



Cost-Benefit Analysis in Reasoning

**Larbi Alaoui
Antonio Penta**

October 2018

Barcelona GSE Working Paper Series

Working Paper n° 1062

Cost-Benefit Analysis in Reasoning*

Larbi Alaoui[†]

Antonio Penta[‡]

UPF and Barcelona GSE

ICREA, UPF and Barcelona GSE

October 8, 2018

Abstract

When an individual thinks about a problem, his decision to reason further may involve a tradeoff between cognitive costs and a notion of value. But it is not obvious that this is always the case, and the value of reasoning is not well-defined. This paper analyzes the primitive properties of the reasoning process that must hold for the decision to stop thinking to be represented by a cost-benefit analysis. We find that the properties that characterize the cost-benefit representation are weak and intuitive, suggesting that such a representation is justified for a large class of problems. We then provide additional properties that give more structure to the value of reasoning function, including ‘value of information’ and ‘maximum gain’ representations. We show how our model applies to a variety of settings, including contexts involving sequential heuristics in choice, response time, reasoning in games and research. Our model can also be used to understand economically relevant patterns of behavior for which the cost-benefit approach does *not* seem to hold. These include choking under pressure and (over)thinking aversion.

Keywords: cognition and incentives – choice theory – reasoning – fact-free learning – sequential heuristics

JEL Codes: D01; D03; D80; D83.

*Alaoui gratefully acknowledges financial support from the Spanish Ministry of Science and Innovation under projects ECO2014-56154-P and SEV-2015-0563. Penta acknowledges the financial support of the ERC Starting Grant #759424. Early versions of this paper have been presented at seminars at Princeton, MIT-Harvard, Northwestern, Columbia, Minnesota, Carnegie-Mellon, UCLA, Boston University, Bocconi, UPF, and many conferences. We are grateful to those audiences for the invaluable comments. Special thanks go to Jose Apesteguia, Andrew Caplin, Gabriel Carroll, Sylvain Chassang, Bart Lipman, Fabio Maccheroni, George Mailath, Paola Manzini, Marco Mariotti, Laurent Mathevet, Efe Ok, Pietro Ortoleva, Andy Postlewaite, Aldo Rustichini, Bill Sandholm, Alvaro Sandroni and Marciano Siniscalchi.

[†]Universitat Pompeu Fabra and Barcelona GSE. E-mail: larbi.alaoui@upf.edu

[‡]ICREA, Universitat Pompeu Fabra and Barcelona GSE. E-mail: antonio.penta@upf.edu

1 Introduction

Making a difficult choice or approaching a hard problem requires thought and introspection. In proving a theorem, all of logic is at our disposal, and yet we rarely reach the solution instantly or are even certain of the truth of the proposition – instead, we reason in steps, pushing the frontier of our understanding, sometimes in leaps of insight and sometimes in small increments, and sometimes never arriving to the solution. Similarly, choosing between complex options often leads people to resort to a deliberation process that relies on heuristics. The reasoning process does not involve receiving new data, but, because of cognitive limitations and failures of logical omniscience, fact-free learning can occur (cf. Aragonés, Gilboa, Postlewaite and Schmeidler (2005)). Furthermore, individuals often stop thinking before achieving full understanding. This need not imply they have reached an absolute cognitive limit. An increase in the academic or financial value of proving a theorem may well provide the researcher with the incentives to reason further. Similarly, higher stakes for making a choice could lead the consumer to be more deliberative. This suggests that individuals trade off the cognitive cost against some possibly vague notion that reasoning is worthwhile. Hence, even though fact-free learning (or ‘reasoning’) spans a wide range of contexts, including computational complexity (e.g., Aragonés et al. (2005) and Gabaix and Laibson (2005)), unawareness (e.g., Schipper (2015)), bounded rationality in games (see Camerer (2003) and Crawford, Costa-Gomes and Iriberry (2013)) and sequential heuristics in choice (e.g., Manzini and Mariotti (2007)), the underlying features are the same. Reasoning is a stepwise process, it is cognitively costly, and individuals may find it worthwhile.

The first instinct to an economist is then to cast the decision to reason further as a cost-benefit tradeoff. But this approach can be problematic, and is often cautioned against in psychology and cognitive science. For instance, there are many well-documented cases in both psychology and economics in which stakes increase but in which performance worsens, which is at odds with basic cost-benefit intuition (e.g., Gneezy and Rustichini (2001)). Moreover, the subjective value that an individual attaches to reasoning is not well-defined: it is not clear how an individual should assess the benefit of reasoning further when he has no sense of what it might bring, especially when he has only a partial understanding of the problem. It is also not obvious how to capture a researcher’s outlook about what thinking more about a proof will lead to. Understanding the relation between incentives and cognition is clearly important for many economic settings, but little is known about when a cost-benefit analysis is justified, about the form it should take, or about which properties of the reasoning process must fail when it does not hold.

These are entirely new questions. To address them, we develop a framework to model the reasoning process and identify the core properties that are necessary and sufficient for reasoning to be captured by a cost-benefit analysis. We find that they are weak and intuitive, suggesting that there is indeed justification for using cost-benefit models for a

large class of problems. But our conditions also identify the limits of the approach: in contexts for which these properties do not seem to hold, reasoning would not follow a cost-benefit criterion. We also provide conditions that give structure to the value of reasoning function, including a ‘value of information’ representation, often assumed in the literature (e.g., Caplin and Dean (2015)) as well as a novel ‘maximum gain’ representation, which is particularly useful for applications.

We model an iterated reasoning process as a machine in which the agent’s reasoning transitions from mental state to mental state. A mental state encodes both the agent’s current understanding of the problem at hand, and his outlook towards what performing an additional step of reasoning may bring. Specifically, each mental state is summarized by two objects: an action that he finds best, given his current understanding, and a preference relation pair. This pair consists of mental preferences, which capture the agent’s view of his understanding at that stage and what he expects to learn, and behavioral preferences, which represents his choice of how deeply to think. We then formalize properties of the reasoning process as restrictions over mental and behavioral preferences.

The key properties in our model can be informally summarized in a simple way: the agent reasons only if it is relevant to his choice; at any mental state, when reasoning has no instrumental value, the agent’s reluctance to think is identical for two problems that only differ by a constant payoff; and he should not have a strict preference for committing to disregard what might be learned. Since these properties are all intuitive, our results show that a cost-benefit approach can be a useful ‘as if’ model for various reasoning processes, thereby justifying the extension of economic methodology to these novel domains. We also discuss additional properties and what they imply for the shape of the value of reasoning function. Under some conditions, the value function takes a ‘maximum gain’ shape – it is as if the value attached to reasoning corresponds to the highest payoff improvement the agent could obtain, relative to his current most-preferred action. Under alternative conditions, it is as if the agent’s outlook is not as extreme, and more measured in what reasoning might bring. For instance, the value function can take a ‘value of information’ form, for which it is as if, at each mental state, the agent has subjective beliefs over the outcome of the next step of reasoning, and the value equals the expected gain of switching from the current action to the future best response. We discuss how these conditions, which relate current understanding and outlook towards future reasoning, might seem plausible in some settings but not in others.

The advantage of a cost-benefit representation is that it allows for comparative statics exercises on the depth of reasoning. But for these to be meaningful, there must be a way of shifting the value without changing the cost or the process itself. We therefore introduce a notion of ‘cognitive equivalence class’, which groups together problems that are equally difficult and approached essentially in the same way. Our representation theorems imply that the costs are uniquely pinned down within each class. This ensures that cognitively equivalent problems only differ in the value of reasoning they entail, thereby allowing for

useful comparative statics on the depth of reasoning. We also obtain a strong uniqueness result in our representations, for both the cost and value functions.

This framework can be applied to diverse theoretical and empirical settings in which cognition is costly. As theoretical applications, we illustrate how the model can be used in domains of research and innovation, to inform a rigorous discussion on alternative cost-benefit criteria for the decision of whether to continue or stop researching, which is often made based on heuristics that are only partially understood. To explain how the model can be taken to the data, and test its joint implications, first we consider an extension to a model of endogenous level- k reasoning (Alaoui and Penta, 2016a), which has strong experimental support and is predictive of well-known behavioral anomalies (see also Alos-Ferrer and Buckenmaier (2018)). Then, we provide an extension of the model to connect our notion of depth of reasoning with measures of response time. We also develop an application to a recent experiment on attention allocation by Avoyan and Schotter (2016), and show that their experimental findings are all consistent with our predictions.

Our analysis is useful not only for understanding the scope of the cost-benefit approach, but also for exploring its limits. In particular, there is a growing literature in economics on contexts where these limits are apparent. These include settings in which individuals display ‘thinking aversion’ (Ortoleva (2013)), and in tasks for which higher financial rewards are sometimes detrimental to performance (e.g. Gneezy and Rustichini (2000) or Camerer and Hogarth (1999)). A number of mechanisms have been proposed in the psychology literature to explain this observed behavior, such as different forms of choking under pressure, anxiety, cognitive overload or distractions. These mechanisms are often intertwined and difficult to disentangle theoretically and empirically. Our model allows a clear separation of these theories, which we show by mapping each mechanism to specific violations of the core properties. Identifying these properties can then serve as a basis for both theoretical and empirical work. Theoretically, the implications of relaxing a specific property can be explored. Empirically, our approach can guide the design of experiments to distinguish the testable implications of competing explanations.

2 The Model of Reasoning

Consider the problem of an agent who chooses from a finite set A of acts taking values in a set of results R and which depend on the realization of a state $\omega \in \Omega$. The realized state is unknown to the agent, but he may reason about it. We model such reasoning processes, where we use the term ‘reasoning’ in the broad sense of general sequential procedures of decision making, including introspection and information acquisition. To fix ideas, however, in this section we adhere to the narrower interpretation of reasoning as a mental process in the strict sense.

2.1 Choice Problems and Mental States

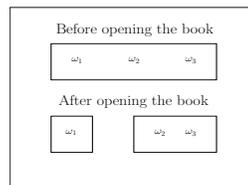
We first define some basic objects. For simplicity, we assume that $R = \mathbb{R}$, and that results are scaled so that acts pay out in utils.¹ Hence, each act $a \in A$ can be written as a function $u(a, \cdot) : \Omega \rightarrow \mathbb{R}$, and a menu A of acts is represented by a utility function $u : A \times \Omega \rightarrow \mathbb{R}$, which the agent seeks to maximize. We let $\mathcal{U} := \mathbb{R}^{|A \times \Omega|}$ denote the set of menu of acts, and refer to $a \in A$ as *actions* or alternatives. Elements of $A \times \Omega$ are referred to as *outcomes*.

Notation: For any $u, v \in \mathcal{U}$ and $\alpha \in [0, 1]$, we denote by $\alpha u + (1 - \alpha)v \in \mathcal{U}$ the utility function that pays $\alpha u(a, \omega) + (1 - \alpha)v(a, \omega)$ for each $(a, \omega) \in A \times \Omega$. It will sometimes be useful to think of utility functions as vectors in $\mathbb{R}^{|A \times \Omega|}$, and write functions that pay $t \in \mathbb{R}$ for all outcomes as $t \cdot \mathbf{1}$, where $\mathbf{1}$ denotes the unit vector. For $t \in \mathbb{R}$, we will abuse notation slightly and write $u + t$ to mean $u + t \cdot \mathbf{1}$. We adopt the following conventions for vectors: for any $x, y \in \mathbb{R}^n$, we let $x > y$ denote the case in which x is weakly larger than y for each component, but not $x = y$; $x \geq y$ means $x > y$ or $x = y$; we let $x \gg y$ denote strict inequality for each component. Finally, for any set X , we let $\Delta(X)$ denote the set of simple (i.e., finite support) probability measures on X .

We provide an example of how standard information can be cast within our framework, although our model need not be limited to information.

Example 1 (Multiple Choice Question) Suppose that the agent answers a multiple choice math question, and obtains utility payoff of 1 if he is correct and 0 otherwise. He chooses between three options, $A = \{a_1, a_2, a_3\}$, and does not know at first which is the correct answer. The agent’s subjective uncertainty is captured by the states $\Omega = \{\omega_1, \omega_2, \omega_3\}$, where a_i is correct in state ω_i . Then, u is described by the following matrix on the left:

	ω_1	ω_2	ω_3
a_1	1	0	0
a_2	0	1	0
a_3	0	0	1



The agent reasons sequentially; we will refer to his steps of reasoning as ‘mental states’ and formalize them below. Suppose that his prior beliefs are $1/2$ of the state being ω_1 , $1/4$ being ω_2 , and $1/8$ being ω_3 . Thinking further could take the form of physically opening a book that reveals whether the true state is ω_1 or not (see the figure above on the right), and nothing more. Alternatively, his reasoning could be introspective and fact-free, but take the same partitional form of confirming or discarding ω_1 . Then, if Bayesian rational, he expects that performing one more step would lead him to choose a_1 if the state is ω_1 and a_2 otherwise. If this information comes at a subjective cognitive cost (e.g., the

¹We abstract away from eliciting the utility function from primitive preferences, which can be done in the classical way and is orthogonal to our objectives.

cost of reading the book, or thinking about the problem), then he may acquire it or not, depending on the magnitude of this cost of cognition. \square

When the reasoning process consists of actual information acquisition (as in Ex. 1), then its unfolding depends not only on u , but also on the *true state* of the world, $\hat{\omega} \in \Omega$. In general, other elements may affect the path of reasoning too, such as framing effects in a choice problem, the opponents' payoffs in a strategic setting, and so on. An *environment* E denotes a full list of all such elements which may affect the agent's reasoning – including the true state $\hat{\omega}$ – and \mathcal{E} denotes the set of possible environments.² A choice problem (u, E) consists of both the menu of acts and the environment. In all the results below, however, E is kept constant, and the representation theorems deliver comparative statics on the agent's reasoning as the payoffs alone are varied. We refer to E nonetheless for conceptual clarity, for instance to accommodate the dependence of an information acquisition process on the true state of nature, which is embedded in E .

Each step in a reasoning process leads to a new 'mental state', meaning a new understanding of the problem and a new outlook over what he might learn. We let S denote the set of all possible mental states, with typical element s . A *reasoning process* is a sequence of mental states $(s^k)_{k \in \mathbb{N}}$ that could potentially be reached. Each mental state $s \in S$ comprises two objects. The first is an action $a^s \in A$ that the agent considers 'best' at s : if the agent had to stop thinking at s , then he would choose a^s . In the example above, at the first step of reasoning $a^1 = a_1$. At the next mental state, it remains a_1 if the true state is ω_1 , otherwise it is a_2 . The second object is a preference relation pair (\succsim_s, \succeq_s) , which we explain below.

For any $(u, E) \in \mathcal{U} \times \mathcal{E}$, we let $\pi(u, E) = (s^k)_{k \in \mathbb{N}}$ denote the reasoning process induced by (u, E) , and let $S(u, E)$ denote the set of states in $\pi(u, E)$. Thus, the function $\pi : \mathcal{U} \times \mathcal{E} \rightarrow S^{\mathbb{N}}$ formalizes the idea that an agent may reason differently in different problems, or in the same problem at a different state $\hat{\omega}$ (the true state is embedded in E).

2.2 Reasoning Preferences

We describe next the elements of the preference relation pair, which are the *behavioral preferences* \succeq_s , and the *mental preferences* \succsim_s .

Behavioral preferences. For every choice problem (u, E) and at every mental state $s \in S(u, E)$, the agent has preferences \succeq_s over whether to stop reasoning and make his choice, given the understanding attained, or to take one extra step. Moreover, we assume that the agent can be compensated with a reward $t \in \mathbb{R}$ for further pursuing his reasoning,

²An environment E encompasses all the information that the agent has about the mechanism that generates ω , although he may not understand it all. For instance, in the problem of proving a theorem, all the required information is at the agent's disposal, but cognitively constrained agents would not instantly understand all the implications. In the language of the unawareness literature (e.g., Modica and Rustichini (1999), Dekel, Lipman and Rustichini (1998, 2001), and Schipper (2015) for a survey), (u, E) is sufficient for the agent's *implicit knowledge*, whereas his *explicit knowledge* is the outcome of the reasoning process.

and that he has preferences over such rewards. Formally, for any (u, E) and $s \in S(u, E)$, \succsim_s is a binary relation which ranks any element of the set $\{(u', 1) : \exists t \in \mathbb{R} \text{ s.t. } u' = u + t \cdot 1\}$ relative to $(u, 0)$: For any $t \in \mathbb{R}$ we write $(u + t, 1) \succsim_s (u, 0)$ if, at mental state s , and given payoff function u , the agent weakly prefers to perform the next step of reasoning, with ‘extra reward’ t , rather than to stop reasoning and choose a^s . (Symmetrically, $(u, 0) \succsim_s (u + t, 1)$ means that the agent prefers to stop reasoning and make his choice rather than think and receive the reward.) Reward $t = 0$ corresponds to the simple choice over whether to think or not.

