Eliciting prior beliefs*

Elias Tsakas†

Maastricht University

First Draft: July 2017
This draft: January, 2018

Abstract

We study elicitation of latent prior beliefs when the agent can acquire information via a costly attention strategy. We introduce a mechanism that simultaneously makes it strictly dominant to (a) not acquire any information, and (b) report truthfully. We call such a mechanism a *robust scoring rule*. Robust scoring rules are crucial for lab experiments, e.g., they are notably needed for testing Bayesian rationality. We prove that a robust scoring rule exists under mild axioms. These axioms are shown to characterize the class of posterior-separable cost functions. Our existence proof is constructive, thus identifying an entire class of robust scoring rules for each posterior-separable cost function. For the most common special case (viz., with entropic attention costs), we characterize the class of robust quadratic scoring rules by means of a simple inequality. Finally, we discuss potential experimental designs for testing Bayesian rationality.

Keywords: Belief elicitation, robust scoring rules, rational inattention, hidden information costs, posterior-separability, Shannon entropy, testing Bayesian rationality.
JEL codes: C91, D81, D82, D83, D87.

1. Introduction

Background and motivation. Subjective beliefs constitute one of the most common latent variables of interest in economics (e.g., Manski, 2004). Having recognized this, statisticians and economists have developed mechanisms, called *(proper) scoring rules*, that incentivize the economic agent to reveal his true latent belief, irrespective of which this belief is. Due to their solid theoretical foundations (i.e., the fact that they are incentive-compatible), proper scoring rules have been extensively used in laboratory experiments and in various applications.

One of the main concerns with scoring rules is that the mechanism itself may affect the very same beliefs it tries to elicit. As Schotter and Trevino (2014, p.109) eloquently put it,

> *"the very act of belief elicitation may change the beliefs of subjects from their true latent beliefs or the beliefs they would hold (respond to) if those beliefs were not elicited (we might have a type of Heisenberg problem)".*

*I am indebted to John Rehbeck, Burkhard Schipper and Peter Wakker for their valuable comments on earlier drafts of the paper. I would also like to thank Chris Chambers, Madhav Chandrasekher, Paul Heidhues, Dorothea Kübler, Andy McKenzie, Marcus Pivato, Mathias Staudigl, Stefan Terstiege, Nikolas Tsakas, Mark Voorneveld and the seminar audience in the Mathematics Department at the University of Athens for fruitful discussions. Finally I would like to thank the Economics Department at UC Davis for its hospitality while working on this project.

†Department of Economics (AE1), Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands; Homepage: www.elias-tsakas.com; E-mail: e.tsakas@maastrichtuniversity.nl

This problem is particularly relevant in the lab, especially in experiments that test Bayesian rationality (see Section 5). The order typically followed in such experiments is to first let the subjects make their choices and subsequently elicit their beliefs about payoff-relevant events. In this case it is crucial to elicit the subjects' prior beliefs (viz., the ones they held when they made their choice), rather than some updated beliefs that they may form after paying extra attention to the problem in an attempt to exploit the incentives provided by the scoring rule. Hence our aim is to identify proper scoring rules that elicit the subjects' prior beliefs, so that we can test whether their choices are rational given the beliefs they held at the moment of their decisions.

Model and results. We consider scoring rules in a model with hidden information costs, which typically emerge as an expression of rationally inattentive preferences (for an overview, see Caplin, 2016). In our formal model, there is a (male) agent – henceforth called the subject – who has a latent (prior) probabilistic belief for some fixed event. A (female) experimenter wants to elicit this belief, and to this end she asks the agent to report it. In order to incentivize him to report truthfully, she designs a scoring rule that rewards the subject on the basis of his report and the realization of the event. Before stating his report, the agent can acquire information through a *costly attention strategy* and then reports his belief after having perhaps updated his prior (see Figure 1 for the timeline).

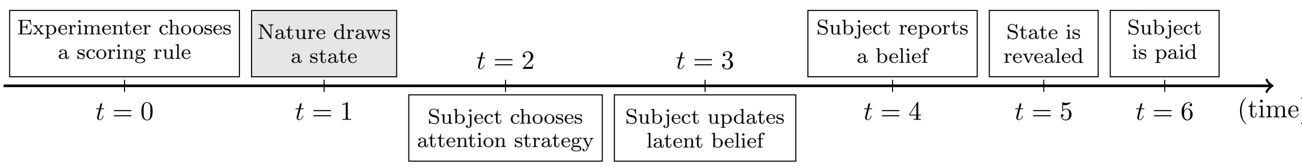

Figure 1: Boxes above the line are observed symmetrically by the subject and the experimenter. Boxes below the line are only observed by the subject. The shaded box is observed with a delay, i.e., it is realized at $t = 1$ and observed at $t = 5$.

In order to elicit the subject's prior belief, the scoring rule must make it simultaneously (a) strictly dominant not to acquire any information (i.e., to choose the zero-attention strategy), and (b) strictly dominant to report truthfully (i.e., the scoring rule is proper). Such a mechanism is called *robust scoring rule*. Two natural questions arise then. *Is there a robust scoring rule?* And if yes, *how does it look like?* Note that, in expectation, every attention strategy yields a benefit (due to the fact that reporting is postponed till after the beliefs have been updated) and a cost (due to the fact that information acquisition is costly). Thus, the experimenter's problem boils down to finding a scoring rule that provides enough incentives for the agent to report truthfully, but not strong enough to offset the costs of acquiring information.

Our main theorem shows that robust scoring rules exist under some mild axioms on the attention costs (Theorem 1). First, as always assumed in the literature, the only costless attention strategy is the one that carries no information. Second, we impose a dynamic consistency axiom, which states that the cost of attention depends only on the distribution of the posteriors and not on the process that yields this distribution. It turns out that our axioms characterize a well-known class of cost functions that has recently attracted attention in the literature, viz., the class of posterior-separable attention costs. Posterior-separability has solid theoretical foundations (Caplin et al., 2017; Zhong, 2017) and is supported by recent experimental evidence (Dean and Neligh, 2017). The proof of our theorem is constructive. Notably, not only do we show existence, but we also explicitly identify an entire class of robust scoring rules for each posterior-separable cost function. In this sense, our theory has strong empirical content.

