



# Proper scoring rules with general preferences: A dual characterization of optimal reports <sup>☆</sup>



Christopher P. Chambers <sup>a</sup>, Paul J. Healy <sup>b</sup>, Nicolas S. Lambert <sup>c,\*</sup>

<sup>a</sup> Dept. of Economics, Georgetown University, 3700 O Street N.W., Washington, DC, USA

<sup>b</sup> Dept. of Economics, The Ohio State University, 1945 North High Street, Columbus, OH, USA

<sup>c</sup> Graduate School of Business, Stanford University, 655 Knight Way, Stanford, CA, USA

## ARTICLE INFO

### Article history:

Received 27 December 2017

Available online 26 July 2019

### JEL classification:

D81

D84

C90

### Keywords:

Scoring rule

Probability forecast

General preferences

## ABSTRACT

Proper scoring rules incentivize truthful probability reports from risk-neutral individuals. For individuals with much more general preferences (including risk and ambiguity aversion) we use duality techniques to characterize the optimal report in any proper scoring rule. We apply this characterization to several well-known families of preferences, and provide several applications. In particular, for the case of CARA and CRRA preferences we (1) can back out true beliefs from optimal reports, (2) show that the quadratic scoring rule minimizes deviations from truthfulness, and (3) provide a two-step procedure in which the agent reports her risk aversion parameter and her belief truthfully.

© 2019 Elsevier Inc. All rights reserved.

## 1. Introduction

A scoring rule is an incentive device for obtaining probabilistic assessments from individuals in the face of subjective uncertainty. It can be represented as a menu of state-contingent payoffs, indexed by the set of possible beliefs of the individual, from which the individual is asked to choose. If the scoring rule is proper, then a risk-neutral individual maximizes her expected utility by choosing, from the menu she is offered, the state-contingent payoff that is associated with her subjective belief.<sup>1</sup>

Although there is a long tradition of using proper scoring rules in experiments,<sup>2</sup> subjects are generally not neutral to risk. They may be risk averse, or may not even be expected utility maximizers, and their decision-making behavior

<sup>☆</sup> An earlier version of this paper circulated under the title “Dual Scoring.” We thank the editor, the associate editor, and three anonymous referees for helpful comments and suggestions. Lambert is thankful to Microsoft Research and the Cowles Foundation at Yale University for their hospitality and financial support.

\* Corresponding author.

E-mail address: [nlambert@stanford.edu](mailto:nlambert@stanford.edu) (N.S. Lambert).

<sup>1</sup> This interpretation is in line with the revealed preference tradition, whereby an individual may not necessarily perceive that she holds a probabilistic assessment, but nevertheless her behavior is consistent with such a belief. To this end, we need not ask the individual to “report” a belief, but rather to choose from a set. It is usually easier to use the language of “reporting” beliefs, however.

<sup>2</sup> See, for example, Nyarko and Schotter (2002). A common alternative to rewarding subjects with proper scoring rules is to use lotteries, also called “probability currencies” (Savage, 1971), as in Roth and Malouf (1979). In theory, lotteries can elicit truthful beliefs from risk-averse expected utility maximizers. In practice however, using lotteries can be worse than using classical proper scoring rules (see, for example, Selten et al., 1999), and experiment designers continue to use both methods.

need not be consistent with probabilistic sophistication (see Machina and Schmeidler, 1992).<sup>3</sup> An important question is how such individuals optimize when facing a given scoring rule, and to what extent their reports may differ from their true belief when they are probabilistically sophisticated. We provide a complete characterization of optimal behavior for individuals with very general preferences, captured by “standard” economic utility functions over state-contingent payoffs. These preferences include individuals who maximize expected utility (with or without risk neutrality), but also ambiguity averse individuals who do not adhere to any concept of “likelihood” or probability. Our goal is to study choice behavior broadly.

The key ingredient of our result is a duality notion that we use to connect a *direct* utility function to an *indirect* utility function.<sup>4</sup> Specifically, suppose there are  $n$  possible states of the world, and that the individual has preferences over state-contingent payoffs captured by a utility function  $U$ , where  $U(x)$  records the individual's utility for state-contingent payoff  $x = (x_1, \dots, x_n) \in \mathbb{R}^n$ . Then, for any such utility function  $U$ —usually understood as a “direct” utility—we define an “indirect utility” over price-wealth pairs as

