity. *American Economic Journal: Economic Policy 13*(2), 467–500.

A Appendix: Some complements to the theory

A.1 The model with sparsity

One can set up

$$Q(L) = \max\left(1 - \left(\frac{L}{C} + 1\right)^{1-\alpha}, 0\right) \tag{52}$$

with $\alpha > 1$, while in the main model $\alpha \in [0, 1)$. This means that attention greater than 0 occurs for $L \in (C, \infty)$. So often we will have 0 attention to a feature, i.e. sparsity: $Q = 0$. Indeed, we have the following.

Lemma 5. (Sparsity threshold for $\alpha > 1$) *Suppose that people use the cognitive production function* (52) *with* $\alpha > 1$. *Then, in the problem* $\max_L pQ(L) - wL$ *the optimal attention is 0 iff:*

$$\frac{p}{wC} < A(\alpha), \qquad A(\alpha) = \frac{1}{\alpha - 1} \tag{53}$$

Proof. We have an interior optimum if the maximand of $V(L) = pQ(L) - wL$ is $L = 0$. As V is concave, this is the case iff $V'(0_+) \leq 0$, i.e. $\frac{p}{C}(1 - \alpha) \leq w$. □

Then, we get the following variant of Proposition 1. Its message is that functional forms do not change – just in all expressions, one replaces Q by $1 - Q$ and m by $1 - m$, and now $\alpha > 1$. In particular, the complexity aggregator (12) is the same.

Proposition 17. (Macro complexity from micro complexity, variant of Proposition *1* with $\alpha > 1$) *Suppose that the "micro" cognitive production function of component i is*

$$m_i(L) = \max\left(1 - \left(\frac{L_i}{c_i} + \phi_i\right)^{1-\alpha}, 0\right) \tag{54}$$

where the complexity of dimension i, c_i, is exogenous for now with $\alpha > 1$, and $\phi_i \geq 0$. If at the optimum of problem (9) *we are in the "interior region" ($m_i\left(L_i^*\right) \in (0,1)$, where L_i^* is the optimal allocation to dimension i), then the macro cognitive production function is*

$$Q(L) = \max\left(1 - \left(\frac{L}{C} + \Phi\right)^{1-\alpha}, 0\right) \tag{55}$$

with "macro complexity" $C = \mathcal{C}\left((s_i, c_i)_{i=1...N}\right)$, using the "complexity aggregator" (12), and with $\Phi = \frac{\sum_i c_i}{C}$. In addition, at the optimum L_i^*

$$1 - m_i(L_i^*) = \left(\frac{C}{c_i/s_i}\right)^{\frac{1-\alpha}{\alpha}} (1 - Q(L)) \tag{56}$$

so that attention to dimension i is larger than the average attention $Q(L)$ if and only if $\frac{c_i}{s_i} \geq C$.

The shifts ϕ_i and Φ may seem a bit odd, but they are largely innocuous. The polar cases are $\phi_i = 0$ (which is notationally simpler), and $\phi_i = 1$,which ensures a concave function $m_i(L)$ with $m_i(0) = 0$.

One can also get a more "intrinsic" formulation of C.

Lemma 6. (Intrinsic characterization of C as a function of marginal effort $L'(Q)$) *Consider the production functions*

$$Q(L) = \begin{cases} \min\left(\left(\frac{L}{C} + \Phi\right)^{1-\alpha}, 1\right) & \text{if } \alpha < 1 \\ \max\left(1 - \left(\frac{L}{C} + \Phi\right)^{1-\alpha}, 0\right) & \text{if } \alpha > 1 \end{cases}$$

with $C > 0$ and Φ is an arbitrary constant independent of L. To achieve a target precision Q, (55), the effort is, in the interior region ($Q \in (0,1)$)

$$L(Q) = C \times \begin{cases} Q^{\frac{1}{1-\alpha}} - \Phi & \text{if } \alpha < 1 \\ (1-Q)^{\frac{1}{1-\alpha}} - \Phi & \text{if } \alpha > 1 \end{cases}$$

so that

$$L'(Q) = C\bar{L}'(Q), \tag{57}$$

where

$$\bar{L}'(Q) = \frac{1}{|\alpha - 1|} \begin{cases} Q^{\frac{\alpha}{1-\alpha}} & \text{if } \alpha < 1 \\ (1-Q)^{\frac{\alpha}{1-\alpha}} & \text{if } \alpha > 1 \end{cases}$$

is independent of C and Φ. Hence, C is the prefactor in the derivative of effort with respect to the target performance.

Optimal effort, for a given shadow cost of effort Proposition 7 shows that in the "shadow cost" version, with a linear cost wL, the performance $Q(L^*(w))$ is independent of the shifter constant Φ.

A.2 Further predictions on actions vs complexity

We next record the attention to each dimension of the problem.

Proposition 18. *(Allocation of attention between components of a problems) Consider the setup of Proposition 1, with $\alpha < 1$. The optimum L_i^* satisfies:*

$$m_i(L_i^*) = \left(\frac{s_i C}{c_i}\right)^{\frac{1-\alpha}{\alpha}} Q(L), \tag{58}$$

so that attention to dimension i is larger than the average attention $Q(L)$ if and only if $\frac{c_i}{s_i} \leq C$.

The following refines Proposition 5

Proposition 19. *In the setup of Proposition 5, regressing $a - a^d$ on its rational components $a^r_{x_i} x_i$*

$$a - a^d = \sum_i m_i a^r_{x_i} \left(x_i - x_i^d\right) + k + \varepsilon$$

yields a regression coefficient equal to (in population)

$$m_i = \left(\frac{s_i C}{c_i}\right)^{\frac{1-\alpha}{\alpha}} M \tag{59}$$

from (58).

A.3 Multi-dimensional actions

In many situations, the action is multidimensional, so is in $a \in \mathbb{R}^K$. We state how the basic analysis essentially carries over.

Proposition 20. (Multidimensional action) *When the action is multidimensional ($a \in \mathbb{R}^K$), the value of more precision to dimension i is*

$$V_i = -\frac{1}{2} a'_{x_i} u_{aa} a_{x_i} \sigma^2_{x_i} \tag{60}$$

so that the relative importance shares are $s_i = \frac{V_i}{\sum_j V_j}$. The measure of complexity of the problem is otherwise the same as in the one-dimensional action case (e.g. Proposition 1).

Note that the dimensions of a_{x_i} is $K \times 1$, and that of u_{aa} (which is the second derivative) is $K \times K$. In the one-dimensional action case, $V_i = -\frac{1}{2} u'' \left(a^d\right) a^2_{x_i} \sigma^2_{x_i}$.

Online Appendix for

"The Complexity of Economic Decisions"

Xavier Gabaix and Thomas Graeber

August 28, 2025

This Online Appendix contains additional proofs, and complements to the theory and to the experimental part.

Table of contents

B Appendix: Review of Related Literature **59**

C Appendix: Omitted proofs **63**

D Appendix: Theory complements **71**

D.1 Endogenizing micro-complexity: Complexity of a problem with two layers in the production function of thought . . . 71
D.2 The mind as a cognitive economy . . . 73
D.3 Outside the interior region . . . 74
D.4 The value of the curvature parameter α . . . 75
D.5 Interaction with learning . . . 75
D.6 First vs. second order complexity aversion . . . 77
D.7 Cognitive-risk adjusted certainty equivalent . . . 78
D.8 Complexity of forecasting: More general version . . . 78
D.9 Ex post simplification of decisions, for instance to communicate an answer . 79
D.10 Application: Complexity of choosing between lotteries . . . 80

E Appendix: Complexity of basic arithmetic operations **81**

E.1 Complexity of a number . . . 81
E.2 Complexity of addition . . . 82
E.3 Complexity of subtraction . . . 82
E.4 Complexity of multiplication . . . 82
E.5 Complexity of the certainty equivalent in a gamble . . . 83
E.6 Complexity of a composite sum, $\sum_i p_i x_i$ . . . 83
E.7 Complexity of choosing between two gambles. . . . 84

F Appendix: Complements to the experimental part **84**
F.1 Survey design 84
F.2 Calibration of utility 85
F.3 Computation of model complexity 86
F.4 Validating complexity elicitations 87

B Appendix: Review of Related Literature

Our paper ties into various literatures. First, on the theoretical side, the central distinguishing feature of our model is that it provides a measure of complexity that applies to canonical economic models in a fairly general way. This is largely owed to the model relying on continuous actions and concepts. Earlier proposals used the number of parts in automata (Rubinstein, 1998), or sequences of discrete decisions (Gabaix et al., 2006), which makes them less tractable: they require either simulation by a computer or rather pain-staking enumeration of cases, and it is unspecified how they could be applied to smooth bread-and-butter problems such as intertemporal consumption. Instead, they are meant for simple games with a few discrete actions. Other models using continuous concepts do not feature an explicit theory of complexity. Sims (2003) builds a theory of Bayesian inattention based on communication through a channel with finite Shannon capacity. There is no theory of different degrees of complexity: how the (typically, entropy) cost of communicating of some things is higher than others (e.g. income). Caplin et al. (2020) uses an analogy with production theory that is similar to ours to propose a methodology for recovering attention costs from choice data in the rational inattention framework, but does not endogenize task complexity. Alaoui and Penta (2022) characterizes the core properties that are necessary and sufficient for the decision to reason about a problem to be captured by a cost-benefit analysis. Ortoleva (2013) models an agent with thinking aversion in risky choice, conceptualized as the cognitive cost required to figure out their preferences in order to make a choice. Gabaix (2014) builds a model of behavioral inattention with sparsity (i.e., not all dimensions are processed), which is tractable enough to apply to basic Arrow-Debreu-style consumer theory and general equilibrium, as well as to public finance and macroeconomics (Farhi and Gabaix, 2020; Gabaix, 2020). But in that line of research, the "cost" of thinking about each dimension is left unmodeled, and typically it's meant to be the same for each dimension. Our work provides such a theory of the "cost" of processing a dimension. Most other behavioral models (e.g. Bordalo et al., 2013) do not feature any explicit notion of complexity. They could be extended, so that their behavioral features would vary with complexity.