Then we focus on entropic attention costs (Sims, 2003; Caplin et al., 2017), which are very common in applications and suitable for parametric tests of experimental data. Entropic costs are posterior-separable, thus guaranteeing, by our previous theorem, that a robust scoring rule exists. Then going a step further, we ask whether a quadratic scoring rule exists under the entropic cost specification. Quadratic scoring rules are widely-used in lab experiments, primarily because they are easily understood by the experimental subjects. Our second theorem then fully characterizes the class of robust quadratic scoring rules when the attention costs are entropic (Theorem 2). This result too has strong empirical content, as our characterization is by means of a simple inequality that contains the (multiplier) parameter of the cost function and the (curvature) parameter of the scoring rule.

Applications. As we have already mentioned, robust scoring rules are particularly important for testing Bayesian rationality in the lab. However, this is not the only application of our theory. In fact, beliefs are often used as an explanatory variable for behavior in a wide range of experiments, such as public good games (Fischbacher and Gächter, 2010), trust games (Costa-Gomes et al., 2014), voting experiments (Duffy and Tavits, 2008) and asset market experiments (Haruvy et al., 2007), just to mention a few. Overall, our theory primarily applies to cases where we are interested in learning the subjects' beliefs, rather than the actual state of the world. This is why we mostly focus on scoring rules being used as an experimental tool, rather than as an incentive scheme.

Related literature. This paper contributes to two different streams of literature, viz., belief elicitation via scoring rules and rational inattention. The following quick overview is by no means exhaustive, but merely provides a taxonomy of the main directions that have been explored so far in the literature.

Scoring rules were originally introduced by meteorologists (Brier, 1950), before being further developed by statisticians (Good, 1952; McCarthy, 1956; Savage, 1971), and eventually being adopted by several disciplines, such as economics, accounting, business, management, psychology, political science and computer science (see Offerman et al., 2009, p. 1462). Within economics, the theory of scoring rules has mostly focused on introducing new mechanisms (Hossain and Okui, 2013; Karni, 2009), on relaxing the underlying assumptions of the standard mechanisms, such as for instance risk-neutrality (Savage, 1971; Offerman et al., 2009; Schlag and van der Weele, 2013) or the expected utility hypothesis (Karni, 1999; Chambers, 2008; Offerman et al., 2009), and on understanding the technical relationships to other economics models (Chambers, Healy and Lambert, 2017). Scoring rules are also used in various economic applications, focusing for instance on incentive schemes in organizations (Thomson, 1979), information markets (Hanson, 2003; Ostrovsky, 2012) and strategic indistinguishability (Bergemann et al., 2017). Finally, there is large experimental literature, focusing primarily on the role of risk-aversion (Offerman et al., 2009; Armantier and Treich, 2013), the comparison of deterministic and stochastic scoring rules (Selten et al., 1999; Harrison et al., 2013, 2014) and the elicitation of beliefs in games (Nyarko and Schotter, 2002; Costa-Gomes and Weizsäcker, 2008; Palfrey and Wang, 2009). For two recent literature reviews, we refer to Schotter and Trevino (2014) and Schlag et al. (2015).

Rational inattention models first appeared in macroeconomics (Sims, 2003, 2006; Maćkowiak and Wiederholt, 2009), before attracting interest of microtheorists. The latter have mostly focused on providing axiomatic foundations (De Oliveira et al., 2017; Ellis, 2018) and on designing revealed-preference tests for identifying the attention costs from choice data (Caplin and Dean, 2015; Chambers, Liu and Rehbeck, 2017; Caplin et al., 2017). Recently, there is interest in dynamic models of rational inattention (Hébert and Woodford, 2016; Morris and Strack, 2017; Zhong, 2017). There are also various economic applications of rational inattention – usually with entropic costs – on topics like discrimination (Bartoš et al., 2016), pricing (Matejka, 2016) and electoral competition (Matejka and Tabellini, 2016). Finally, there is recent work on experimentally testing models of rational inattention (Dean and Neligh, 2017). For an overview of this literature, see Caplin (2016).

Of particular interest is the relationship between our paper and the one of Chambers and Lambert (2017) in that they are among the handful of papers that study dynamic belief elicitation. The only other paper is the one by Karni (2017).[1] In their paper, Chambers and Lambert (2017) consider an agent who has a latent prior belief and receives new information over time based on an exogenously given dynamic process. Then, they construct a mechanism which makes it incentive-compatible for the agent to simultaneously reveal his prior, his anticipated information flow and his realized posteriors. The conceptual difference to our paper is that the agent does not strategically choose the process of his information flow (viz., the attention strategy in our terminology). Moreover, the two papers differ in the formal approaches that they employ, viz., as opposed to our paper, their mechanism does not rely on the usual subgradient characterization, but rather on a randomization technique originally introduced by Allais (1953). On the other hand, a major similarity is that both our paper and the one of Chambers and Lambert (2017) truthfully elicit the agent's prior beliefs.

Structure of the paper. In Section 2 we introduce our model. In Section 3 we introduce our axioms, and we state and prove our main result. In Section 4 we study the special case with entropic attention costs. In Section 5 we discuss potential experimental designs to test Bayesian rationality in the lab. Section 6 contains a discussion. All proofs are relegated to the Appendix.

2. Robust scoring rules

Proper scoring rules. Consider a binary state space $\Omega = \{\omega_0, \omega_1\}$. A risk-neutral (male) experimental subject has a latent subjective belief $\mu_0 \in [0,1]$ of ω_0 occurring, which is not observed by the (female) experimenter. The subject is asked to state μ_0 and reports some $r \in [0,1]$, which is not necessarily equal to μ_0. A *scoring rule* is a function

$$S : [0,1] \times \Omega \to \mathbb{R},$$

chosen by the experimenter, which takes the subject's report (r) and the realized state (ω) as an input and returns a monetary payoff $(S_r(\omega))$ as an output. In economics we sometimes consider stochastic scoring rules where the subject is paid in probabilities of winning a fixed prize. Stochastic scoring rules are used to elicit the subject's belief for arbitrary risk attitudes. In statistics on the other hand the image of S is often allowed to take values in $\overline{\mathbb{R}} = [-\infty, \infty]$, in order to deal with a common subdifferentiability issue that often appears with ordinary scoring rules like the ones described above. For a discussion on such scoring rules, see Section 6.

