A robust measure of complexity

Egor Bronnikov* Elias Tsakas†

Maastricht University

January, 2025

Preliminary draft

Latest version of the paper.

Abstract

We introduce a robust belief-based measure of complexity. The idea is that task A is deemed more complex than task B if the probability of solving A correctly is smaller than the probability of solving B correctly regardless of the reward. The corresponding complexity order over the set of tasks is incomplete, being represented by a vector-valued function over the two-dimensional space of difficulty and ex ante uncertainty. Then, we aggregate the individual measures in a group of agents to obtain an objective measure of complexity. Whenever the group is sufficiently large, the resulting objective complexity order is complete and ranks the tasks lexicographically, comparing them first with respect to difficulty and then with respect to ex ante uncertainty. The contribution of these results is twofold: on the one hand, we identify ex ante uncertainty as a novel dimension of complexity; on the other hand, we provide microeconomic foundations for belief-based measures of complexity.

Keywords: complexity, measure, difficulty, ex ante uncertainty, incomplete relation, effort.
JEL codes: D83, D90.

1. Introduction

Complexity is a fundamental concept across numerous scientific domains, including computer science, cognitive sciences, neuroscience, etc. More recently, its importance has also been

*Homepage: egorbronnikov.github.io; E-mail: egor.bronnikov@maastrichtuniversity.nl
†Homepage: www.elias-tsakas.com; E-mail: e.tsakas@maastrichtuniversity.nl

recognized by (behavioral) economists, as a key determinant of decisions in many settings, with the potential to explain mistakes that people systematically make in such decisions (Banovetz and Oprea, 2023; Enke et al., 2024a; Oprea, 2024b).

At the same time, it is also the case that complexity has been traditionally used in a casual way, without consensus on what exactly it means (Gabaix and Graeber, 2024). This lack of a widely accepted precise definition can possibly explain why scholars have also focused on measurements of complexity, which can in turn serve as proxies for different types of complexity (Oprea, 2024a).[1] Common ways to measure complexity include direct metrics (Oprea, 2020), behavioral metrics (Banovetz and Oprea, 2023), and most importantly for this paper belief-based metrics (Enke and Graeber, 2023; Enke et al., 2024a,b).

The key idea within this last stream of literature is to use an agent's belief about their own accuracy as a proxy for complexity. This is also consistent with theoretical results which show that people are more accurate when solving simpler tasks (Goncalves, 2024). However, as appealing as this approach is, there is a serious caveat. Namely, expected accuracy depends both on complexity as well as on the effort to handle the task. And since effort depends non-linearly on the underlying reward (for correctly solving the task), it often happens that expected accuracy is sensitive to the reward. And this naturally gives rise to the question: if different rewards yield different accuracy orders over the set of tasks, which one should we use as a measure of complexity?

In this paper, we take a robust approach, proposing the following way to rank tasks: we say that task A is deemed more complex than task B, whenever the chances of correctly solving A are smaller than the chances of correctly solving B, *for every reward*. Obviously, this is a rather conservative criterion, as it imposes a strong dominance condition. But at the same time, whenever satisfied, it provides quite convincing evidence that the tasks are indeed ranked in this way, thus indirectly providing an arguably necessary condition for every reasonable definition of complexity.

The first main result of the paper provides a full characterization of the complexity order that the aforementioned criterion induces (Theorem 1). Not surprisingly, this complexity order is incomplete. But they key insight is that it depends on two distinct parameters, viz., the difficulty of the task and the ex ante uncertainty. More specifically, higher difficulty is only a necessary condition for higher complexity. In order to also become sufficient, the agent cannot be ex ante much more uncertain about the state realization. For example, in order for a student to find exam A more complex than exam B, it must be both the case that A is more difficult than B, and moreover that the student is not much better prepared about B than about A.

[1] Measurements and formal definitions of complexity are closely linked with each other, in an analogous way belief elicitation mechanisms (Brier, 1950; Savage, 1971) are linked with definitions of subjective probability (Savage, 1954; Anscombe and Aumann, 1963).

The fact that complexity has these two distinct dimensions allows us to establish a link with the decision-theoretic literature on incomplete preferences. In particular, we show that our complexity order is represented by a vector (utility) function, like in Ok (2002).

Then, we aggregate individual complexity measure of a group of agents to obtain an objective measure of complexity. Accordingly, we will say that task A is objectively more complex than task B if the two tasks are comparable by at least one agent in this group, and moreover everyone who ranks them deems A more complex than B. The interesting fact is that the larger the group becomes, the more complete the objective complexity order will be. Then, our second main result shows that, once the group becomes sufficiently large, the objective complexity order becomes lexicographic, in that it first compares tasks with respect to difficulty and then with respect to ex ante uncertainty (Theorem 2).

Overall, the contribution of the paper is twofold. First, we identify ex ante uncertainty as a novel dimension of complexity. This dimension is intuitively appealing: as we have already mentioned above, an exam is deemed complex by a student, mainly because it is difficult, but also because the student is not well prepared. This novel channel of complexity has been previously overlooked, as complexity is typically associated only with difficulty (Oprea, 2024a, and references therein). At the same time, our results are not entirely at odds with the existing literature in the sense that, although difficulty is not the only dimension of complexity, it is still the primary one. And crucially, this is not something we assume, but rather a property that we naturally follows from our measure of complexity.

The second main contribution of the paper is to provide foundations for belief-based measures of complexity. Such measures have been recently surging (Enke and Graeber, 2023; Enke et al., 2024a,b), as they provide a portable and simple way to elicit individual perceptions of complexity. Unfortunately, solid microeconomic foundations have been missing. Therefore, by filling this gap, we make this type of belief-based measures not only empirically appealing, but also theoretically sound, and a fortiori more likely to be widely used in practice.

The literature on complexity is vast, and as such we are de facto forced to make a selection of what in our view is the most relevant subset.

Early work focused primarily on the role of strategy complexity within game theory (Rubinstein, 1986; Abreu and Rubinstein, 1988). More recently, the focus has shifted towards explaining mistakes and irrationalities, e.g., Oprea (2024b) study the effect of complexity on risk preferences, and Enke et al. (2024a) the respective effect on time preferences. There is also interest in formalizing definitions of complexity, e.g., Gabaix and Graeber (2024) build a general model of production within a cognitive economy in order to operationalize complexity, whereas Oprea (2024a) borrows insights from computer science to introduce a framework within which complexity reflects the cost for handling a task. Others, define it as the signal-to-noise ratio (Goncalves, 2024), similarly what is often done in psychometrics. The common denominator throughout most of this literature is that definitions of complexity are typically input-based, i.e., they somehow reflect the underlying difficulty to process and handle a task.

What is closer to our work is the literature on measuring complexity. As Oprea (2024a) elegantly points out, this literature can be classified into three large streams, depending the measurement tool. Within the first stream, we encounter measurement with direct metrics, such as willingness to pay in order to avoid dealing with a certain task (Oprea, 2020), response times (Goncalves, 2024), and biometrics (van der Wel and van Steenbergen, 2018). The second stream leverages behavioral metrics, such as procedural measurements (Banovetz and Oprea, 2023), and choice inconsistencies (Woodford, 2020).

Finally, the third stream, within which our paper belongs, uses belief-based metrics. These include subjective rankings, like for instance in Gabaix and Graeber (2024) where subjects are simply asked to rank tasks with respect to complexity, and beliefs about optimality of the subject's own accuracy (Enke and Graeber, 2023; Enke et al., 2024a,b), like in our paper. Of course, in all aforementioned papers, rewards are fixed. The effects of varying rewards are discussed in (Alaoui and Penta, 2022).