Behavioral preferences can be elicited directly. In the multiple choice example, for instance, the agent can be asked whether or not he wishes to read the book before giving the answer, possibly compensated with an extra reward $t \in \mathbb{R}$. For choice problems in which learning is not observed in as simple a manner, a similar exercise can be conducted: at any step s , the agent can be asked whether he prefers to keep thinking $((u, 1) \succeq_s (u, 0))$ or not $((u, 0) \succeq_s (u, 1))$, and similarly if $t \neq 0$.

Mental preferences. The agent’s choice of thinking or not, i.e., his behavioral preferences \succeq_s , rely on what he believes to have understood about the problem (his current understanding), and what he believes he might understand if he were to think another step (his outlook over his future understanding). The behavioral preferences defined above do not capture the full richness of this process. For this reason, we define mental preferences, which are a complete and transitive binary relation \succcurlyeq_s over $\mathcal{U} \times \{0, 1\}$, with asymmetric and symmetric parts \succ_s and $\dot{\succ}_s$, respectively. These mental preferences serve as an input into the decision to think further or not and will be linked to the behavioral preferences.

The projection of \succcurlyeq_s on $\mathcal{U} \times \{0\}$ and $\mathcal{U} \times \{1\}$, denoted by \succcurlyeq_s^0 and \succcurlyeq_s^1 , represent respectively the agent’s *current understanding* and his *outlook over his future understanding*. To understand how \succcurlyeq_s^0 can express the former, consider an agent at mental state $s \in S(u, E)$. If his reasoning stops or is interrupted at that step, then he chooses action a^s . Now suppose that this agent is asked whether, given his choice, he prefers to receive an extra transfer of utils according to a *side transfer* $\tau' : A \times \Omega \rightarrow \mathbb{R}$ or $\tau'' : A \times \Omega \rightarrow \mathbb{R}$. Intuitively, since the agent has already chosen a^s , he should be indifferent between the two if $\tau'(a^s, \omega) = \tau''(a^s, \omega)$ for all ω . But if $\tau'(a^s, \omega) \neq \tau''(a^s, \omega)$ for some ω , and the agent strictly prefers one over the other, then we can infer which states he regards as more likely at s , in a manner analogous to the approach in De Finetti (1937) (see ex.2 below).³

Note that, in this thought experiment, the agent is effectively being asked whether he prefers to be paid according to $u', u'' \in \mathcal{U}$, where $u' = u + \tau'$ and $u'' = u + \tau''$. In the following, it will be convenient to work with the ‘total payoffs’ formulation, with no explicit reference to transfers. For any $u', u'' \in \mathcal{U}$, we thus write $u' \succcurlyeq_s^0 u''$ if and only if transfer $\tau' = u' - u$ is preferred to $\tau'' = u'' - u$ in the thought experiment above.

³De Finetti (1937) uses preferences over acts with monetary payoffs to elicit subjective probabilities over states. Similarly, our preferences embed the agent’s assessment, probabilistic or not, over outcomes (a, ω) .

Example 2 (Multiple Choice Question (continued)) Continuing with Example 1, the agent's reasoning process $\pi(u, E) = (s^1, s^2, s^3, \dots)$ starts at an initial mental state s^1 . Consider his current understanding, \succcurlyeq_s^0 . If he must choose without thinking further, then his choice is observed to be $a^1 = a_1$. Since the agent's choice is already made and the payoffs for action a_1 are not changed at any state, the agent should be indifferent between the two; i.e., $u \doteq_s^0 u'$. We will assume properties of this type.

	ω_1	ω_2	ω_3
a_1	1	0	0
a_2	0	1	0
a_3	0	0	1

 $u =$

	ω_1	ω_2	ω_3
a_1	1	0	0
a_2	1	0	0
a_3	1	0	0

 $u' =$

	ω_1	ω_2	ω_3
a_1	1/2	1/2	1/2
a_2	0	1	0
a_3	0	0	1

 $u'' =$

To see how these comparisons may be used to elicit beliefs, suppose that the agent's mental preferences are such that $u \doteq_s^1 u''$, where u and u'' are as in the table above. Under some assumptions, this indifference could allow us to infer that his beliefs at mental state s^1 are that ω_1 has likelihood 1/2. \square

The \succcurlyeq_s^1 relation is analogous to \succcurlyeq_s^0 , but used to represent the agent's outlook over his *future understanding*. The only difference is that the agent has not yet made his choice. In essence, while \succcurlyeq_s^0 is a ranking of different problems in terms of how well-understood they are in the current state, \succcurlyeq_s^1 is an *anticipated* ranking of problems in terms of how well-understood they will be if the individual decides to think further. The \succcurlyeq_s^1 relation can be used to elicit his beliefs not only about the state, but also about his future choice. In the previous example, for instance, a Bayesian agent at state s^1 would not be indifferent between receiving u or u' *after* having opened the book. This is because he thinks it possible that he might learn that the correct state is not ω_1 . In that case, he would take action a_2 , which leads to a higher payoff, and hence we would have $u \succ_s^1 u'$. Conversely, preferences $u \succ_s^1 u'$ indicate to an outside observer that the agent is uncertain over whether he will choose a_1 at the next step, otherwise his preferences would be $u \doteq_s^1 u'$.

Thus, the two preferences \succcurlyeq_s^0 and \succcurlyeq_s^1 will in general have different properties, but they share the same logic. Jointly, they provide a window into understanding the rationales behind the agent's choice to think further or to stop, which involves the 0-1 comparisons, i.e. the agent's comparison of whether or not to reason another step. The comparison between $(u, 1)$ and $(u, 0)$ in the example is straightforward, as no side transfers are needed. But suppose instead that we wish to elicit the comparison of $(u', 1)$ and $(u, 0)$. Then, the side transfers $\tau' = u' - u$ would be added after the agent has reasoned about the problem. In other words, the agent is asked to think about whether he would open the book or not for exactly the same multiple choice question, but with the understanding that he would then receive side transfers if he does. While there is no conceptual issue with this idea, its implementation may seem more challenging, since in principle side transfers may change the agent's reasoning process. The elicitation method provided in Section 2.5, however, resolves this apparent difficulty.

As already mentioned, the *behavioral preferences* represent the agent’s choice of whether to stop thinking or not. The *mental preferences* describe the agent’s current and future understanding, and in that sense they are best thought of as a representation of the criteria which inform that choice. We connect the two with by the following condition, which we maintain throughout:

Condition 1 (Reasoning Consistency) *For every (u, E) and for every $s \in S(u, E)$, $(u + t, 1) \succ_s (u, 0)$ if and only if $(u + t, 1) \succeq_s (u, 0)$.*

This condition formalizes the idea that the criteria expressed by the mental preferences effectively guide the actual choice to think further. Formally, this implies that the binary relation \succ_s is a subset of \succeq_s , and hence the distinction may seem unnecessary. Conceptually, however, the two objects are distinct. First, because agents may sometimes have a compulsion to keep reasoning even though they find it no longer beneficial. Such psychological phenomena can be thought of precisely as an inconsistency between mental and behavioral preferences (Section 5), akin to other phenomena of inconsistency in choice. The second reason is that one could imagine extensions of the model to express richer criteria behind an agent’s decision to reason. Being mainly interested in modeling boundedly rational agents, we consider decisions based on at most one-step ahead criteria, and hence we take $\mathcal{U} \times \{0, 1\}$ as the space of mental preferences. But our framework could be extended to accommodate more sophisticated forward looking agents, or agents whose reasoning criteria are affected by other factors. To express such richer criteria, one would need to extend the space of mental preferences, and different reasoning processes (e.g., with different degrees of myopia) would correspond to different ways of connecting mental and behavioral preferences. Spelling out the conceptual distinction between the two binary relations therefore provides a better guidance for future extensions of the model.

2.3 Reasoning Processes and Cognitive Equivalence

The model can accommodate diverse reasoning processes, and the same agent may reason about different problems in very different ways. In those cases, comparing depth of reasoning is of unclear significance. We expect, however, that the agent approaches cognitively similar problems using the same form of reasoning. For this reason, we introduce a notion of cognitive equivalence, which is a partition on $\mathcal{U} \times \mathcal{E}$ that groups together choice problems that the agent approaches in essentially the same way. The basic notion of mental states suggests a natural criterion to define this cognitive partition:

Definition 1 (Cognitive Equivalence) *Two mental states $s, \hat{s} \in S$ are cognitive equivalent (c.e.) if $a^s = a^{\hat{s}}$ and $\succeq_s = \succeq_{\hat{s}}$. Two choice problems (u, E) and (\hat{u}, \hat{E}) are cognitively equivalent if they induce sequences of pairwise c.e. states.*

We let \mathcal{C} denote the *cognitive partition* on $\mathcal{U} \times \mathcal{E}$ induced by the cognitive equivalence relation, and let $C \in \mathcal{C}$ denote a generic cognitive equivalence class. We also let $C(u, E)$

denote the equivalence class that contains (u, E) . Hence, by Definition 1, any (u, E) and $(\hat{u}, \hat{E}) \in C(u, E)$ induce sequences $(s^k)_{k \in \mathbb{N}}$ and $(\hat{s}^k)_{k \in \mathbb{N}}$ that correspond to the same action a^k and mental preferences \succsim_k for each k . Since we maintain Condition 1 throughout, with slight abuse of notation we also use the notation \succeq_k in the same manner.

Intuitively, choice problems that are in the same cognitive equivalence class are equally difficult to think about. In all representation theorems that follow, cognitively equivalent problems have the same associated cost of reasoning at any step s . As an illustration, consider Example 1, and let (u', E) consist of the same math problem with all the utils increased by 1 (that is, $u'(a_i, w_j) = 2$ if $i = j$ and 1 otherwise), or (u'', E) , where all utils are multiplied by 10 ($u''(a_i, w_j) = 10$ if $i = j$ and 0 otherwise). Aside from the higher payoffs, these problems are exactly the same and they are framed the same way. Note also that since all payoffs are in utils, there are no changing risk considerations to account for. We therefore expect problems (u, E) , (u', E) and (u'', E) to be cognitively equivalent. This is made explicit in our next condition, which contains the main substantive restrictions imposed on the determinants of the reasoning process:

Condition 2 *For any E , if u and v are such that, for some $\alpha \in \mathbb{R}_+$ and $\beta : \omega \rightarrow \mathbb{R}$, $v(a, \omega) = \alpha u(a, \omega) + \beta(\omega)$ for all (a, ω) , then $(v, E) \in C(u, E)$.*

This condition restricts the dependence of the reasoning process on the payoff function. It says that, holding E fixed, positive affine transformation of payoffs do not change the c.e. class. This condition therefore implies that c.e. classes cannot be too small.

2.4 Cost-Benefit Representation of Depth of Reasoning

We now define what we mean by a cost-benefit representation.

Definition 2 (Cost-Benefit Representation) *The agent's reasoning admits a cost-benefit representation if, for any cognitive equivalence class $C \in \mathcal{C}$, there exist cost and value functions $c : \mathbb{N} \rightarrow \mathbb{R}_+ \cup \{\infty\}$ and $V : \mathcal{U} \times \mathbb{N} \rightarrow \mathbb{R}_+$ such that, for any $k \in \mathbb{N}$, for any $(u, E) \in C$, and $t \in \mathbb{R}$: $(u + t, 1) \succsim_{k-1} (u, 0)$ if and only if $V(u, k) + t \geq c(k)$, and $V(u', k) = 0$ whenever u' is constant in a .*

We say that such a representation is unique if, for any cognitive equivalence class $C \in \mathcal{C}$, whenever (V, c) and (V', c') represent the agent's depth of reasoning, then $c = c'$ and $V(\cdot, k) = V'(\cdot, k)$ for all k such that $c(k) \neq \infty$.

In words, for any choice problem (u, E) , the agent keeps thinking as long as the value of reasoning $V(u, k)$ at step k is higher than the cost $c(k)$. In addition, if the agent receives an extra t to think about u , then our representation linearly separates $V(u, k)$ from t , so that the agent reasons if and only if $V(u, k) + t > c(k)$.

The separation between the V and c functions is what makes the representation meaningful. The role of the cognitive equivalence classes (Def. 1) is that they allow for simple

comparative statics. Specifically, the representation maintains that the cost function is the same for any choice problems (u, E) and (u', E') within a c.e. class, while the value of reasoning need not be. That is, cognitively equivalent problems may only differ in the value of reasoning they entail. The manner in which the value of reasoning may vary with u will be discussed in Section 3, but note that the value of reasoning in the definition above reaches its minimum of zero for payoff functions which are constant in the action. This is because our focus is on consequentialist reasoning: for such problems, reasoning is not instrumental to decision-making, and so it has no value.

Finally, note that the notion of uniqueness provided in Definition 2 is the strongest possible in this context: if the cost of reasoning is infinite at some step, then the agent would choose not to reason for any specification of the value of reasoning.

2.5 Dataset, elicitation and testability

Our framework and axioms serve two purposes. The first is to understand at a foundational level what cost-benefit analysis in reasoning entails, and what the different representations mean in terms of the fundamental properties of reasoning. The second concerns testability of these properties. In this subsection, we summarize the requirements on the datasets needed to test our axioms and elicitation methods for the reasoning preferences.

The data we require are as follows. We assume that the agent's menu of choices over which he is thinking is observed (recall that we are abstracting from the elicitation of utils, which we take to be achieved in the standard way). From there, the agent's reasoning can be interrupted at a step. One way to proceed at this point is with the side-transfer method we considered above, which implicitly assumes that side transfers can be made conditional on: (i) the agent's current reasoning, to obtain preferences \succcurlyeq_s^0 ; (ii) his future reasoning, to obtain preferences \succcurlyeq_s^1 ; and (iii) reasoning one more step or not, to obtain the 0-1 comparisons.

While the side-transfer method is conceptually sound, especially as a way to express fundamental properties of reasoning, for elicitation purposes it may be difficult to operationalize asking the agent to think about u' and u'' from the perspective of mental state $s \in S(u, E)$ – that is, through the reasoning he would use to think about (u, E) – so as to elicit \succcurlyeq_s^1 or the 0-1 comparisons. For those cases, the following lottery method circumvents the issue. Suppose that the agent is interrupted at mental state s . To elicit \succcurlyeq_s^1 , the following procedure is used. The agent is told that if he were to reason an extra step, he will receive (i) u with probability $1 - \epsilon$ and u' with probability ϵ , or (ii) u with probability $1 - \epsilon$ and u'' with probability ϵ . For sufficiently low ϵ , the probability of receiving u' or u'' should be negligible enough that the agent would not change how he thinks about the problem, and maintains the reasoning as if he were thinking about u . At the same time, since he would receive u in either case, the choice between (i) and (ii) only differs in the comparison of u' and u'' (using the reasoning associated with u), and so a preference for (i) over (ii) indicates that $u' \succcurlyeq_s^1 u''$. This method is akin to commonly used experimental

procedures such as the widely used Multiple Price List method (e.g., Kahneman, Knetsch and Thaler (1990), Holt and Laury (2002), Andersen et al. (2006), Azrieli et al., (2018)), which relies on making choices with the knowledge that only a subset of the choices will be used for actual payments.⁴

When used to elicit 0-1 comparisons, the lottery method requires a subtle adjustment. Now, in option (i*) the subject would have to perform an extra step of reasoning. Then, after having reasoned, with probability $1 - \epsilon$ he would receive u , while with probability ϵ he would receive u' . In option (ii*) he commits to performing an extra step about u with probability $1 - \epsilon$, and to *not* performing and receiving u'' – but with his chosen action a^s being already locked in – with probability ϵ . As for the elicitation used for \succsim_s^1 , for low enough ϵ , he would reason in (i*) about the problem as if he were reasoning about u . And as before, since he reasons about u with probability $1 - \epsilon$ in both cases, the difference is in the ϵ event of reasoning about u'' compared to the ϵ event not reasoning and receiving u' . Hence, option (i*) should be preferred to option (ii*) if and only if $(u'', 1) \succsim (u', 0)$.

3 Properties of the Reasoning Process and Representations

In this section we introduce the properties of the reasoning process – formally, axioms on the mental preference relation \succsim_s – and provide various representation theorems. Conditions 1 and 2 will be maintained throughout, as well as the assumption that \succsim_s is complete and transitive. In Section 3.1 we focus on the core properties that characterize the possibility of a cost-benefit representation. Section 3.2 provides an ‘expected value’ representation of the value of reasoning. The important special case of the ‘value of information’ representation is presented in Section 3.3. Section 3.4 introduces a tractable ‘maximum gain’ representation. As will be discussed, this representation is *not* consistent with the value of information representation, but it is nonetheless plausible and particularly convenient for applications. Further variations and extensions are discussed in Section 4.1.