The subject is assumed to maximize Subjective Expected Utility (SEU), i.e., given the scoring rule (S) and his actual belief (μ_0), he chooses the report (r) that maximizes

$$\mathbb{E}_{\mu_0}(S_r) := \mu_0 S_r(\omega_0) + (1-\mu_0) S_r(\omega_1).$$

A scoring rule is called proper whenever it is strictly dominant for the subject to report his true latent belief, irrespective of what this belief is. Formally, S is a *proper scoring rule*, whenever

$$\mathbb{E}_{\mu}(S_{\mu}) > \mathbb{E}_{\mu}(S_r) \tag{1}$$

for every $r \neq \mu$ and every $\mu \in [0,1]$ (Brier, 1950; Good, 1952). The most commonly used proper scoring rule is the quadratic (QSR), which is defined by $S_r(\omega_0) := \alpha - \beta(1-r)^2$ and $S_r(\omega_1) := \alpha - \beta r^2$,

[1] I am indebted to Chris Chambers for pointing out these connections.

where $\alpha \in \mathbb{R}$ and $\beta > 0$. For a review of the standard proper scoring rules, we refer to Schlag et al. (2015, Section 2).

It is well-known that a proper scoring rule is characterized by a class of strictly convex functions (McCarthy, 1956; Savage, 1971). In particular, the scoring rule S is proper if and only if there exists a strictly convex and subdifferentiable function $\phi : [0,1] \to \mathbb{R}$ such that S_r is a subtangent line at r evaluated 1 and 0 respectively. Formally, let ϕ be such that, for each $r \in [0,1]$ there are $a_r, b_r \in \mathbb{R}$ such that $\phi(s) \geq a_r + b_r s$ for all $s \in [0,1]$, with equality holding if and only if $s = r$. In this case, $S_r(\omega_0) := a_r + b_r$ and $S_r(\omega_1) := a_r$ is a proper scoring rule. Strict convexity of ϕ guarantees that (1) holds, while subdifferentiability at the boundary guarantees that the subtangent of ϕ is not vertical and therefore S is well-defined. This last condition can be dispensed with if we allow S to take values in $\overline{\mathbb{R}}$ instead of $\mathbb{R}$ as often done in statistics. Whenever S is characterized by ϕ, it is the case that

$$\phi(\mu) = \mathbb{E}_\mu(S_\mu),$$

i.e., $\phi(\mu)$ is the subject's expected utility when he reports his true belief. As an example, the QSR is characterized by the function $\phi_\beta(\mu) = \alpha - \beta\mu(1-\mu)$. For an overview of the subgradient characterization of proper scoring rules, see Gneiting and Raftery (2007).

Costly attention. We now enrich the agent's preferences to allow for information acquisition by means of costly attention (Sims, 2003). An *attention strategy* is a signal, designed by the subject in an attempt to form updated "more precise" subjective beliefs. Given his prior μ_0, each attention strategy is identified by a (Bayes-plausible) distribution over posteriors, chosen from the set $\Pi(\mu_0) := \{\pi \in \Delta([0,1]) : \int_0^1 \mu d\pi = \mu_0\}$. We define the degenerate zero-attention strategy, $\hat{\mu}_0 \in \Pi(\mu_0)$, that puts probability 1 to the prior μ_0. For notation simplicity, we henceforth denote by $\hat{\Pi}(\mu_0) := \Pi(\mu_0) \setminus \{\hat{\mu}_0\}$ the set of non-degenerate attention strategies. If $\mu \in \{0,1\}$ then $\hat{\Pi}(\mu) = \emptyset$. Given the prior μ_0 and a scoring rule ϕ, the (expected) benefit of an attention strategy $\pi \in \Pi(\mu_0)$ is equal to

$$B_\phi(\pi) := \langle \phi, \pi \rangle - \phi(\mu_0),$$

where $\langle \phi, \pi \rangle := \mathbb{E}_\pi(\phi)$ denotes the usual inner product duality. Since ϕ is strictly convex, we obtain $B_\phi(\pi) \geq 0$, with equality holding if and only if $\pi = \hat{\mu}_0$. That is, attention always has strictly positive benefits when the scoring rule is proper. However, attention is also costly. In particular, there is a non-negative-valued cost function,

$$C : \Delta([0,1]) \to \mathbb{R}_+$$

assigning a cost to each attention strategy. Obviously the cost does not depend on the scoring rule, but only on the attention strategy. Attention costs can be identified from choice data (Caplin and Dean, 2015; Chambers, Liu and Rehbeck, 2017). They are part of the standard axiomatizatic characterizations of rationally inattentive preferences (De Oliveira et al., 2017; Ellis, 2018). The common entropic cost function is discussed in Section 4.

Cost-benefit analysis. Given a prior μ_0 and a proper scoring rule ϕ, the subject will choose an attention strategy in $\Pi(\mu_0)$ that maximizes the value

$$V_\phi(\pi) := B_\phi(\pi) - C(\pi).$$

After (optimally) choosing some π, the subject will first update his beliefs to some – also latent – posterior $\mu \in \text{supp}(\pi)$, and then – as ϕ is proper – he will truthfully report his posterior belief μ. Therefore, in order to guarantee that the agent will report his prior belief μ_0, it must be the case that

$\hat{\mu}_0$ is a strictly dominant attention strategy. Whenever this is the case for every prior, we say that the scoring rule is robust. Formally, ϕ is a *robust scoring rule*, whenever it is proper and satisfies

$$V_\phi(\hat{\mu}) > V_\phi(\pi) \tag{2}$$

for every $\pi \in \hat{\Pi}(\mu)$ and every $\mu \in [0,1]$. Then, we naturally ask: *is there a robust scoring rule?* And if yes, *how does it look like?*

3. Existence of robust scoring rules

As turns out, a robust scoring rule exists under some mild regularity conditions, imposed on the cost function. In what follows in this section, we prove existence constructively, thus identifying an entire family of robust scoring rules.

3.1. Axioms

We begin with two standard conditions, postulating that the zero-attention strategy is costless, whereas every other attention strategy is costly. Formally:

(C_1) NORMALIZATION: $C(\hat{\mu}) = 0$ for all $\mu \in [0,1]$.

(C_2) ATTENTION IS COSTLY: $C(\pi) > 0$ for all $\pi \in \hat{\Pi}(\mu)$ and all $\mu \in [0,1]$.

The crucial restriction imposed by (C_1) is that every zero-attention strategy induces the same cost irrespective of the prior μ. The fact that this cost is set equal to 0 is merely a normalization. Then, (C_2) postulates that new information is always costly, in the sense that the cost of paying attention is higher than the normalized cost of the zero-attention strategy. This last condition is necessary for the existence of a robust scoring rule: indeed, if there is some $\pi \in \hat{\Pi}(\mu)$ with $C(\pi) \leq C(\hat{\mu})$, then for every strictly convex ϕ we obtain $V_\phi(\hat{\mu}) < V_\phi(\pi)$, implying that (2) is violated, and therefore the subject will update his prior belief.