This entire literature is part of a surging field of Cognitive Economics (Caplin, 2025; Enke, 2024), which also incorporates topics such as rational inattention, cognitive uncertainty, etc.

The paper is structured as follows: In Section 2 we introduce our measure of complexity and prove our main characterization result. In Section 3 we introduce an objective measure of complexity, by aggregating the individual measures across many individuals. In Section 4 we study the relationship between our complexity measure and the induced effort. In Section 5 we discuss assumptions that we have imposed throughout the paper, as well as possible extensions.

2. A belief-based measure of complexity

An agent's task is to guess the realization of a binary state space $\Theta = \{\theta_0, \theta_1\}$. Let $Y = \{0, 1\}$ denote the possible results of this guessing task, i.e., 0 denotes a wrong guess and 1 denotes a correct guess. Let $X := [0, \infty)$ be a convex set of material rewards for guessing correctly. The agent's preferences over the set of acts $(X \times Y)^\Theta$ admit a SEU representation with Bernoulli utility function $u : X \times Y \to \mathbb{R}$, which is often abbreviated by $u_0(x) := u(x, 0)$ and $u_1(x) := u(x, 1)$. Let u_1 be continuously increasing and unbounded in X. We assume that

$$u_1(x) - u_0(x) \geq 0.$$

That is, the agent may care (viz., > 0) or may not care (viz., $= 0$) intrinsically about being correct, but she will never intrinsically want to fail. Other than this, do not impose any condition on the difference $u_1(x) - u_0(x)$ across $x \in X$. Neither do we assume separability of u

with respect to X and Y. Finally, without loss of generality, the utility function is normalized so that $u(0,0) = 0$.

Going back to the guessing task, for some fixed extrinsic reward $x \in X$, the guess $r \in \Theta$ is an act which yields utility $u_1(x)$ at state $\theta = r$ and utility 0 at $\theta \neq r$. Thus, the agent's indirect expected utility, as a function of the probability $q \in [0,1]$ that she attaches to θ_1, is given by

$$g(q) := u_1(x) \max\{q, 1-q\}, \tag{1}$$

which is obviously proportional to the probability she attaches to her best guess being correct.

Suppose that the agent holds a prior belief which assigns probability $p \in [0,1]$ to θ_1. This belief incorporates her prior knowledge/experience, and therefore represents the degree of her ex ante uncertainty in terms of proximity to the (maximally uncertain) uniform belief:

$$\phi := 1 - 2\big|p - {}^1\!/_2\big|. \tag{2}$$

That is, the larger ϕ, the more ex ante uncertain the agent is, meaning that $\phi = 0$ whenever $p \in \{0,1\}$, and respectively $\phi = 1$ whenever $p = {}^1\!/_2$. This notion of uncertainty resembles the usual measures of uncertainty from information theory (Cover and Thomas, 2006). Then, before making a guess, the agent decides how much attention to pay. Attention is modelled with a Bayesian signal, which is uniquely identified by a mean-preserving distribution of posterior probabilities (Kamenica and Gentzkow, 2011). The set of all signals is denoted by

$$\Pi(p) = \Big\{\ \pi \in \Delta\big([0,1]\big) : \mathbb{E}_\pi(q) = p\ \Big\}\ .$$

For any signal $\pi \in \Pi(p)$, define the agent's ex ante indirect expected utility,

$$G(\pi) := \mathbb{E}_\pi\big(g(q)\big). \tag{3}$$

It is not difficult to see that $G(\pi)/u_1(x)$ is the probability that she attaches ex ante (i.e., before π is realized) to her best guess being eventually correct.

Aligned with the rational inattention literature, we assume that attention is costly. The cost function is assumed to be uniformly posterior separable (Caplin et al., 2022), i.e., there is a strictly convex function $c : [0,1] \to \mathbb{R}$ such that the cost of signal $\pi \in \Pi(p)$ is given by

$$C(\pi) = \kappa\Big(\mathbb{E}_\pi\big(c(q)\big) - c(p)\Big), \tag{4}$$

where $c(q)$ represents the agent's marginal cost for acquiring information, and $\kappa > 0$ is a parameter of the task's difficulty. Note that consistently with the complexity literature (Oprea, 2024a), the cost consists an objective part (viz., the parameter κ) and a subjective part (viz., the function c). Such cost functions have solid foundations (Denti, 2022) and are supported by experimental evidence (Dean and Neligh, 2023). Throughout the paper, we will focus on symmetric cost functions, which include the Shannon entropy (Sims, 2003), the Shorrocks entropy (Shorrocks, 1980), the Tsallis entropy (Caplin et al., 2022) as special cases. Formally, this means that for every $q \in [0,1]$, we have $c(q) = c(1-q)$. For an axiomatization of symmetric cost functions, see Hébert and Woodford (2021). We further discuss the symmetry assumption in Section 5.2.

The agent faces a tradeoff, in that more informative signals help her to achieve higher expected utility, but at the same time are also more costly. That is, formally speaking, the agent solves the following optimization problem:

$$\max_{\pi \in \Pi(p)} \Big(G(\pi) - C(\pi) \Big). \tag{5}$$

It is not difficult to verify that for every tuple of parameters (x, ϕ, κ) there is a unique optimal signal, henceforth denoted by $\pi(\cdot|x, \phi, \kappa)$. It follows from standard arguments (e.g., Kamenica

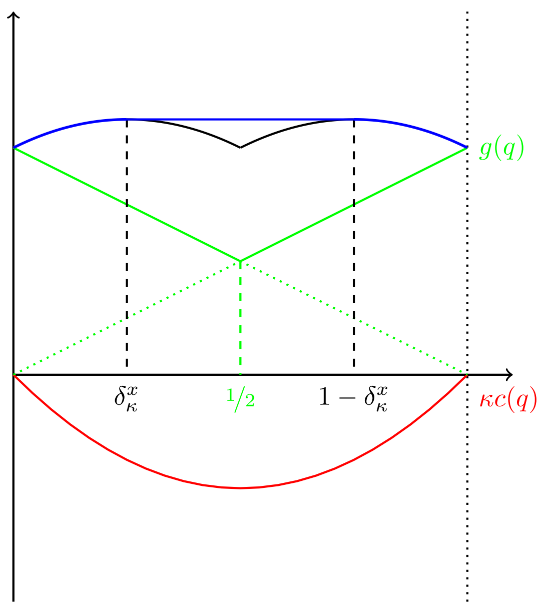

Figure 1: The green solid piecewise linear function is the net expected utility from guessing correctly, and the red line is the marginal cost for acquiring information. The blue curve is the concave closure of the difference $g(q) - \kappa c(q)$. The optimal signal distributes all the probability between the posteriors δ^x_κ and $1-\delta^x_\kappa$ whenever ex ante uncertainty is large (i.e., whenever the prior lies between δ^x_κ and $1-\delta^x_\kappa$), and it is completely uninformative otherwise.

and Gentzkow, 2011; Matějka and McKay, 2015) that $\pi(\cdot|x, \phi, \kappa)$ is given by concavifying the function $g(q) - \kappa c(q)$, as illustrated in Figure 1 below. In particular, there is some $\delta^x_\kappa \in \big[0, 1/2\big]$ such that, for small levels of ex ante uncertainty (i.e., $\phi \leq 2\delta^x_\kappa$) the agent does not exert any effort in acquiring information about the task and maintains her prior beliefs, whereas for large levels of ex ante uncertainty (i.e., $\phi > 2\delta_\kappa^x$) her optimal signal mixes between the posteriors δ_κ^x and $1-\delta_\kappa^x$. Note that δ_κ^x is continuously decreasing in the reward x, and continuously increasing in the difficulty κ. Throughout the paper, we will denote by

$$\phi_\kappa := 2\delta_\kappa^0 \tag{6}$$

the lower bound of ex ante uncertainty for which the agent will acquire information, in the absence of a reward. By monotonicity of δ_κ^x, it follows that ϕ_κ is continuously increasing in κ.