Notions of complexity also have a long tradition in behavioral economics. Simon (1955) introduces the notion of bounded rationality. Lipman (1995) gives a literature review on a model of bounded rationality stemming from limited information processing. Aragones et al. (2005) studies "fact-free learning", i.e. discovering new patterns among already known facts, and show that this amounts to an algorithmically NP-complete variable selection task. The analysis of fact-free learning, or *reasoning*, has been conducted for

a variety of settings, including unawareness (Schipper, 2015), sequential use of heuristics (Manzini and Mariotti, 2007) and strategic interactions (Crawford and Iriberri, 2007).

While general and precise notions of complexity have proven elusive, there is much work on the type of behaviors that are plausibly induced by complex decision environments, especially noisy or imprecise choices (see, e.g., Woodford, 2020) and heuristic or rule-based decision-making (Nielsen and Rehbeck, 2022; Halevy and Mayraz, 2022; Lacetera et al., 2012a; Tversky and Kahneman, 1974; Enke et al., 2023).

Various notions of complexity have received substantial attention in game theory. For example, Rubinstein (1986) explores strategic complexity in two-person supergames where players' strategies have to be carried out by finite automata. Abreu and Rubinstein (1988) extend Rubinstein (1986) by deriving more general properties for games with limited strategic complexity, and show that it results in discontinuities of strategies and payoffs with respect to game parameters.[31] More recently, Li (2017) introduces the concept of "obvious strategy-proofness" for a set of simple auctions and matching algorithms. The concept is binary in nature: when a mechanism is not obviously strategy-proof, there is no remaining notion of lower or greater complexity. Chatterjee and Sabourian (2020) provide an extensive literature review on strategic complexity in game theory.

In the field of computer science, there is a vast literature on complexity that mostly revolves around the prominent P vs. NP problem (see e.g. Lloyd, 2001). This work typically looks at asymptotic notions of complexity, focusing on the number of operations (for instance) needed to solve a problem with N components, and in particular considering whether this number is bounded by a polynomial. These concepts are of limited usefulness to economists, who have mostly been interested in problems that feature bounded rationality even with a very small number N of components, e.g., one or few elements. Hence, many of these insights from computer science are potentially less directly relevant for economic decision-making.

In contrast, to this literature, we try to propose a smooth, tractable, and psychologically intuitive notion of complexity. A few papers are already using some of the modeling ideas in this paper. Fenig and Petersen (2024) studies analytically and experimentally a life-cycle planning model related to the life-cycle we present, with a rich series of variants. Kurtz-David et al. (2024) model the complexity strategic interactions. Gabaix (2025) derives the optimal simplicity in a number of situations, e.g. the optimal simplicity of a tax system, or a contract.

Second, on the empirical side, we connect to several literatures. A recent line of experimental work on bounded rationality studies the effects of specific notions of complexity.

[31] See Rubinstein (1998) for a review of game theory with *finite automata*.

Enke and Shubatt (2023) study the complexity of binary choice between lotteries. Oprea (2020) asks respondents how much they would be willing to pay to avoid implementing different sets of rules, motivated by an automata-based concept of algorithmic complexity. Kendall and Oprea (2024) ask respondents to predict the output of data-generating processes, and show that they can rarely explain their mental model and seem attracted to common, simple models; "partition complexity" (adapted from Lipman (1995)) and "sparsity complexity" (adapted from Gabaix (2014)) perform best at predicting extraction. Concurrent work shows that a large set of classical decision anomalies in risky (Oprea, 2024b) and intertemporal choice (Enke et al., 2023) occur in variants of the decision problems that rule out the motivation explanations. These papers conclude that valuation errors as a consequence of complexity, i.e. the cognitive cost of information processing, are the source of these anomalies, but do not provide a theory of such costs. Martínez-Marquina et al. (2019) experimentally show that the difficulty of contingent reasoning through multiple possible states compromises maximization in the acquiring-a-company game. Another line of work explores the neural foundations and decision-making signatures of computational complexity (Bossaerts et al., 2019; Bossaerts and Murawski, 2017). For example, Murawski and Bossaerts (2016) make subjects solve the knapsack problem and find that human performance decreases with algorithmic complexity as defined in computer science. Jin et al. (2021a) design an experiment in which senders are forced to write a truthful report to receivers, and find that senders strategically shroud information by making their reports intentionally complex. A different notion of complexity is studied in work on choice overload (Iyengar and Lepper, 2000; Iyengar and Kamenica, 2010), which suggests that a larger set of options induces various behaviors, including a preference for simple-to-understand options. Fenig and Petersen (2024) provide an experimental investigation of complexity in a planning experiment.

Abeler and Jäger (2015) conduct an experiment in which subjects work for a piece rate and face taxes of varying complexity but identical marginal incentives, and show that more complex tax schedules generate underreaction to changes in incentives. Dean and Neligh (2023) experimentally examine the predictions of rational inattention models and find that people adjust attention when there are changes in incentives, but note that results are qualitatively inconsistent with information costs that are linear in Shannon entropy. Salant and Spenkuch (2022) propose a model of satisficing with evaluation errors that features complexity of individual alternatives. They test their model on chess data and document that complex moves are chosen less frequently, among other findings.

Our experimental measurement approach leverages insights from psychometrics. In particular, we adapt techniques proposed by Bradley and Terry (1952). Puri (2024), build-

ing on Huck and Weizsäcker (1999), Sonsino et al. (2002) and the discussion of complexity in Bernheim and Sprenger (2020), proposes lotteries are less favorably evaluated if they are more complex, and uses the number of outcomes as a measure of the complexity. Evidence for this effect is provided in Fudenberg and Puri (2023) and Goodman and Puri (2022). This evidence indirectly supports our measure (we develop lotteries in Section D.10), which is richer than the number of outcomes (and admits that as a particular case), as each outcome is weighted by its importance.

Third, our proposed theory and experimental findings speak to a more applied empirical literature on the effects of complexity. Aghion et al. (2024) analyze taxes on the self-employed in France, who can choose one of three possible complex systems; they find that people take time to switch to the most favorable regime and only learn over time about changes. You and Zhang (2009) analyze companies's annual SEC 10-K filings and show that markets react more sluggishly to more complex reports. Célérier and Vallée (2017) find that more complex financial products (of the equity protection type) have higher markups. Relatedly, Carvalho and Silverman (2023) study people's sophistication in opting out of complex financial situations. Colliard and Georg (2025) find that under more complex financial regulations, students make more mistakes in computing the risk-weighted asset. Carlin (2009) builds a model in which firms increase their market power by adopting more complex rules in retail financial markets: our framework would endogenize the noise in people's mind while considering financial products. Molavi et al. (2024) argue that constraints on the complexity of agents' models of asset prices can generate return predictability. Zwick (2021) studies how the complexity of the tax code affects corporate filing behavior.

Fourth, our paper relates to a long line of work on task complexity in the cognitive sciences. Campbell (1988) synthesizes this literature on complexity in psychology and identifies four key determinants: multiple paths to a desired end-state, multiple desired end-states, conflicting interdependence and uncertain or probabilistic linkages. Byrne and Johnson-Laird (1989) conduct experiments about about how people reason about spatial relations between objects, and find that the number of mental models required for a task (rather than the number of steps in each model) determine its difficulty of a task. Liu and Li (2012) provide a more recent review of definitions of task complexity in psychology and suggest a framework with six components (goal, output, input, process, presentation, and time). Gershman et al. (2015) build bridges with computer science to chart potential advances on the question of "computational rationality", i.e. identifying optimal decisions while taking into consideration the costs of computation. Griffiths et al. (2015) argue that bounded rationality is also found in the algorithmic level of analysis, i.e.

how we come up with approximately correct decision rules because of cognitive resource constraints even when the mind knows problem determinants.

Finally, we relate to the project of behavioral macroeconomics. There is now much evidence that inattention, or the hardness to think about the future, has important implications for macro, e.g. fiscal and monetary policy (Gabaix (2020)) and general equilibrium thinking (Angeletos and Lian (2023)). But how to endogenize attention? Doing so is important, in part to make models robust to the Lucas critique. Much has been done on how to model the benefit of attention, modelled as some gains in the traditional utility from more optimal decisions (Gabaix and Laibson (2002), Sims (2003), Gabaix (2014), Woodford (2020)). But endogenizing meaningfully the cost of attention (e.g. based on the complexity of the subparts) has proven much more challenging. Doing this is a key contribution of our paper.

C Appendix: Omitted proofs

Proof of Lemma 1 This is quite standard (see e.g. Gabaix (2014), Lemma 2). We use the fact that (5) implies

$$\begin{aligned}
\mathbb{E}\left[(y_i^s - y_i)^2\right] &= var\left((m_i - 1)\left(y_i - y_i^d\right) + \sqrt{m_i\left(1 - m_i\right)}\varepsilon_i\right) \\
&= var\left((m_i - 1)\left(y_i - y_i^d\right)\right) + var\left(\sqrt{m_i\left(1 - m_i\right)}\varepsilon_i\right) \\
&= \left(1 - m_i\right)^2 \sigma_{y_i}^2 + m_i\left(1 - m_i\right)\sigma_{y_i}^2 \\
&= \left(1 - m_i\right)\sigma_{y_i}^2.
\end{aligned}$$

We call $\sigma := \max_i \sigma_i$. By Taylor expansion, $a - a^r = \sum_i \left(y_i^s - y_i\right) + O\left(\sigma^2\right)$. As under the subjective expectations $y_i^s - y_i$ are uncorrelated across i's,

$$\begin{aligned}
\mathbb{E}\left[(a - a^r)^2\right] &= \left(\sum_i \left(y_i^s - y_i\right) + O\left(\sigma^2\right)\right)^2 = \sum_i \mathbb{E}\left[(y_i^s - y_i)^2\right] + o\left(\sigma^2\right) \\
&= \sum_i \sigma_{y_i}^2\left(1 - m_i\right) + o\left(\sigma^2\right)
\end{aligned}$$

Next, by Taylor expansion around a^r,

$$\begin{aligned}\Delta u &:= u(a,x) - u(a^r,x)\\ &= u_a(a^r,x)(a-a^r) + \frac{1}{2}u_{aa}(a^r,x)(a-a^r)^2 + O\left((a-a^r)^3\right)\\ &= \frac{1}{2}u_{aa}(a^r,x)(a-a^r)^2 + O\left(\sigma^3\right) \text{ as } u_a(a^r,x) = 0\\ &= \frac{1}{2}u_{aa}\left(a^d,x\right)(a-a^r)^2 + O\left(\sigma^3\right) \text{ as } u_{aa}\left(a^d,x\right) - u_{aa}(a^r,x) = O(\sigma)\end{aligned}$$