Our next axiom is relatively new to the literature, postulating that the cost of an attention strategy is only a function of the distribution of posteriors, and not of the underlying process that yields this distribution. Formally:

(C_3) DYNAMIC CONSISTENCY: If $\sigma : [0,1] \to \Delta([0,1])$ satisfies $\sigma(\mu) \in \Pi(\mu)$ for all $\mu \in [0,1]$, then

$$C(\mathbb{E}_\pi(\sigma)) = C(\pi) + \mathbb{E}_\pi(C \circ \sigma) \tag{3}$$

for all $\pi \in \Delta([0,1])$.

Intuitively, if the subject chooses a sequential attention strategy, according to which he first picks π (first-period attention strategy) and then conditional on observing some posterior $\mu \in \text{supp}(\pi)$ he picks a new attention strategy $\sigma(\mu)$ (second-period attention strategy), the total cost that he incurs is equal to the cost of his first-period strategy ($C(\pi)$) plus the expected cost of his second-period strategies ($\mathbb{E}_\pi(C \circ \sigma)$). The distribution of posteriors at the end of the second period is then $\mathbb{E}_\pi(\sigma)$. Dynamic consistency postulates that the cost of an attention strategy that directly yields this distribution of posteriors (viz., $C(\mathbb{E}_\pi(\sigma))$) is equal to the total cost of the aforementioned sequential attention strategy (viz., $C(\pi) + \mathbb{E}_\pi(C \circ \sigma)$).

Our two axioms, (C_1) and (C_3), impose some basic coherency on the costs across different priors, similarly to recent work on dynamic information acquisition (Hébert and Woodford, 2016; Morris and

Strack, 2017; Zhong, 2017). This is in contrast to the standard decision-theoretic models of rational inattention, which specify a cost function $C : \Pi(\mu) \to \mathbb{R}_+$ for each prior $\mu \in [0,1]$ but remain silent on the relationship of the costs across the different priors (De Oliveira et al., 2017). Finally, as we show later in the paper, cost functions that satisfy $(C_1) - (C_3)$ are canonical in De Oliveira et al.'s (2017) sense, i.e., they also satisfy Blackwell monotonicity and convexity (see Section 6).

3.2. Main result

The following result answers our first question affirmatively, for the rather large class of cost functions that satisfy our axioms.

Theorem 1. *If the cost function satisfies $(C_1) - (C_3)$, there exists a robust scoring rule.*

The overall idea behind the previous result is to find a scoring rule that provides strong enough incentives to induce truth-telling (viz., ϕ must be strictly convex), but not so strong that lead the subject to update his beliefs (viz., ϕ should not be "too convex"). The proof is constructive, thus allowing us not only to prove that a robust scoring rule exists, but also to identify its functional form. Let us sketch the main steps here, while the full proof is relegated to Appendix A.

We begin with the following intermediate result, which provides a characterization of the cost functions that satisfy our axioms by means of a property that has recently attracted interest in the rational inattention literature (Caplin et al., 2017). A similar result has been proven by Zhong (2017) in a somewhat different context, relying on standard properties of mutual information (e.g., Cover and Thomas, 2006).

Lemma 1. *The cost function satisfies $(C_1) - (C_3)$ if and only if it satisfies:*

Posterior-Separability: *There is a strictly concave function $K : [0,1] \to \mathbb{R}$ such that*

$$C(\pi) = K(\mu) - \langle K, \pi \rangle \tag{4}$$

for every $\pi \in \Pi(\mu)$ and every $\mu \in [0,1]$.

One can interpret $K(\mu)$ as the cost of the most informative attention strategy when the prior belief is μ, i.e., it is the cost that the subject must incur in order to learn the true state with certainty. The curvature of K puts a bound on the incentives that the scoring rule can give. Loosely speaking, "ϕ must be less convex than $-K$", i.e., formally, *a proper scoring rule ϕ is robust if and only if $\phi + K$ is strictly concave*. Therefore, we must focus entirely on proper scoring rules that satisfy this last property. The most obvious such candidate is

$$f := \gamma - \lambda K \tag{5}$$

which is obviously strictly convex for every $\gamma \in \mathbb{R}$ and every $\lambda \in (0,1)$. In fact, if K is subdifferentiable at the boundary of $[0,1]$, so is f, and therefore we can simply set $\phi := f$.

So let us focus on the case where K is not subdifferentiable, and a fortiori f is not either. One such example is the entropic cost function that we study in detail in the next section.

Lemma 2. *Consider a strictly convex function $f : [0,1] \to \mathbb{R}$. Then, there exists a strictly convex and subdifferentiable function $g : [0,1] \to \mathbb{R}$ such that $f - g$ is convex.*

The previous result, takes f as a benchmark and introduces the strictly convex function g which provides weaker incentives than f, i.e., formally, $B_g(\pi) \le B_f(\pi)$ for every $\pi \in \Delta([0,1])$. Therefore, since $B_f(\pi) < C(\pi)$, it will also be the case that $B_g(\pi) < C(\pi)$, thus guaranteeing that (2) will be satisfied by g. Finally, since g is subdifferentiable, we can set $\phi := g$, thus completing the proof of our theorem. Notice that our proof of Lemma 2 is constructive, implying that not only do we prove existence, but we also identify an entire family of robust scoring rules.

4. Entropic attention cost

The most common functional form of attention costs within the rational inattention literature is the *entropic cost* specification (Sims, 2003, 2006; Caplin et al., 2017), which among other nice properties, allows us to provide microeconomic foundations to the multinomial logit model (Matejka and McKay, 2015). Accordingly, the cost of an arbitrary $\pi \in \Pi(\mu_0)$ is equal to

$$C_\kappa(\pi) = \kappa\big(H(\mu_0) - \langle H, \pi\rangle\big), \tag{6}$$

where $H(\mu) = -\mu \log \mu - (1-\mu)\log(1-\mu)$ is the Shannon entropy (Shannon, 1948), and $\kappa > 0$ is a multiplier parameter. It is straightforward to verify that C_κ is posterior-separable with $K := \kappa H$. Therefore, C_κ satisfies (C_1)–(C_3), and by Theorem 1, there exists a robust scoring rule. Since entropic attention costs are widely-used in applications and empirical studies, we naturally ask whether there is a common robust scoring rule. In particular, we ask: *is there a robust quadratic scoring rule when the attention costs are entropic?* The following result answers the previous question affirmatively.