Note that three parameters enter the picture when computing the optimal signal:

- Difficulty (κ) : It is a feature of the task itself, and it is in general unobservable. It is treated as an objective task characteristic, i.e., one that does not vary across agents.
- Ex ante uncertainty (ϕ) : It is an unobservable individual characteristic of the agent in relation to the task, as it reflects the agent's prior experience about the specific task. It is treated as a subjective characteristic of the task, i.e., one that may be perceived differently across agents, but which nonetheless is inherently linked to the task.
- Reward (x) : It is the only parameter which is directly observable, and can be controlled by the analyst. It is treated as an experimental parameter/treatment, and it is not a task characteristic as the utility from x is not affected by the task itself.

Given the previous classification, for a given agent, we will henceforth identify each task with a pair (ϕ, κ), and we will denote the set of all tasks by

$$\mathcal{T} := [0,1] \times (0, \infty).$$

Remark 1. Of course, the optimal signal also depends on individual characteristics of the agent, and in particular of the utility function u and the cost function c. However, given that these characteristics are exogenously fixed, we will take them for granted. We come back to the role of these parameters in Section 3 when we aggregate our soon-to-be-defined complexity measure over a set of agents. ◁

Assuming that the agent is rational (in the sense that she picks the optimal signal), her expected accuracy is given by

$$F(x, \phi, \kappa) := \frac{G\big(\pi(\cdot|x, \phi, \kappa)\big)}{u_1(x)}. \tag{7}$$

This is equal to the expected probability of guessing correctly, as illustrated in Figure 2. Note that whenever her ex ante uncertainty is small (i.e., $\phi \leq 2\delta_\kappa^x$), her expected accuracy will be $\max\{p, 1-p\}$. On the other hand, whenever her ex ante uncertainty is large (i.e., $\phi > 2\delta_\kappa^x$), her expected accuracy becomes $1 - \delta_\kappa^x$.

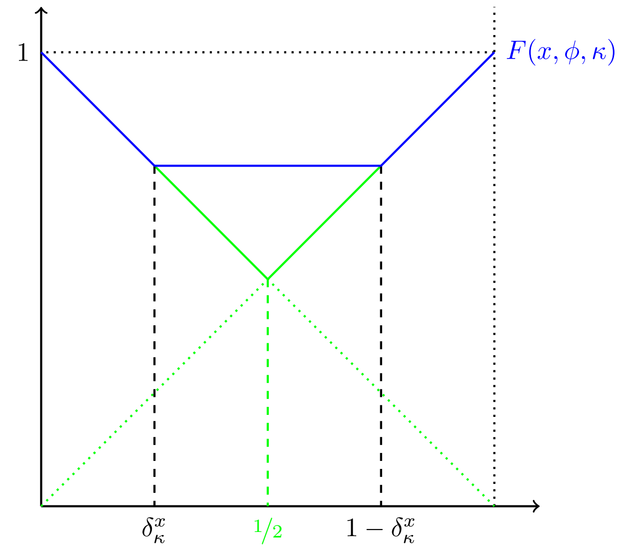

Figure 2: The blue piecewise linear function is the expected accuracy (as a function of the agent's prior), assuming that the optimal signal $\pi(\cdot|x, \phi, \kappa)$ has been used.

In the literature on belief-based complexity measures, expected accuracy is used as a measure of task complexity (Enke and Graeber, 2023; Enke et al., 2024a,b; Oprea, 2024a). This is also consistent with theoretical results that show certain definitions of complexity to be positively correlated with accuracy (Goncalves, 2024). However, there are is an important caveat: Expected accuracy depends both on task characteristics and on an experimental parameter. This leads to an undesirable situation where expected accuracy leads to a task complexity order that varies with respect to characteristic that is not linked with the task itself, e.g., it may very well be the case that task A is labelled more complex than task B under the reward x, while at the same time A is labelled simpler than B under the reward x'. And this naturally makes us wonder: does the agent deem A or B more complex?

In this paper, we address this question by proposing a robust belief-based measure of complexity, which labels A more complex than B whenever the expected accuracy of A is always smaller than the expected accuracy of B, for every extrinsic reward. This idea is formalized in the following definition.

Definition 1. We say that task (ϕ, κ) is deemed ***more complex*** than task (ϕ', κ'), and we write $(\phi, \kappa) \succeq (\phi', \kappa')$, whenever

$$F(x, \phi, \kappa) \leq F(x, \phi', \kappa') \tag{8}$$

for all $x \in X$. ◁

The asymmetric and the symmetric parts of $\succeq$ are defined as usual. That is, we have

$(\phi, \kappa) \succ (\phi', \kappa')$ whenever $(\phi, \kappa) \succeq (\phi', \kappa')$ and $(\phi, \kappa) \npreceq (\phi', \kappa')$, and respectively $(\phi, \kappa) \sim (\phi', \kappa')$ whenever $(\phi, \kappa) \succeq (\phi', \kappa')$ and $(\phi, \kappa) \preceq (\phi', \kappa')$.

A task (ϕ, κ) is said to be trivial if the optimal signal reveals the true state with certainty for every $x \geq 0$. Obviously, if a task (ϕ, κ) is trivial, then it is simpler than any other task, i.e., $(\phi', \kappa') \succeq (\phi, \kappa)$ for all $(\phi', \kappa') \in \mathcal{T}$. The set of non-trivial tasks is henceforth denoted by $\mathcal{T}_0 \subseteq \mathcal{T}$, and it is characterized by a difficulty threshold $\kappa_0 \geq 0$.

Proposition 1. *There is some $\kappa_0 \geq 0$, such that the following are equivalent:*

(i) $\kappa > \kappa_0$.

(ii) $(\phi, \kappa) \in \mathcal{T}_0$ *for every* $\phi \in [0, 1]$.

The idea is quite simple: in order for a task to be non-trivial, the information costs must be sufficiently large to guarantee that the intrinsic incentives alone are not strong enough to always lead to a perfectly informative signal. In our main result below, we characterize how non-trivial tasks are ranked in terms of complexity if we use our robust belief-based measure.

Theorem 1. *For any pair (ϕ, κ) and (ϕ', κ') of non-trivial tasks, the following are equivalent:*

(i) $(\phi', \kappa') \succeq (\phi, \kappa)$.

(ii) $\kappa' \geq \kappa$ *and* $\phi' \geq \min\{\phi_\kappa, \phi\}$.

Recall that ϕ_κ is the information-acquisition constraint in the absence of any reward, i.e., the agent will acquire information if and only if her ex ante uncertainty satisfies the following inequality:

$$\phi > \phi_\kappa. \tag{9}$$

Moreover, ϕ_κ is continuously increasing in κ, with $\phi_\kappa \to 0$ as $\kappa \to \kappa_0$, and $\phi_\kappa \to 1$ as κ grows arbitrarily large. So, without sufficiently large costs, the agent will always acquire the perfectly informative signal regardless of her prior, whereas when the cost becomes infinitely large she will not acquire information regardless of her prior belief.