3.1 Core Properties and Representations

The first three properties are standard monotonicity and continuity properties:

Property 1 (Monotonicity) *For each $s \in S$, and $x = 0, 1$: If $u \geq v$, then $u \succsim_s^x v$. If $u \gg v$ then $u \succ_s^x v$.*

Property 2 (Archimedean) *For each $s \in S$, $x = 0, 1$ and $v, u, w \in \mathcal{U}$ such that $u \succ_s^x v \succ_s^x w$, there exist $0 \leq \beta \leq \alpha \leq 1$ such that $\alpha u + (1 - \alpha)w \succ_s^x v \succ_s^x \beta u + (1 - \beta)w$.*

Property 3 (Continuity) *For each $s \in S$ and $u \in \mathcal{U}$, if there exists $t \in \mathbb{R}$ s.t. $(u + t, 1) \succsim_s (u, 0)$, then there exists $t^* \leq t$ such that $(u + t^*, 1) \dot{\succ}_s (u, 0)$.*

⁴The formalization of this method, along with discussions of the assumptions required, is available upon request.

Since payoffs are already scaled in utils, the first two properties are natural in this context. The third property is also particularly weak. For instance, it allows lexicographic preferences such as $(u, 0) \succ_s (u + t, 1)$ for all u and t , which means that the agent cannot be incentivized to perform the next step of reasoning, no matter how high t is.

The next property defines the scope of our theory. Part (S.1) states that the reasoning process is purely instrumental in informing the player's choice. Thus, if i 's payoffs are constant in his own action, the agent would never strictly prefer to think harder. Part (S.2) pushes the idea further, requiring that the incentives to reason are completely driven by the payoff differences between actions: if u and v are such that $u(a, \omega) - u(a', \omega) = v(a, \omega) - v(a', \omega)$ for all a, a' and ω , then u and v provide the agent with the same incentives to reason.

Property 4 (Scope) *For each $s \in S$:*

S.1 If u is constant in a , then $(u, 0) \succcurlyeq_s (u, 1)$.

S.2 If $u - v$ is constant in a , then $(u, 0) \succcurlyeq_s (u, 1)$ if and only if $(v, 0) \succcurlyeq_s (v, 1)$.

An agent who is observed strictly preferring to think further for a constant u is in violation of S.1 (see Section 5.2). For S.2, notice that since $u - v$ is constant in a , (u, E) and (v, E) are cognitively equivalent for any environment E (Condition 2). Hence, Axiom S.2 will be violated whenever the agent thinks further in one of these problems than in the other. For instance, suppose, that in the multiple choice example (Example 1), for payoff function u the agent receives 1 util if the answer is correct and 0 otherwise, while in v he receives 2 utils if the answer is correct and 1 otherwise. If an agent thinks further in v than in u (or vice versa), then S.2 is violated.

The next axiom is a ‘cost-independence’ condition, and will be used as a calibration axiom to pin down the cost of reasoning:

Property 5 (Cost-Independence) *For each $s \in S$, for any u, v that are constant in a , and for any $t \in \mathbb{R}$, $(u + t, 1) \dot{\succ}_s (u, 0)$ if and only if $(v + t, 1) \dot{\succ}_s (v, 0)$.*

To understand this property, first note that (S.1) implies that $(u, 0) \succcurlyeq_s (u, 1)$ whenever u is constant in a , since thinking has no instrumental value. Suppose the agent can be made indifferent between thinking and not, if reasoning were accompanied by an extra reward $t \geq 0$ (that is, $(u + t, 1) \dot{\succ}_s (u, 0)$). Then, that reward would also make him indifferent between reasoning or not given any other constant function. This condition is weak because it is only required for constant payoff functions. Yet, together with the other axioms, it will ensure a separate identification of the cost of reasoning for all choice problems, including those with non-constant payoff functions. Note that there may not be a reward t that can induce the agent to think, as would be the case for an absolute cognitive bound at mental state s .

For each $u \in \mathcal{U}$ and k , let u^k be such that $u^k(a, \omega) = u(a^k, \omega)$ for all (a, ω) , and similarly define u^s such that $u^s(a, \omega) = u(a^s, \omega)$ for all (a, ω) . For instance, in the multiple choice example with matrix u replicated below, if $a^s = a^k = a_1$, then

$$u = \begin{array}{c|ccc} & \omega_1 & \omega_2 & \omega_3 \\ \hline a_1 & 1 & 0 & 0 \\ a_2 & 0 & 1 & 0 \\ a_3 & 0 & 0 & 1 \end{array} \quad u^s = u^k = \begin{array}{c|ccc} & \omega_1 & \omega_2 & \omega_3 \\ \hline a_1 & 1 & 0 & 0 \\ a_2 & 1 & 0 & 0 \\ a_3 & 1 & 0 & 0 \end{array} .$$

The next property states that, given that the agent would choose a^s if he stopped reasoning at s , at that mental state he is indifferent between the original payoff function u and u^s . We thus call this property ‘consequentialism’:

Property 6 (Consequentialism) *For any $s \in S$ and for any u : $u \stackrel{0}{\cong}_s u^s$.*

Any agent who only cares about his final payoffs should satisfy this property. Given that he chooses a^s , he should be indifferent between u and u^s , because the payoffs for the unchosen actions are irrelevant.

The next property requires that, for any problem and for any mental state reached by the corresponding reasoning process, the agent does not strictly prefer to commit to disregarding what might be learned:

Property 7 (No Improving Obstinance) *For any (u, E) and $s \in S(u, E)$: $u \not\geq_s^1 u^s$.*

As will be illustrated in Example 3, this is an important property for a cost-benefit representation in that it guarantees that the value of reasoning is non-negative. If Property 7 is violated then at some mental state s the agent feels that, were he to think further, he would be better off committing to action a^s (equivalently: being rewarded by u^s) rather than following his future understanding. In this sense, violations of this property entail a preference for committing to disregard what might be learned.

Any standard model would trivially satisfy this property. Violations of this axiom entail an extreme lack of confidence in one’s own reasoning process, and may be associated to phenomena of ‘thinking aversion’. For instance, it may be due to fear of being unable to ignore what is learned when this knowledge is not beneficial. These violations are discussed in Section 5.2.

It will be convenient to introduce the following ‘payoff differences’ notation: for any $u \in \mathcal{U}$ and $\hat{a} \in A$, let $D(u, \hat{a}) : A \times \Omega \rightarrow \mathbb{R}$ be defined as

$$D(u, \hat{a})(a, \omega) = u(a, \omega) - u(\hat{a}, \omega) \text{ for all } (a, \omega) \in A \times \Omega. \quad (1)$$

In words, given u and \hat{a} , and for each $(a, \omega) \in A \times \Omega$, $D(u, \hat{a})(a, \omega)$ represents how much the agent gains if he chooses a over \hat{a} in state ω . For instance, considering again the same u from the multiple choice example, then $D(u, a_1)$ is as follows.

$$u = \begin{array}{c|ccc} & \omega_1 & \omega_2 & \omega_3 \\ \hline a_1 & 1 & 0 & 0 \\ a_2 & 0 & 1 & 0 \\ a_3 & 0 & 0 & 1 \end{array} D(u, a_1) = \begin{array}{c|ccc} & \omega_1 & \omega_2 & \omega_3 \\ \hline a_1 & 0 & 0 & 0 \\ a_2 & -1 & 1 & 0 \\ a_3 & -1 & 0 & 1 \end{array}$$

Note that, for any $\hat{a} \in A$, $D(u, \hat{a}) = \mathbf{0}$ if and only if u is constant in \hat{a} . Hence, for $\hat{a} = a^s$, $D(u, a^s) : A \times \Omega \rightarrow \mathbb{R}$ can be seen as an array of all the potential gains (losses, if the entry is negative) from switching from the current choice a^s to some other a , at every ω .

Theorem 1 (Core Representation) *The agent's reasoning satisfies Properties 1-7 if and only if, for any cognitive equivalence class $C \in \mathcal{C}$:*

1. *There exists a unique cost-benefit representation (c, V) of the agent's depth of reasoning. Moreover, for any $k \in \mathbb{N}_+$, $V(u, k) \geq V(u', k)$ if $D(u, a^{k-1}) \geq D(u', a^{k-1})$, with $V(u, k) = 0$ if $D(u, a^{k-1}) = \mathbf{0}$ and $V(u, k) \geq 0$ for all $(u, E) \in C$.*
2. *Moreover, for any $k \in \mathbb{N}$ and $x = 0, 1$ there exist continuous functions $W^x(\cdot, k) : \mathcal{U} \rightarrow \mathbb{R}$ that represent \succcurlyeq_k^x , are increasing in u , satisfy $W^1(u, k) \geq W^1(u^k, k)$, $W^0(u^k, k) = W^0(u, k)$ and such that $V(u, k+1) = W^1(u - u^k, k)$ for all u .*

Part 1 provides a minimal cost-benefit representation. It requires that $D(u, a^{k-1}) = \mathbf{0}$ within a cognitive equivalence class, which is desirable for a minimal representation because when u is constant in a , nothing could be learned that is relevant for one's choice. Moreover, the value of reasoning is (weakly) increasing in the payoff differences $D(u, a^{k-1})$, because $D(u, a^{k-1})$ represents a minimal notion of payoff gains. Moreover, notice that since $V(u, k) \geq 0$ for all u and $D(u, \hat{a}) \leq \mathbf{0}$ if $u(\hat{a}, \omega) \geq u(a, \omega)$ for all $(a, \omega) \in A \times \Omega$, this monotonicity in $D(u, a^{k-1})$ also implies that $V(u, k) = 0$ if a^{k-1} is a best response at all states. In other words, if the action viewed to be best is dominant, then there cannot be any perceived instrumental value in thinking further. A representation that does not obey these desiderata would arguably not capture a basic notion of cost-benefit in reasoning.

Part 2 represents \succcurlyeq_k^0 and \succcurlyeq_k^1 and provides the connection between the decision to stop thinking and the rationale behind it. It also implies that it is possible to rank payoff functions by the incentives to reason they provide using the \succcurlyeq_k^1 relation alone. In particular, for any cognitive equivalence class, say that u provides 'stronger incentives to reason at' than v at step k if $(v+t, 1) \succcurlyeq_k (v, 0)$ implies $(u+t, 1) \succcurlyeq_k (u, 0)$ for all $t \in \mathbb{R}$. By part 1, this is possible if and only if $V(u, k+1) \geq V(v, k+1)$, which (by part 2) is equivalent to $W^1(u - u^k, k) \geq W^1(v - v^k, k)$. But since W^1 represents \succcurlyeq_k^1 , we have:

Corollary 1 *Fix a cognitive equivalence class, and let k be such that $c(k) < \infty$. Then, u provides stronger incentives to reason than v at step k if and only if $(u - u^k) \succcurlyeq_k^1 (v - v^k)$.*

As previously mentioned, Property 7 is key to a meaningful notion of value, for which the value of reasoning is non-negative and zero at constant payoff functions:

Example 3 Let $A = \{T, B\}$ and $\Omega = \{\omega\}$, so that we can omit ω in the following. Let the choice problem (\hat{u}, E) be such that $\hat{u}(T) = 0$ and $\hat{u}(B) = -1$, and let the reasoning process be such that, for each k : (i) $a^k = T$; (ii) \succsim_k^0 and \succsim_k^1 are represented, respectively, by functions $W^0(u, k) = u(a^k)$ and $W^1(u, k) = \max_{a' \neq a^k} u(a')$; (iii) preferences \succsim_k are such that $(u, 1) \succsim_k (u, 0)$ if and only if $V(u, k+1) \equiv W^1(u, k) - W^0(u, k) = V(u, k+1) = W^1(u - u^k, k) \geq c(k+1) > 0$. In this case, V does not reach its minimum at payoff functions that are constant in i 's action, where it is equal to 0. To see this, let u be constant and identically equal to 0. Then, we have $V(\hat{u}, k) = -1 < V(u, k) = 0$. We now check that the pair (V, c) satisfies all the axioms and properties used in Theorem 1, except Property 7: (i) Property 4.1 is satisfied, since $V(u, k) = 0$ whenever u is constant, and hence $(u, 0) \succ_k (u, 1)$; 4.2 is satisfied, since $V(v, k) = V(u, k)$ if $v - u$ is constant in a ; (ii) For any constant u , $(u + c(k+1), 1) \sim_k (u, 0)$, so Property 5 is satisfied; (iii) Function V is continuous, hence Property 3 is satisfied; (iv) Properties 1-6 are clearly satisfied. Property 7, however, fails: for $a^k = T$, $W^1(\hat{u}, k) = -1$ and $W^1(\hat{u}^k, k) = 0$, hence $\hat{u}^k \succ_k^1 \hat{u}$. \square

As emphasized in Section 2.4, the cost-benefit representation in Theorem 1 is uniquely pinned down by each equivalence class. Comparative statics exercises are therefore possible within each class, in that two cognitively equivalent problems (u, E) and (u', E') are associated to the same cost of reasoning, and may only differ in the values $V(u, k)$ and $V(u', k)$. The little structure on V , however, limits the applicability of Theorem 1: instances in which payoffs can be transformed to ensure that $D(\cdot, a^{k-1})$ increases uniformly are rare. In practice, an intuitive notion of varying the stakes is to transform a problem (u, E) into another $(\alpha u, E)$ where $\alpha \in \mathbb{R}_+$: if $\alpha > 1$ (respectively, $\alpha < 1$), it is natural to think of the transformation as increasing (resp., decreasing) the stakes. But if $D(u, a^{k-1})$ has both positive and negative components (as in the example above), then $D(\alpha u, a^{k-1})$ is not uniformly above or below $D(u, a^{k-1})$. Hence, the representation in Theorem 1 does not record this payoff transformation as increasing the incentives to reason. The next property characterizes a representation which allows these easy-to-implement comparative statics exercises.

Property 8 (Payoff Magnification) *For any $s \in S$, if $u \dot{\succ}_s^1 t \cdot \mathbf{1}$ for some $t \geq 0$, then $\alpha u \dot{\succ}_s^1 t \cdot \mathbf{1}$ for all $\alpha \geq 1$.*

This is a natural restriction. If the agent expects that applying the insights of future understanding to payoff function u is as valuable as receiving a certain payoff $t \geq 0$, then it would be at least as valuable if those payoffs were magnified by a constant (recall that payoffs are already in utils). As further discussed in Section 5, violations of this property are associated to phenomena of ‘choking’ (e.g., Ariely et al. (2009)) and to ‘fear of failure’.

Theorem 2 (Core* Representation) *Under the maintained assumptions of Theorem 1, Property 8 is satisfied if and only if V is such that, for any $(u, E) \in C$ and $\alpha \geq 1$, $V(\alpha u, k) \geq V(u, k)$.*

Since the conditions for this representation are weak, it reveals that increasing incentives to reason by magnifying payoffs, which is easily implemented in experimental settings, is often justified.

3.2 Expected Value of Reasoning

In this section we investigate properties which lead to an ‘expected value’ representation of the value of reasoning. From a formal viewpoint, the first such property is a standard independence axiom for the binary relation \succcurlyeq_s^1 :

Property 9 (1-Independence) *For each $s \in S$: For all $u, v, w \in \mathcal{U}$, $u \succcurlyeq_s^1 v$ if and only if $\alpha u + (1 - \alpha) w \succcurlyeq_s^1 \alpha v + (1 - \alpha) w$ for all $\alpha \in (0, 1)$.*

Joint with monotonicity (Property 1), this property implies Property 8. Hence, at a minimum, adding this property to those of Theorem 1 ensures a value of reasoning with the properties stated in Theorem 2. More broadly, independence as usual ensures that beliefs can be represented in probabilistic terms. There is, however, one important difference relative to standard applications of independence: The beliefs to be elicited here are not only over the states ω , but also over the action a that the agent expects to choose at the next step of reasoning. This is because, for any $s \in S(u, E)$, the mental preferences \succcurlyeq_s^1 are over real functions of (a, ω) . Independence ensures that, when comparing ‘counterfactual’ payoff functions $u' = u + \tau'$ and $u'' = u + \tau''$, payoffs at different outcomes (a, ω) are traded-off in a linear way, so that the agent’s beliefs can be expressed by a probabilistic assessment (cf. De Finetti, 1937).

Intuitively, one would expect a certain consistency between an agent’s view on his future beliefs over ω and his outlook over his future action. This consistency does not come from independence, but it will be the subject of the next property, which imposes a minimal notion of aptness for the problem at hand. For any $u \in \mathcal{U}$ and $a' \in A$, let

$$R(u, a') := \left\{ v \in \mathcal{U} : \begin{array}{l} \text{for all } a \neq a', v(a, \omega) = u(a, \omega) \text{ for all } \omega, \text{ and} \\ \text{for some } a'' \in A, v(a', \omega) = u(a'', \omega) \text{ for all } \omega \end{array} \right\}.$$

In words, the set $R(u, a')$ comprises all payoff functions that can be obtained from u just by replacing act $u(a', \cdot) : \Omega \rightarrow \mathbb{R}$ with some other act $u(a'', \cdot) : \Omega \rightarrow \mathbb{R}$ in the feasible set A , leaving everything else unchanged. Effectively, it is as if a' is dropped from the set of alternatives and replaced by a copy of one of the other alternatives.