Taking expectations,

$$U(m) - U(\iota) = \mathbb{E}\left[u(a,x) - u(a^r,x)\right] = \sum_i V_i(1-m_i) + O\left(\sigma^3\right), \qquad V_i = \frac{1}{2}\left|u''\left(a^d\right)\right|\sigma^2_{y_i}.$$

Applying this to $m = 0$ gives $U(0) - U(\iota) = \sum_i V_i + O\left(\sigma^3\right)$, so that

$$U(m) - U(\iota) = \sum_i V_i m_i + O\left(\sigma^3\right).$$

Proof of Proposition 1 Consider the dual problem, with a cost of effort $(1-\alpha)w$:

$$\max_{L_1,\ldots L_N} \sum_i v_i m_i(L_i) - (1-\alpha)w\sum_i L_i \tag{61}$$

with $v_i = \sigma^2_{y_i}$, and $m_i(L) = \left(\frac{L_i}{c_i}\right)^{1-\alpha}$. The FOC is $\frac{v_i L_i^{-\alpha}}{c_i^{1-\alpha}} = w$, i.e.

$$L_i = w^{-\frac{1}{\alpha}} v_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}} \tag{62}$$

and the attention is

$$m_i = w^{1-\frac{1}{\alpha}} v_i^{\frac{1}{\alpha}-1} c_i^{1-\frac{1}{\alpha}} \tag{63}$$

and the cost incurred for dimension i is: $(1-\alpha)wL_i = (1-\alpha)v_i m_i$.

The aggregate attention is, with $v = \sum_i v_i$ and $s_i = \frac{v_i}{v}$,

$$M = \frac{\sum_i v_i m_i}{v} = w^{1-\frac{1}{\alpha}}\frac{1}{v}\sum_i v_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}} = w^{1-\frac{1}{\alpha}} v^{\frac{1}{\alpha}-1}\sum_i \left(\frac{v_i}{v}\right)^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}}$$

so, using the definition of C in (12):

$$M = w^{1-\frac{1}{\alpha}} v^{\frac{1}{\alpha}-1} C^{1-\frac{1}{\alpha}} \tag{64}$$

This also gives (58).

Finally, the constraint $\sum_i L_i = L$ gives w:

$$L = \sum_i L_i = w^{-\frac{1}{\alpha}} \sum_i v_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}} = w^{-\frac{1}{\alpha}} v^{\frac{1}{\alpha}} \sum_i s_i^{\frac{1}{\alpha}} c_i^{1-\frac{1}{\alpha}} = w^{-\frac{1}{\alpha}} v^{\frac{1}{\alpha}} C^{1-\frac{1}{\alpha}}$$

so

$$w = vL^{-\alpha} C^{\alpha-1} \tag{65}$$

and (64) gives $M = \left(\frac{L}{C}\right)^{1-\alpha}$.

We note

$$\frac{\partial C}{\partial s_i} = \frac{1}{1-\alpha} \left(\frac{c_i}{s_i C}\right)^{1-\frac{1}{\alpha}} \tag{66}$$

So, to maximize complexity subject to $\sum_i s_i = 1$, all the $\frac{\partial C}{\partial s_i}$ should be equal, which means that s_i should be proportional to c_i.

We also record for future reference.

$$\frac{\partial C}{\partial c_i} = \left(\frac{c_i}{s_i C}\right)^{-\frac{1}{\alpha}} \tag{67}$$

Proof of Proposition 4 We shall prove the following slightly more general result.

Lemma 7. Calling $\beta = \frac{1}{\alpha}$, the complexity aggregator $\tilde{\mathcal{C}}_\beta := C_{1/\beta} = \left(\sum_i s_i^\beta c_i^{1-\beta}\right)^{\frac{1}{1-\beta}}$ is weakly decreasing in $\beta \in [-\infty, \infty]$, and we have:

$$\tilde{\mathcal{C}}_{-\infty} = \max_i \frac{c_i}{s_i} \geq \tilde{\mathcal{C}}_0 = \sum_i c_i \geq \tilde{\mathcal{C}}_1 = \exp\left(\sum_i s_i \ln \frac{c_i}{s_i}\right) \geq \tilde{\mathcal{C}}_{+\infty} = \min_i \frac{c_i}{s_i}. \tag{68}$$

Proof. With $\gamma := 1 - \beta$, and $Y_i := \frac{c_i}{s_i}$ we have

$$\tilde{\mathcal{C}}_\beta = \Gamma(\gamma) := \left(\sum_i s_i Y_i^\gamma\right)^{\frac{1}{\gamma}} = \mathbb{E}\left[Y_i^\gamma\right]^{1/\gamma}$$

where the Y_i are drawn with probability s_i. Hence, we need to show that $\Gamma(\gamma)$ is weakly increasing in γ for γ from $-\infty$ to ∞. The proof is well-known. Suppose for instance $\gamma' > \gamma > 0$. Then, set $\delta = \frac{\gamma'}{\gamma}$ and $Z_i = Y_i^\gamma$. Then, we have $\mathbb{E}\left[Z_i^\alpha\right] \geq \mathbb{E}\left[Z_i\right]^\alpha$ by Jensen's inequality. Taking the power $\frac{1}{\gamma'}$ gives $\Gamma(\gamma') \geq \Gamma(\gamma)$. The other cases (including negative γ's) are similar. □

Proof of Proposition 7 This is by calculation. For $\alpha > 1$, we solve,

$$\max_L 1 - \left(\frac{L}{C} + \Phi\right)^{1-\alpha} - \frac{w}{p}L$$

so, assuming an interior solution:

$$\frac{\alpha - 1}{C}\left(\frac{L}{C} + \Phi\right)^{-\alpha} = \frac{w}{p}$$

and

$$L = C\left(\left(\frac{wC/p}{\alpha - 1}\right)^{-\frac{1}{\alpha}} - \Phi\right)$$

and the resulting attention is

$$Q = 1 - \left(\frac{L}{C} + \Phi\right)^{1-\alpha} = 1 - \left(\frac{wC/p}{\alpha - 1}\right)^{1-\frac{1}{\alpha}}$$

This also implies:

$$\frac{1}{\alpha - 1}\frac{w}{p}\left(L + C\Phi\right) = \left(\frac{wC}{\alpha - 1}\right)^{1-\frac{1}{\alpha}} = 1 - Q$$

so that the cost of attention paid is:

$$wL = (\alpha - 1)\, p\, (1 - Q) - wC\Phi$$

So an increase in p increases attention Q and effort L. But controlling for p, there is a negative relation between attention Q and the cost wL.

In the case $\alpha > 1$, effort is increasing in C iff $\Phi \leq (wC/p)^{-1/\alpha}\,(\alpha - 1)^{1+\frac{1}{\alpha}}$.

Proof of Proposition 10, and complements The rational answer to the problem is

$$a_t = \frac{w_t}{T - t} + \left(1 - \frac{1}{T - t}\right) h_t + \sum_{s=t}^{T-1} \frac{\hat{e}_s}{T - t} = \sum_i e_{it}$$

where $e_{0,t} = \frac{w_t}{T-t}$, and for $i \geq 1$, the e_{it} are as in the Proposition. Then, we apply the general machinery of the paper.

For Figure 6 we use the following parameters. People start working at age 25 and life ends at age 85. $A = \frac{\sigma_H^2}{\sigma_e^2} = 0.005$ is the relative variance from taste vs. income shocks and $\alpha = \frac{1}{2}$. With $t = \text{age} - 25$, and $T = 85 - 25 = 60$,

$$C(t) = \frac{\left((T-t)\left(\frac{1}{t-T}\right)^{\frac{2}{\alpha}} + A^{\frac{1}{\alpha}}\left(1-\frac{1}{t-T}\right)^{\frac{2}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}}}{\left((T-t)\left(\frac{1}{t-T}\right)^{2} + A\left(1-\frac{1}{t-T}\right)^{2}\right)^{\frac{1}{\alpha-1}}}. \tag{69}$$

Proof of Lemma 13 Before proving Lemma 13, we record how actions depend on parameters.

Lemma 8. (Discrete choice: sensitivity of actions to parameters) *Take the discrete choice problem. A change in information dx_i creates a change in the probability of choosing action k equal to*

$$da^k = \frac{1}{\beta} a^k \sum_i \left(v^k_{x_i} - \bar{v}_{x_i} \right) dx_i \tag{70}$$

where $\bar{v}_{x_i}$ is the average value of valuation changes created by the change in variable x_i:

$$\bar{v}_{x_i} = \sum_k a^{k,d} v^k_{x_i}. \tag{71}$$

In addition, $u_{aa} = -\beta Diag\left(\frac{1}{a^k}\right)$.