Remark 2. The function ϕ_κ depends on the cost function c and the value $u_1(0)$. Therefore, the utilities assigned to any other $x > 0$ is inconsequential for $\succeq$. $\triangleleft$

Graphically, the previous result is illustrated in the Figures below which correspond to two cases, i.e., the one where the information-acquisition constraint of (9) holds, and the one where it does not. In both figures, we have taken $c(q) = q^2 - q$, which is subdifferentiable at the boundaries of the unit interval, and therefore for small cost parameters the agent will optimally acquire the perfectly informative signal regardless of the size of the reward. This explains why

ϕ_κ is initially constant at 0, and as a result there is a grey region which contains the trivial tasks. Such horizontal part would not have appeared if the cost function was for instance entropic, in which case a perfectly informative signal would have never been optimal.

Starting with the first case (Figure 3), we take a task (ϕ, κ) such that $\phi > \phi_\kappa$. Then, the red region contains the tasks (ϕ', κ') that are demmed more complex than (ϕ, κ). On the other hand, the green region contains the tasks (ϕ', κ') that are deemed simpler than (ϕ, κ). How we obtained the red region is obvious given our previous theorem. So, let us elaborate on how the green region arises. First of all, $\kappa \geq \kappa'$ and $\phi \geq \phi'$ jointly imply $(\phi, \kappa) \succeq (\phi', \kappa')$, again by our theorem. Then, let us consider some (ϕ', κ') such that $\kappa' < \kappa$ and $\phi' \geq \phi$. By ϕ_κ being increasing, it follows that $\phi_{\kappa'} < \phi_\kappa < \phi$. Hence, once again by our previous theorem, we obtain $(\phi, \kappa) \succeq (\phi', \kappa')$.

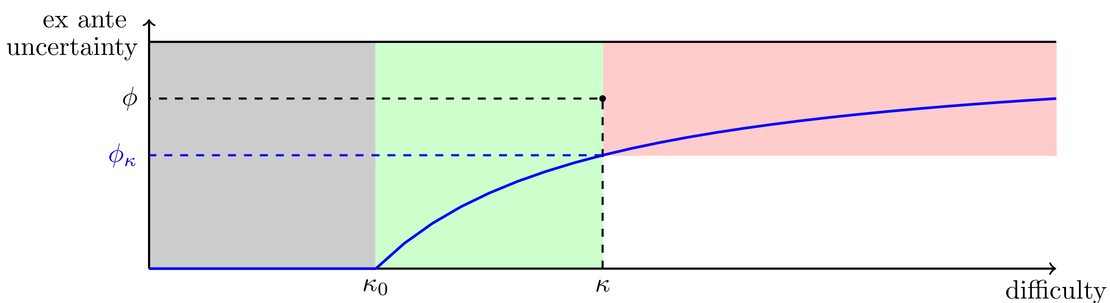

Figure 3: This is the case when the ex ante uncertainty is large and therefore it is optimal to acquire information before making a guess. The red area contains the tasks that are deemed more complex than (ϕ, κ), and the green area are the tasks that are deemed simpler than (ϕ, κ). Finally, the grey area contains the trivial tasks, meaning that they are deemed simpler than every other task, including (ϕ, κ).

Let us now focus on the second case (Figure 4), where we take a task (ϕ, κ) such that $\phi \leq \phi_\kappa$. Once again, by our previous theorem, the red region is the set of tasks that are deemed more complex than (ϕ, κ). And once again, the green region is the set of tasks that are deemed simpler than (ϕ, κ). Regarding this last part, let us elaborate only on the non-obvious case where $\kappa' < \kappa$ and $\phi' > \phi$. If κ' is such that $\phi_{\kappa'} > \phi$, then by our theorem (ϕ, κ) and (ϕ', κ') are not comparable via $\succeq$.

Remark 3. Note that in most of the existing literature, complexity is taken as a synonym of difficulty. Here we show that this is not the case, as complexity depends both on difficulty and ex ante uncertainty. This is consistent with the idea that a task may be deemed complex because it has not been previously encountered by the agent. For instance, students often complain that they did not have enough practice questions similar to the ones they faced in their final exam, even though their actual exam was not particularly difficult. Nevertheless, even though increased difficulty is not sufficient for increased complexity, it is still a necessary

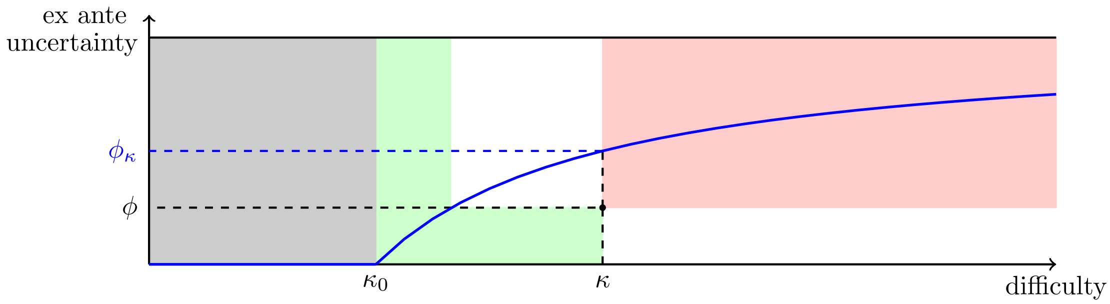

Figure 4: This is the case when the ex ante uncertainty is small and therefore it is optimal not to acquire any information before making a guess. Once again, the red area contains the tasks that are deemed more complex than (ϕ, κ), the green area the tasks that are deemed simpler than (ϕ, κ), and the grey area the trivial tasks.

condition, i.e., unless both (ϕ, κ) and (ϕ', κ') are both trivial, $(\phi', \kappa') \succeq (\phi, \kappa)$ will necessarily imply $\kappa' \geq \kappa$. Therefore, even though neither difficulty nor ex ante uncertainty are directly observable, by observing that task A is deemed more complex than task B, we can conclude that A is more difficult than B. We come back to this point in Section 3. $\triangleleft$

As conceptually appealing as robustness is, it comes at the cost of not being able to rank all tasks. Below, we provide the formal statement:

Proposition 2. *The relation $\succeq$ is transitive and incomplete.*

The previous result can be easily illustrated in the two figures above, i.e., there is always a white region which includes the tasks that are subjectively neither more complex nor simpler than any given task.

Note that within $\mathcal{T}_0$ the relation $\succeq$ has the same structure as the incomplete preference relations in Ok (2002). To understand the connection, recall that in that paper preferences are represented by a vector-valued utility function. In the context of our paper this would mean that there exists some

$$\boldsymbol{u} : \mathcal{T}_0 \to \mathbb{R}^2$$

such that $(\phi, \kappa) \succeq (\phi', \kappa')$ if and only if $\boldsymbol{u}_1(\phi, \kappa) \geq \boldsymbol{u}_1(\phi', \kappa')$ and $\boldsymbol{u}_2(\phi, \kappa) \geq \boldsymbol{u}_2(\phi', \kappa')$, where

$$\boldsymbol{u}_1(\phi, \kappa) := \kappa \text{ and } \boldsymbol{u}_2(\phi, \kappa) := \min\{\phi_\kappa, \phi\}.$$

The incompleteness property also distinguishes our measure of complexity from most definitions in the literature which typically induce a complete order, e.g., the signal-to-noise ratio (Callander, 2011; Fehr and Rangel, 2011; Goncalves, 2024) or willingenss to pay for avoiding a task (Oprea, 2020) or the subjective-ranking metric (Gabaix and Graeber, 2024). Nonetheless, as we will discuss in the next section, completeness can be restored once we move to an objective measure of complexity that aggregates the subjective measures from many agents.