$$G(p, w) = \sup\{U(x) : p \cdot x \leq w\},$$

where  $w$  is a scalar,  $p = (p_1, \dots, p_n) \in \mathbb{R}_+^n$ , and  $p \cdot x$  denotes the dot product  $\sum_{i=1}^n p_i x_i$ . In words, the value  $G(p, w)$  is the maximal utility achievable by an individual with utility function  $U$  when market prices for the primitive securities are given by vector  $p$ , and the wealth available for expenditure is  $w$  (see, for example, Mas-Colell et al., 1995). Our main result shows that for any decision maker who has a quasiconcave, weakly increasing, and continuous utility function over state-contingent payoffs, the unique optimal announcement  $p^*$  in scoring rule  $f$  coincides with the unique  $p^*$  which minimizes  $G(p, V(p))$ , where  $V(p) = p \cdot f(p)$  is the value function associated with  $f$ , i.e., the function that gives the expected payoff or score, under  $p$ , to the individual who announces  $p$ .

Our work is most related to Grünwald and Dawid (2004), which describes two classical approaches to the problem of robust statistics. Given a convex and compact set of probability measures, an individual is asked to choose a probability measure from this set. Assuming that state-contingent payoffs are evaluated according to the minimal expected value according to all probabilities in the set, Grünwald and Dawid show that the individual facing a proper scoring rule can be viewed as if she is minimizing a convex function on the set of probabilities, with each proper scoring rule being associated with its own convex function. For example, the authors observe the duality between the logarithmic scoring rule (Good, 1952) and the entropy function.<sup>5</sup> While the result of Grünwald and Dawid captures the behavior of an individual who is neutral to risk but ambiguity averse with maxmin preferences, the duality characterization of this paper extends to arbitrary “well-behaved” utilities in an environment with finitely many states, including risk neutral ambiguity averse individuals as a special case.<sup>6</sup>

Besides extending Grünwald and Dawid (2004) to a broad class of preferences, the dual characterization captures transparently the trade-offs that individuals face when rewarded with a proper scoring rule. This interpretation of the individual's optimization problem facilitates comparative statics—for example, to understand how behavior is expected to change as the individual becomes more averse to uncertainty—and makes it possible to derive results about the individual's behavior and its implications more simply and more naturally than with a direct approach. It is also worth noting that many preference specifications in economics are defined only via their indirect utility functions—chief among these preference classes is the *Gorman polar form* (Gorman, 1961), commonly used in applied modeling. The dual formulation offers the possibility to work directly with these functions. We stress that our main result itself does not provide new tools or methods for practitioners. Rather, we provide a framework that makes certain analyses of scoring rules more tractable, which in turn can lead to innovations.

The paper is organized as follows. We present our main result in Section 2. We illustrate the result in Section 3 with several examples for commonly used preferences. In Section 4 we put the result to work in simple applications. We review the literature and conclude in section 5.

## 2. Main result

Let  $\Omega = \{1, \dots, n\}$  be a finite set of *states* and  $\Delta(\Omega)$  the set of probability measures on  $\Omega$ . A *scoring rule* is a mapping  $f : \Delta(\Omega) \rightarrow \mathbb{R}^\Omega$  that associates a state-contingent payoff (or score) to every possible announcement of a probability distribution over states. It is *proper* if for all  $p, p' \in \Delta(\Omega)$ ,  $p \cdot f(p) \geq p \cdot f(p')$  and *strictly proper* if for all  $p, p' \in \Delta(\Omega)$  for which  $p \neq p'$ , we

<sup>3</sup> Although such individuals do not “hold probabilities” as part of their preferences over uncertainty, they can still make probability assessments and their report can still be useful. One can also use the report to make inferences about their preferences.

<sup>4</sup> These duality notions go back to Konüs (1939) and Ville and Newman (1951–1952), which are translations of earlier foreign language works. Roy (1947) popularized the concept. There are several related dualities. The ones studied in Lau (1969); Shepherd (1970); Weymark (1980) and Cornes (1992) require all payoffs to be nonnegative. In this work, we focus on the duality exploited in Cerreia-Vioglio et al. (2011b), which allows for negative payoffs.

<sup>5</sup> See also Chambers (2008), which independently proves the result in a simpler environment.