Proof of Lemma 8 We have

$$u_{a_k} = v^k - \beta \ln \frac{a^k}{a^{d,k}}$$

which gives

$$u_{aa} = -\beta Diag\left(\frac{1}{a^k}\right) \tag{72}$$

and

$$u_{a_k x_i} = v^k_{x_i}$$

Hence, $a_x = -u_{aa}^{-1} u_{ax}$ gives:

$$\frac{da_k}{dx_i} = \frac{1}{\beta} a^k v^k_{x_i}$$

and the compensated change is (so that $\sum_k a_k = 1$; see Gabaix (2014) section V.D)

$$\frac{d\bar{a}_k}{dx_i} = \frac{1}{\beta} a_k \left(v^k_{x_i} - \bar{v}_{x_i} \right) \tag{73}$$

□

We can now finish the proof of Proof of Lemma 13. The LQ approximation from a higher precision is:

$$dU = -\frac{1}{2}\sum_{i,j} dx_j \bar{a}'_{x_j} u_{aa} \bar{a}_{x_i} dx_i$$

As the agent imagines that the dx_i are uncorrelated across i's, we keep only the diagonal terms, so that

$$dU = -\frac{1}{2}\sum_{i} \bar{a}'_{x_i} u_{aa} \bar{a}_{x_i} (dx_i)^2$$

Using the value of $\frac{d\bar{a}_k}{dx_i}$ and u_{aa}, we get:

$$\delta U = \frac{1}{2\beta}\sum_{k} \bar{a}^k_{x_i} \left(v^k_{x_i} - \bar{v}_{x_i}\right)^2 \sigma^2_{x_i} m_i \tag{74}$$

In the case where information i only affects option k. Then $v^\ell_{x_i} = 1_{k=\ell} v^k_{x_i}$, $\bar{v}_{x_i} = a^k v^k_{x_i}$, and

$$\sum_\ell a^\ell \left(v^\ell_{x_i} - \bar{v}_{x_i}\right)^2 = \left(v^k_{x_i}\right)^2 \sum_\ell a^\ell \left(1_{k=\ell} - a^k\right)^2$$

where

$$\begin{aligned}\sum_\ell a^\ell \left(1_{k=\ell} - a^k\right)^2 &= \left[\sum_\ell a^\ell \left(a^k\right)^2\right] + a^k\left(\left(1-a^k\right)^2 - \left(a^k\right)^2\right) \\ &= \left(a^k\right)^2 + a^k\left(1-2a^k\right) = a^k\left(1-a^k\right)\end{aligned}$$

so

$$\sum_\ell a^\ell \left(v^\ell_{x_i} - \bar{v}_{x_i}\right)^2 = \left(v^k_{x_i}\right)^2 a^k\left(1-a^k\right).$$

Proof of Proposition 11 By Definition 2 of the complexity of discrete choice, the complexity is that of $u(a,x)$. Lemma 4 gives $\Delta U = \sum_i \frac{V_i^0}{2\beta} m_i$. Hence, we apply Proposition 1, with $s_i \propto \frac{V_i^0}{2\beta} \propto V_i^0$.

Proof of Proposition 12 We follow Definition 1. The "naive" answer $a^{d,0}$, the agent randomized over actions with equal probability, yields average utility $v(0) = \bar{\mu} = \frac{1}{K}\sum_{k=1}^K \mu_k$. When the anchor is a^d (obtained after a small amount of thinking), the baseline utility is $U(0^+) = \sum_k a^{d,k}\mu^k$, and by Lemma 4, the utility is

$$U(m) = \sum_i \frac{V_i^0}{2\beta} m_i + U(0^+)$$

Hence, the output-based utility is

$$\begin{aligned} 1-Q(L) &= \frac{v^r - v(L)}{v^r - v(0_-)} = \frac{\left(\sum_i \frac{V_i^0}{2\beta} 1 + U(0^+)\right) - \sum_i \frac{V_i^0}{2\beta} m_i + U(0^+)}{\sum_i \frac{V_i^0}{2\beta} 1 + U(0^+) - v(0_-)} \\ &= \frac{\sum_i \frac{V_i^0}{2\beta}(1-m_i)}{\sum_i \frac{V_i^0}{2\beta} 1 + \sum_k a^{d,k}\mu^k - \bar{\mu}} \end{aligned}$$

Proof of Proposition 14 First, we derive a^d and β, following Assumption 1. The naive initial action distribution is $a^{d,0,k} = \frac{1}{K}$. We also develop the case with uncorrelated information across goods, but different variances $\sigma_{kd}^2 = var(b_d x_{dk})$. By (44) we have $V_{i=(k,\delta)}^0\left(a^{d,0}\right) = \frac{1}{K}\left(1-\frac{1}{K}\right)\sigma_{k\delta}^2$. So, by (45),

$$\beta^0 = \left(\sum_{i=(k,d)} V_i^0\left(a^{d,0}\right)\right)^{1/2}\bar{\beta} = \left(\frac{1}{K}\left(1-\frac{1}{K}\right)\sum_{k,d}\sigma_{kd}^2\right)^{1/2}\bar{\beta}. \tag{75}$$

This gives $a^{d,k} \propto e^{\beta^0 \mu_k}$. Then, (45) gives (in this expression, the superscript d means default and the underscript δ means dimension)

$$\beta = \bar{\beta}\left(\sum_{k,\delta} a_k^d\left(1-a_k^d\right)\sigma_{k\delta}^2\right)^{1/2} \tag{76}$$

Then, $V_{k\delta}^0 = a_k^d\left(1-a_k^d\right)\sigma_{k\delta}^2$, and $s_{k\delta} = \frac{V_{k\delta}^0}{\sum_{k',\delta'} V_{k'\delta'}^0}$. The complexity is then $\mathcal{C}(s_{k\delta}, c_{k\delta})$.

Particular case where all goods have similar variances across dimensions. We next suppose that $\sigma_{k\delta} = \sigma_\delta$. Then, (45) gives

$$\beta = \bar{\beta}\sqrt{\left(\sum_k a_k^d\left(1-a_k^d\right)\right)\left(\sum_d \sigma_d^2\right)} \tag{77}$$

and

$$V_{k\delta}^0 = a_k(1-a_k)\sigma_\delta^2 \tag{78}$$

This implies $s_{kd} = \frac{V_{kd}^0}{\sum_{k',d'} V_{k'd'}^0} = s_k s_d$ where $s_k = \frac{a_k(1-a_k)}{\sum_{k'} a_{k'}\left(1-a_{\prime_k}\right)}$ is the importance share of good k, and $s_d = \frac{\sigma_d^2}{\sum_{d'}\sigma_{d'}^2}$ is the importance share of dimension d. Hence, the complexity of is $\mathcal{C}(s_{kd}, c_{kd})$ with $c_{kd} = c_d$, as it is independent of the good. Hence, the complexity C

satisfies

$$\begin{aligned} C &= \mathcal{C}\left(s_{kd}, c_{kd}\right) = \left(\sum_{k,d}\left(s_k s_d\right)^{\frac{1}{\alpha}} c_d^{1-\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}} = \left(\left(\sum_k s_k^{\frac{1}{\alpha}}\right)\left(\sum_d s_d^{\frac{1}{\alpha}} c_d^{1-\frac{1}{\alpha}}\right)\right)^{\frac{\alpha}{\alpha-1}} \\ &= \left(\sum_k s_k^{\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}}\left(\sum_d s_d^{\frac{1}{\alpha}} c_d^{1-\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}} = K^f C^f. \end{aligned}$$

Proof of Proposition 13

Proof. The optimal attention is $m_i = Q\left(\frac{L_i}{c_i}\right)$ where $L_i = \text{argmax}_{L_i} V_i Q\left(\frac{L_i}{c_i}\right) - wL_i$, i.e. $\frac{L_i}{c_i} = \text{argmax}_{L_i} \frac{V_i}{wc_i} Q\left(\frac{L_i}{c_i}\right) - \frac{L_i}{c_i}$. □

Proof of Proposition 15 The meta-objective function is (50). So,

$$p^{x_i,Meta} = \frac{1}{2}\left|V_{LL}\right|\left(L_p p\right)^2 = \frac{1}{2}\left|V_{LL}L^2\right|\left(\frac{L_p p}{L}\right)^2$$

Because the FOC is $(1-\alpha)\, pC^{\alpha-1}L^{-\alpha} = w$, we have $\frac{L_p p}{L} = \frac{1}{\alpha}$. Also,

$$\begin{aligned} -V_{LL}L^2 &= -p^x Q''(L) L^2 = p^x(\alpha-1)\alpha L^{-\alpha-1}C^{\alpha-1}L^2 = p^x(\alpha-1)\alpha\left(\frac{L}{C}\right)^{1-\alpha} \\ &= p^x(\alpha-1)\alpha\left(1-m^d\right) \end{aligned}$$

for $\alpha > 1$. So, still for $\alpha > 1$,

$$p^{x_i,Meta} = \frac{1}{2}p^x\left(\frac{\alpha-1}{\alpha}\right)\left(1-m^d\right)$$

For $\alpha < 1$, $-V_{LL}L^2 = p^x(\alpha-1)\alpha\left(-m^d\right)$ for $\alpha < 1$, and

$$p^{x_i,Meta} = \frac{1}{2}p^x\left(\frac{1-\alpha}{\alpha}\right)m^d$$

Proof of Proposition 16 Indeed,

$$\begin{aligned} C &= \mathcal{C}\left(s_i\left(\mathcal{M}^s\right), c_i\left(\mathcal{M}^c\right)\right) = \left(\sum_i \left(\frac{s_i^{\mathcal{M}^s}}{\sum_j s_j^{\mathcal{M}^s}}\right)^{\frac{1}{\alpha}} \left(c_i^{\mathcal{M}^c}\left(c^d\right)^{1-\mathcal{M}^c}\right)^{1-\frac{1}{\alpha}}\right)^{\frac{\alpha}{\alpha-1}} \\ &= \left(\sum_i s_i^{\frac{\mathcal{M}^s}{\alpha}} c_i^{\mathcal{M}^c\left(1-\frac{1}{\alpha}\right)}\right)^{\frac{\alpha}{\alpha-1}} \left(\sum_j s_j^{\mathcal{M}^s}\right)^{\frac{'1}{\alpha-1}} \left(c^d\right)^{1-\mathcal{M}^c} \end{aligned}$$

Proof of Proposition 20 The proof is as for Lemma 1, but with multidimensional action a. The main arguments are as follows. The expected utility gains are as in (6): with $y_i := a_{x_i} x_i$ is now a vector, we have

$$2\delta U = -\sum_{i,j} y_j' u_{aa} y_i \tag{79}$$

Assuming that the expected innovations are uncorrelated, we get

$$\delta U = -\frac{1}{2}\sum_{i,j} a_{x_i}' u_{aa} a_{x_i} \sigma_{x_i}^2 m_i$$

Hence, the value of information — of knowing more about i — is $V_i = -\frac{1}{2} a_{x_i}' u_{aa} a_{x_i} \sigma_{x_i}^2$.

D Appendix: Theory complements

D.1 Endogenizing micro-complexity: Complexity of a problem with two layers in the production function of thought

So far, the "atomic" complexities c_i of the components x_i were exogenous. Now, we endogenize them. To do so, we use the same technique as above, but model the cognitive hierarchy down to one level deeper. That gives us a theory of the complexity of parts as a function of the complexity of the subparts, and sub-subparts, etc., extending the predictive power of the theory to the extent that it can then be based on assumptions about the complexity of the smaller subparts (instead of the larger parts). We will leverage this in the experimental Section 4.