3. Towards an objective measure

The incomplete complexity order that we obtained through our complexity measure reflects an agent's subjective assessment. In this section, we will study the question of aggregating such measures in an attempt to obtain an objective measure of complexity.

Consider a group of agents I. Each $i \in I$ is identified by a utility function u_i and a cost function c_i, which respectively satisfy the conditions we introduced in the previous section. These two functions induce an information-acquisition constraint ϕ^i_κ for every $i \in I$, which in turn yields agent i's incomplete complexity relation $\succeq^i$ (via Theorem 1).

Let us then introduce a new complexity order which aggregates the different $\succeq^i$ for all the agents $i \in I$. This new order will label task A more complex than task B whenever at least one agent is able to rank the two tasks, and moreover none of the agents deems B strictly more complex than A. Formally, this criterion is stated as follows:

Definition 2. Within group I, we say that task (ϕ, κ) is ***objectively more complex*** than task (ϕ', κ'), and we write $(\phi, \kappa) \succeq^I (\phi', \kappa')$, if the following conditions hold:

(a) There is at least one $i \in I$ such that (ϕ, κ) and (ϕ', κ') are $\succeq^i$-comparable.

(b) For every $i \in I$, either (ϕ, κ) and (ϕ', κ') are not $\succeq^i$-comparable or $(\phi, \kappa) \succeq^i (\phi', \kappa')$.

The strict relation $\succ^I$ and the indifference relation $\sim^I$ are defined in the usual way, i.e., as the asymmetric and symmetric part of $\succeq^I$ respectively. ◁

Before moving forward, let us point out that $\succeq^I$ is well-defined in the following (desired) sense: two agents will never completely contradict each other, i.e., it will never be the case that one agent deems A more strictly complex than B, and another agent will deem B strictly more complex than A (see Lemma A1 in the Appendix). And this, in turn, yields the following result.

Proposition 3. *For any $i \in I$ and any pair of tasks $(\phi, \kappa), (\phi', \kappa') \in \mathcal{T}$, we have:*

(a) If $(\phi, \kappa) \succeq^i (\phi', \kappa')$, then (ϕ, κ) and (ϕ', κ') are $\succeq^I$-comparable.

(b) If $(\phi, \kappa) \succ^i (\phi', \kappa')$, then $(\phi, \kappa) \succ^I (\phi', \kappa')$.

Note that $(\phi, \kappa) \succeq^i (\phi', \kappa')$ does not directly imply $(\phi, \kappa) \succeq^I (\phi', \kappa')$. This is because, if we have $(\phi, \kappa) \sim^i (\phi', \kappa')$ and $(\phi', \kappa') \succ^j (\phi, \kappa)$, then it is the case that $(\phi', \kappa') \succ^I (\phi, \kappa)$.

So now, let us focus on the effect of adding more agents into a group. The idea is that the moment two tasks are $\succeq^i$-comparable for some $i \in I$, then they are also $\succeq^I$-comparable. That is, the more agent we are adding to I, the more complete $\succeq^I$ becomes. Then, the question is whether we will eventually have sufficiently many agents in I so that $\succeq^I$ becomes complete.

To answer the question, let us first formalize a notion of richness. Recall that $\succeq^i$ is derived from the information-acquisition constraint ϕ^i_κ, which depends on agent i's individual characteristics (Remark 2). We will say that I is rich if for every task (ϕ, κ) there is some $i \in I$ such that $\phi^i_\kappa = \phi$. In other words, there is always an agent for whom the information-acquisition constraint is binding. Graphically, this corresponds to covering the entire set $\mathcal{T}$ with the graphs of the functions ϕ^i_κ for the different agents.

Theorem 2. *For every rich group of agents I, and any pair (ϕ, κ) and (ϕ', κ') of tasks, the following are equivalent*

(i) $(\phi', \kappa') \succeq^I (\phi, \kappa)$.

(ii) Either $\kappa' > \kappa$, or simultaneously $\kappa' = \kappa$ and $\phi' > \phi$.

The previous result has several direct implications. First of all, quite obviously, $\succeq^I$ is the lexicographic order that first ranks tasks with respect to difficulty, and then —in case of a tie— ranks them with respect to ex ante uncertainty. In this sense, the complexity order that our objective measure induces as two dimensions: the usual difficulty dimension that is identified throughout the literature as almost synonymous to complexity (Oprea, 2024a), as well as the novel dimension of ex ante uncertainty that we identify in this paper. Nonetheless, conceptually, our conclusion is not that different from the existing literature. The reason is that, although we identify this new dimension of complexity, the primary dimension remains the same, viz., difficulty.

We say that $\succeq^I$ is robust with respect to adding more agents, if for every $J \supseteq I$ it is the case that $\succeq^I = \succeq^J$. As it turns out, the only relation $\succeq^I$ which is robust with respect to adding more agents in the lexicographic relation of Theorem 2, as formally shown below.

Proposition 4. *For a group I, the following are equivalent:*

(i) $\succeq^I$ is robust with respect to adding more agents.

(ii) $\succeq^I = \succeq^{I'}$ for some rich group I'.

By the previous result, the lexicographic order is the only complete and asymmetric $\succeq^I$ that can be obtained for some group I. The intuition is quite simple. If $\succeq^I$ is already complete and asymmetric and we add more agents to I, nothing is going to change (Proposition 3). Hence, $\succeq^I$ is robust with respect to adding more agents, which in turn implies that I is rich (Proposition 4). Therefore, $\succeq^I$ must necessarily be the lexicographic relation (Theorem 2).

4. Subjective complexity and effort

Let us go back to the case where there is a single agent. In this context, the relationship between complexity and effort has been studied by several papers in the literature (e.g., Goncalves, 2024, and references therein). Similarly to these earlier papers, we naturally assume that effort is positively correlated with the cost $C\big(\pi(\cdot|x,\phi,\kappa)\big)$ that the agent incurs for optimally acquiring information about the task.

For starters, it is not difficult to see that for any fixed $x \in X$, the induced effort is not always monotonic in complexity. This finding is not surprising, and it has been already pointed out in the literature. Here we will focus on the the problem from a different angle, focusing on robustness with respect to the reward, i.e., we ask whether the non-monotonicity can arise for every $x \in X$.

Proposition 5. *Take an arbitrary $(\phi,\kappa) \in \mathcal{T}_0$ such that $\phi > \phi_\kappa$. Then, there exists a task $(\phi',\kappa') \succ (\phi,\kappa)$ such that $C\big(\pi(\cdot|x,\phi',\kappa')\big) < C\big(\pi(\cdot|x,\phi,\kappa)\big)$ for all $x \in X$.*

The previous result shows that for all tasks for which the agent optimally exerts some effort, we can always find a more complex task which will induce less effort regardless of the size of the reward. In this sense, non-monotonicity of effort with respect to complexity is both generic and robust with respect to the reward. Graphically this is illustrated in Figure 5 below. Of course, for reward $x = 0$ it is pretty obvious, as the agent will not acquire any information when facing (ϕ',κ'), as opposed to when facing (ϕ,κ). The interesting part arises when the extrinsic reward increases to any $x > 0$. In this case, although (ϕ',κ') is harder, it involves lower ex ante uncertainty. And this is exactly what makes it cheaper.