The next property requires that, if mental state s belongs to the reasoning process about (u, E) , it is expected that the future understanding – whatever it might be – would be at least as apt to the actual problem than to a similar one, in which one of the alternatives has been replaced in the sense above. Formally:

Property 10 (No Improving Replacement) *For any (u, E) and $s \in S(u, E)$, $u \succcurlyeq_s^1 v$ for all a and $v \in R(u, a)$.*

This property is extremely weak and intuitive. To see this, consider the problem in Example 1 of Section 2, and let $v \in R(u, a_3)$ have a_3 replaced with a copy of a_1 :

u				$v \in R(u, a_3)$			
	ω_1	ω_2	ω_3		ω_1	ω_2	ω_3
a_1	1	0	0	a_1	1	0	0
a_2	0	1	0	a_2	0	1	0
a_3	0	0	1	a_3	1	0	0

If $v \succ_s^1 u$, so that Property 10 is violated, then the agent expects his future understanding to yield higher payoffs in v than in u . But since v is the same as u but with a_3 replaced by a_1 , it is as if the agent expects to regret following his own reasoning, should it suggest a_3 , because he would prefer a_1 instead. Property 10 ensures that no such regret is expected, for any action. Hence, Property 10 imposes a very weak notion of aptness, and does not directly require optimality of the reasoning process. Yet, together with the other properties, it yields a representation in which, at each stage of reasoning k , it is as if the agent expects that if he were to further pursue his reasoning, with some probability $p^k(\mu)$ he would form new beliefs $\mu \in \Delta(\Omega)$, and that in each case he would choose an optimal response $a^*(\mu)$ to those beliefs. Formally, for any $u \in \mathcal{U}$ and $\mu \in \Delta(\Omega)$, we let $BR^u(\mu) := \arg \max_{a \in A} \sum \mu(\omega) u(a, \omega)$. Then:

Theorem 3 (EV-representation) *Under the maintained assumptions of Theorem 1,⁵ Properties 9 and 10 are satisfied if and only if V in Theorem 1 is such that*

$$V(u, k) = \sum_{\mu \in \Delta(\Omega)} p^k(\mu) \sum_{\omega} \mu(\omega) \left[u(a^*(\mu), \omega) - u(a^{k-1}, \omega) \right], \quad (2)$$

where $p^k \in \Delta(\Delta(\Omega))$, $\mu \in \Delta(\Omega)$, and $a^*(\mu) \in BR^u(\mu)$ for all $(u, E) \in C$.

Since, for every k , the beliefs over the future understanding are pinned down by \succ_s^1 , and the mental preferences at the k -th step of reasoning are the same within each cognitive equivalence class, it follows that cognitively equivalent problems generate the same beliefs over outcomes at every k . This is the reason why the p^k 's and the $a^*(\mu)$ in this representation are common within each equivalence class, as is the cost of reasoning. Thus, these properties jointly entail a bound on the richness of the c.e. classes.

3.3 Value of Information

The functional form of V in Theorem 3 is reminiscent of the standard notion of ‘expected value of information’. But the beliefs in the representation are obtained from the relation \succ_s^1 , which need not relate to the determinants of the current action a^{k-1} . In contrast,

⁵In fact, Property 7 is not needed for this result, as it is implied by the other properties jointly.

standard models of information require that agents are Bayesian in the stronger sense that they use a single prior over everything that affects choices (for a discussion, see Gilboa (2009)). As such, for a value of information representation, the beliefs that describe the outlook on future steps of reasoning should be consistent with current choice a^{k-1} . Formally, let the mean beliefs $\hat{\mu}^k \in \Delta(\Omega)$ be defined from the representation in Theorem 3 so that $\hat{\mu}^k(\omega) = \sum_{\mu \in \Delta(\Omega)} p^k(\mu) \mu(\omega)$ for each $\omega \in \Omega$. Then, in a standard Bayesian model, $a^{k-1} \in BR^u(\hat{\mu}^k)$ for each $(u, E) \in C$.

To obtain this kind of dynamic consistency, we add one other property of the reasoning process, which requires that, from the viewpoint of the *future* understanding relation \succcurlyeq_s^1 , it is best to commit to the *current* action a^s than to any other $a \in A$.

Property 11 (1-0 Consistency) *For all (u, E) , $s \in S(u, E)$ and $v \in R(u, a^s)$, $u^s \succcurlyeq_s^1 v^s$.*

Theorem 4 (Value-of-Information Representation) *Under the maintained assumptions of Theorem 3, Property 11 is satisfied if and only if the representation in Theorem 3 is such that, for each $C \in \mathcal{C}$, for each $(u, E) \in C$ and for each k , $a^{k-1} \in BR^u(\hat{\mu}^k)$.*

The representation in this theorem says that it is as if the agent holds current beliefs $\hat{\mu}^k$, to which the current action is a best response, and regards the next step of reasoning as conclusive and consisting of an information structure that generates signals which he processes via Bayesian updating. Moreover, the agent anticipates that, for each signal realization, he would best respond to the posterior beliefs induced from the signal.

This is a completely standard value of information for a Bayesian agent who expects the next step of reasoning to be conclusive. This representation, however, does not accommodate agents who consider that the reasoning process might continue beyond the next step.⁶ In that case, a standard forward-looking Bayesian agent may be willing to think further even if he expects to learn little at the next step, if he thinks that he may learn a lot from the subsequent steps. To account for this possibility one needs to expand the domain of the mental preferences and modify the reasoning consistency condition, to express the idea that the choice to continue reasoning (i.e., the behavioral preferences) need not reflect only the 0-1 comparisons, but possibly more complex relations with the outlook over future understandings at different horizons (see Alaoui and Penta, 2018a).

Property 11 requires that the agent's attitude towards the next step of reasoning is the same as that entailed by his current understanding. While natural in standard information problems, this property is too narrow to accommodate general reasoning processes, since it rules out that the agent may not be fully aware of the determinants of the reasoning he has yet to perform. A disconnect between current behavior and future understanding is typical of several domains of research, in which the sheer complexity of certain problems makes it difficult for a researcher to anticipate what he might learn.

⁶The model by Gabaix and Laibson (2005), Gabaix et al. (2006), which mainly focuses on the problem of directed cognition, also maintains a similar assumption of myopia.

In conjunction with these factors, ethical constraints, prudence or lack of confidence about the current understanding often prevent the process of discovery from exhibiting the kind of dynamic consistency entailed by Property 11. In medical research, for instance, current practices are not abandoned even when researchers are very optimistic about new experimental techniques. As a simple example, suppose that a medical doctor can choose one of three treatments, $\{a_1, a_2, a_3\}$, whose effectiveness depends on the state of the world $\omega \in \{\omega_1, \omega_2, \omega_3\}$. Action a_1 , the *status quo* treatment, yields a payoff of 0.4 in every state, while a_2 and a_3 are experimental treatments which perform better but only in certain states of the world: a_2 yields a payoff of 0.5 in states ω_2 and ω_3 , but 0 otherwise; a_3 yields a payoff of 1 if the state is ω_3 , and 0 otherwise. The doctor has not yet ruled out ω_1 , but he is very optimistic that, upon further research, he will prove that ω_3 is the true state. He may still recommend the conventional treatment a_1 and at the same time assess the value of further research as close to 0.6. That is because he expects to confirm his belief that ω_3 is the true state, in which case he would switch from a_1 to a_3 .⁷ Hence, it is *as if* his current $a^s = a_1$ is dictated by one set of beliefs, but his value of reasoning is dictated by another. This is at odds with Property 11. Note that, with these beliefs, a Bayesian researcher would choose a_3 right away and have no incentive to research further.

In summary, for settings in which Property 11 is appealing, the Value of Information representation of Theorem 4 is justified. For settings in which it is not, such as situations in which unawareness, circumspection or sensitivity to confidence are relevant factors, alternative representations for the value of reasoning should be explored. These include Theorem 3, or the ‘maximum gain’ representation (Theorem 5), which we introduce next.

3.4 Circumspection in Deliberation

In this section we provide a representation which is particularly useful in applications, and which results from appending one property to those of Theorem 3. This new property expresses the idea that the agent is particularly deliberative or ‘attentive’ when making his choice. We first introduce the following notation: For any u and (a', ω') , let

$$\Gamma(u, (a', \omega')) = \left\{ v \in \mathcal{U} : \begin{array}{l} v(a, \omega) = u(a, \omega) \text{ for all } (a, \omega) \neq (a', \omega'), \text{ and} \\ v(a', \omega') = u(a'', \omega'') \text{ for some } (a'', \omega'') \in A \times \Omega \end{array} \right\};$$

then define $\Gamma(u) := \bigcup_{(a', \omega') \in A \times \Omega} \Gamma(u, (a', \omega'))$.

In words, $\Gamma(u)$ comprises all payoff functions which can be obtained from u by replacing a single outcome (a', ω') with some other (a'', ω'') , leaving everything else unchanged. Payoff functions in $\Gamma(u)$ thus represent situations that are very similar to u . The next property requires that, for any mental state involved in the reasoning about (u, E) , the

⁷We argue that this is descriptively plausible, whether or not it is considered normatively compelling. For discussions on why rationality need not entail Bayesianism when there is insufficient understanding, see Gilboa, Postlewaite and Schmeidler (2012) and Gilboa, Samuelson and Schmeidler (2013).

agent regards further introspection as weakly more valuable (in the sense of the ‘stronger incentives to reason’ of Corollary 1) in (u, E) than in any such similar situation:

Property 12 (Circumspection) *For any (u, E) and for any $s \in S(u, E)$, $u - u^s \geq_s^1 v - v^s$ for all $v \in \Gamma^s(u)$.*

This property formalizes a weak notion of ‘attentiveness’ of the decision maker. Despite its weakness, this property has remarkable consequences for the representation:

Theorem 5 (Maximum Gain) *Under the maintained assumptions of Theorem 1, Properties 9 and 12 are satisfied if and only if V in Theorem 1 is such that, for any $(u, E) \in C$,*

$$V(u, k) = \max_{\omega \in \Omega} u(a^*(\omega), \omega) - u(a^{k-1}, \omega).$$

In words, it is *as if* the agent values the next step of reasoning by looking at the highest opportunity cost that playing according to the current understanding, a^{k-1} , may entail. This is a simple and intuitive heuristic, which has the advantage that it only depends on the payoffs, with no extra subjective parameters such as the beliefs in Theorems 3 and 4. This representation therefore is easy to use and limits the degrees of freedom of the model, which is especially desirable in applications (see Alaoui and Penta (2016a)).

The representation in Theorem 5 is a special case of the one in Theorem 3, in which the decision maker has the strongest incentives to reason. In this sense, the agent is fully optimistic about the value of pursuing his reasoning further (or, equivalently, he is fully pessimistic about his *current* understanding). However, this representation is inconsistent with that of Theorem 4: in Theorem 5, it is *as if* the agent is certain that the result of the next step of reasoning will be to learn that the state is $\omega^* \in \arg \max_{\omega \in \Omega} [u(a^*(\omega), \omega) - u(a^{k-1}, \omega)]$. But with these beliefs, the representation in Theorem 4 implies that a^{k-1} is itself a best response to ω^* , hence $u(a^*(\omega^*), \omega^*) = u(a^{k-1}, \omega^*)$, which means that the value of reasoning is 0. Given the definition of ω^* , however, this is only possible if a^{k-1} is a best response to every $\omega \in \Omega$. Thus, Properties 11 and 12 are inconsistent with each other in non-trivial problems, in which the optimal action varies with the state.⁸

4 Extensions and Applications

In this section we discuss different ways in which our framework can be applied or extended. Section 4.1 illustrates how the axioms and representations above may be varied to accommodate specific forms of reasoning or applications. Section 4.2 instead focuses on how the model can be taken to the data, to test its joint implications within specific settings. The section also explains how to connect our model to the important literature on response time, and provides an application to a recent experiment on attention allocation by Avoyan and Schotter (2018).

⁸Note that Property 10 is implied by the representation in Theorem 5, hence the gap between the representations in Theorems 4 and 5 is due to Properties 11 and 12 alone.

4.1 Theoretical Variations

A useful feature of our framework is that it is amenable to variations, which may be desirable in specific settings. This section illustrates the point through extensions to accommodate specific reasoning processes (e.g., based on categorizing by conditional dominance relations), or to develop cost-benefit criteria for applied settings (e.g., R&D problems).

4.1.1 Categorizing by Dominance and Challenge

Consider an agent who, at any point in the reasoning process, perfectly understands for which states the current action is optimal, and whose decision about whether to continue reasoning is solely driven by the set of states within which the current action is not dominant. This kind of reasoning can be easily accommodated in the model through a simple adaptation of the Scope Property 4 (part S.2). First, let $D(u) = (D(u, a))_{a \in A}$ denote the collection of payoff differences, and note that (S.2) may be rewritten as follows:

$$\text{If } D(u) = D(v), \text{ then: } (u, 1) \succsim_s (u, 0) \text{ if and only if } (v, 1) \succsim_s (v, 0) \quad (\text{S.2})$$

Now, for each $u \in \mathcal{U}$ and action a^s , let $\hat{D}(u, a^s) \in \mathcal{U}$ be such that, for all $(a, \omega) \in A \times \Omega$,

$$\hat{D}(u, a^s)(a, \omega) = \begin{cases} u(a, \omega) - u(a^s, \omega) & \text{if } a^s \notin \arg \max_a u(a, \omega) \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

Replacing $D(u)$ with $\hat{D}(u, a^s)$ in Property 4-(S.2) requires that only the payoff differences at states in which a^s is *not* optimal matter. This captures the desired notion that the agent understands in which states the current action is optimal, namely $\Omega(a^s) := \{\omega : a^s \in \arg \max_a u(a, \omega)\}$, and evaluates whether to continue reasoning according to the consequences he anticipates when his current action is not optimal.

With this substitution, a ‘Core Representation’ analogous to Theorem 1 obtains, with \hat{D} replacing D in the statement. Moreover, since $\hat{D}(u, a^{k-1}) = \mathbf{0}$ if a^{k-1} is a best response at all states, the representation resulting from this modification effectively entails that $V(k) = 0$ if a^{k-1} is dominant (cf. Alaoui and Penta (2016a)). Similarly, if Property 9 is added to this modified version of Theorem 1 (without Property 10), then

$$V(u, k) = \sum_{\mu \in \Delta(\Omega)} p^k(\mu) \sum_{\omega} \mu(\omega) \left[u(a(\mu), \omega) - u(a^{k-1}, \omega) \right], \quad (4)$$

where the outlook on future understanding takes a partitional form, in the sense that $\text{supp}(p^k) \subseteq \Delta(\Omega(a^{k-1})) \cup \Delta(\Omega \setminus \Omega(a^{k-1}))$, and $a(\mu)$ does not necessarily maximize $E_\mu[u(\cdot, \omega)]$ (as in Theorem 3), but it is a ‘better response’: for $(u, E) \in C$, it ensures $\mathbb{E}_\mu[u(a(\mu), \omega)] \geq \mathbb{E}_\mu[u(a^{k-1}, \omega)]$.

4.1.2 Research and Development

In many domains of research and innovation, there is no a priori obvious way of specifying the value of conducting an additional step of reasoning. When attempting to improve on a computer algorithm or develop a new product, it is difficult to assess what investing extra resources in research might bring in insight, how useful it would be, or with what likelihood. It is thus unclear how the value of reasoning should be assessed. These cost-benefit decisions are of enormous relevance for a large number of firms and high-tech startups, effectively determining their success or bankruptcy. Yet these important choices are often determined by heuristics at best partially understood, and the standard approach is not always applicable. Our framework can be used to inform a rigorous discussion on alternative criteria that lead to new representations. In this respect, the theorems in Section 3 serve as templates for obtaining appropriate representations by suitably adjusting the baseline properties to match the desiderata of the problem at hand. We provide next two illustrative applications.