We still have $a^r = \sum_i a_{x_i} x_i$, but now we study what happens when a_{x_i}depends on parameters z_{ij}, as in:

$$a_{x_i} = \sum_{j \geq 0} a_{x_i, z_{ij}} z_{ij},$$

where $z_{i0} = 1$ is used for the "default" value $a_{x_i,z_{i0}} = a^d_{x_i}(0)$. Figure 2 (right panel) shows the corresponding schema. Instead of stopping the analysis at x_i (as in the left panel), we study their construction with the z_{ij} (right panel): their cognitive production builds $a_{x_i}x_i$ (bottom arrows pointing up), which in turn builds a (top arrows pointing up). Then, the cognitive production function is

$$F^a = \sum_{i,j} \left| a_{x_i,z_{ij}} x_i z_{ij} \right|^2 m_{ij}\left(L_{ij}\right)$$

where $m_{ij}\left(L_{ij}\right) = \mathcal{M}\left(\frac{L_{ij}}{c_{ij}}\right)$, and $\mathcal{M}$ is our Cobb-Douglas function (3). We obtain the same structure as in the prior one-layer case. Also, the production function for $a_{x_i}x_i$ is

$$F^{a_{x_i}x_i} = \sum_{j} \left| a_{x_i,z_{ij}} x_i z_{ij} \right|^2 m_{ij}\left(L_{ij}\right)$$

So, with $s_{ij} = \frac{\left(a_{x_i,z_{ij}} z_{ij}\right)^2}{\sum_{j'}\left(a_{x_i,z_{ij'}} z_{ij'}\right)^2}$ as the importance share of z_{ij} for x_i, and, as before,

$$c = \mathcal{C}\left(\left(s_i, c_i\right)_{i=1...N}\right),$$

we see how we get production functions from underlying parts. In the two-layer setup above, the macro-complexity of decision a is the following function of the micro-complexity c_i of the dimensions i, as was true in Proposition 1

$$c = \mathcal{C}\left(\left(s_i, c_i\right)_{i=1...N}\right). \tag{80}$$

But in addition, the complexity c_i of dimension i is endogenized as:

$$c_i = \mathcal{C}\left(\left(s_{ij}, c_{ij}\right)_{j=0...N_i}\right), \tag{81}$$

where $s_{ij} = \frac{\left(a_{x_i,z_{ij}} z_{ij}\right)^2}{\sum_{j'}\left(a_{x_i,z_{ij'}} z_{ij'}\right)^2}$ is the importance of z_{ij} for dimension i. Furthermore, the complexity of a can be also expressed as a function of the nano-components:

$$c = \mathcal{C}\left(\left(s_i s_{ij}, c_{ij}\right)_{i=1...N, j=0...N_i}\right) \tag{82}$$

Clearly, one could go deeper, and endogenize the complexity of each nano component z_{ij}, in turn.

D.2 The mind as a cognitive economy

In the following, we make a number of speculative remarks on how the model proposed in this paper may map to actual cognitive process in the human brain. More precisely, if we were consulted during the design of the brain (which, sadly, we were not), this is how we would propose to implement the "cognitive economy" of this paper. We would implement it like an economy, with not just production, but also price signals circulating in the mind.

Recall that for a, the problem is (from (9))

$$\max_L p^a Q(L) - wL, \tag{83}$$

where $p^a = \frac{1}{2}\left|u''\left(a^d\right)\right|\sigma_a^2$ is the benefit (in utils) of going from zero to full precision, and w is the shadow cost of effort.

Then, the precision problem for x_i is also

$$\max_L p^{x_i} m_i(L_i) - wL_i, \tag{84}$$

with

$$p^{x_i} := p^a \frac{a_{x_i}^2 \sigma_{x_i}^2}{\sigma_a^2} = p^a s_i. \tag{85}$$

This has the interpretation that the price p^a was sent "down the production chain" as a transformed price p^{x_i}, to indicate to the module thinking about x_i how much effort to invest in x_i.

This can be generalized to the two-layer production function. Then, we form

$$p^{z_{ij}} = p^{x_i} s_{ij} \tag{86}$$

so, the price p^a generates p^{x_i}, which in turn generates $p^{z_{ij}}$, which instructs how much to think about $p^{z_{ij}}$.

We see that prices fall as we proceed down the network. At some point, this can stop. Indeed, when $\alpha > 1$, by Lemma 5, we stop when

$$\frac{p^{z_{ij}}}{wc_{ij}} < A(\alpha) \tag{87}$$

for a value of $A(\alpha)$ specified in the Lemma.

One could contemplate that this is how the mind allocates prices throughout its archi-

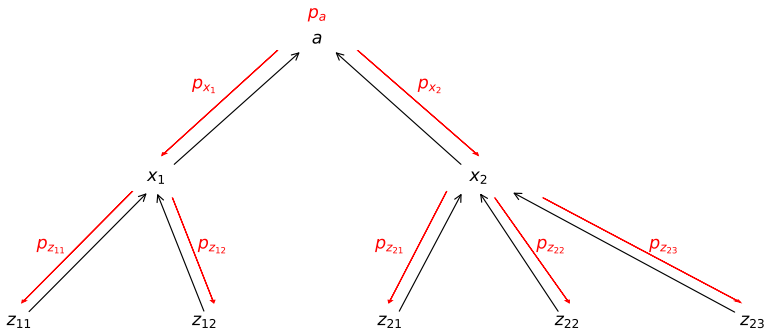


Figure D.1: Cognitive architecture with quantities and prices. Notes. Inputs contribute to the outcome, i.e. flow upwards on the tree, while prices inform about the value of different inputs, i.e. flow downwards on the three.

tecture. One could also speculate that one day, programming computers might explicitly use this kind of architecture. For many calculations, one does need only a finite, and potentially small, precision. The value of each calculation could be circulated in a program with this sort of "prices", in the sense of value.

D.3 Outside the interior region

We now generalize Proposition 1, to the case where the m_i can be equal to 1 at the optimum.

Proposition 21. *Given the assumption of Proposition 1, in the general case where we may be outside the "interior" region, so that we allow $m_i(L_i)$ to be equal to 1, we have*

$$Q\left(L,(s_i,c_i)\right)=\min_{w}\left[\sum_{i:\frac{c_i}{s_i}\leq w}s_i+\left(\frac{L-\sum_{i:\frac{c_i}{s_i}\leq w}c_i}{\mathcal{C}\left((s_i,c_i)_{i:\frac{c_i}{s_i}>w}\right)}\right)^{1-\alpha}\right] \tag{88}$$

Then, the attention $m_i=1$ for $\frac{c_i}{s_i}\leq w$, and $m_i<1$ for $\frac{c_i}{s_i}>w$.

The corresponding intuition is simple. Suppose that there are just two components, and let us say that $\frac{c_1}{s_1}<\frac{c_2}{s_2}$. For low total cognitive budget L, both components get some partial attention $m_i<1$. Then, for L above a value L_*, the attention to dimension 1 is full

($m_1 = 1$), so that all the residual effort is invested in L_2. Then, the production function is $Q(L) = s_1 + s_2 \left(\frac{L-c_1}{c_2}\right)^{1-\alpha}$.

D.4 The value of the curvature parameter α

The parameter $\alpha > 0$ enters in the formula. We submit that $\alpha = 1/2$ is a good "Laplacian" prior when $\alpha \in [0, 1]$. Fortunately, the outcomes are generally not very sensitive to its value (for instance, in the graph of Figure, the values at $s = 0$, $\frac{1}{N}$ and 1 do not depend on $\alpha > 0$). This has a good impact: the predictions are relatively robust to the value of α. On the other hand, that makes establishing α delicate. In a revision of this paper, we plan to offer a more systematic analysis of its value. In the meantime, we propose to stay with $\alpha = 1/2$.

Figure D.2 illustrates this. The complexity weakly increases in α (as per Proposition 4); for very large α, it is always 5, as long as s is neither 0 or 1. So high α tends to lower the sensitivity to shares, when they're not very extreme.

In that sense, a high value $\alpha > 1$ is useful when modelling agents who perceive fuzzily the exact values of the shares, as in the model of imperfect metacognition of Section 6. This is why, provisionally, we took $\alpha = 5$ in the numerical analysis of the experiment.

D.5 Interaction with learning

We sketch how to incorporate learning in the model. Suppose that people learned the optimal action $a^k = a^*\left(x^k\right)$ for some parameters x^k, with $k \in \{1, \ldots, K\}$. Then, given a new vector x, a sensible strategy it so find an prior situation x^k close to x, anchor in the optimal action a^k, and adjust for the small difference between x and x^k.

We formalize this idea, and enrich it in the process. Call Ω^k the utility loss given L cost if the default is taken to to be a^k, x^k:

$$V^k(L) = \max_{L_1, \ldots L_N} \frac{-1}{2} |u''\left(x^k\right)| \sum_i a_{x_i}^2 \left(x_i^k - x_i\right)^2 (1 - m_i(L_i)) \text{ s.t. } \sum_i L_i \leq L$$

Then, the agent chooses the best situation according to:

$$\min_k V^k(L)$$

And the associated complexity is low is a paradigmatic situation x^k is close to the situation x faced by the agent. We state the following simple result, whose proof is immediate.