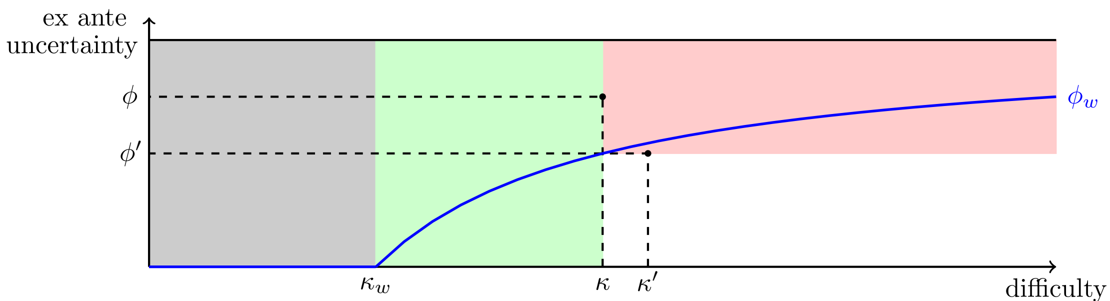

Figure 5: THere is some $\kappa' > \kappa$ such that task (ϕ',κ') is strictly more complex than (ϕ,κ), but at the same time it always induces less effort, regardless of how large the reward is.

In fact, it is not difficult to show that if complexity is entirely driven by ex ante uncertainty, more complex tasks will always induce more effort. That is, whenever $\phi > \phi'$, it will be the case that $(\phi,\kappa) \succ (\phi',\kappa)$, and furthermore $C\big(\pi(\cdot|x,\phi,\kappa)\big) > C\big(\pi(\cdot|x,\phi',\kappa)\big)$ for all $x \geq 0$.

Interestingly, while the effect of ex ante uncertainty is clear, the same is not true for difficulty.

Proposition 6. *For every ϕ, there exists some $\kappa_\phi > \kappa_0$ such that, for any two $\kappa' > \kappa > \kappa_\phi$:*

$$\begin{aligned} C\big(\pi(\cdot|x,\phi,\kappa)\big) &> C\big(\pi(\cdot|x,\phi,\kappa')\big), \\ C\big(\pi(\cdot|x',\phi,\kappa)\big) &< C\big(\pi(\cdot|x',\phi,\kappa')\big), \end{aligned}$$

for rewards $x' > x \geq 0$.

The previous result suggests that effort rankings are never robust along the difficulty dimension. Thus, besides the fact that effort is not monotonic with respect to complexity, the strength of the extrinsic reward also matters, in the sense that for relatively small rewards the simpler task induces more effort, whereas for high rewards the more complex task induces more effort. This result can be seen as the robust counterpart of Goncalves (2024, Prop. 3), who shows that for a fixed x, effort is not monotonic with respect to difficulty.

5. Discussion

5.1. Eliciting subjective complexity

Let us focus on a simple —albeit important— aspect regarding the practical implementation of our complexity measure. Suppose that we elicit a subject's expected accuracy after she has undertaken the task, like in Enke and Graeber (2023). The problem is that in this case, we can no longer vary the reward, and therefore our measure will no longer be robust. Hence, the only option would be to elicit expected accuracy before the task is undertaken, using the strategy method across different rewards. However, in this case, the subject would have incentives to hedge for each reward level, thus potentially misreporting her expected accuracy (Blanco et al., 2010).

As a result, elicit beliefs about one's own expected accuracy would not be feasible in practice. In order to overcome this experimental hurdle, we can rely on an idea similar to the one Bayesian markets (Baillon, 2017), where we use beliefs about other individuals to proxy one's own beliefs about themselves. Accordingly, in our setting, we can elicit our subject's beliefs about the expected accuracy of another individual for different reward levels. Then, under the assumption that the subject considers this other individual being similar to themselves, the elicited accuracy can be used to construct the complexity measure.

5.2. Cost specification

There is a natural question regarding the cost specification that we have assumed throughout the paper: how restrictive is the symmetry assumption?

Starting with symmetry of c, note that in the literature this assumption has been critizised mostly because distinguishing between two states might be more difficult than distinguishing between two other states (Hébert and Woodford, 2021). In other words, asymmetries enter the picture primarily in cases where the state space has some underlying distance, and more similar states are harder to tell apart. However, given that here we focus on binary tasks, it is reasonably justified to maintain symmetric costs.

A. Proofs

Proof of Proposition 1. By a standard concavification argument, there exists some $\delta_\kappa^x \in [0, 1/2]$ such that the optimal signal distributes all probability between δ_κ^x and $1-\delta_\kappa^x$ whenever $p \in (\delta_\kappa^x, 1-\delta_\kappa^x)$, and it is completely uninformative otherwise. Note that δ_κ^x is increasing in κ and decreasing in x. Then, define

$$\kappa_0 := \sup\{\kappa > 0 : \delta_\kappa^0 = 0\}. \tag{A.1}$$

$(i) \Rightarrow (ii)$: By construction, for every $\kappa \leq \kappa_0$ and every $x \geq 0$, it is the case that $\delta_\kappa^x \leq \delta_\kappa^0 \leq \delta_{\kappa_0}^0 = 0$. This means that (ϕ, κ) is trivial for every $\phi \in [0, 1]$.

$(ii) \Rightarrow (i)$: Suppose that $\kappa > \kappa_0$. Then, by construction $\delta_\kappa^0 > 0$. Therefore, there is some $\phi > 0$ sufficiently close to 0 such that the optimal signal is completely uninformative. This means that $F(0, \phi, \kappa) < F(0, \phi, \kappa_0) = 1$, i.e., (ϕ, κ) is not trivial. □

Proof of Theorem 1. Take an arbitrary pair (ϕ, κ). Using the definition of δ_κ^x from the proof of Proposition 1, we obtain

$$F(x, \phi, \kappa) = \begin{cases} 1 - \phi/2 & \text{if } \phi \leq \phi_\kappa^x, \\ 1 - \delta_\kappa^x & \text{if } \phi > \phi_\kappa^x, \end{cases} \tag{A.2}$$

where $\phi_\kappa^x := 1 - 2\left|\delta_\kappa^x - 1/2\right|$ is the threshold of the degree of ex ante uncertainty for acquiring information. Then, define $\phi_\kappa := \phi_\kappa^0$.

$(ii) \Rightarrow (i)$: Take $\kappa' \geq \kappa$ and $\phi' \geq \min\{\phi_\kappa, \phi\}$, and consider two cases:

First, let $\phi' \geq \phi$. Note that Equation (A.2) can be equivalently rewritten as

$$F(x, \phi, \kappa) = \max\{1 - \phi/2, 1 - \delta_\kappa^x\}. \tag{A.3}$$

Letting $\delta_\kappa^{x'}$ be the low posterior we obtain from the concavification exercise given the parameters (ϕ', κ'), it is not difficult to verify that $\delta_\kappa^{x'} \leq \delta_\kappa^x$. Hence, it follows directly from (A.3) that $F(x, \phi', \kappa') \leq F(x, \phi, \kappa)$ for all $x \geq 0$, and a fortiori we obtain $(\phi', \kappa') \succeq (\phi, \kappa)$.

Second, let $\phi > \phi' \geq \phi_\kappa$. By definition, ϕ_κ^x is decreasing in x, and hence $\phi > \phi_\kappa^x$. So, combined with (A.2), we obtain $F(x,\phi,\kappa) = F(x,\phi',\kappa) = 1 - \delta_\kappa^x$. Moreover, as shown in the first case above, F is decreasing in κ and therefore $F(x,\phi',\kappa') \leq F(x,\phi',\kappa)$. Combining the two yields $F(x,\phi',\kappa') \leq F(x,\phi,\kappa)$ for all $x \geq 0$, and a fortiori we obtain $(\phi',\kappa') \succeq (\phi,\kappa)$.