Research as Filtration. In many cases, research consists of progressively ruling out states of the world. This process of elimination may be accommodated within the model by enriching the notion of mental state: for each step of reasoning k , the mental state also specifies a set $\Omega^k \subseteq \Omega$ of states that are still regarded as possible, and such that $\Omega^{k+1} \subseteq \Omega^k$ for all k . In this case it is natural to restrict the suitability relation to only depend on states in Ω^k (formally: for $x = 0, 1$, $u \doteq_x v$ if $u(a, \omega) = v(a, \omega)$ for all $(a, \omega) \in A \times \Omega^k$). Under this additional property, the beliefs in the representation of Theorems 3 and 4 would be concentrated on Ω^k . Similarly, if the notion of attentiveness formalized by Property 12 is compelling in a research context, then it can be adapted by replacing the set $\Gamma(u)$ in that property to allow duplication and replacement of outcomes only in $A \times \Omega^k$.⁹ With this change, under the assumptions of Theorem 5, we obtain a representation in which the ‘maximum gain’ heuristic is applied solely to the set of states still viewed as possible:

$$V(u; k) = \max_{\omega \in \Omega^{k-1}} u(a^*(\omega), \omega) - u(a^{k-1}, \omega). \quad (5)$$

As with Theorem 5, this representation may be particularly useful in applications, since the utility function, the filtration and the path of reasoning fully characterize the value of reasoning, with no extra degrees of freedom.

R&D and Consideration sets. In firms’ R&D processes, research is often focused on understanding the properties of particular ‘actions’ (such as products, organizational

⁹Formally, $\Gamma(u)$ in Property 12 should be replaced by $\Gamma^k(u) := \bigcup_{(a', \omega') \in A \times \Omega^k} \Gamma^k(u, (a', \omega'))$, where

$$\Gamma^k(u, (a', \omega')) := \left\{ v \in \mathcal{U} : \begin{array}{l} v(a, \omega) = u(a, \omega) \text{ for all } (a, \omega) \neq (a', \omega'), \text{ and} \\ v(a', \omega') = u(a'', \omega'') \text{ for some } (a'', \omega'') \in A \times \Omega^k \end{array} \right\}.$$

procedures, etc.) in different states of the world. Suppose, for instance, that the firm is interested in assessing a particular set $A' \subseteq A$ of alternatives relative to a default a^s . We refer to $A' \cup \{a^s\}$ as the *consideration set*. Then, it may be normatively appealing to modify the axioms and properties as follows: (i) first, assume that the mental preferences are only affected by the payoffs of actions in the consideration set: for $x = 0, 1$, $u \stackrel{x}{\succeq}_s v$ if $u(a, \omega) = v(a, \omega)$ for all $(a, \omega) \in A' \cup \{a^s\} \times \Omega$; (ii) second, strengthen part (S.2) of Axiom 4 by requiring that the incentives to reason are completely pinned down by the extent by which the alternatives in the consideration set improve on the current action at different states, if at all. Formally, for any u let

$$D^+(u, a^s)(a, \omega) := \max\{0, u(a, \omega) - u(a^s, \omega)\}; \quad (6)$$

then, (S.2) should be replaced by the following: if $D^+(u, a^s)(a, \omega) = D^+(v, a^s)(a, \omega)$, for all $(a, \omega) \in A' \times \Omega$, then: $(u, 1) \succsim_s (u, 0)$ if and only if $(v, 1) \succsim_s (v, 0)$.

With these two changes, under the maintained assumptions of Theorem 1, plus Property 9, the value of reasoning takes the following form,

$$V(u, k) = \sum_{(a, \omega) \in A' \times \Omega} p^k(a, \omega) \max\{0; u(a, \omega) - u(a^{k-1}, \omega)\}, \quad (7)$$

which effectively represents the value of the option of switching from a^{k-1} to the alternatives $a \in A'$ in different states, weighted by the probability that a pair alternative-state is considered.¹⁰

4.2 From the model to the data

The most immediate way to test the model's prediction is to look at the effect that changing the stakes of a problem – holding constant the cognitive equivalence class – has on the number of steps of reasoning undertaken by the agent. This is best done in settings in which individuals' reasoning processes are well-understood and in which the steps of reasoning are easily identified. One such example is provided by level- k reasoning in games. Alaoui and Penta (2016a, AP hereafter) test experimentally the implications of the cost-benefit approach in the context of a model of *endogenous level- k reasoning*. As in standard models of level- k reasoning (e.g., Nagel (1995), Crawford and Iriberri (2007), Costa-Gomes and Crawford (2006)), the path of reasoning is determined by iterating players' best-responses. It is thus assumed that games with the same (pure action) best-response functions are cognitively equivalent.¹¹ AP consider two versions of the model.

¹⁰Manzini and Mariotti (2014) study consideration sets in a stochastic choice setting. Whereas the focus is very different, one interesting similarity between their model and the representation in (7) is that different actions in the consideration sets may be chosen with positive probability. In our model, however, it is *as if* the agent expects to make this choice only if it is effectively superior to the default action a^{k-1} . This effectively entails a *satisficing* criterion, in which the critical threshold is set by the default a^{k-1} .

¹¹To apply the model of reasoning to strategic settings, let $A \equiv A_i$ denote the set of the agent's actions in the game; the context E specifies a set of opponents N , with actions $A_{-i} = \times_{j \neq i} A_j \equiv \Omega$ and payoff

In the first, ‘detail free’ version, the only restrictions on the cost and benefit functions are those entailed by the core representation of Theorem 1.¹² The experimental results are consistent with the model’s predictions across a variety of treatments, thereby providing strong support to the empirical validity of the cost-benefit approach. The second version of the model applies the maximum gain representation of Theorem 5 to obtain sharper predictions. With this added structure, AP apply the model to the experiments in the seminal paper by Goeree and Holt (2001, GH), and show by means of a calibration exercise that it delivers quantitatively accurate predictions across different games.

But in many settings individuals’ reasoning processes are not as well understood, or the mapping from choice to the step reached may not be as clear. For these reasons, a large part of the literature on reasoning processes has focused on more indirect measures of individuals’ reasoning, such as *response time* and *attention allocation*. Such measures of time provide an intuitive and practical way to measure subjects’ level of attention or their cognitive effort. It seems thus tempting to identify changes in the depth of reasoning in our model with changes of response-time (or of the attention allocated) of an equal direction. This connection, however, is not straightforward, since the model above does not refer directly to time. In particular, the unit of measure of a step of reasoning in our model is not time, but a ‘unit of understanding’, and the amount of time needed to achieve a certain ‘unit of understanding’ may vary from case to case: if it is harder to reason about problem (u, E) than (u', E') , then any given ‘unit of understanding’ may take longer in (u, E) , and hence (u', E') may induce a lower response time even if it were associated with a larger depth of reasoning (see, e.g., Alos-Ferrer and Buckenmaier, 2018).¹³ Based on this logic – and abstracting from the possibility of endogenous changes in the individual’s *focus* (cf. footnote 13) – one should expect that if two problems are cognitively equivalent, and as such equally difficult to think about, then every step of reasoning would take an equal amount of time in the two problems. Hence, if (u, E) and (u', E') are cognitively equivalent, and the former induces a higher depth of reasoning than the latter, then we expect a higher response time in (u, E) than (u', E') .

The connection between response time and depth of reasoning is less straightforward when problems are compared across cognitive equivalence classes, for the following reason: if (u, E) and (u', E') induce the same value of reasoning, but (u, E) has a higher cost of rea-

functions $u_{-i} : A_i \times \Omega \rightarrow \mathbb{R}^{|\mathcal{N}|}$ and $u_i : A_i \times \Omega \rightarrow \mathbb{R}$. In standard models of level- k reasoning, agents go through a sequence of ‘increasingly sophisticated’ actions: for each i , and for given ‘level-0 profile’ $a^0 = (a_i^0, a_{-i}^0)$, i ’s path of reasoning $\{a_i^k\}_{k \in \mathbb{N}}$ is such that, for each k and j , a_j^k is the best response to a_{-j}^{k-1} . The path of reasoning is thus a sequence of mental states $\{s_{\pi(u, E)}^k\}_{k \in \mathbb{N}}$ with associated actions $\{a_i^k\}_{k \in \mathbb{N}}$.

¹²More precisely, we used the version of Theorem 1 in which D is replaced by \hat{D} (eq. (3), Section 4.1).

¹³Alos-Ferrer and Buckenmaier (2018) study the impact of changing stakes on individuals’ deliberation time in an experimental setting similar to Alaoui and Penta’s (2016a). They find that the effect of higher incentives on the depth of reasoning, as identified by the level- k action, is consistent with Alaoui and Penta’s (2016a) findings and with the cost-benefit model; but the effect on response time is unclear, and if anything there is evidence of a shorter response time when stakes are higher. They suggest endogenous changes in *focus* as a possible explanation, and propose a model which reflects the idea that higher payoffs may induce more focus from the individual, and hence shorter response time.

soning than (u', E') , then even though the number of steps would be lower in (u, E) , each of them would take longer than in (u', E') , and the overall effect is indeterminate. Hence, our model suggests special caution in drawing inferences on subjects' depth of reasoning from response time data, when non-cognitively equivalent problems are compared.

4.2.1 Response Time and Attention Allocation

In this section we formalize the logic above and connect our model of endogenous depth of reasoning to response time. We also propose an application to Avoyan and Schotter's (2016) experiment on attention allocation. For simplicity, throughout this section we maintain the maximum gain representation from Theorem 5:

$$V(u, k) = \max_{\omega \in \Omega} u(a_i^*(\omega), \omega) - u(a^{k-1}, \omega). \quad (8)$$

Relying on the representation theorem, for any problem (u, E) and cost of reasoning $c : N \rightarrow \mathbb{R}_+$, we let $K(u, E; c)$ denote the associated depth of reasoning:

$$\mathcal{K}(u, E; c) = \min\{k \in \mathbb{N}_+ : c(k) \leq V(u, k) \text{ and } c(k+1) > V(u, k+1)\}. \quad (9)$$

As already discussed, the maximum gain representation is particularly convenient in applications, because the value of reasoning at every step is completely pinned down by the sequence $\{a^k\}_{k \in \mathbb{N}}$. For later reference, we say that two decision problems are *path-equivalent* if they induce the same sequences $\{a^k\}_{k \in \mathbb{N}}$; they are *cost-equivalent* if they are associated with the same cost-of-reasoning. Cognitively equivalent problems are always both path- and cost-equivalent; under the maximum gain representation, the converse holds too, and it is also possible to obtain unambiguous comparative statics on the depth of reasoning between path-equivalent problems (even if they are not cognitively equivalent), whenever the associated costs of reasoning are ranked uniformly across steps.

For any problem (u, E) , let $\tau : \mathbb{N}_+ \rightarrow \mathbb{R}$ denote the *time function*, where $\tau(k)$ represents the time it takes to perform the k -th step of reasoning. Then, the *total response time* in problem (u, E) , given the associated cost and time functions (c, τ) , is equal to:

$$T(u, E; c, \tau) := \sum_{k=1}^{\mathcal{K}(u, E; c)} \tau(k).^{14}$$

It seems reasonable to assume that if it is harder to attain a certain understanding in one problem than in another, then not only is the cost of reasoning (weakly) higher, but it

¹⁴The $T(\cdot)$ function above corresponds to the so-called *chronometric function* in the neuroeconomics and psycho-metric literature. This specification abstracts from the possibility that changes in stakes may have a direct impact on individual's *focus* (cf. footnote 13). This possibility could be accommodated in the model by letting the $\tau(\cdot)$ function be directly affected by the payoff function, independent of the cognitive equivalence class. Such an extension, however, would limit the possibility of unambiguous comparative statics on response time to the subset of cognitively equivalent problems which also induce the same focus.

also takes (weakly) more time to attain that understanding. We thus connect our baseline model of reasoning to response time by means of the following simple assumption:

Assumption 1 (Difficulty) *Let (u, E) and (u', E') be path-equivalent, with associated costs of reasoning c and c' , and time functions τ and τ' , respectively. Then, for every $k \in \mathbb{N}$: $c(k) \geq c'(k)$ implies $\tau(k) \geq \tau'(k)$.*

Together with Theorem 1, this assumption implies the following:

Remark 1 *If (u, E) and (u', E') are cognitively equivalent, and (u', E') has larger payoff differences, then $T(u', E') \geq T(u, E)$. If the two games are not cognitively equivalent, then the comparison between $T(u', E')$ and $T(u, E)$ is indeterminate, even if they are path-equivalent and $c(k) \geq c'(k)$.*

The model of response time thus obtained can be easily applied to derive testable predictions for Avoyan and Schotter’s (2016, AS hereafter) experiment on attention allocation in games. In their *attention allocation task*, subjects are presented a pair of two-player games in matrix form, (G^1, G^2) , and they are asked the fraction $\alpha^l \in [0, 1]$ of a total (unspecified) time X to allocate to game G^l ($\alpha^1 + \alpha^2 = 1$). In AS’s experiments, such games consist of variations of archetypal two-person games such as battle of the sexes, pure coordination, constant-sum games, etc. (see Appendix B). In the following we explain how all the experimental results in AS are consistent with a straightforward extension of the model above, appended with one simple assumption on the cognitive equivalence classes which formalizes one of AS’s working hypotheses.

In particular, consider a two-stage procedure in which first the agent is presented a pair (G^1, G^2) , and chooses the fraction $\alpha^l \in [0, 1]$ to allocate to game G^l . Then, in the second stage, the agent plays the game G^l according to the model from the previous section, with induced response time $T(G^l)$. In the first stage, the agent chooses the time allocation without having really reasoned about the two games, and hence without knowing how much time he would actually need, whether a given time allocation will be binding or not, nor having a precise understanding of how he will play in the game. Suppose, however, that the agent’s choice in the first stage is ‘consistent’ with his future reasoning process in the following (weak) sense:

Assumption 2 (Attention Consistency) *for any pair (G^1, G^2) , α^1 is increasing in $T(G^1)$, decreasing in $T(G^2)$, and such that $\alpha^1 > \alpha^2$ if $T(G^1) > T(G^2)$.*

Concerning the equivalence classes, we assume that:

(CE.1) whenever two games in matrix form induce the same pure-action best responses, then they are *path equivalent*; if they also have the same zero payoffs, then they are also *cognitively equivalent*; and **(CE.2)** if a game in matrix form is changed turning some zero payoff into non-zero, but without changing the strategic structure, then the cost of reasoning (weakly) increases at every step.

Note that (CE.1) is the same assumption made in Alaoui and Penta (2016a), and it is implied by the maintained assumption on cognitive equivalence classes from Section 2. (CE.2) instead merely formalizes one of Avoyan and Schotter’s (2016) working hypothesis in terms of a restriction on our cognitive equivalence classes. It is motivated by the idea that the presence of zeros (weakly) simplifies the identification of the payoff differences between actions and states, by sparing the individuals’ an arithmetic operation.

Example 4 Consider the following games:

G^1	L_2	R_2	G^2	L_2	R_2	G^3	L_2	R_2
T_1	3, 1	0, 0	T_1	6, 2	0, 0	T_1	3, 1	0, 1
B_1	0, 0	1, 3	B_1	0, 0	3, 6	B_1	1, 0	1, 3

According to Assumption 2 (and, hence, CE.1), games G^1 and G^2 are cognitively equivalent, and hence they are associated with the same path and costs of reasoning. Moreover, the payoff differences are higher in G^2 , and hence this game entails a higher value of reasoning than G^1 (cf. eq. 8). It follows that G_1 induces a (weakly) higher depth of reasoning than G^2 . In contrast, G^3 is not cognitively equivalent to G^1 , because some zero entries were turned into non-zeros, leaving everything else unchanged. Then, by Assumption CE.2, G^3 has a higher cost of reasoning than G^1 . In this case, it also has a lower value of reasoning, so it will induce a lower depth of reasoning. Based on Remark 1, however, the effect on response time is ambiguous in this case. \square

Under these mild assumptions, our model generates a rich set of predictions for AS’s attention allocation task, all of which are consistent with AS’s findings. The tables in Appendix B summarize all the model’s predictions for the binary comparisons in attention allocation in AS’s tasks. The theoretical predictions of our model are all in terms of weak ordering, and they are all confirmed. In fact, with the only exception of two comparisons, for which AS find that the time allocation is no different from uniform at a statistically significant level, all the other data are consistent with the theoretical predictions in the strict ordering sense, at a statistically significant level (the vast majority at the 5% level). Hence, all data from AS’s experiment are consistent with all the predictions of the model above for their attention allocation task.

AS’s experimental results therefore support the joint implications of the model of this section, which combines the assumptions on the individuals’ reasoning processes contained in Theorem 1, with the mild assumptions on the connection with response time (the Difficulty Assumption 1), attention allocation (the Attention Consistency Assumption 2), and on the cognitive equivalence classes (which played the role of identification restrictions, by determining the domain of applicability of the comparative statics). The same joint assumptions – and particularly those on the cost-benefit representation – are also consistent with the experimental findings of Alaoui and Penta (2016a), their applications to Goeree and Holt’s (2001) experiments, and with Alos-Ferrer and Buckenmaier’s (2018) results on

depth of reasoning. AS’s results therefore further strengthen a rich set of experimental evidence in support of our model.