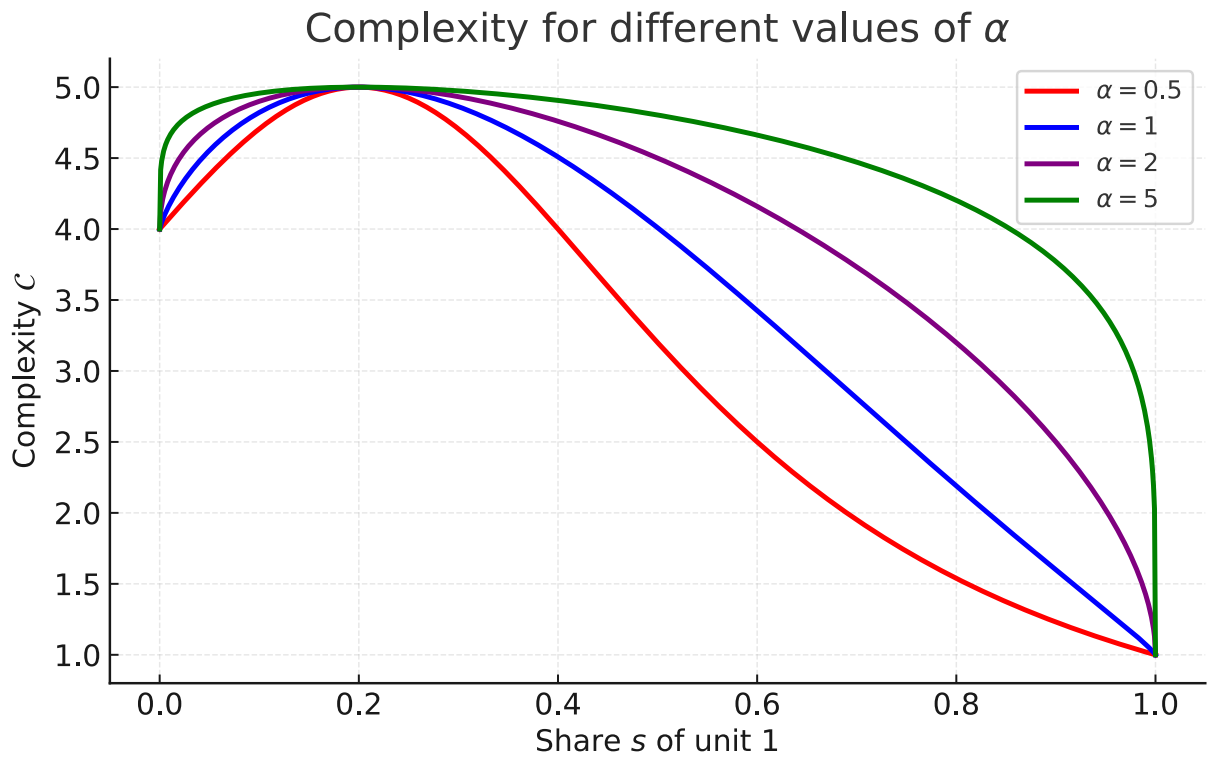


Figure D.2: Graph of Figure 3, for different values of the curvature parameter α. We see that a high α tends to lower the sensitivity to shares, when they're not very extreme.

Lemma 9. (Learning decreases complexity) *Hence, a richer set of learned paradigms $x^k$is leads to less complexity: a lower value of* $\min_k V^k(L)$.

One could imagine thinking of an "optimal" learning, e.g. an optimal spacing of paradigmatic situations x^k.

D.6 First vs. second order complexity aversion

So far, the errors affect precision via second order terms. But one could well imagine that they are first order terms. This is analogous to the corresponding issue in risk aversion. This theme is developed in great detail in Gabaix (2025). We show a bridge between the two approaches here.

Indeed, if the errors $e_i := y_i^s - y_i$ are uncorrelated, $a = \sum_i y_i$, $a^s = \sum_i y_i^s$, then the error in a is $a^s - a = \sum_i e_i$. So, if the errors are all uncorrelated, the quadratic error is

$$\mathbb{E}\left[(a^s - a)^2\right] = \sum_i \sigma_{e_i}^2 \tag{89}$$

where $\sigma_{e_i}^2$ is the variance of error e_i in dimension i . However, suppose that the errors have the maximal correlation of 1. Then

$$\mathbb{E}\left[(a^s - a)^2\right] = \left(\sum_i \sigma_{e_i}\right)^{1/2} \tag{90}$$

So each σ_{e_i} enters with a power 1 rather than 2. Then, the utility loss is in (6) is:

$$U = U^r - \frac{1}{2}\left|u''\left(a^d\right)\right|\left(\sum_i \sigma_{y_i}\left(1 - m_i\right)^{1/2}\right)^2$$

So, defining $\tilde{m}_i = 1 - (1 - m_i)^{1/2}$, and the shares (8) become:

$$\tilde{s}_i = \frac{\sigma_{y_i}}{\sum_j \sigma_{y_j}} \tag{91}$$

the problem (9) becomes:

$$\tilde{Q}(L) := \max_{L_1,...L_N} \sum_i \tilde{s}_i \tilde{m}_i(L_i) \text{ s.t. } \sum_i L_i \leq L \tag{92}$$

Moreover, we can imagine that the production function in $\tilde{m}$ space is the one we had used for m in (10).Then, the problem is formally the same, except that the share $\tilde{s}_i$ depend on

standard deviations rather than their squares—first order rather than second order. Mechanically, first order subjective losses from complexity will lead to a even more powerful force to avoid complexity.

D.7 Cognitive-risk adjusted certainty equivalent

When performing $\max u(a, x)$, the agent might use a certainty equivalent, which we posit to be as follows:

$$x_i^{s,C} = x_i^s - \lambda \operatorname{sign}\left(u_{ax_i} x_i \left(a - a^d\right)\right)(1 - m_i)\,\sigma_{x_i} \tag{93}$$

where $\lambda \geq 0$ is akin to a loss aversion coefficient, with as a rough guesstimate $\lambda \simeq 0.2$. Then, the agent takes the action:

$$\max_a u\left(a, x^{s,C}\right) \tag{94}$$

For instance, if $u(a, x) = -\left(a - a^r(x)\right)^2$ with $a^r(x) = x_1 - x_2$ with $x_1 > 0$ and $x_2 > 0$, and $a^d = 0$. Then the action using the certainty equivalent is:

$$a = x_1 - \lambda(1 - m_1)\,\sigma_{x_1} - x_2 - \lambda(1 - m_2)\,\sigma_{x_2} \tag{95}$$

Here is the interpretation of (93). Suppose that $u_{ax_i} x_i > 0$, so that x_i is a "good news" for a (like $x_1 > 0$ in our example), i.e. increases the optimal a compared to the case where $x_i = 0$. Then, the perceived certainty equivalent is a little dampened, but a quantity proportional to the uncertainty about x_i. Likewise, if $u_{ax_i} x_i < 0$, x_i is "bad news" for a (like x_2 in our example) then the agent, mindful of cognitive noise, is even more pessimistic, or distrustful of the situation. This capture the fact that willingness to pay is generally lower than willingness to accept.

D.8 Complexity of forecasting: More general version

We want to capture that some forecasts are "more complex" than others. For instance, if $x_t = \rho x_{t-1} + \varepsilon_t$, then $\mathbb{E}_t[x_{t+h}] = \rho^h x_t$ is reasonably simple. But suppose that the process is multidimensional,

$$x_t = A x_{t-1} + \varepsilon_t$$

we want to calculate for the value of a variable $y_t = b' x_t$ in h periods, $f_{t,h} \coloneqq \mathbb{E}_t[y_{t+h}]$. This is more complex, as we pretty much need to forecast every component. Formally, $f_{t,h} = \mathbb{E}_t[b' x_{t+h}] = b' A^h x_t$ is rather complex. We formalize and quantify that.

We suppose a default matrix $A^d = D$, typically $D = \rho I_N$, for some ρ. Then, the default prediction is $f^d_{t,h} = b' D^h x_t$, and the discrepancy is

$$\delta(A) := b' \left(A^h - D^h \right) x_t$$

The marginal impact of an difference $A_{ij} - D_{ij}$ between the true and default transition matrix is:

$$v_{ij} = \mathbb{E}\left[\delta^2_{A_{ij}} \left(A_{ij} - D_{ij} \right)^2 \right]$$

and it's share is: $s_{ij} = \frac{v_{ij}}{\sum_{i'j'} v_{i'j'}}$. Then, the complexity is just $\mathcal{C}(s)$.

For instance, suppose that $A = Diag(\rho_i)$ and $D = \bar{\rho} I$, for some $\bar{\rho}$, with ε_{it} uncorrelated across i and t. Then,

$$v_{ii} = \left(b_i h \rho_i^{h-1} (\rho_i - \bar{\rho}) \right)^2 var(x_i) \tag{96}$$

while $v_{ij} = 0$ if $i \neq j$. Hence, the complexity is higher there the variable y is "effectively" composed of more variables x_i with different autocorrelations ρ_i.

D.9 Ex post simplification of decisions, for instance to communicate an answer

We often wish to have a simple decision rule — for instance, one that does only use the simplest of numbers, 0 and 1. For instance, if the correct answer is $a^r = 1.1x_1 + 0.2x_2$, we might replace it by $a^s = x_1$. Similarly, if the answer is $a^r = 3.12$, we might just say $a^s = 3$. Saying otherwise would be pedantic — uselessly burdening the recipient of the information with largely superfluous extra information. We formalize this basic fact here.

Suppose that the answer is $a^r = \sum_i b_i x_i$.[32] The metric of complexity of an answer $a = \sum_i b^s_i x_i$ is $C(b^s) = \sum_i C\left(b^s_i\right)$ with

$$C\left(b^s_i\right) = c_0 1_{b^s_i \notin \{0,1\}} + c_1 1_{b^s_i = 1} \tag{97}$$

which means that there a penalty of c_1 for a b^s_i equal to 1, and a penalty of $c_0 \geq c_1$ for a b^s_i different from the "simple" values of 0 and 1. We could have other "simple" numbers, e.g. $\frac{1}{2}$, but we leave the generalization there.

[32] For instance, if $a^r = 3.12$, we express this as $a^r = x_1 + 0.1x_2 + 0.01x_3$, with the x_i's equal to 3, 2 and 1.

Then, the problem of the "simplifier" is

$$\min_{\left(b_i^s\right)_{i=1...N}} \frac{\left(b_i - b_i^s\right)^2 \|x_i\|^2}{\|a\|^2} + C\left(b_i^s\right) \tag{98}$$

The first term is the error in the recommended answer (as a fraction of the modulus of the correct answer); the second term is the cost of using this.