<sup>6</sup> To see why Grünwald and Dawid (2004) is a special case of our setup, observe that when  $U(x) = \min_{p \in \mathcal{P}} p \cdot x$ , then  $G(p, w) = w$  when  $p \in \mathcal{P}$  and  $+\infty$  otherwise. In other words, the uniquely optimal announcement for such a utility function is equivalently given by  $\min_{p \in \mathcal{P}} V(p)$ . While Grünwald and Dawid (2004) and Chambers (2008) both rely on versions of the minimax theorem, the proof of this paper is simpler and leverages the separating hyperplane theorem.

have  $p \cdot f(p) > p \cdot f(p')$ . In words, (strictly) proper scoring rules (strictly) induce probabilistically sophisticated individuals who care about maximizing their expected payoff to announce their honest assessment of the state probabilities. For a strictly proper scoring rule  $f$ , we define the associated *value function*  $V : \Delta(\Omega) \rightarrow \mathbb{R}$  by  $V(p) = p \cdot f(p) = \sup_{p' \in \Delta(\Omega)} p \cdot f(p')$ .<sup>7</sup> That is,  $V(p)$  is the average score of individuals who announce  $p$  when states are distributed according to  $p$ .

We assume that preferences over state-contingent payoffs are represented by a utility function  $U : \mathbb{R}^\Omega \rightarrow \mathbb{R}$  that is weakly increasing, quasiconcave, and continuous.<sup>8</sup> We call such a utility *standard*, as these assumptions are the basis for classical demand and general equilibrium theory.<sup>9</sup> We stress that this utility is defined *over the multidimensional domain of state-contingent payoffs*, which, of course, is richer than the perhaps more commonly used notion of utility for money used to represent preferences based on expected utility theory. To distinguish between the two, we use upper case  $U$  for utility over state-contingent payoffs which we refer to as utility function, and lower case  $u$  for utility over monetary amounts which we refer to as utility for money. In the case of an individual who has beliefs or tastes, these will be reflected in behavior, which is captured by  $U$ . In particular, any belief of a probabilistically sophisticated individual is included as part of the utility function, which itself defines the individual's preferences in full. The *indirect utility*  $G : \Delta(\Omega) \times \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$  is then defined by  $G(p, w) = \sup\{U(x) : p \cdot x \leq w\}$ . Quite generally (under monotonicity, upper semicontinuity, and convexity), we have  $U(x) = \inf_{p \in \Delta(\Omega)} G(p, p \cdot x)$ .<sup>10</sup> Indirect utility captures the concept of duality that we use for our results.

For example, this model can incorporate decision makers with utility functions of the form

$$U(x) = u^{-1} \left( \sum_{i=1}^n p_i u(x_i) \right),$$

where  $u$  is an increasing and concave utility for money, and  $p$  is a probability distribution, with  $p_i$  the subjective probability of occurrence of state  $i$ . This utility is a classical certainty-equivalent representation of a risk-averse expected utility maximizer. The representation of these preferences in terms of indirect utility is given in Cerreia-Vioglio et al. (2011b). One benefit of indirect utilities is that the class of utility functions that can be accommodated is much broader than the risk-averse expected utility class, and includes ambiguity averse individuals who are not probabilistically sophisticated. It is important to note that individuals can still participate meaningfully in a scoring rule even if they do not make use of probabilities in their own decision calculus, since the reported probabilities still leads to meaningful inferences about their own preferences and subjective beliefs.

We now state our main result.

**Theorem 1.** *Suppose that  $f$  is a continuous strictly proper scoring rule. For a standard utility with indirect utility function  $G$ , there exists a unique solution to the problem  $\max_{p \in \Delta(\Omega)} U(f(p))$ , given by  $\arg \min_{p \in \Delta(\Omega)} G(p, V(p))$ , where  $V$  is the value function associated with  $f$ .*

Theorem 1 claims that for any standard utility and any strictly proper scoring rule, there is a unique optimal announcement  $p^* \in \Delta(\Omega)$ , and further, this unique announcement can be obtained by duality: it solves the problem of minimizing  $G(p, V(p))$  over  $p \in \Delta(\Omega)$ . The latter problem is often easier to describe. Theorem 1 can be extended to include scoring rules which take infinite-valued payoffs (such as the classical logarithmic scoring rule). In this case, existence of an optimal announcement is not guaranteed, but when there is such an announcement, the duality will hold.