$(i) \Rightarrow (ii)$: Suppose that $\kappa_0 < \kappa' < \kappa$. Then, there exists some $x \geq 0$ such that $\phi > \phi_\kappa^x$, and therefore $F(x,\phi,\kappa) = 1 - \delta_\kappa^x$. But then, by (A.3), it is also the case that $F(x,\phi',\kappa') \geq 1 - \delta_\kappa^{x'}$. And by $\kappa' < \kappa$, we have $1 - \delta_\kappa^{x'} > 1 - \delta_\kappa^x$, and a fortiori $F(x,\phi',\kappa') > F(x,\phi,\kappa)$, meaning that $(\phi',\kappa') \not\succeq (\phi,\kappa)$.

Finally, suppose that $\kappa_0 < \kappa \leq \kappa'$ and $\phi' < \min\{\phi,\phi_\kappa\}$. The latter implies that $\phi' < \phi$. Since both (ϕ,κ) and (ϕ',κ') are non-trivial, it means that $F(x,\phi,\kappa) = 1 - {}^{\phi}/_{2} < 1 - {}^{\phi'}/_{2} = F(x,\phi,\kappa)$, and therefore $(\phi',\kappa') \not\succeq (\phi,\kappa)$. □

Proof of Proposition 2. TRANSITIVITY: It follows directly from Theorem 1. Namely, begin with $(\phi'',\kappa'') \succeq (\phi',\kappa') \succeq (\phi,\kappa)$. Hence, we have $\kappa'' \geq \kappa' \geq \kappa$ and $\phi'' \geq \min\{\phi',\phi_{\kappa'}\}$ and $\phi' \geq \min\{\phi,\phi_\kappa\}$. Then, we consider two cases:

(i) $\phi'' \geq \phi'$: It follows directly $\phi'' \geq \min\{\phi,\phi_\kappa\}$, and therefore $(\phi'',\kappa'') \succeq (\phi,\kappa)$.

(i) $\phi' > \phi'' \geq \phi_{\kappa'}$: By monotonicity, we have $\phi_{\kappa'} \geq \phi_\kappa \geq \min\{\phi,\phi_\kappa\}$. Hence, we obtain $\phi'' \geq \min\{\phi,\phi_\kappa\}$, and a fortiori $(\phi'',\kappa'') \succeq (\phi,\kappa)$.

INCOMPLETENESS: By the fact that ϕ_κ approaches 1 in the limit, there exists some $\kappa > \kappa_0$ such that $\phi_\kappa > 0$. Then, consider $\phi > \phi_\kappa > \phi'$, and $\kappa' > \kappa$. By Theorem 1, we obtain $(\phi,\kappa) \not\succeq (\phi',\kappa')$ and $(\phi',\kappa') \not\succeq (\phi,\kappa)$, meaning that $\succeq$ is incomplete. □

Lemma A1. *There is no pair of agents $i,j \in I$ and tasks $(\phi,\kappa),(\phi',\kappa') \in \mathcal{T}$ such that $(\phi,\kappa) \succ^i (\phi',\kappa')$ and $(\phi',\kappa') \succ^j (\phi,\kappa)$.*

Proof. By Theorem 1, if $(\phi,\kappa) \succ^i (\phi',\kappa')$, then it is either the case that $\kappa > \kappa'$, or simultaneously $\kappa = \kappa'$ and $\phi \geq \phi_\kappa > \phi'$. But then, in either of these cases, again by Theorem 1, we cannot have $(\phi',\kappa') \succ^j (\phi,\kappa)$. □

Proof of Proposition 3. (a) : There are two possible cases. First, let $(\phi,\kappa) \succ^i (\phi',\kappa')$. This means that condition (a) in Definition 2 is satisfied. Moreover, by Lemma A1, for every $j \in I$ for which (ϕ,κ) and (ϕ',κ') are $\succeq^j$-comparable, we will have $(\phi,\kappa) \succeq^j (\phi',\kappa')$, meaning that condition (b) in Definition 2 is also satisfied. Hence, we obtain $(\phi,\kappa) \succeq^I (\phi',\kappa')$, and therefore (ϕ,κ) and (ϕ',κ') are $\succeq^I$-comparable.

Turning to the second case, let $(\phi,\kappa) \sim^i (\phi',\kappa')$. This means that condition (a) in Definition 2 is satisfied. If there is some $j \neq i$ such that either $(\phi,\kappa) \succ^j (\phi',\kappa')$ or $(\phi',\kappa') \succ^j (\phi,\kappa)$, then following the same steps as in the previous case, we conclude that (ϕ,κ) and (ϕ',κ') are

$\succeq^I$-comparable. If on the other hand, for every $j \in I$ for which (ϕ, κ) and (ϕ', κ') are $\succeq^j$-comparable, we have $(\phi, \kappa) \sim^j (\phi', \kappa')$, then $(\phi, \kappa) \sim^I (\phi', \kappa')$, meaning that (ϕ, κ) and (ϕ', κ') are $\succeq^I$-comparable.

(b) : By the first case in the proof of part (a) above, $(\phi, \kappa) \succ^i (\phi', \kappa')$ implies $(\phi, \kappa) \succeq^I (\phi', \kappa')$. Then, suppose contrary to what we want to prove that $(\phi', \kappa') \succeq^I (\phi, \kappa)$. Since the two tasks are $\succeq^i$-comparable, part (b) of Definition 2 would then imply $(\phi', \kappa') \succeq^i (\phi, \kappa)$, which would in turn contradict $(\phi, \kappa) \succ^i (\phi', \kappa')$. $\square$

Proof of Theorem 2. $(ii) \Rightarrow (i)$: If $\kappa' = \kappa$ and $\phi' > \phi$, then it follows directly from Theorem 1 that for every $i \in I$, we will have $(\phi', \kappa') \succeq^i (\phi, \kappa)$, and therefore $(\phi', \kappa') \succeq^I (\phi, \kappa)$.

So, let us focus on the case where $\kappa' > \kappa$. Again by Theorem 1, if $\phi' \geq \phi$ then $(\phi', \kappa') \succeq^i (\phi, \kappa)$ for every $i \in I$, and a fortiori $(\phi', \kappa') \succeq^I (\phi, \kappa)$. So, suppose that together with $\kappa' > \kappa$ we have $\phi' < \phi$. Since I is rich, there exists some $i \in I$ such that $\phi^i_\kappa = \phi'$, meaning that $(\phi', \kappa') \succeq^i (\phi, \kappa)$. Furthermore, by $\kappa' > \kappa$, there is no $j \in I$ such that $(\phi, \kappa) \succeq^j (\phi', \kappa')$, as explained in Remark 3. Hence, $(\phi', \kappa') \succeq^I (\phi, \kappa)$.

$(i) \Rightarrow (ii)$: We will proceed by contraposition, starting with the premise that (ii) does not hold. Then, there are two cases. First, let $\kappa' < \kappa$. Thus, by Remark 3, there is no $j \in I$ such that $(\phi', \kappa') \succeq^j (\phi, \kappa)$, and therefore $(\phi', \kappa') \not\succeq^I (\phi, \kappa)$.