D.10 Application: Complexity of choosing between lotteries

What's the complexity of evaluating a lottery? Let's say that a lottery has Ω events, $\omega = 1 \ldots \Omega$, and event ω has probability p_ω and payoff r_ω. The expected value is $\sum_\omega p_\omega r_\omega$. To represent the imprecise perception, we draw on a rich psychometric literature (e.g., Dehaene, 2003; Zhang and Maloney, 2012; Gabaix, 2019; Khaw et al., 2021; Enke and Graeber, 2023) documenting that (i) probabilities are processed in log odds (i.e. agents use noisy versions of $\tilde{p}_\omega = \ln \frac{p_\omega}{1-p_\omega}$) and (ii) absolute values of payoffs are processed as logs (i.e. agents use noisy versions of $\tilde{r}_\omega = \ln |r_\omega|$), while preserving the sign $S_\omega = sign\left(r_\omega\right)$. Accordingly, we posit that the agents processes values using the function:

$$V\left(\left(\tilde{p}_\omega, \tilde{r}_\omega, S_\omega\right)_{\omega \in \Omega}\right) = \frac{\sum_\omega P\left(\tilde{p}_\omega\right) R\left(\tilde{r}_\omega, S_\omega\right)}{\sum_\omega P\left(\tilde{p}_\omega\right)} \tag{99}$$

with the inverting functions (mapping the log odds $\tilde{p}$ into probabilities P) $P\left(\tilde{p}\right) = \frac{1}{1+e^{-\tilde{p}}}$ and $R\left(\tilde{r}, S\right) = e^{\tilde{r}} \Lambda\left(S\right)$. We also allow for loss aversion, where $\Lambda\left(S\right) = 1$ if $S = 1$, and $\Lambda\left(S\right) = -\lambda$ if $S = -1$, with $\lambda \geq 1$ the loss aversion coefficient. If $\lambda = 1$, the agent is simply risk neutral.

The complexity of the lottery is simply the complexity of calculating this V in (99). We assume that the cost are the same for $\tilde{p}_\omega$ and $\tilde{r}_\omega$, except when the corresponding value of $\tilde{p}_\omega$ or $\tilde{r}_\omega$ is infinite – then the cost is 0 as the underlying value is 0 or 1.

When $\alpha = \infty$, the complexity of a gamble is just twice the number of non-zero outcomes, as in Puri (2024).[33] However, when $\alpha < \infty$, our complexity measure is more nuanced than simply the number of outcomes, as it is importance-weighted. For instance, consider the lottery with payoffs: $\left(0.8, \$3\right), \left(\varepsilon, \$2\right), \left(0.2 - \varepsilon, \$1\right)$. As the small probability ε goes to 0, this converges to the lottery $\left(0.8, \$3\right), \left(0.2, \$1\right)$. In our model, the complexity of the first lottery does indeed converge to that of the second. But this would not be the case in a "number of non-zero outcomes" model.

[33] For simplicity, we only discuss lotteries with non-zero dollar outcomes.

Several lotteries If there are K lotteries, we proceed similarly. Each lottery k is characterized by p_ω^k, r_ω^k, and $v^k = V\left(\left(\tilde{p}_\omega^k, \tilde{r}_\omega^k\right)_{\omega\in\Omega^k}\right)$. We use the complexity of choosing between K lotteries from Proposition 11.

E Appendix: Complexity of basic arithmetic operations

It is clear that computing $3+2$ is easier than $3.989+6.2$. For many issues in economics, this can be forgotten. However, for others (e.g. when thinking about the "cognitive noise" in laboratory evidence), this may be important. Accordingly, we next explore a modeling of the complexity of basic arithmetic operations. This section is still exploratory, and likely to change.

E.1 Complexity of a number

Suppose that we need to calculate xy. What is the complexity of that? We stylize the mental operations in the following way.

For instance, 1 or 10 is simpler than 3 or 0.37. People first choose a simplification of x, which is the nearest number number of the type 10^n or $\frac{1}{2}\times 10^b$, perhaps sign s that can be $+1$ or -1:

$$\min_{a\in\mathbb{Z},b\in\left\{0,1,\frac{1}{2}\right\},s\in\{-1,1\}} |x - sa\times 10^b|\pi(a)$$

where $\pi(0)=\pi(1)=1$ and $\pi\left(\frac{1}{2}\right)=1.5$, is a penalty capturing that $\frac{1}{2}$ is a bit more complex than 0 and 1. We also call $x' = sa\times 10^b$ the resulting value.

Then, we calculate the relative error of replacing x by x':

$$r\left(x,x'\right) = \frac{2|x-x'|}{|x|+|x'|}\in[0,1]$$

Then, the difficulty of x is

$$\delta(x) = r\left(x,x'\right)^\gamma, \qquad \gamma = \frac{1}{2} \tag{100}$$

This captures that difficulty is higher when there is an error, and the function is concave. The value of $\gamma=\frac{1}{2}$ is just a simple value yielding a concave function.

We could refine all this, e.g. to see that 2 is a simpler number than 3: that would be refining the function $\pi(a)$. We defer that to future research.

E.2 Complexity of addition

Suppose that we want to perform $\sum_i x_i$. We posit that the first number is "free". So, the complexity of that is

$$\mathcal{C}(s_i, c_i) \tag{101}$$

with $c_i = c_+$ for all numbers, except the largest x_i, in absolute value and $s_i = \frac{x_i^2}{\sum_j x_j^2}$. Here c_+ is the complexity of an addition.

As an enrichment, there can be a penalty for handling negative numbers, so that $c_i = c_{+,>0}$ for $x_i > 0$, $c_i = c_{+,<0}$ for $x_i < 0$, and $c_i = 0$ for $x_i = 0$:

$$c_i = c_+ \left(1 + \lambda 1_{x_i<0}\right) \tag{102}$$

As another enrichment, we can use the penalty for the complexity of the number x_i. So, we

$$c_i = c_+ \delta(x_i) \tag{103}$$

so that simpler number are easier to add.

E.3 Complexity of subtraction

It is the complexity of addition, with the penalty for negative numbers (102).

E.4 Complexity of multiplication

Suppose that we want to perform px. We use the "difficulty" of p and x, $\delta(x)$,$\delta(p)$, given in given in (100). We observe that the difficulty is 0 if p or x is equal to 0 or 1. So, we say

$$c^{p\times x} = c_\times \delta(x)\, \delta(p) \tag{104}$$

with for instance $c_\times = 5c_+$ is the complexity of a multiplication.

There is a connection with Khaw et al. (2021). They show convincingly that for "complex" products, people perform a noisy operation, so instead of px returns something like (see also Weber, Dehaene, Zhang and Maloney (2012); Gabaix (2014, 2019); Enke and Graeber (2023))

$$p^s x^s = (px)^m \left(p^d x^d\right)^{1-m} e^{\sqrt{m(1-m)}\tilde{u} - \frac{1}{2}m(1-m)\sigma_u^2} \tag{105}$$

i.e. there is shrinkage $m < 1$, anchoring on a default $p^d x^d$, and noise $\tilde{u}$. A refinement is the probability is perceived in log odds space; then, defining $L(p) = \ln \frac{p}{1-p}$ the log odds,

the perceived log odds are $mL(p) + (1-m)\, L\left(p^d\right)$, and after inversion,

$$p^s = L^{-1}\left(mL(p) + (1-m)\, L\left(p^d\right) + \sqrt{m(1-m)}\tilde{\varepsilon}\right)$$

and

$$p^s x^s = p^s x^m \left(x^d\right)^{1-m} e^{\sqrt{m(1-m)}\tilde{u} - \frac{1}{2}m(1-m)\sigma_u^2}$$

What we bring here is the endogeneization of the shrinkage and noise, via $m = Q\left(\frac{L}{c^{p\times x}}\right)$, where $c^{p\times x}$ is the cost of computing that product. In particular, when either p or x is has a low difficulty $\delta(x)$, the noise is 0: plainly, when asked to perform 1×3, people answer 3 with no difficulty.

E.5 Complexity of the certainty equivalent in a gamble

When forming the certainty equivalent of (p, X), things are easy when $p = 0$ or 1. They're hard where p is different from those values. In part, because we need to simulate. So, we say that the complexity is

$$c^{(p,X)} = c_* \delta(p)\,(1 + \delta(x)) \tag{106}$$

where $\delta(p)$ is the complexity of the probability number, and c_* is a constant to be calibrate.

E.6 Complexity of a composite sum, $\sum_i p_i x_i$

Let us call $v_i = p_i x_i$. We have c^{v_i} given above in (106). We set

$$c_i = c^{v_i} + c_+\left(1 + \lambda 1_{x_i<0}\right)$$

which is the complexity of a multiplication, then an addition (enriched by the penalty λ for negative numbers); except for the largest $p_i x_i$, which has a complexity of simply $c_i = c^{v_i}$.

Then, we use $\mathcal{C}(s_i, c_i)$ for all this, with $s_i = \frac{(p_i x_i)^2}{\sum_j \left(p_j x_j\right)^2}$

E.7 Complexity of choosing between two gambles.

Suppose that we need to evaluate a gamble $p_1 x_1$ vs a certainty equivalent. We use the complexity of discrete choice in Section 5. We get

$$q^{\text{binary}}\left(\ln\frac{p_A X_A}{p_B X_B}, \sigma_1 Q^p \left(\frac{L}{C}\right)^{1/2}, \sigma_1\right)$$

F Appendix: Complements to the experimental part

F.1 Survey design

Figure F.3 shows the decision screen respondents face in the survey. Figure F.4 shows the complexity comparison screen from the survey using the "Your subjective complexity" elicitation. Figure F.5 shows the complexity comparison screen using the "Previous respondents' subjective complexity" elicitation.