5 Core Violations and Psychological Phenomena

The previous section has focused on the applicability of this framework when the cost-benefit representation holds; this section focuses on its usefulness in understanding patterns of choice for which it may fail. These are *choking under pressure* (Section 5.1), *(over)thinking aversion* and *ruminatio*n (Section 5.2). These behavioral patterns are well-known in psychology research, and their relevance to economic settings has been increasingly acknowledged. Using our approach, we show that different psychological mechanisms that have been proposed to explain these phenomena map to distinct violations of the assumptions of our model. Our framework can serve to unknot these hypotheses and design tests to contrast their testable implications empirically.

5.1 Choking Under Pressure

Increasing financial incentives does not always improve performance in cognitive tasks, and it is sometimes even detrimental, as with “choking under pressure” (see Camerer and Hogarth (1999) and Ariely, Gneezy, Loewenstein and Mazar (2009)).¹⁵ The recent literature has shown a detrimental effects of increasing incentives in choice problems and in cognitive tasks (see, for instance, the ‘packing quarters’ and ‘adding’ tasks in Ariely et al.(2009), or the experiments in Gneezy and Rustichini (2001)). In this type of problems, it is reasonable to maintain that a deeper understanding would bring the subject closer to the objective true state of the world, and is therefore not detrimental to performance. Hence, from the viewpoint of our model, a worse performance is associated to a lower depth of reasoning.¹⁶

A lack of impact of financial incentives is not at odds with cost-benefit per se, since the costs for the task at hand may be steep, and possibly infinite. But a detrimental effect is not consistent with our representation. Different mechanisms have been offered to explain performance declining with stakes, and are summarized in Camerer and Hogarth (1999) and Ariely et al. (2009). Using the key elements of our model, we classify these competing mechanisms into three categories: *Cost-based*, *value-based* and *(deflection of the) path of reasoning*. The precise mapping between each mechanism and the assumption of the model that it violates is discussed below, and summarized in Table 1. We will then discuss how an experiment can be designed to separate these different mechanisms.

¹⁵The papers above contain detailed discussions of this topic and the proposed mechanisms in the literature. There is also a vast psychology literature which we do not survey; see also Baumeister and Showers (1986) and Beilock, Kulp, Holt and Carr (2004).

¹⁶This assumption does not necessarily apply to tasks that do not involve actual reasoning or introspection, such as tasks based on motor skills (e.g., Neiss (1988)) and choices in which there is a detrimental effect of reflection interrupting automatic processes - also referred to as the “centipede effect” (see Camerer, Loewenstein and Prelec (2005) and references therein).

Mechanism	Property Violated
Cost-based explanations: Level effect (Anxiety) Differences effect (Fear of failure)	Condition 2 (c.e., cost) β α
Value-based explanations: Level effect (Distraction) Differences effect (Self-handicapping)	Scope Axiom 4-(S.2) Property 8 (Payoff magnification)
Path of Reasoning: Deflection	Condition 2 (c.e., path)

Table 1: *Choking under pressure*. From psychological mechanisms to violations.

5.1.1 Cost-based Explanations: Anxiety and Fear of Failure

One view is that the increased pressure of higher rewards takes up working memory, which makes reasoning more difficult. Within our model, an increase in reasoning difficulty translates to higher costs. Hence, even though the task to be completed remains exactly the same, the cost of reasoning increases as rewards increase. This implies that the reward itself changes the cognitive equivalence class, which is in violation of Condition 2.

But, by this condition, two choice problems (u, E) and (v, E) are cognitively equivalent if they differ only by an additive constant ($v = u(a, \omega) + \beta(\omega)$), or only by a multiplicative factor ($v(a, \omega) = \alpha u(a, \omega)$). This suggests that the general mechanism of increased pressure can be separated into two factors. A higher cognitive load can be due to the pure level effect of higher rewards, which would lead to a reduced performance even when payoffs are increased by a constant. Alternatively, it can be due to larger payoff differences between succeeding and failing. These explanations can therefore be distinguished with the appropriate experimental design. The level effect mechanism would lead to a reduced performance when payoffs $u(a, \omega) + \beta(\omega)$ are increased by a sufficiently large $\beta(\omega)$, while the differences effect mechanism would lead to reduced performance when payoffs $\alpha u(a, \omega)$ are magnified by a sufficiently large α . Whereas the terms are sometimes used with different meanings in the psychology literature, it is natural to associate the level effect mechanism to phenomena of *anxiety*, and the differences effect mechanism to *fear of failure*.

5.1.2 Value-based Explanations: Distractions and Self-handicapping

Another view is that higher incentives distract the individual from the task at hand, thereby diminishing his motivation. In our model, a reduction in motivation translates to a decrease in the value of reasoning. Paralleling the discussion on costs, this value-based explanation can also be decomposed into differences and pure level effects.

The first mechanism, in which the value is lower when the incentives should be higher, is a failure of the Payoff Magnification Property 8: under this mechanism, reasoning about αu when $\alpha > 1$ has less value than reasoning about u , even though there is more to gain. We note that violations of this nature may be non-monotonic in the increase in α , which

would be consistent with empirical findings. Gneezy and Rustichini (2000), for instance, find that performance first worsens from awarding no monetary payoffs to low rewards, but then improves as rewards are higher. In contrast, in the Yerkes-Dodson law (see Ariely et al., (2009)) the violation of Property 8 is in the opposite direction: performance first improves and then diminishes.

The second mechanism, in which adding a constant payoff reduces the value of reasoning (V is higher for u than for $u(a, \omega) + \beta(\omega)$, $\beta(\omega) > 0$), is a failure of Axiom 4 (S.2). Psychological mechanisms such as *self-handicapping* and *distraction* in this context map naturally to the differences and pure level effects, respectively.

Experimental design to identify violations

It is conceptually straightforward to disentangle level effects from difference effects, in that level effects concern increasing incentives by a constant, while difference effects concern changing the differences between the payoffs. Distinguishing between cost-based and value-based explanations, however, may seem more challenging. In the absence of our framework, the parallel between the cost and value mechanisms creates difficulty in disentangling them, even conceptually. Our model can be used to achieve this objective, because these different components (Condition 2 (parts α and β), Axiom 4 and Property 8) are clearly distinct. We now discuss an experimental design that identifies separately which of the cost or value-based explanations (i.e., a violation of Condition 2 (α), Condition 2 (β), Scope Axiom 4, or Property 8) would explain the observed departure from behavior consistent with cost-benefit analysis in reasoning.

First, since the different explanations map to different axiom violations, these can be directly tested, for instance using the lottery elicitation method discussed in Section 2.5. But testing each relevant axiom requires a rich dataset, and so we provide here a simpler method which relies on fewer direct tests. Consider a cognitive task, similar to the experiment of Dean and Neligh (2017), in which a subject is asked whether an urn with n balls has more red or green balls. Suppose, for simplicity, that the states of the world are $\omega \in r, g$, which corresponds to more red or green balls, respectively. Let u be a payoff matrix which pays 1 if the decision is correct, 0 otherwise. Now suppose that we increase all the payoffs to $u' = \alpha u$ or to $u'' = u + \beta$, where $\alpha > 1$ and $\beta > 0$, while leaving the task (and hence the environment E) constant:

$$u = \begin{array}{|c|c|c|} \hline & r & g \\ \hline a_1 & 1 & 0 \\ \hline a_2 & 0 & 1 \\ \hline \end{array} \quad u' = \begin{array}{|c|c|c|} \hline & r & g \\ \hline a_1 & \alpha & 0 \\ \hline a_2 & 0 & \alpha \\ \hline \end{array} \quad u'' = \begin{array}{|c|c|c|} \hline & r & g \\ \hline a_1 & 1 + \beta & \beta \\ \hline a_2 & \beta & 1 + \beta \\ \hline \end{array}$$

We maintain here that performance increases with depth of reasoning, as we would expect with such a task.¹⁷ Since our aim is to disentangle them, we assume that violations,

¹⁷For simplicity, suppose that there are no learning effects from task repetition across treatments, if the

if they occur, would be due to either the cost-based or the value-based mechanisms, but not both. Given this assumption, if the performance increases going from u to u' but has no change going from u to u'' , then we cannot rule out that the individual is consistent with all the properties. But if his performance decreases from u to u' then he is violating either C.E. Condition 2 (α), or Property 8. If his performance changes from u to u'' , then he is violating either C.E. Condition 2 (β) or the Scope property 4-S.2 (see Table 1).

Consider the case in which performance decreases from u to u' . To further separate whether it is C.E. Condition 2 (α) or Property 8 that is violated, the payoff magnification property can then be tested through the elicitation method discussed in Section 2.5. If it holds, then it is C.E. Condition 2 that is violated.

For the case in which performance decreases from u to u'' , an additional treatment can be used to separate C.E. Condition 2 (β) from the Scope property 4-S.2. The experiment consists of maintaining u fixed, but letting the environment vary, so that (u, E) can be compared to (u, E') , where environment E' is identical in every way to E except that the task is more complex. For instance, suppose that there are now m balls to count, with $m > n$, but otherwise maintaining the same composition of the urn. We make the identification restriction that an increase in complexity increases the cost, and leaves all else identical. In that case, if the agent's performance increases (in addition to his performance having changed from u to u'' previously), then it must be that it is the Scope property 4-S.2 that is violated.

5.1.3 Changing the path of reasoning: deflection

The discussions above focused on cost and value-driven explanations, and on designing experiments to separate between them and identifying which of the relevant properties would be violated. Here we focus on a distinct explanation, which is that “increased motivation tends to narrow individuals’ focus of attention on a variety of dimensions [...] This can be detrimental for tasks that involve insight or creativity, since both require a kind of open-minded thinking that enables one to draw unusual connections between elements.” (Ariely et al. (2009), p. 453). That is, the additional motivation affects how individuals think, which translates in our model to higher rewards changing the reasoning process itself. Formally, the reasoning process $(s^k)_{k \in \mathbb{N}}$ associated with the lower rewards is not the same as the reasoning process $(\hat{s}^k)_{k \in \mathbb{N}}$ associated with higher rewards. This is a separate violation of Condition 2, in that it is not necessarily the costs that change, but the path of reasoning itself.

Our framework is useful for distinguishing this mechanism from the previous ones. For instance, with the value-based explanations, the path of reasoning is *not* affected by the reward. In particular, consider the value-based mechanism, and suppose that gradually

comparison is made within groups. Alternatively, these treatments can be compared across groups. In this discussion we have also abstracted from the issue of mapping monetary payments to utils, which should obviously be addressed with care in an actual experiment.

Psychological Phenomenon	Property Violated
Thinking aversion	Property 7 (No Improving Obstnacy)
Overthinking	Condition 1 (Reasoning Consistency)
Rumination and worry	Scope Axiom 1-(S.1)

Table 2: From psychological phenomena to violation of properties

increasing payoffs has the non-monotonic effect of first decreasing and then increasing the value of reasoning, which would translate to the performance first (weakly) worsening and then improving. But since, according to that mechanism, it is the value that changes and not the path of reasoning, high stakes may induce the individual to revert to solutions from much lower rewards.¹⁸ If, instead, the path itself has changed, then such a reversion need not be observed.

Concretely, suppose that the sequence of actions in the path of reasoning is a^1, a^2, \dots both for u and for higher αu (for $\alpha > 1$), and that incentives to reason decrease in α up to some value $\alpha^* > 1$ and increase afterwards. Then, if the path does not change, behavior for high $\alpha > \alpha^*$ would be equal to the behavior for some low $\alpha < \alpha^*$. For instance, the agent may stop at a^2 for low reward, at a^1 for intermediary reward, and again at a^2 for high reward. In contrast, if the reasoning process associated with higher stakes induces a different sequence $\hat{a}^1, \hat{a}^2, \dots$, as implied by the change in path mechanism, such a reversion need *not* be observed. For instance, the agent may stop at a^2 for low rewards, and then (once rewards are sufficiently high to *deflect* the path of reasoning) at \hat{a}^3, \hat{a}^4 , and so on.

5.2 Thinking aversion, overthinking and rumination

Other violations of cost-benefit analysis include thinking aversion, overthinking and rumination (Table 2). *Thinking aversion* refers to the notion that the agent may find thinking to be detrimental. This is the case if, for instance, the agent fears he could not neglect a harmful outcome of reasoning. He might then want to commit not to reason, i.e., $u^s >_s^1 u$. But this is a direct violation of Property 7, and is therefore inconsistent with the representation theorems in which the value of reasoning is never negative.¹⁹

A related phenomenon, *overthinking*, refers to the agent having a compulsion to keep reasoning even though he finds it no longer beneficial. That is, he effectively fails to stop thinking even though the cost of reasoning exceeds the benefit. There is thus an inconsistency between the agent’s behavioral preferences over thinking and the mental preferences, since the agent effectively fails to stop thinking even though he would prefer committing to. Therefore, overthinking violates Condition 1 (Reasoning Consistency).

¹⁸As previously discussed, this non-monotonicity in performance would be consistent with Gneezy and Rustichini (2000). This reversion would still be observed for the range of intermediary rewards if, instead, incentives to reason first increase and then decrease, as is consistent with the Yerkes-Dodson Law.

¹⁹An alternative view of *thinking aversion* is that considered by Ortoleva (2013), in which a decision maker may prefer smaller menus so as to avoid the cognitive costs associated with larger menus.

Notice that overthinking is separate from thinking aversion, in that a violation of Property 7 would not, in itself, induce overthinking. If Condition 1 holds but not Property 7, then the value of reasoning may be negative. But then the agent would not perform the next step of reasoning, and therefore he would not be at risk of overthinking. If instead Property 7 holds but not Condition 1, then the value of reasoning would never be negative, but the agent could still perform the next step of reasoning even when this (non-negative) value is lower than the cost. Hence, although the two phenomena likely occur together, they are distinct. Overthinking, by violating Condition 1, entails thinking when it is no longer worth the cost; thinking aversion, by violating Property 7, entails a wish to disregard what would be learned, even if the cost were sunk.

Lastly, *ruminatio*n and *worry* refer to reasoning about problems in a way that is not useful for choice, either because they are in the past (ruminatio)n or because nothing can be done to alter the outcome (worry; see Nolen-Hoeksema, Wisco and Lyubomirsky (2008)). In other words, the agent keeps thinking even though reasoning is not instrumental. This means that even for some choice problem with a u that is constant in actions, $(u, 1)$ could be strictly preferred to $(u, 0)$. This is a violation of Scope property (S.1). Hence, although here as well it is plausible that ruminatio)n and overthinking often occur simultaneously, within our model they map to specific properties, and can be disentangled.

The Scope property (S.1) can be tested directly by comparing whether the agent would strictly prefer to think about a payoff function u that is constant in the action. Testing whether reasoning consistency holds is more subtle, as it requires separating the criterion used, \succcurlyeq_s , from the outcome of that criterion, \succeq_s . Our discussion of overthinking hints at the method, as it concerns the distinction between wanting to stop and actually stopping. Consider a choice problem, and suppose first that the agent is interrupted at a step of reasoning s and asked whether he wishes to continue to reason or commit to stop, in which case he chooses a^s and immediately receives the payoff associated with it.²⁰ If he commits not to reason further, then we conclude that $(u, 0) \succcurlyeq_s (u, 1)$. Now suppose instead that he is asked at that step whether he wishes to continue to reason or to stop, but without the commitment device being offered. That is, in this latter case he is not ‘forced’ to immediately stop thinking even if he has expressed that preference when the commitment was offered. If he remains with the preference not to think an extra step then $(u, 1) \succeq_s (u, 0)$. But if he thinks further when he cannot commit to stop, then $(u, 1) \succeq_s (u, 0)$, which, given $(u, 0) \succcurlyeq_s (u, 1)$, entails a violation of the reasoning consistency condition.

²⁰In practice it is of course difficult to force a subject to stop thinking about a problem; a proxy would be to let the uncertainty resolve immediately and give him his earnings, and move on to the next experimental task so as to shift his attention away. Conducting an experiment to overcome such practical limitations would be a valuable contribution for future research.

6 Conclusions

In this paper we have provided a foundation for a cost-benefit analysis in reasoning, and we have analyzed different properties of the reasoning process together with the representations that they entail. Focusing on the scope of a cost-benefit representation, we have shown that the model delivers new testable predictions in a number of applications. We have also illustrated that our framework can inform rigorous discussions of alternative criteria in the specific context of interest. Focusing on the limits of the approach, we have shown that our model provides a framework to formally disentangle competing theories of important psychological phenomena. Here we discuss some directions of future research.

One direction is to explore further representations of the value of reasoning, following the route suggested in the paper. This agenda, while guided by applied considerations and empirical findings, is particularly well-suited to decision theoretic research. This is made easier by our approach, which adapts standard tools to this new domain. Concerning future experimental research, we note that while our axioms serve to understand the primitive properties underlying cost-benefit in reasoning, a practical limitation that arises in separately testing them is that mental states are not directly observable. Future empirical research requires a method to circumvent this issue; the lottery elicitation method provided in this paper provides one such method.