The proof of Theorem 1 is in Appendix A. Fig. 1 illustrates the intuition of the proof in a simple case. In this figure,  $\Omega = \{1, 2\}$ , so  $n = 2$ . By strict properness, the range of the scoring rule  $f$  forms the upper boundary of a strictly convex set. Intuitively, this fact owes to the continuity of  $f$  and the observation that each  $p \in \Delta(\Omega)$  induces a hyperplane with a unique tangency to the range (in  $\mathbb{R}^\Omega$ ) of the scoring rule. The problem of maximizing  $U$  over this set results in an optimal vector of payments  $f(p^*)$ , which is achieved by announcing the distribution  $p^*$ .<sup>11</sup> The payments of  $f(p^*)$  give an expected value (under  $p^*$ ) of  $V(p^*) = p^* \cdot f(p^*)$ . Now consider the indirect utility maximization program: for any  $p$  we calculate the expectation of the scoring rule under  $p$ —which equals  $V(p)$ —and ask what point  $x \in \mathbb{R}^2$  maximizes  $U$  subject to the constraint that  $p \cdot x \leq V(p)$ . This constraint is shown for  $p^*$  and  $p'$  as dashed lines. For  $p^*$  the maximizing point is again  $f(p^*)$ , which gives indirect utility  $G(p^*, V(p^*)) = U(f(p^*))$ . But for  $p'$  the constraint includes points strictly better than  $f(p^*)$  for the decision maker, so the maximum indirect utility (which obtains at  $x'$ ) is  $G(p', V(p')) > U(f(p^*)) = G(p^*, V(p^*))$ . By the strictly convex shape of  $f(\Delta(\Omega))$  this is true for any  $p' \neq p^*$ . Thus, the original utility-maximizing point  $f(p^*)$  is the unique minimizer of the indirect utility function  $G(p, V(p))$ .

<sup>7</sup> For a general, nonproper scoring rule, one can define the value function via  $V(p) = \sup_{p' \in \Delta(\Omega)} p \cdot f(p')$ .

<sup>8</sup> Weakly increasing means that if  $x_\omega > y_\omega$  for all  $\omega \in \Omega$ , then  $U(x) > U(y)$ .

<sup>9</sup> In applications,  $U$  is known to the individual but not to the experimenter. Rather, the experimenter observes the reported  $p$  and makes an inference about  $U$ .

<sup>10</sup> This duality is related to, but distinct from, the duality presented in Lau (1969); Shepherd (1970); Weymark (1980), and Cornes (1992). That notion of duality is for functions on the nonnegative orthant. See Cerreia-Vioglio et al. (2011b) for details.

<sup>11</sup> We reiterate that the decision maker need not have “true” beliefs  $p^*$ —or any probabilistic beliefs at all. Here,  $p^*$  is interpreted only as the decision maker's optimal announcement given  $f$ .




Fig. 1. An illustration of Theorem 1.

According to Theorem 1, the announced probability  $p^*$  separates the convex hull of the image of payoffs of the scoring rule, and the upper contour set of the preferences. By exploiting Roy's identity (Roy, 1947) and the fact that a scoring rule is a subdifferential of its homothetic extension, we get the following corollary.

**Corollary 1.** Suppose  $U$  is standard and  $f$  is a continuous and strictly proper scoring rule. If  $p^* \in \arg \max_{p \in \Delta(\Omega)} U(f(p))$ , then  $f(p^*) \in \arg \max_{x \in \mathbb{R}^n : p^* \cdot x \leq V(p^*)} U(x)$ .

Corollary 1 states that a standard decision maker will choose to announce the  $p^*$  at which her Walrasian demand (when given wealth  $V(p^*)$ ) includes  $f(p^*)$ . This result illustrates the connection between the optimal choice of a decision maker from a scoring rule, and the equilibrium price in a Robinson Crusoe economy where the output of the scoring rule plays the role of a technology.

*Interpretation*

Let us consider a utility function  $U$  which is a *certainty equivalent* representation, namely, one which has the property that for all  $x \in \mathbb{R}$ ,  $U(x, x, \dots, x) = x$ .<sup>12</sup> This utility function measures utility in monetary units. In this case,  $G(p, w)$  has a simple interpretation: it measures the monetary value to the individual with wealth  $w$  and who faces state prices represented by vector  $p$ . Clearly,  $G(p, w) \geq w$ . In addition, there exists  $p^* \in \Delta(\Omega)$  such that  $G(p^*, w) = w$  (see Cerreia-Vioglio et al., 2011b). A high value of  $G(p, w)$  relative to  $w$  reflects a willingness to stake a “bet against” the odds  $p$  with wealth  $w$ .