So, let us turn to the second case where $\kappa' = \kappa$ and $\phi' < \phi$. Then, by I being rich, there is some $i \in I$ such that $\phi' < \phi^i_{\kappa'} = \phi^i_\kappa < \phi$. As a result, by Theorem 1, we obtain $(\phi, \kappa) \succeq^i (\phi', \kappa')$. Hence, by Proposition, it is the case that $(\phi, \kappa) \succeq^I (\phi', \kappa')$, and a fortiori $(\phi', \kappa') \not\succeq^I (\phi, \kappa)$. $\square$

Proof of Proposition 4. $(ii) \Rightarrow (i)$: Take some rich I' such that $\succeq^I = \succeq^{I'}$. Then, by Theorem 2, the order $\succeq^I$ is complete and asymmetric, i.e., for any two tasks $(\phi, \kappa), (\phi', \kappa') \in \mathcal{T}$, we will have either $(\phi, \kappa) \succ^I (\phi', \kappa')$ or $(\phi', \kappa') \succ^I (\phi, \kappa)$. Without loss of generality, let $(\phi, \kappa) \succ^I (\phi', \kappa')$. Now, suppose that we take some $J \supsetneq I$. Then, by Lemma A1, for every $j \in J \setminus I$, we will have either $(\phi, \kappa) \succeq^j (\phi', \kappa')$ or the tasks will not be $\succeq^j$-comparable. Hence, by definition, it will be the case that $(\phi, \kappa) \succ^J (\phi', \kappa')$, i.e., $\succeq^I = \succeq^J$.

$(i) \Rightarrow (ii)$: Since I is robust to adding more agents, we can take some $J \supseteq I$ such that for every pair (u, c) there exists some $j \in J$ such that $(u_j, c_j) = (u, c)$. This means that J will be rich. But then, by $\succeq^I$ being robust with respect to adding more agents, we get $\succeq^I = \succeq^J$. $\square$

Proof of Proposition 5. Let $\phi' := \phi_\kappa$. Denote by p and p' two arbitrary priors that correspond to ϕ and ϕ' respectively. Then, for every $x \geq 0$, we obtain

$$C\big(\pi(\cdot|x, \phi, \kappa)\big) - C\big(\pi(\cdot|x, \phi', \kappa)\big) = \kappa\big(c(p) - c(p')\big) = \varepsilon > 0. \tag{A.4}$$

Take some $\kappa' > \kappa$ such that $(\kappa' - \kappa)\big(c(0) - c(1/2)\big) = \varepsilon$. Then, for any $x \geq 0$, we obtain

$$C\big(\pi(\cdot|x, \phi', \kappa')\big) - C\big(\pi(\cdot|x, \phi', \kappa)\big) = \kappa'\big(c(\delta_{\kappa'}x) - c(p')\big) - \kappa\big(c(\delta_\kappa^x) - c(p')\big) < \varepsilon, \tag{A.5}$$

where δ_κ^x and $\delta_{\kappa'}x$ are the posteriors that we obtain from the concavification exercise, like in the proof of Theorem 1. The last inequality follows from $c(\delta_\kappa^x), c(\delta_{\kappa'}^x) \leq c(0)$ and $c(p') > c(1/2)$. Hence, by (A.4) and (A.5), we obtain $C\big(\pi(\cdot|x, \phi', \kappa')\big) - C\big(\pi(\cdot|x, \phi, \kappa)\big) < 0$. □

Proof of Proposition 6. Let $\phi < 1$, and define $\kappa_0 := \phi_\kappa^{-1}(\phi)$. Fix $\kappa' > \kappa > \kappa_0$, and observe that there is some $x > 0$ such that $\pi(\cdot|x, \phi, \kappa')$ is completely uninformative as opposed to $\pi(\cdot|x, \phi, \kappa)$ which provides some information. Therefore, we obtain $C\big(\pi(\cdot|x, \phi, \kappa)\big) > C\big(\pi(\cdot|x, \phi, \kappa')\big)$. Moreover, for each $x'' \geq 0$, denote by δ_κ^x and $\delta_{\kappa'}^x$ the low posteriors in the support of $\pi(\cdot|x'', \phi, \kappa)$ and $\pi(\cdot|x'', \phi, \kappa')$ respectively, like in the proof of Theorem 1. Note that, as x'' becomes arbitrarily large, both δ_κ^x and $\delta_{\kappa'}^x$ converge to 0. Thus, by continuity of c, there exists some x' such that $C\big(\pi(\cdot|x', \phi, \kappa)\big) < C\big(\pi(\cdot|x', \phi, \kappa')\big)$. □

References

Abreu, D. and Rubinstein, A. (1988). The structure of nash equilibria in repeated games with finite automata. *Econometrica*, 56:1259–1281.

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

Anscombe, F. and Aumann, R. (1963). A definition of subjective probability. *Annals of Mathematical Statistics*, 34:199–205.

Baillon, A. (2017). Bayesian markets to elicit private information. *Proceedings of the National Academy of Sciences*, 114:7958–7962.

Banovetz, J. and Oprea, R. (2023). Complexity and procedural choice. *American Economic Journal: Microeconomics*, 15:3913–3951.

Blanco, M., Engelmann, D., Koch, A., and Normann, H.-T. (2010). Belief elicitation in experiments: is there a hedging problem? *Experimental Economics*, 13:412–438.

Brier, G. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:1–3.

Callander, S. (2011). Searching and learning by trial and error. *American Economic Review*, 111:2277–2308.

Caplin, A. (2025). *An introduction to cognitive economics: The science of mistakes*. Palgrave.

Caplin, A., Dean, M., and Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing Shannon entropy. *Journal of Political Economy*, 130:1676–1715.

Cover, T. and Thomas, J. (2006). *Elements of Information Theory*. Wiley-Interscience.

Dean, M. and Neligh, N. (2023). Experimental tests of rational inattention. *Journal of Political Economy*, 131:3415–3461.

Denti, T. (2022). Posterior separable cost of information. *American Economic Review*, 112:3215–3259.

Enke, B. (2024). The cognitive turn in behavioral economics. *Working Paper*.

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

Enke, B., Graeber, T., and Oprea, R. (2024a). Complexity and time. *Journal of the European Economic Association*.

Enke, B., Graeber, T., Oprea, R., and Yang, J. (2024b). Behavioral attenuation. *Working Paper*.

Fehr, E. and Rangel, A. (2011). Neuroeconomic foundations of economic choice – recent advances. *Journal of Economic Perspectives*, 25:3–30.

Gabaix, X. and Graeber, T. (2024). The complexity of economic decisions. *Working Paper*.

Goncalves, D. (2024). Speed, accuracy, and complexity. *Working Paper*.

Hébert, B. and Woodford, M. (2021). Neighborhood-based information costs. *American Economic Review*, 111:3225–3255.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101:2590–2615.

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

Ok, E. (2002). Utility representation of an incomplete preference relation. *Journal of Economic Theory*, 104:429–449.

Oprea, R. (2020). What makes a rule complex? *American Economic Review*, 110:3913–3951.

Oprea, R. (2024a). Complexity and its measurement. *Handbook of Experimental Methods in the Social Sciences.*

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

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

Savage, L. (1954). *The foundations of statistics.* Wiley.

Savage, L. (1971). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66:783–801.

Shorrocks, A. (1980). The class of additively decomposable inequality measures. *Econometrica*, 48:613–625.

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

van der Wel, P. and van Steenbergen, H. (2018). Pupil dilation as an index of effort in cognitive control tasks: A review. *Psychonomic Bulletin and Review*, 25:2005–2015.

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