Theorem 2. *For an entropic cost function with multiplier parameter $\kappa > 0$, the quadratic scoring rule ϕ_β is robust if and only if $\beta \leq 2\kappa$.*

Note that only the parameter β is relevant for robustness. This is not surprising, given that the incentives of a scoring rule are measured in terms of its convexity, and the constant α does not affect the degree of convexity of ϕ_β, but rather it simply rescales the payments by adding a constant.

The proof of the previous result exploits the fact that H is twice differentiable. Indeed, the condition $\beta \leq 2\kappa$ is equivalent to $\phi''_\beta(\mu) + \kappa H''(\mu) \leq 0$ for every $\mu \in [0,1]$ with equality holding in at most countably many points. This last condition corresponds to $\phi_\beta + K$ being strictly concave, which as we have already discussed in the previous section is equivalent to (2).

5. Testing Bayesian rationality

Suppose that our main aim is to test Bayesian rationality in the lab. We consider a subject that makes choices under uncertainty about a payoff-relevant event. The subject is said to be rational if his choice maximizes his subjective expected utility given his (prior) beliefs. Hence, in order to test Bayesian rationality we need to have data about both his choices and his prior beliefs.

Let us first define what exactly we mean by "prior beliefs". In our theoretical model, the prior beliefs are those derived from the agent's Anscombe-Aumann preferences over acts, viz., over singleton menus in De Oliveira et al. (2017). That is, "prior beliefs" do not refer to the subject's beliefs before entering the lab, but rather to his beliefs at the moment that he makes his choice. In fact, it is plausible that before encountering the decision problem that he faces in the lab, he may not even have beliefs about the payoff-relevant event of interest. Hence, all we care about is the beliefs that are relevant for his choice.

The most obvious experimental test for Bayesian rationality would rely on comparing direct with indirect measurements of beliefs, i.e., comparing the subject's reported belief with the ones that rationalize his actual choice. For the extensive experimental literature on this topic, see Schotter and Trevino (2014, Section 3.2). A discrepancy between the two measurements could theoretically be attributed to two different factors, viz., belief updating due to the experimental procedure (experimental effect) or failure of Bayesian rationality. Hence, in order to test the subject's rationality, an experimental effect must be ruled out. For starters, it is necessary to elicit the subject's beliefs after the choice has been made, to avoid possible belief updating between elicitation and choice. Then, assuming that first we observe choices and then we elicit beliefs, we must necessarily use a robust scoring rule, to make sure that no updating takes place between choice and elicitation.

Recall that a scoring rule ϕ is robust if and only if $K + \phi$ is strictly concave. Therefore, to guarantee that the subject reports his prior beliefs, we must first know the cost function. There are two ways to practically deal with this problem. One solution is to calibrate the attention costs using a revealed-preference test (Caplin and Dean, 2015; Chambers, Liu and Rehbeck, 2017). Alternatively, we may rely on distributional assumptions, e.g., we may use past data to estimate the distribution of the multiplier parameter of the entropic cost in the population of our subjects. In this case, we can define a scoring rule which is robust with sufficiently high probability.

Concluding, while there are many open questions regarding the experimental design that we should use, the aforementioned discussion provides a roadmap for testing Bayesian rationality. In particular, we illustrate that robustness of the scoring rule is necessary, but not always sufficient for testing Bayesian rationality.

6. Discussion

Other experimental currencies. Throughout the paper we have focused exclusively on scoring rules that pay in monetary payoffs, i.e., formally, S takes values in $\mathbb{R}$. However, as we have already mentioned, there are large literatures dealing with scoring rules that pay either in probabilities over a fixed prize or allowing for infinite rewards/losses, i.e., formally, S takes values in $[0, 1]$ in the former and in $\overline{\mathbb{R}}$ in the latter case, respectively. Let us first briefly present the motivation and then discuss our results for each of these alternative mechanisms.

Scoring rules that pay in probability currencies are called stochastic and have been introduced in economics in order to deal with subjects who are not risk-neutral (Savage, 1971; Schlag and van der Weele, 2013). Formally, $S_r(\omega) \in [0, 1]$ is the objective probability of the subject winning the prize, when he reports r and the realized state is ω. In this case the subject's expected utility is linear in the probability of winning the prize irrespective of his risk preferences, and our analysis follows verbatim except for one small detail, viz., in order for a function ϕ to characterize a stochastic scoring rule, not only should it be subdifferentiable, but it should also have at every point a subtangent that takes values in $[0, 1]$ both when evaluated at 0 and at 1. The latter holds whenever ϕ satisfies the following two conditions: $0 \leq \phi(0) + \phi'(0) \leq 1$ and $0 \leq \phi(1) - \phi'(1) \leq 1$, for some $\phi'(\mu) \in \partial\phi(\mu)$, for each $\mu \in \{0, 1\}$. Hence, for a robust ϕ there exists some $\beta \in \mathbb{R}$ and a sufficiently small $\gamma \in (0, 1)$ such that $\psi(\mu) := \gamma(\phi(\mu)+\beta\mu)$ satisfies the previous two inequalities, thus implying that ψ is a robust stochastic scoring rule. It is important to mention that stochastic scoring rules have been criticized based on experimental evidence (Selten et al., 1999), although such criticism is not unanimous (Harrison et al., 2013, 2014). For an in-depth discussion on the role of risk preferences, we refer to Offerman et al. (2009).

Scoring rules that allow for infinite rewards/losses are common in statistics (Gneiting and Raftery, 2007). The main reason behind such generalization is in order to be able to dispense with the subdifferentiability of ϕ. More specifically, when S is allowed to take infinite values, every strictly convex function ϕ characterizes a proper scoring rule, even if it is not subdifferentiable at the boundary. In this last case, the respective subtangents are infinitely sloped, which is why S needs to be unbounded. In fact, under such generalized scoring rules, the proof of our main result becomes straightforward, viz., the function $f = \gamma - \lambda K$ (see (5)) is always robust. However, it would still be very difficult to practically implement such a scoring rule.