For instance, a risk-neutral individual with subjective probability distribution  $\pi$  will be willing to bet against any odds  $p \neq \pi$ , meaning that with state prices  $p$ , she is willing to short at least one state (namely, the states  $\omega$  for which  $p_\omega > \pi_\omega$ ), and to long at least one state. Further, because of risk-neutrality, she will be willing to invest arbitrarily large amounts of money in such a bet. Hence, for such an individual,  $G(p, w) = +\infty$  while, of course,  $G(\pi, w) = 0$ . A risk-averse individual with subjective probability  $\pi$  is unwilling to bet against the odds of  $\pi$ , and hence we continue to have  $G(\pi, w) = w$ , but, if sufficiently risk averse, she may also not be willing to bet much against odds  $p \neq \pi$  due to risk aversion. She would prefer instead to keep her wealth (relatively) state-independent. In this case,  $G(p, w)$  remains close to  $w$  for  $p \neq \pi$ . Finally, at the extreme end of the spectrum, an individual who wants to maximize her payoff in the worst state will have the completely flat indirect utility  $G(p, w) = w$ : this individual is unwilling to deviate from certainty. Roughly, the more the individual cares about the worst-case scenario, the flatter her corresponding indirect utility function, with respect to  $p$ .

The following lemma is elementary and proved in Appendix A for completeness. We denote by  $\Delta_{++}(\Omega)$  the set of full-support probabilities, i.e., the interior of the simplex.

<sup>12</sup> Existence of such a utility function is easy to establish for any preference relation in our domain.

**Lemma 1.** Let  $f$  be continuous and proper. If  $p^* \in \Delta(\Omega)$  has the property that for all  $\omega, \omega' \in \Omega$ ,  $f(p^*)(\omega) = f(p^*)(\omega')$ , then  $p^* \in \operatorname{argmin}_{p \in \Delta(\Omega)} V(p)$ . Conversely, if  $p^* \in \Delta_{++}(\Omega)$  and  $p^* \in \operatorname{argmin}_{p \in \Delta(\Omega)} V(p)$ , then  $p^*$  has the property that for all  $\omega, \omega' \in \Omega$ ,  $f(p^*)(\omega) = f(p^*)(\omega')$ .

In words, the value function of a continuous proper scoring rule is minimized exactly at a point which, when interior, yields the same payoff in every state, i.e., the riskless state-contingent payoff. Therefore, when the indirect utility  $G$  is applied to the value function  $V$  as in Theorem 1, the individual who is risk-neutral with subjective probability  $\pi$  is best off reporting  $\pi$  independently of  $V$ , and so independently of which strictly proper scoring rule is being applied. As the individual becomes increasingly risk averse, the indirect utility becomes flatter which means that the individual increasingly cares about  $V$ , and so the individual's reported probability distribution reflects a trade-off between the minimization of the indirect utility for a fixed wealth, obtained at  $\pi$  as in the risk-neutral case, and the minimization of  $V$ , which minimizes risk or the dispersion of payoffs. In the extreme situation of an individual who maximizes her payoff in the worst-case scenario, the individual entirely disregards  $\pi$  and simply chooses the report that minimizes  $V$  and so yields the same payoff in every state.

### 3. Examples

In this section, we provide examples of characterizations.

**Example 1 (Translation-invariant utility functions).** Consider the variational preferences model of Maccheroni et al. (2006); applied to our setting, these preferences are those which can be written as utility functions  $U$  that are translation invariant in the sense that for all  $x \in \mathbb{R}^\Omega$  and all  $\lambda \in \mathbb{R}$ ,  $U(x + \lambda(1, \dots, 1)) = U(x) + \lambda$ . In this case, we can write  $G(p, w) = w + c(p)$ , for some proper convex (possibly infinite-valued) lower semicontinuous function  $c$  (Cerrea-Vioglio et al., 2011b). Applying Theorem 1, the individual with such a utility function announces the probability distribution  $p^*$  that solves

$$\min_{p \in \Delta(\Omega)} \{V(p) + c(p)\}.$$

The function  $c$  has the interpretation of a *certainty equivalent*:  $c(p) - c(q)$  measures the sure amount a decision maker holding only a riskless asset would pay to move from state prices  $q$  to state prices  $p$ . To understand this claim, consider a decision maker maximizing  $U(x)$  subject to  $p \cdot x \leq w$ , with  $U$  as above. Note that, by the translation-invariant property,  $U(x)$  can be interpreted as the certainty equivalent of the state-contingent payoff  $x$ , up to an additive constant. Also observe that  $c(p) = \max_{x: p \cdot x \leq 0} U(x)$ . Hence, measured in terms of monetary units,  $c(p)$  is the value of facing state prices  $p$  when endowment is 0. To move from state prices  $q$  to state prices  $p$ , the individual would offer to pay  $c(p) - c(q)$  (by the translation invariant property, this willingness to pay is independent of wealth). The translation-invariant utility functions include common classes of preferences, some are described in the examples that follow.