Another direction of research is more centered around *costs*. Understanding the determinants of the costs of reasoning, which are constant within each cognitive equivalence (c.e.) class, requires a deeper analysis of the cognitive partition. Definition 1 and Condition 2 provide basic restrictions for cognitive equivalence, but these conditions can be extended, along the lines suggested in Section 4.1. The extent to which these restrictions hold could also be investigated empirically, as can the shape of cost functions associated with each c.e. class. Specifically, for a given representation of the value of reasoning, the shape of the cost function may be identified by continuously varying the value of reasoning through suitable changes in the stakes. A partial ranking of cost functions across c.e. classes, as well as the extent to which such a ranking is robust across individuals, could also be identified. This ranking would then be of use to compare difficulties of choice problems, and to shed light on properties of the cost functions.²¹

Yet another direction of research includes linking this framework with a neuroeconomics approach. Brain activity and response times can be useful to elicit the degree of cognitive effort and deliberation, and our model can inform the design of experiments inspired by comparative statics on deliberation as incentives vary. In this vein, it would be interesting to explore the connection between our approach and neuroeconomics models of choice (e.g., Bogacz et al. (2006), Rustichini and Padoa-Schioppa (2014), and Padoa-Schioppa and Rustichini (2015), among others). Combining neuroeconomics with our model, and

²¹If cost functions c and c' associated with c.e. classes C and C' , respectively, are such that $c(k) \geq c'(k)$ for all $k \in \mathbb{N}$, then we can say that choice problems in C are more difficult than problems in C' . See Mathevet (2014) for a classification of games by degree of complexity, based on an axiomatic approach.

the axiomatic approach more broadly, also has the potential to identify c.e. classes and to analyze different dimensions of cognitive sophistication.

From a more applied perspective, ideas of bounded depth of reasoning have recently been applied to market settings as well. For instance, Hortacsu et al. (2016) apply Camerer et al.’s (2003) cognitive hierarchy model to firms’ bidding behavior in electricity markets. They document that larger firms typically exhibit higher strategic sophistication, and they also find evidence that mergers endogenously increase firms’ sophistication. Our model provides a framework to investigate the causes of both phenomena. For instance, an obvious explanation of the positive correlation between size and sophistication is that larger firms have lower ‘costs of reasoning’. While this might be plausible, it is also possible that larger firms indeed may have higher costs. Our model suggests a more nuanced view. In particular, the larger firms’ higher sophistication may be due precisely to the larger stakes of their choices which, given our representations, increase the value of reasoning. Different representations of the value of reasoning would thus provide different ways of connecting the underlying market structure to firms’ strategic sophistication. This in turn can provide the extra structure to identify how much of the increase in strategic sophistication is due to internal characteristics of the firm (as in the ‘lower cost’ explanation), or to the change in the market structure, which affects the value of reasoning.

The research agenda discussed here is non-exhaustive, as the tradeoff between cognitive costs and a notion of value of reasoning is pervasive. There are various important economic environments in which individuals reason in a stepwise fashion, and in which tradeoffs exist between cognitive costs and a value of reasoning. Our model provides a language for analyzing these diverse settings in a unified way.

Appendix

A Proofs

Proof of Theorem 1:

Fix the cognitive equivalence class $C \in \mathcal{C}$. For any k and for any u , let u^k be such that $u^k(a, \omega) = u(a^k, \omega)$ for all (a, ω) . By Condition 1, without loss of generality we use, for $x = 0, 1$, $(u, x) \succcurlyeq_k (u, 1 - x)$ to denote $(u, x) \succcurlyeq_k (u, 1 - x)$.

Part 1 (Cost Identification):

Step 1.1: If $(u, 0) \succcurlyeq_k (u + t, 1)$ for all $u \in \mathcal{U}$ and $t \in \mathbb{R}$, then set $c(k + 1) = \infty$. If not, let $\hat{u}_i \in \mathcal{U}$ be such that $(\hat{u}_i + t, 1) \succcurlyeq_k (\hat{u}_i, 0)$ for some $t' \in \mathbb{R}$. We show next that

$$\exists c(k + 1) \in \mathbb{R}_+ \text{ s.t. } u \in \mathcal{U}, \left(u^k + c(k + 1), 1 \right) \dot{\succ}_k \left(u^k, 0 \right) \text{ for all } u \in \mathcal{U}. \quad (10)$$

To this end, we first show that for the \hat{u}_i above, $\exists c \in \mathbb{R}_+$ s.t. $(\hat{u}_i^k + c, 1) \succ_k (\hat{u}_i^k, 0)$: suppose not, then (using, in order, the hypothesis on \hat{u} , Property 6 and the absurd hypothesis),

$(\hat{u}_i + t, 1) \succcurlyeq_k (\hat{u}_i, 0) \dot{=} (\hat{u}_i^k, 0) \succcurlyeq_k (\hat{u}_i^k + c, 1)$ for all $c \in \mathbb{R}_+$, but $(\hat{u}_i + t, 1) \succcurlyeq_k (\hat{u}_i^k + c, 1)$ for all $c \in \mathbb{R}_+$ contradicts the monotonicity of \succcurlyeq_k^1 (Property 1). Now, since $\exists c \in \mathbb{R}_+$ s.t. $(\hat{u}_i^k + c, 1) \succ_k (\hat{u}_i^k, 0)$, Axiom 3 implies $\exists c^* \in \mathbb{R}_+$ such that $(\hat{u}_i^k + c^*, 1) \dot{=} (\hat{u}_i^k, 0)$. We then set $c(k+1) \equiv c^*$. Note that, by Axiom 4.2, $(\hat{u}_i^k + c^*, 1) \dot{=} (\hat{u}_i^k, 0)$ implies $(u^k + c^*, 1) \dot{=} (u^k, 0)$ for all u .

The following characterization follows: for any $u \in \mathcal{U}$,

$$(u, 1) \succcurlyeq_k (u, 0) \text{ if and only if } (u, 1) \succcurlyeq_k (u^k + c(k+1), 1). \quad (11)$$

(The only if follows by transitivity, using the hypothesis, result (10) and Property 6.)

Step 1.2: Note that, for any $u \in \mathcal{U}$ and for any k , $D(u) = D(u - u^k)$. Hence, Axiom 4.2 implies

$$(u, 1) \succcurlyeq_k (u, 0) \text{ if and only if } (u - u^k, 1) \succcurlyeq_k (u - u^k, 0) \quad (12)$$

Using (11) and (12), it follows that, for any $u \in \mathcal{U}$,

$$(u, 1) \succcurlyeq_k (u, 0) \text{ if and only if } (u - u^k, 1) \succcurlyeq_k \left((u^k - u^k)^k + c(k+1), 1 \right),$$

where $(u - u^k)^k$ is defined such that, for every (a, ω) , $(u - u^k)^k(a, \omega) = u(a^k, \omega) - u^k(a^k, \omega) \equiv 0$. Hence, we conclude that, for any $u \in \mathcal{U}$,

$$(u, 1) \succcurlyeq_k (u, 0) \text{ if and only if } (u - u^k) \succcurlyeq_k^1 c(k+1). \quad (13)$$

Part 2 (Value Identification): From part 1, we have two cases: (1) $(u, 0) \succcurlyeq_k (u + t, 1)$ for all $u \in \mathcal{U}$ and $t \in \mathbb{R}$, in which case we set $c(k+1) = \infty$. In this case, any $V : \mathcal{U} \rightarrow \mathbb{R}$ vacuously represents preferences \succcurlyeq_k : $(u + t, 0) \succcurlyeq_k (u, 1)$ if and only if $V(k+1) + t < c(k+1) = \infty$. (2) The complementary case is such that there exists $\hat{u}_i \in \mathcal{U}$ and t such that $(\hat{u}_i + t, 1) \succcurlyeq_k (\hat{u}_i, 0)$ and $c(k+1)$ such that $(u, 1) \succcurlyeq_k (u, 0)$ if and only if $(u - u^k) \succcurlyeq_k^1 c(k+1)$ for all $u \in \mathcal{U}$. We thus focus here on this second case.

Step 2.1: First we show that, for each $u \in \mathcal{U}$, there exists $t^u \in \mathbb{R}$ such that $(u + t^u, 1) \sim_k u^0$. We consider different cases: (i) if $(u, 1) \dot{=} (u, 0)$ then $t^u = 0$; (ii) if u is constant in a , then $t^u = c(k+1)$; (iii) If $(u, 1) \succ_k (u, 0)$, then $t_i^u < 0$ by Axiom 3; (iv) If $(u, 0) \succ_k (u, 1)$, then there exists sufficiently high $t^* \in \mathbb{R}_+$ such that, $(u + t^*, 1) \succcurlyeq_k (u, 0)$: suppose not, then (by transitivity, Property 6 and the absurd hypothesis) we have $(u^k, 0) \succ_k (u + t, 1)$ for all $t \in \mathbb{R}$. But since u^k is constant, we know from the previous step that, for $t' > c(k+1) \geq 0$, $(u^k + t', 1) \succcurlyeq_k (u^k, 0)$. Hence, by transitivity, $(u^k + t', 1) \succ_k (u + t, 1)$ for all $t \in \mathbb{R}$, which contradicts monotonicity (Property 1). Now, since $\exists t^* \in \mathbb{R}_+$ s.t. $(u + t^*, 1) \succ_k (u, 0)$, property 3 implies $\exists t^u \in \mathbb{R}_+$ such that $(u + t^u, 1) \dot{=} (u, 0)$.

Step 2.2: Thus, $t^u \geq 0$ if and only if $u^0 \succcurlyeq_k u^1$. Furthermore, by axiom 4.2, $t^u = t^{u-u^k}$, because $D(u) = D(u - u^k)$. Consider the following specification of the V function: for each $u \in \mathcal{U}$, let $V(u, k+1) = c(k+1) - t^u$. Now notice that $V(u, k+1) \geq V(v, k+1)$

if and only if $t^u \leq t^v$, but since $t^u = t^{u-u^k}$, $t^u \leq t^v$ if and only if $t^{u-u^k} \leq t^{v-v^k}$, that is if and only if $u - u^k \geq_k^1 v - v^k$. Since $t^{u^k} = c(k+1)$ (because u^k is constant in a), we also have $V(u^k, k+1) = 0$. We show next that $V(u, k+1) \geq 0$ whenever $u \in \mathcal{U}(C)$: to see this, note that Property 7 implies that $u \geq_s^1 u^k$ whenever $u \in \mathcal{U}(C)$, and since $u \in \mathcal{U}(C)$ implies $u - u^k \in \mathcal{U}(C)$ (by Condition 2), Property 7 also implies $u - u^k \geq_k^1 u^k - u^k = \mathbf{0}$ (this is because $u^k - u^k = (u - u^k)^k$). Hence $V(u, k) \geq V(u^k, k) = 0$ for all $u \in \mathcal{U}(C)$.

Step 3 (Asymmetric Rewards): It remains to show that, for any t and u , $(u+t, 1) \succ_k (u, 0)$ if and only if $V(u, k+1) + t \geq c(k+1)$. Given V defined above, clearly $V(u, k+1) + t \geq c(k+1)$ if and only if $t > t^u$. But monotonicity and the definition of t^u imply that $t > t^u$ if and only if $(u+t, 1) \succ_k (u+t^u, 1) \dot{=} (u, 0)$.

Part 4 (Necessity of Axioms 4-3): Axiom 4 follows trivially from the representation; Axiom 3 is trivial; Axiom 5 follows in that for u and v constant, in the representation we have $t^u = t^v = c(k+1)$. by definition from continuity.

Part 5 (Necessity of Properties 1-7): The existence of continuous functions (W^0, W^1) representing $(\geq_k^x)_{x=0,1}$ clearly implies that such binary relations are weak orders (maintained assumption) and satisfy the Archimedean Property 2. Property 1 also follows directly from the monotonicity of W^0 and W^1 . The necessity of Property 7 is due to the restriction $V(u, k+1) \geq 0$ for $u \in \mathcal{U}(C)$ in the theorem, as illustrated by Example 3. Property 6 follows trivially from the stated property of W^0 : $W^0(u) = W^0(u^k)$.

Part 6 (Uniqueness) Suppose that (V', c') also represent the preferences. Then, for any $u \in \mathcal{U}$ and $t \in \mathbb{R}$,

$$\begin{aligned} V'(u, k+1) - c'(k+1) \geq t &\iff V(u, k+1) - c(k+1) \geq t \\ \text{and } V'(u^k, k+1) &= V(u^k, k+1) = 0 \end{aligned}$$

It follows that $c(k+1) = c'(k+1)$, and if $c(k+1) < \infty$, then $V'(u, k+1) = V(u, k+1)$ for every u .

Proof of Theorem 2:

Sufficiency: Relative to Theorem 1, we only need to show that, for any $u \in \mathcal{U}(C)$, $V(\alpha u_i, k) \geq V(u, k)$ whenever $\alpha \geq 1$. First note that Property 7 implies that $D(u, k) = u - u^k \geq_k^1 u^k - u^k = \mathbf{0}$ (cf. Step 2.1 of the proof of Theorem 1). Then, by Properties 2 and 1, there exist $c^u \in \mathbb{R}_+$ such that $c^u \cdot \mathbf{1} \dot{=}^1_k D(u, k) \geq_k^1 \mathbf{0}$ (the positiveness of c^u is due to monotonicity). Property 8 then implies that for $\alpha \geq 1$, $\alpha D(u, k) \geq_k^1 D(u, k)$.

Now, let V be defined as in Step 2.1 of the proof of Theorem 1, then $V(\alpha u_i, k) \geq V(u, k)$ if and only if $t^u \leq t^{\alpha u_i}$ which is the case if and only if $t^{D(u, k)} \leq t^{D(\alpha u_i, k)} = t^{\alpha D(u, k)}$ (the latter equality follows trivially from the definition of D). This in turn is the case if and only if $\alpha D(u, k) \geq_k^1 D(u, k)$, which is satisfied for $\alpha \geq 1$.

Necessity: By contradiction, suppose there exist $u \in \mathcal{U}(C)$, $c \in \mathbb{R}_+$, $\alpha \geq 1 : c \cdot \mathbf{1} \dot{=}^1_k u$ and $u \not\geq_k^1 \alpha u_i$. Now define $\hat{u}_i = u + u^k$, so that $D(\hat{u}_i, k) = u$, hence we have $D(\hat{u}_i, k) \geq_k^1 \alpha D(\hat{u}_i, k)$. By Theorem 1, $V(u, k) = W^1(u - u^k, k)$, that is $V(v, k) > V(u, k)$ if and

only if $W^1(D(v, k)) > W^1(D(u, k))$. Hence, $D(\hat{u}_i, k) \succsim_k^1 \alpha D(\hat{u}_i, k)$ implies $V(u, k) > V(\alpha u_i, k)$, contradicting the stated property of the value of reasoning.

Proof of Theorem 3:

Sufficiency: Given Theorem 1, we only need to show that for any k , the function $V(\cdot, k) : \mathcal{U} \rightarrow \mathbb{R}$ in that representation has the form in eq. 2. To this end, notice that with the addition of Property 9, the suitability relation \succsim_k^1 satisfies the conditions for the mixture space theorem. Hence, for any $k \in \mathbb{N}$, there exists a function $\hat{W}^1(\cdot, k) : \mathcal{U} \rightarrow \mathbb{R}$ that represents \succsim_k^1 and satisfies $\hat{W}^1(\alpha u_i + (1 - \alpha)v, k) = \alpha \hat{W}^1(u, k) + (1 - \alpha) \hat{W}^1(v, k)$ for all $\alpha \in [0, 1]$ and $u, v \in \mathcal{U}$. Moreover, $\hat{W}^1(\cdot, k)$ is unique up to positive affine transformations. Because $\hat{W}^1(\cdot, k)$ is linear in $u \in \mathcal{U} = \mathbb{R}^{|A \times \Omega|}$, there exist $(\hat{\rho}(a, \omega))_{(a, \omega) \in A \times \Omega} \in \mathbb{R}^{|A \times \Omega|}$ s.t. $W^1(u, k) = \sum_{a, \omega} \hat{\rho}(a, \omega) \cdot u(a, \omega)$. By monotonicity, $\hat{\rho}(a, \omega) \geq 0$ for each (a, ω) , and we define $\rho^k \in \Delta(A \times \Omega)$ normalizing such weights in the unit simplex, so that

$$\rho^k(a, \omega) = \begin{cases} \hat{\rho}(a, \omega) / \sum_{(a', \omega')} \hat{\rho}(a', \omega') & \text{if } \sum_{(a', \omega')} \hat{\rho}(a', \omega') > 0 \\ \frac{1}{|A \times \Omega|} & \text{otherwise} \end{cases}$$

Since this is a positive affine transformation, $W^1(u, k) = \sum_{a, \omega} \rho^k(a, \omega) \cdot u(a, \omega)$ also represents \succsim_k^1 by the uniqueness part of the mixture space theorem. For any such ρ^k , define $p^k \in \Delta(A)$ and $\mu = (\mu^a)_a \in \Delta(\Omega)^A$ as follows: for any $a \in A$, let $p^k(a) = \sum_{\omega} \rho^k(a, \omega)$ and define $\mu^a \in \Delta(\Omega)$ such that, for any ω ,

$$\mu^a(\omega) = \begin{cases} \frac{\rho^k(a, \omega)}{p^k(a)} & \text{if } p^k(a) > 0 \\ \frac{1}{|\Omega|} & \text{otherwise} \end{cases}.$$

It is immediate to see that, for any (a, ω) , $\rho^k(a, \omega) = p^k(a) \cdot \mu^a(\omega)$. Hence, without loss of generality we can represent $W^1(\cdot, k)$ as follows:

$$W^1(u, k) = \sum_a p^k(a) \sum_{\omega} \mu^a(\omega) \cdot u(a, \omega). \tag{14}$$

We show next that for each $u \in \mathcal{U}(C)$, $a \in BR_i^u(\mu^a)$ whenever there exists $\omega \in \Omega$ such that $\rho(a, \omega) > 0$.