Canonical attention costs. In the rational inattention literature, there are two natural regularity properties that we typically require the cost function to satisfy, viz., Blackwell monotonicity and convexity (De Oliveira et al., 2017). First, let us define the (partial) Blackwell order in $\Pi(\mu)$. For two attention strategies, $\pi, \rho \in \Pi(\mu)$, we say that π is Blackwell more informative than ρ, and we write $\pi \trianglerighteq \rho$, whenever $\langle f, \pi \rangle \geq \langle f, \rho \rangle$ for every convex function $f : [0,1] \to \mathbb{R}$ (Blackwell, 1953). The two axioms postulate that, for every $\mu \in [0,1]$, the following hold respectively:

(C_4) Blackwell monotonicity: $C(\pi) \geq C(\rho)$ for all $\pi, \rho \in \Pi(\mu)$ with $\pi \trianglerighteq \rho$.

(C_5) Convexity: $C(\lambda\pi + (1-\lambda)\rho) \leq \lambda C(\pi) + (1-\lambda)C(\rho)$ for all $\pi, \rho \in \Pi(\mu)$ and all $\lambda \in (0,1)$.

The obvious question is whether our axioms $(C_1)-(C_3)$ imply $(C_4)-(C_5)$, i.e., are the cost functions that we consider canonical?

Starting with Blackwell monotonicity, take $\pi, \rho \in \Pi(\mu)$ such that $\pi \trianglerighteq \rho$. Then, by setting $f := -K$, we obtain

$$\begin{aligned} -\langle K, \pi \rangle \geq -\langle K, \rho \rangle \;\; &\Rightarrow \;\; K(\mu) - \langle K, \pi \rangle \geq K(\mu) - \langle K, \rho \rangle \\ &\Rightarrow \;\; C(\pi) \geq C(\rho), \end{aligned}$$

with the second implication following from posterior separability. Hence, (C_4) holds as desired. Switching now to convexity, take $\pi, \rho \in \Pi(\mu)$ and $\lambda \in (0,1)$. Then, we obtain

$$\begin{aligned} C(\lambda\pi + (1-\lambda)\rho) &= K(\mu) - \langle K, \lambda\pi + (1-\lambda)\rho \rangle \\ &= \lambda K(\mu) - \lambda\langle K, \pi \rangle + (1-\lambda)K(\mu) - (1-\lambda)\langle K, \rho \rangle \\ &= \lambda C(\pi) + (1-\lambda)C(\rho), \end{aligned}$$

with the first and third equations following from posterior separability, and the second one following from the linearity of the inner product. Hence, our axioms imply an even stronger axiom than (C_5), viz., C is linear in $\Pi(\mu)$.

Quadratic scoring rules. As we have shown above, for every entropic cost function there is a robust quadratic scoring rule. Does this result extend to every posterior-separable cost functions? It is not difficult to see that this is not the case. For instance, if $K(\mu) = 1 - \mu^3$, the second derivative is not bounded away from 0, implying that the cost function becomes arbitrarily flat close to 0. On the other hand, the second derivative of an arbitrary ϕ_β is bounded away from 0. That is, formally $\phi''_\beta(\mu) + K''(\mu) = 2\beta - 6\mu$, implying that $\phi_\beta + K$ is (strictly) convex in $[0, \frac{\beta}{3})$. Hence, ϕ_β is not robust for any β.

Multinomial beliefs. Throughout the paper we have focused on binary state spaces, thus eliciting the probability of a single event. The difficulty to directly extend our main result to the multinomial case lies on the extension of Lemma 2 to higher-dimension euclidean spaces not being straightforward. Nevertheless, for practical purposes, it can be shown that approximately robust scoring rules exist. That is formally, for every $\varepsilon > 0$ there is a scoring rules that elicits a belief within ε distance from the subject's actual prior belief. Interestingly, our Theorem 2 extends verbatim to the multinomial case, i.e., a quadratic scoring rule is robust if and only if $\beta \leq 2\kappa$.

A. Proofs

A.1. Intermediate results

Proof of Lemma 1. Sufficiency. Let C be posterior separable. Then, it is obvious that $C(\hat{\mu}) = 0$ for every $\mu \in [0,1]$, thus proving (C_1). By strict concavity of K it follows that $C(\pi) = K(\mu) -$

$\langle K, \pi \rangle > 0$ for every $\pi \in \hat{\Pi}(\mu)$ and every $\mu \in [0,1]$, thus proving (C_2). Finally, for an arbitrary $\sigma : [0,1] \to \Delta([0,1])$ satisfying $\sigma(\mu) \in \Pi(\mu)$ for every $\mu \in [0,1]$, and an arbitrary $\pi \in \Pi(\mu_0)$,

$$\begin{aligned} C(\pi) + \mathbb{E}_\pi(C \circ \sigma) &= K(\mu_0) - \langle K, \pi \rangle + \mathbb{E}_\pi(K - \langle K, \sigma \rangle) \\ &= K(\mu_0) - \langle K, \pi \rangle + \langle K, \pi \rangle - \langle K, \mathbb{E}_\pi(\sigma) \rangle \\ &= C(\mathbb{E}_\pi(\sigma)), \end{aligned}$$

with the first and the third equation following from posterior-separability, and the second one following from the linearity of the expectation ($\mathbb{E}_\pi$) and the inner product ($\langle K, \cdot \rangle$). Hence, (C_3) is also proven.

Necessity. Assume that C satisfies $(C_1) - (C_3)$, and let $K : [0,1] \to \mathbb{R}_+$ be the cost of learning the state with certainty, i.e., $K(\mu) := C(\sigma(\mu))$ for each $\mu \in [0,1]$, where $\sigma(\mu) \in \Pi(\mu)$ induces with probability μ the posterior that puts probability 1 to ω_0 and with probability $1 - \mu$ the posterior that puts probability 1 to ω_1. Now, for an arbitrary $\pi \in \Pi(\mu_0)$,

$$\begin{aligned} C(\pi) &= C(\mathbb{E}_\pi(\sigma)) - \mathbb{E}_\pi(C \circ \sigma) \\ &= K(\mu_0) - \langle K, \pi \rangle, \end{aligned}$$

with the first equation following directly from rearranging (C_3), and the second one following from the definition of K. Hence, it suffices to prove that K is strictly concave. Take arbitrary $0 \leq \mu_1 < \mu_2 \leq 1$ and $\theta \in (0,1)$, and let $\pi_0 \in \hat{\Pi}(\theta\mu_1 + (1-\theta)\mu_2)$ be the attention strategy that assigns probability θ to μ_1 and probability $1 - \theta$ to μ_2. Then,

$$\begin{aligned} K(\theta\mu_1 + (1-\theta)\mu_2) &= C(\pi_0) + \theta K(\mu_1) + (1-\theta)K(\mu_2) \\ &> \theta K(\mu_1) + (1-\theta)K(\mu_2), \end{aligned}$$

with the equation above following from (C_3), and the inequality following from (C_2). Hence, K is strictly concave, thus completing the proof. □

Proof of Lemma 2. If f is subdifferentiable in $[0,1]$ then the result follows trivially by setting $g := f$. Therefore, we assume that there exists $x \in \{0,1\}$ such that the subderivative

$$\partial f(x) := \{t \in \mathbb{R} : f(y) \geq f(x) + t(y - x) \text{ for all } y \in [0,1]\}$$

is empty.