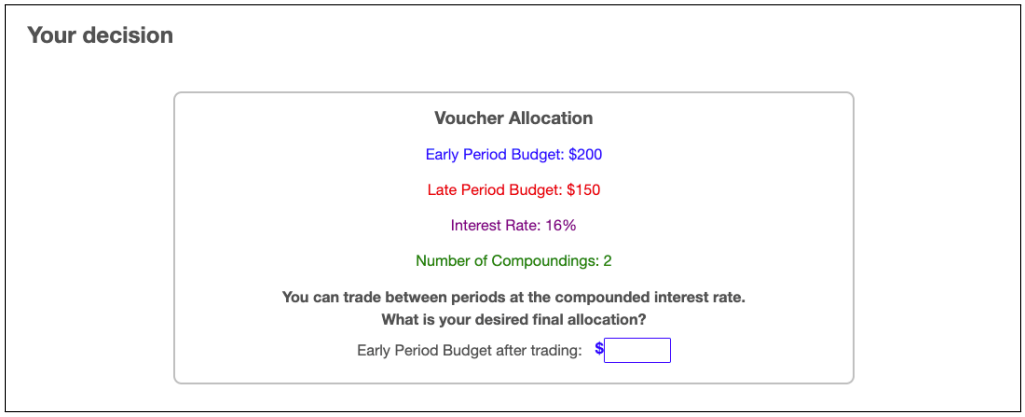


Figure F.3: Decision Screen in the Intertemporal Consumption Task

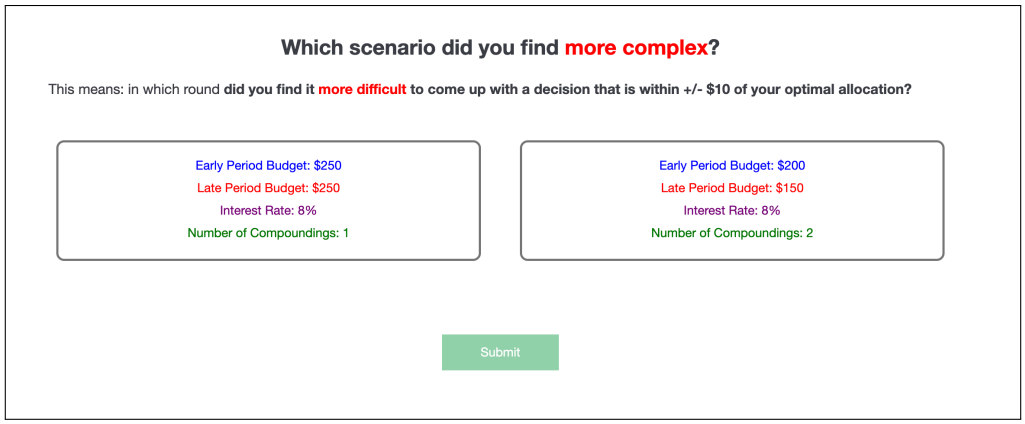


Figure F.4: Complexity comparison screen in the "Your subjective complexity" elicitation

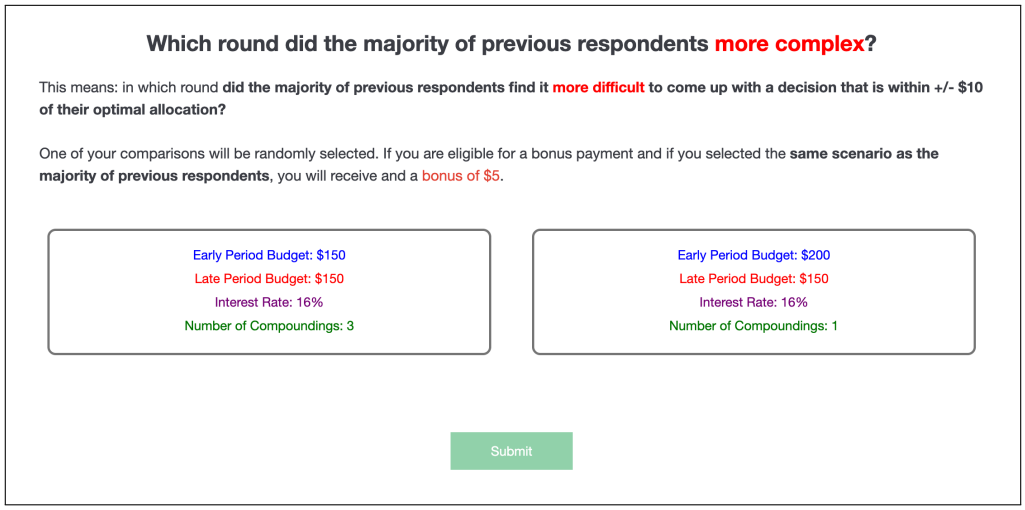


Figure F.5: Complexity comparison screen in the "Previous respondents' subjective complexity" elicitation

F.2 Calibration of utility

As described in the main text, we find no discounting in our settings. To calibrate the utility parameter ψ, we asked respondents to choose quasi-certainty equivalents. After completing all tasks, respondents are shown a bundle of *Early Period* and *Late Period* budgets (e_1, e_2) and asked which symmetric bundle (x, x) would make them indifferent between the two. To make subjects think about the bundles more concretely, this is elicited using a price list mechanism, where respondents have to indicate their switching point in a list of increasing symmetric bundles, as shown in Figure F.6. They are asked to repeat this choice for 6 different bundles, with the random parameters $e_1, e_2 \in \{150, 200, 250\}$.

To calibrate an average ψ, we then minimize the loss $\mathcal{L}(\psi) = \sum_{i,b} \left(\log \frac{x_{i,b}}{x_b^r(\psi)} \right)^2$, where $x_{i,b}$ is the observed choice for subject i and bundle $b = (e_1^b, e_2^b)$ while $x_b^r(\psi)$ is the rational

choice for bundle b given ψ. We obtain $\psi = 0.85$. Using a quadratic loss functions yields a very similar estimate of $\psi = 0.89$. To calibrate a subject-level ψ_i, we perform the same loss-minimization exercise on each respondent's choices The median ψ_i is 0.89. To limit the impact of outliers, given the low number of observations per subject, we enforce the reasonable parameter bounds $\psi \in [0.01, 10]$.



Figure F.6: Calibration of utility via price-list choice between bundles

F.3 Computation of model complexity

Call e_1, e_2 the *Early Period* and *Late Period* budgets, r the interest rate and T the number of compoundings. Denote $R = (1+r)^T - 1$ the effective interest rate between periods. The rational *Early Period* consumption is, as in (32),

$$c_1^r = \frac{e_1 + \frac{e_2}{(1+R)}}{1 + (1+R)^{\psi-1}} \tag{107}$$

We take the default action to be allocating the *Early Period* budget, i.e. $c_1^d = e_1$. Taking the Taylor expansion around $R = 0$ and $e_1 = e_2$, this becomes, up to $o\left(\|(R, e_2 - e_1)\|^2\right)$:

$$c_1^r = e_1 + \frac{1}{2}(e_2 - e_1) - \frac{\psi}{2} e_1 R = c_1^d + y_1 + y_2,$$

where the inputs of the problem are the two terms:

$$y_1 = \frac{1}{2}(e_2 - e_1), \qquad y_2 = -\frac{\psi}{2} e_1 R \tag{108}$$

We assign them "micro" complexities $c_1 = 1$ and $c_2 = 3$ as 1 and 3 values are respectively needed to compute them. This is justified in Section 3.6. Their shares are given by $s_i = \frac{y_i^2}{y_1^2 + y_2^2}$. We then use the usual complexity aggregator

$$C = \left(s_1^{\frac{1}{\alpha}} c_1^{1-\frac{1}{\alpha}} + s_2^{\frac{1}{\alpha}} c_2^{1-\frac{1}{\alpha}} \right)^{\frac{\alpha}{\alpha-1}}, \tag{109}$$

with $\alpha = 5$, a value that seems natural as it induces sparsity. For example, this yields the following complexities:

$$C(200, 200, 0\%, 1) = 0.00 \qquad C(200, 250, 8\%, 1) = 2.98$$
$$C(150, 200, 0\%, 1) = 1.00 \qquad C(200, 250, 8\%, 3) = 3.83$$

F.4 Validating complexity elicitations

We find that asking respondents about their subjective feeling of complexity produces credible results. To validate this unincentivized measure, we compare it to an incentivized version. One half of subjects were asked to predict which parameter combination respondents in a previous survey found more complex,. They were also informed that after the survey, some participants would be selected for a bonus, that one of their comparisons would be selected randomly, and that if it aligned with the majority of previous respondents they would receive a bonus of \$5. See Figure F.5 for the precise formulation.

Column 1 of Table F.1 shows a regression of task-level complexity using the "Previous respondents" elicitation on task-level complexity using the "You" elicitation. We only include tasks that have been compared by at least three different respondents in each elicitation mode and apply the MLE separately on each set of comparisons. The regression shows a significantly positive coefficient quite close to 1, as well as a high R^2: the former reflects attenuation due to measurement error in the dependent variable and the latter reflects measurement error in both. We conclude that there is no reason to believe the two differ systematically.

Column 2 shows a regression of task-level complexity using the "Previous respondents" elicitation on model complexity. We only include tasks that have been compared by at least three different respondents in the "Previous respondents" elicitation. We also

obtain our main finding, namely a highly significant coefficient and high R^2, in this incentivized subset. Overall, we conclude that the incentivized elicitation of complexity yields estimates that are not substantially different from the unincentivized one. Our methodological recommendation to researchers is therefore to use the simpler, unincentivized elicitation. In all analyses except Columns 1 and 2 of Table F.1, we pool observations across both modes.

Columns 3 compares our headline measure of complexity with an alternative measure using sliders. After playing all rounds, subjects are shown slider where each task has a numeric complexity between 1 and 100, set by default at a location corresponding to its complexity ranking in the main round. Respondents are the given the opportunity to adjust and refine their complexity judgements, and we stress that they are allowed to reverse orders. We find that respondents rarely reverse their judgements, and that MLE complexity and mean slider complexity are highly correlated. Finally, Column 4 compares subjective complexity with the average certainty respondents have about their answer, in line with Enke and Graeber (2023): the regression coefficient is negative and significant, as predicted by Proposition 6.

Table F.1: Validating Complexity Elicitations

	Dependent variable: "Previous" Complexity			
	"Previous"-subjective complexity		Subjective complexity	
"You"-subjective complexity	0.837***			
	(0.076)			
Model complexity		1.380***		
		(0.104)		
Slider complexity			0.101***	
			(0.005)	
Certainty				-0.109***
				(0.014)
Constant	1.805***	2.768***	1.063***	12.963***
	(0.432)	(0.260)	(0.270)	(0.921)
Observations	71	76	81	81
R^2	0.574	0.683	0.860	0.327

Notes. Column 1 regresses task-level complexity using the "Previous respondents" elicitation on task-level complexity using the "You" elicitation, keeping only tasks with three different respondents in each elicitation mode, and applying the MLE separately on each set of comparisons. Column 2 regresses task-level complexity using the "Previous respondents" elicitation on model complexity, keeping only tasks that have been compared by at least three different respondents in the "Previous respondents" elicitation. Columns 3 regresses pooled subjective complexity with the mean complexity slider value. Columns 4 regresses subjective complexity with the average certainty respondents have about their answer. Columns 3 & 4 only keep tasks that have been compared by at least three respondents in either elicitation.