Suppose not. Then $\exists \hat{a}$ s.t. $\rho(\hat{a}, \omega) > 0$ for some ω s.t. $\hat{a} \notin BR_i^u(\mu^{\hat{a}})$. Then, let $a^* \in BR_i^u(\mu^{\hat{a}})$ and consider payoff function $u^{\hat{a} \rightarrow a^*}$ (that is, the payoff function identical

to u except that the payoffs associated to action \hat{a} are replaced by those of a^* . Then,

$$\begin{aligned}
W^1(u^{\hat{a} \rightarrow a^*}, k) &= \sum_{a, \omega} \rho^k(a, \omega) \cdot u^{\hat{a} \rightarrow a^*}(a, \omega) \\
&= \sum_{(a, \omega): a \neq \hat{a}} \rho^k(a, \omega) \cdot u(a, \omega) + \sum_{\omega \in \Omega} \rho^k(\hat{a}, \omega) \cdot u(a^*, \omega) \\
&> \sum_{(a, \omega): a \neq \hat{a}} \rho^k(a, \omega) \cdot u(a, \omega) + \sum_{\omega \in \Omega} \rho^k(\hat{a}, \omega) \cdot u(\hat{a}, \omega) \\
&= W^1(u, k).
\end{aligned}$$

But this conclusion contradicts Property 10. It follows that

$$W^1(u, k) = \sum_a p^k(a) \sum_{\omega} \mu^a(\omega) \cdot u(a^*(\mu), \omega) \text{ for all } u \in \mathcal{U}(C). \quad (15)$$

Notice next that functional $V(\cdot, k+1)$ in Theorem 1 represents the binary relation \succsim^* defined as $u \succsim^* v$ if and only if $u - u^k \succcurlyeq_k^1 v - v^k$. Since $W^1(\cdot, k)$ represents \succcurlyeq_k^1 , it follows that $u \succsim^* v$ if and only if $W^1(u - u^k, k) \geq W^1(v - v^k, k)$. Thanks to Property 9 we can thus set $V(u, k+1) = W^1(u - u^k, k)$. Since

$$W^1(u - u^k, k) = \sum_a p^k(a) \sum_{\omega} \mu^a(\omega) \cdot \left[u(a^*(\mu^a), \omega) - u(a^k, \omega) \right],$$

the representation follows from Theorem 1, noticing that $W^1(u^k - u^k, k) = 0$, and that (by Theorem 1) $u^1 \succcurlyeq_k u^0$ if and only if $u - u^k \succcurlyeq_k^1 c(k)$.

Necessity: Assume that

$$\begin{aligned}
V(u, k+1) &= W^1(u - u^k, k) \\
&= \sum_a p^k(a) \sum_{\omega} \mu^a(\omega) \cdot \left[u(a^*(\mu^a), \omega) - u(a^k, \omega) \right].
\end{aligned}$$

Then, for any u , functions W^1 represents \succcurlyeq_k^1 and is such that

$$W^1(u) = \sum_{a \in A} p^k(a) \sum_{\omega \in \Omega} \mu^a(\omega) \cdot u(a^*(\mu), \omega).$$

It is immediate to verify that \succcurlyeq_k^1 satisfies Properties 9 and 10. ■

Proof of Theorem 4:

Sufficiency: By contradiction, suppose there exists a' s.t. $\sum_{\omega} \hat{\mu}^k(\omega) u(a', \omega) > \sum_{\omega} \hat{\mu}^k(\omega) u(a^k, \omega)$. Then, substituting the definition of $\hat{\mu}^k$, we obtain

$$\sum_{\mu \in \Delta(\Omega)} p^k(\mu) \sum_{\omega} \mu(\omega) \left[u(a', \omega) - u(a^k, \omega) \right] > 0 \quad (16)$$

Define $v \in R(u, a^k)$ such that $v(a, \omega) = u(a, \omega)$ for all $a \neq a^k$ and $v(a^k, \omega) = u(a', \omega)$, we have $v^k \succ_k^1 u^k$ because eq. 14 represents \succ_k^1 . But $v^k \succ_k^1 u^k$ contradicts Property 11.

Necessity: it holds by construction, for the definition of $\hat{\mu}^k$. ■

Proof of Theorem 5:

Sufficiency: From the proof of Theorem 3, it follows that under the assumptions of Theorem 1 plus Property 9, the value of reasoning takes the following form:

$$V(u, k+1) = W^1(u - u^k) = \sum_{(a, \omega)} \rho^k(a, \omega) \cdot [u(a, \omega) - u(a^k, \omega)].$$

We need to show that, adding Property 12, we have $\rho^k(a, \omega) = 0$ whenever $(a, \omega) \neq (a^*(\omega^*), \omega^*)$, where $\omega^* \in \arg \max_{\omega \in \Omega} u(a^*(\omega), \omega) - u(a^k, \omega)$. By contradiction, suppose that there exists $(a', \omega') \neq (a^*(\omega^*), \omega^*)$ s.t. $\rho^k(a, \omega) > 0$. Now, let v be such that

$$v(a, \omega) = \begin{cases} u(a^*(\omega^*), \omega^*) & \text{if } (a, \omega) = (a', \omega') \\ u(a, \omega) & \text{otherwise} \end{cases}.$$

Then, by construction:

$$\begin{aligned} V(v, k+1) &= W^1(v - v^k) = \sum_{(a, \omega)} \rho^k(a, \omega) \cdot [v(a, \omega) - v(a^k, \omega)] \\ &= \rho^k(a', \omega') [u(a^*(\omega^*), \omega^*) - u(a^k, \omega^*)] + \sum_{(a, \omega) \neq (a', \omega')} \rho^k(a, \omega) \cdot [u(a, \omega) - u(a^k, \omega)] \\ &> \rho^k(a', \omega') [u(a', \omega') - u(a^k, \omega')] + \sum_{(a, \omega) \neq (a', \omega')} \rho^k(a, \omega) \cdot [u(a, \omega) - u(a^k, \omega)] \\ &= V(u, k+1) = W^1(u - u^k). \end{aligned}$$

Hence, $v - v^k \succ_k^1 u - u^k$. But since $v \in \Gamma(u)$ (by construction), this contradicts Property 12.

Necessity: trivial. ■

B AS's attention allocation experiment

The following matrices summarize all the predictions of the model's in Section 4.2.1 for the allocation tasks in AS's experiment: in each matrix, different rows and columns correspond to different games considered by AS (the descriptions of the games associated to the labels in the rows and columns is left to the appendix.) Each cell in the matrices yields the relative ranking between α^1 and α^2 , where G^1 denotes the game in the corresponding row and G^2 in the corresponding column. Blank cells correspond to comparisons for which the

model is indeterminate.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
(1) <i>Chance</i>	=	=	<	=	<						
(2) <i>PC</i> ₅₀₀	=	=	<	=	<						
(3) <i>PC</i> ₈₀₀	>	>	=	>	=						
(4) <i>BS</i> ₅₀₀	=	=	<	=	<						
(5) <i>BS</i> ₈₀₀	>	>	=	>	=						
(6) <i>CS</i> ₅₀₀						=	<				
(7) <i>CS</i> ₈₀₀						>	=				
(8) <i>CS</i> ₄₀₀								=			
(9) <i>PD</i> ₈₀₀									=	>	>
(10) <i>PD</i> ₅₀₀									<	=	>
(11) <i>PD</i> ₃₀₀									<	<	=

	<i>PC</i> ₈₀₀	<i>PC</i> _{$\frac{800}{500}$}	<i>PC</i> _{$\frac{800}{500}$} ₁₀₀		<i>BS</i> ₅₀₀	<i>BS</i> _{$\frac{800}{0}$}	<i>BS</i> _{$\frac{800}{100}$}
<i>PC</i> ₅₀₀	<	<			<i>BS</i> ₈₀₀	>	
<i>PC</i> ₈₀₀	=	>			<i>BS</i> ₅₀₀	=	<
<i>PC</i> _{$\frac{800}{500}$}	<	=			<i>BS</i> _{$\frac{800}{0}$}	=	
	<i>PD</i> ₅₀₀	<i>PD</i> ₃₀₀	<i>PD</i> _{$\frac{800}{500}$}				
<i>PD</i> ₈₀₀	>	>					
<i>PD</i> ₅₀₀	=	>					
<i>PD</i> ₃₀₀	<	=					

We note that the theoretical predictions are in terms of weak ordering, and they are all confirmed. In fact, with the exception of the $(BS_{800}, Chance)$ and (PC_{800}, BS_{500}) comparisons, for which AS find that the time allocation is no different from uniform at a statistically significant level, all the other data are consistent with the theoretical predictions in the strict ordering sense, at a statistically significant level (the vast majority at the 5% level).

References

1. Alaoui, L. and A. Penta. 2016a. “Endogenous Depth of Reasoning,” *Review of Economic Studies*, forthcoming.
2. Alaoui, L. and A. Penta. 2016b. “Endogenous Depth of Reasoning and Response Time, with an application to the Attention-Allocation Task ,” *mimeo* , UPF
3. Alaoui, L. and A. Penta. 2018. “Cost-Benefit Analysis in Reasoning, Supplementary Appendix: The Value of Information Case with Forward Looking Agent,” *mimeo* , UPF

4. Aragonès, E., I. Gilboa, A. Postlewaite and D. Schmeidler. 2005. "Fact-free Learning," *American Economic Review*, 95(5): 1355-1368.
5. Ariely, D., U. Gneezy, G. Loewenstein, N. Mazar. 2009. "Large Stakes and Big Mistakes," *Review of Economic Studies*, 76: 451-469.
6. Bernheim, D. 1984. "Rationalizable Strategic Behavior," *Econometrica*, 52(4): 1007-1028.
7. Bogacz R., Brown E., Moehlis J., Holmes P. and JD. Cohen. 2006. "The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks," *Psychological Review*, 113: 700-765.
8. Camerer, C. F. 2003. *Behavioral Game Theory*, Princeton Univ. Press, Princeton, NJ.
9. Camerer, C. F., G. Loewenstein and D. Prelec. 2005. "Neuroeconomics: How neuroscience can inform economics," *Journal of Economic Literature*, 34(1): 9-65.
10. Camerer, C. F. and R. Hogarth. 1999. "The effects of financial incentives in economics experiments: A review and capital-labor-production framework," *Journal of Risk and Uncertainty*, 19(1): 7-42
11. Caplin, A. and M. Dean. 2008. "Axiomatic Neuroeconomics," in: *Neuroeconomics: Decision Making and the Brain*, Chapter 3, P.W. Glimcher, C.F. Camerer, E. Fehr and R.A. Poldrack (Eds.), Academic Press, New York, 21–32.
12. Caplin, A. and M. Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105(7): 2183-2203
13. Costa-Gomes, M.A., and V. P. Crawford. 2006. "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, 96(5): 1737-1768.
14. Crawford, V. P., M. A. Costa-Gomes, and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications," *Journal of Economic Literature*, 51(1): 5–62.
15. 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.
16. Dean, M. and N. Neligh. 2017. "Experimental Tests of Rational Inattention," *mimeo*.
17. De Finetti, B.. 1937. "La Prévision: Ses Lois Logiques, Ses Sources Subjectives," *Annales de l'Institut Henri Poincaré*, 7: 1-68.
18. Dekel, E., B. L. Lipman, and A. Rustichini. 1998. "Standard state-space models preclude unawareness," *Econometrica*, 66(1): 159–173.

19. Dekel, E., B. L. Lipman, and A. Rustichini. 2001. "Representing preferences with a unique subjective state space," *Econometrica*, 69(4): 891–934.
20. Dekel, E. and M. Siniscalchi. 2014. "Epistemic Game Theory" in: *Handbook of Game Theory*, vol. 4.
21. Gabaix, X. and D. Laibson. 2005. "Bounded rationality and directed cognition," *Mimeo*.
22. Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg. 2006. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model," *American Economic Review*, 96(4): 1043–1068.
23. Gigerenzer, Gerd and Wolfgang Gaissmaier. 2011. "Heuristic Decision Making," *Annual Review of Psychology*, 62:451–82.
24. Gilboa, I. 2009. *Theory of Decision under Uncertainty*, Econometric Society Monographs, Cambridge Univ.Press, Cambridge, MIT.
25. Gilboa, I., A. Postlewaite, and D. Schmeidler. 2012. "Rationality of Beliefs or: Why Savage's Axioms are neither Necessary nor Sufficient for Rationality," *Synthese*, 187, 11–31.
26. Gilboa, I., L. Samuelson and D. Schmeidler. 2013. "Dynamics of Inductive Inference in a Unified Framework," *Journal of Economic Theory*, 148, 1399–1432.
27. Gneezy, U. and A. Rustichini. 2000. "Pay Enough or Don't Pay At All," *Quarterly Journal of Economics*, 115(3), 791–810.
28. Goeree, J. K., and C. A. Holt. 2001. "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5): 1402-1422.
29. Hortacsu, A., Luco, F., Puller, S. L. and D. Zhu. 2016. "Strategic Ability and Productive Efficiency in Electricity Markets," *mimeo*, Chicago.
30. Manzini, P. and M. Mariotti. 2007. "Sequentially Rationalizable Choice," *American Economic Review*, 97(5) 1824-1839.
31. Manzini, P. and M. Mariotti. 2014. "Stochastic Choice and Consideration Sets," *Econometrica*, 82(3): 1153-1176.
32. Mathevet, L. 2014. "Axiomatic Behavior in Repeated Games," *mimeo*, NYU.
33. Modica, S. and A. Rustichini. 1999. "Unawareness and Partitional Information Structures," *Games and Economic Behavior* 27: 265-298.
34. Nagel, R. 1995. "Unraveling in Guessing Games: An Experimental Study," *American Economic Review*, 85(5): 1313-1326.

35. Neiss, R. 1988. "Reconceptualizing Arousal: Psychological States in Motor Performance," *Psychological Bulletin*, 103(3): 345–366.
36. Nolen-Hoeksema, S., B.E. Wisco, S. Lyubomirsky. 2008. "Rethinking Rumination," *Perspectives on Psychological Science*, 3(5): 400–424.
37. Ortoleva, P. 2012. "Modeling the Change of Paradigm: Non-Bayesian Reaction to Unexpected News," *American Economic Review* 102(6): 2410–2436.
38. Ortoleva, P. 2013. "The price of flexibility: Towards a theory of Thinking Aversion," *Journal of Economic Theory* 148: 903–934.
39. Padoa-Schioppa, C. and A. Rustichini. 2015. "A neuro-computational model of economic decisions," *Journal of Neurophysiology*, 114: 1382–1398.
40. Pearce, D. 1984. "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52, 1029–1050.
41. Perea, A. 2012. *Epistemic Game Theory* Cambridge Univ. Press.
42. Rubinstein, A. and Y. Salant. 2008. "Some Thoughts on the Principle of Revealed Preference," in: *The Foundations of Positive and Normative Economics: A Handbook*, Chapter 5, A. Caplin and A. Schotter (Eds.), Oxford University Press, Oxford.
43. Padoa-Schioppa, C. and A. Rustichini. 2014. "Rational Attention and Adaptive Coding: A Puzzle and a Solution," *American Economic Review: Papers & Proceedings*, 104(5): 507–513.
44. Rustichini, A. and Padoa-Schioppa, C. 2015. "A neuro-computational model of economic decisions," *Journal of Neurophysiology*, 114: 1382–1398.
45. Schipper, B. 2015. "Awareness," in: *Handbook of Epistemic Logic*, Chapter 3, H. van Ditmarsch, J.Y. Halpern, W. van der Hoek and B. Kooi (Eds.), College Publ., London.
46. Tversky, A. 1972. "Elimination by Aspects: A Theory of Choice," *Psychological Review*, 79: 281–99.