Step 1: By convexity, f is continuous in $(0,1)$. Let $\hat{f} : [0,1] \to \mathbb{R}$ be the continuous extension of $f : (0,1) \to \mathbb{R}$ to $[0,1]$. It is straightforward that $\hat{f}$ exists and is strictly convex. Let us now prove that $f - \hat{f}$ is convex. Take arbitrary $0 \leq x_1 < x_2 \leq 1$ and $\theta \in (0,1)$. Since $(\theta x_1 + (1-\theta)x_2) \in (0,1)$, we trivially obtain $(f - \hat{f})(\theta x_1 + (1-\theta)x_2) = 0$. Moreover, by $f(x) \geq \hat{f}(x)$, we obtain $\theta(f - \hat{f})(x_1) + (1-\theta)(f - \hat{f})(x_2) \geq 0$. Hence, $f - \hat{f}$ is convex, as claimed. Therefore, it suffices to prove that there is a strictly convex and subdifferentiable g such that $\hat{f} - g$ is convex.

Step 2: For each $x \in [0,1]$ define the left $a_x := \hat{f}'_-(x)$ and right $b_x := \hat{f}'_+(x)$ derivative respectively. We adopt the notational convention that $a_0 = -\infty$ and $b_1 := \infty$. It follows from (strict) convexity of $\hat{f}$ that $\partial \hat{f}(x) = [a_x, b_x]$, with $a_x = b_x$ whenever $\hat{f}$ is differentiable at x. Moreover, $\partial \hat{f}$ is strictly increasing, i.e., $x < y$ if and only if $a_x \leq b_x < a_y \leq b_y$. Obviously, $\hat{f}$ is subdifferentiable if and only if $-\infty < b_0 < a_1 < \infty$, in which case we simply set $g := \hat{f}$. Hence, we henceforth focus on the case where $\partial \hat{f}(x) = \emptyset$ for some $x \in \{0,1\}$, i.e., $b_0 = -\infty$ or $a_1 = \infty$. Let $x_0 \in [0,1]$ be the unique minimizer of $\hat{f}$, and define the strictly increasing function $F : [0,1] \to \overline{\mathbb{R}}$ as follows: $F(x) := a_x > 0$ for all $x \in (x_0, 1]$, $F(x) := b_x < 0$ for all $x \in [0, x_0)$, and $F(x_0) = 0$.

STEP 3: Since $\hat{f}$ is continuous in a closed interval, it is also absolutely continuous, and therefore by the Fundamental Theorem of Calculus, F is Lebesgue integrable and

$$\hat{f}(x) = \hat{f}(0) + \int_0^x F(t)dt. \tag{A.1}$$

Take a strictly increasing Lipschitz function $h : \overline{\mathbb{R}} \to [-1, 1]$ (with Lipschitz constant $c \leq 1$), and let $G := h \circ F$. Since F is Lebesgue integrable, so is G. Thus, we can define $g : [0,1] \to \mathbb{R}$ by

$$g(x) := \hat{f}(0) + \int_0^x G(t)dt. \tag{A.2}$$

Since G is strictly increasing, g is strictly convex, and therefore subdifferentiable in $(0,1)$. Moreover, since G takes values in $[-1,1]$, it is the case that $\int_0^x G(t)dt \geq -2x$, implying that $g(x) \geq g(0) - 2x$ for every $x \in [0,1]$, i.e., g is subdifferentiable at 0. We prove identically that g subdifferentiable also at 1, implying that it is subdifferentiable in $[0,1]$.

STEP 4: Let us finally prove that $\hat{f} - g$ is convex. Consider arbitrary $0 \leq x_1 < x_2 \leq 1$. Since h is Lipschitz with constant $c \leq 1$, it is the case that $F(x_2) - F(x_1) \geq G(x_2) - G(x_1)$, implying that $F - G$ is increasing. Moreover, by (A.1) and (A.2), we obtain $(\hat{f} - g)(x) = \int_0^x (F(t) - G(t))dt$, implying that $\hat{f} - g$ is convex, which completes the proof. □

A.2. Proof of Theorem 1

First, observe that a proper scoring rule ϕ is robust if and only if $\phi + K$ is strictly concave. Indeed, for arbitrary $\pi \in \hat{\Pi}(\mu)$ and $\mu \in [0,1]$,

$$\begin{aligned} V_\phi(\pi) &= \langle \phi, \pi \rangle - \phi(\mu) - K(\mu) + \langle K, \pi \rangle \\ &= \langle \phi + K, \pi \rangle - (\phi + K)(\mu), \end{aligned}$$

with the first equation following from Lemma 1 and the definition of V_ϕ. Now for arbitrary $\gamma \in \mathbb{R}$ and $\lambda \in (0,1)$, define the strictly convex function $f := \gamma - \lambda K$ (see Equation (5)). Then, by Lemma 2, there exists some strictly convex and subdifferentiable function g, such that $f - g$ is convex. Set $\phi := g$. By strict convexity and subdifferentiability of g, it follows that ϕ is a proper scoring rule. Finally, notice that $g + K$ is strictly concave, as it is the sum of a strictly concave function (viz., $K + f$) and a concave function (viz., $g - f$). Therefore, by our first argument, ϕ is robust.

A.3. Proof of Theorem 2

Recall from the proof of Theorem 1 that ϕ_β is robust if and only if $v := \phi_\beta + \kappa H$ is strictly concave. Observe that $v''(\mu) = 2\beta - \frac{\kappa}{\mu(1-\mu)}$, implying that $v''(\mu) \leq 0$ if and only if $2\beta\mu^2 - 2\beta\mu + \kappa \geq 0$. Take $\Delta := 4\beta^2 - 8\beta\kappa$ and observe that there are three possible cases. If $\Delta < 0$ then $v''(\mu) < 0$ for all $\mu \in [0,1]$. If $\Delta = 0$ then $v''(\mu) \leq 0$ for all $\mu \in [0,1]$ with equality holding only for $\mu = 1/2$. Finally, if $\Delta > 0$ then $v''(\mu) > 0$ for all $\mu \in [0,1] \cap (\frac{1}{2} - \frac{\sqrt{\Delta}}{4\beta}, \frac{1}{2} + \frac{\sqrt{\Delta}}{4\beta})$. Hence, v is strictly concave in $[0,1]$ if and only if $\Delta \leq 0$, which is the case if and only if $\beta \leq 2\kappa$.

References

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. *Econometrica* 21, 503–546.

Armantier, O. & Treich, N. (2013). Eliciting beliefs: proper scoring rules, incentives, stakes and hedging. *European Economic Review* 62, 17–40.

Bartoš, V., Bauer, M., Chytilová, J. & Matejka, F. (2016). Attention discrimination: theory and field experiments with monitoring information acquisition. *American Economic Review* 106, 1437–1475.

Bergemann, D., Morris, S. & Takahashi, S. (2017). Interdependent preferences and strategic distinguishability. *Journal of Economic Theory* 168, 329–371.

Blackwell, D. (1953). Equivalent comparisons of experiments. *Annals of Mathematical Statistics* 24, 265–272.

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

Caplin, A. (2016). Measuring and modeling attention. *Annual Review of Economics* 8, 379–403.

Caplin, A. & Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review* 105, 2183–2203.

Caplin, A., Dean, M. & Leahy, J. (2017). Rationally inattentive behavior: characterizing and generalizing Shannon entropy. *Working Paper.*

Chambers, C. (2008). Proper scoring rules for general decision models. *Games and Economic Behavior* 63, 32–40.

Chambers, C. & Lambert, N. (2017). Dynamic belief elicitation. *Working Paper.*

Chambers, C., Healy, P. & Lambert, N. (2017). Dual scoring. *Working Paper.*

Chambers, C., Liu, C. & Rehbeck, J. (2017). Nonseparable costly attention and revealed preference. *Working Paper.*

Costa-Gomes, M., Huck, S. & Weizsäcker, G. (2014). Beliefs and actions in the trust game: creating instrumental variables to estimate the causal effect. *Games and Economic Behavior* 88, 298–309.

Costa-Gomes, M. & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies* 75, 729–762.

Cover, T.M. & Thomas, J.A. (2006). *Elements of information theory.* John Wiley & Sons.

De Oliveira, H., Denti, T., Mihm, M. & Ozbek, K. (2017). Rationally inattentive preferences and hidden information costs. *Theoretical Economics* 12, 621–654.

Dean, M. & Neligh, N. (2017). Experimental tests of rational inattention. *Working Paper.*

Duffy, J. & Tavits, M. (2008). Beliefs and voting decisions: a test of the pivotal voter model. *American Journal of Political Science* 52, 603–618.

Ellis, A. (2018). Foundations for optimal inattention. *Journal of Economic Theory* 73, 56–94.

Fischbacher, U. & Gächter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public good experiments. *American Economic Review* 100, 541–556.

Gneiting, T. & Raftery, A. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102, 359–378.

Good, I.J. (1952). Rational decisions. *Journal of the Royal Statistical Society, Series B*, 14, 107–114.

Hanson, R. (2003). Combinatorial information market design. *Information Systems Frontiers* 5, 107–119.

Harrison, G., Martìnez-Correa, J. & Swarthout, J.T. (2013). Inducing risk neutral preferences with binary lotteries: a reconsideration. *Journal of Economic Behavior and Organization* 94, 145–159.

——— (2014). Eliciting subjective probabilities with binary lotteries. *Journal of Economic Behavior and Organization* 101, 128–140.

Haruvy, E., Lahav, Y. & Noussair, C. (2007). Traders' expectations in asset markets: experimental evidence. *American Economic Review* 97, 1901–1920.

Hébert, B. & Woodford, M. (2016). Rational inattention with sequential information sampling. *Working Paper.*

Hossain, T. & Okui, R. (2013). The binarized scoring rule. *Review of Economic Studies* 80, 984–1001.

Karni, E. (1999). Elicitation of subjective probabilities when preferences are state-dependent. *International Economic Review* 40, 479–486.

——— (2009). A mechanism for eliciting probabilities. *Econometrica* 77, 603–606.

——— (2017). A mechanism for the elicitation of second-order belief and subjective information structure. *Working Paper.*

Maćkowiak, B. & Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review* 99, 769–803.

Manski, C.F. (2004). Measuring expectations. *Econometrica* 72, 1329–1376.

McKelvey, R. & Page, T. (1990). Public and private information: an experimental study of information pooling. *Econometrica* 58, 1321–1339.

McCarthy, J. (1956). Measures of the value of information. *Proceedings of the National Academy of Sciences* 42, 654–655.

Matejka, F. (2016). Rationally inattentive seller: sales and discrete pricing. *Review of Economic Studies* 83, 1156–1188.

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

Matejka, F. & Tabellini, G. (2016). Electoral competition with rational inattentive voters. *Working Paper.*

Morris, S. & Strack, P. (2017). The Wald problem and the equivalence of sequential sampling and static information costs. *Working Paper.*

Nyarko, Y. & Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica* 70, 971–1005.

Offerman, T., Sonnemans, J., van de Kuilen, G. & Wakker, P. (2009). A truth serum for non-Bayesians: correcting proper scoring rules for risk attitudes. *Review of Economic Studies* 76, 1461–1489.

Ostrovsky, M. (2012). Information aggregation in dynamic markets with strategic traders. *Econometrica* 80, 2595–2647.

Palfrey, T. & Wang, S. (2009). On eliciting beliefs in strategic games. *Journal of Economic Behavior and Organization* 71, 98–109.

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

Schlag, K., Tremewan, J. & van der Weele, J. (2015). A penny for your thoughts: a survey of methods for eliciting beliefs. *Experimental Economics* 18, 457–490.

Schlag, K. & van der Weele, J. (2013). Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters* 3, 38–42.

Schotter, A. & Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics* 6, 103–128.

Selten, R., Sadrieh, A. & Abbink, K. (1999). Money does not induce risk neutral behavior, but binary lotteries do even worse. *Theory and Decision* 46, 211–249.

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423, 623– 656.

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

——— (2006). Rational inattention: beyond the linear-quadratic case. *American Economic Review* 96, 158–163.

Thomson, W. (1979). Eliciting production possibilities from a well-informed manager. *Journal of Economic Theory* 20, 360–380.

Zhong, W. (2017). Optimal dynamic information acquisition. *Working Paper.*