**Example 2 (Ambiguity aversion).** The case of risk-neutral, multiple-prior maxmin utility is nested in the model of translation-invariant utility, with  $c = 0$  on the set of priors  $\mathcal{P}$ , and  $c = +\infty$  otherwise (Cerrea-Vioglio et al., 2011b). Hence, the ambiguity-averse individual announces the probability in the set of priors which uniquely minimizes the value function  $V$ ; this is one of the main results of Grünwald and Dawid (2004). Thus, if the value function  $V$  is strictly increasing along some direction in  $\Delta(\Omega)$ , then the individual reports the prior in her set that is highest in that direction. By varying the direction of increase, it is possible to uncover the entire set of priors.

**Example 3 (Subjective expected value maximizers).** If the set of priors in the previous example contains only the single prior  $\pi$ , i.e., the individual wants to maximize expected payments under belief  $\pi$ , then  $c$  is infinite-valued everywhere except at  $\pi$ . Hence, such an individual always strictly prefers to announce her true belief  $\pi$ , which indeed defines what it means for a scoring rule to be strictly proper.

**Example 4 (Subjective expected CARA utility maximizers).** Consider the subjective expected utility maximizer with CARA utility for money  $u(x) = -\exp(-\alpha x)$ , where  $\alpha > 0$ , and subjective probability distribution  $\pi$ . We can consider the certainty equivalent utility representation of these preferences, written

$$U(x) = -\frac{1}{\alpha} \log \sum_{i=1}^n \pi_i e^{-\alpha x_i}.$$

This utility function is translation invariant, and the indirect utility has a very special form: the function  $c(p)$  is given by a scaling of the *relative entropy* or *Kullback Leibler (KL) divergence*. Formally, if probability distribution  $q$  is absolutely continuous with respect to probability distribution  $p$ , i.e.,  $p_\omega > 0$  whenever  $q_\omega > 0$ , define  $\text{KL}(q \parallel p)$  as

$$\text{KL}(q \parallel p) = \sum_{\omega \in \Omega: q(\omega) > 0} q_\omega \log \frac{q_\omega}{p_\omega},$$

and otherwise let  $KL(q \parallel p) = +\infty$ . Then, the indirect utility is  $G(w, p) = w + \alpha^{-1} KL(p \parallel \pi)$ . Hence, by Theorem 1, the expected utility maximizer with CARA utility chooses to announce the probability distribution  $p^*$  that solves

$$\min_{p \in \Delta(\Omega)} \{ \alpha V(p) + KL(p \parallel \pi) \}.$$

Observe that  $KL(p \parallel \pi) = 0$  only when  $p = \pi$ , so that when  $\alpha \rightarrow 0$ , this individual behaves as a risk-neutral expected utility maximizer with subjective probability distribution  $\pi$ . The observation that the relative entropy function leads to CARA-style preferences in this context is made in Strzalecki (2011).

**Example 5 (Subjective expected CRRA utility maximizers).** Let us now consider the subjective expected utility maximizer with CRRA utility for money  $u(x) = x^{1-\eta}/(1-\eta)$ , where  $\eta \geq 0, \eta \neq 1$  is the coefficient of relative risk aversion, and subjective probability distribution  $\pi$ . (The case  $\eta = 1$  corresponds to the CRRA utility  $u(x) = \log x$  and is treated separately below.) Here, the preferences of the individual are defined over vectors of positive payoffs, and we suppose that there are no binding nonnegativity constraints in the proceeding. The certainty equivalent utility representation of these preferences is written

$$U(x) = \left( \sum_{i=1}^n \pi_i x_i^{1-\eta} \right)^{1/(1-\eta)}.$$

The indirect utility  $G(p, w)$  is obtained by maximizing  $U(x)$  subject to the constraint  $\sum_{i=1}^n p_i x_i \leq w$ . To do so, we internalize the constraint and maximize the Lagrange function

$$U(x) - \lambda \left( \sum_{i=1}^n p_i x_i - w \right),$$

where  $\lambda$  is the Lagrange multiplier. The first-order conditions yield, for every state  $\omega$ , the equality

$$\pi_\omega \left( \frac{U(x)}{x_\omega} \right)^\eta = \lambda p_\omega.$$

Hence, the constraint is binding. Solving for  $x_\omega$ , we get

$$x_\omega = U(x) \left( \frac{\pi_\omega}{\lambda p_\omega} \right)^{1/\eta},$$

and using  $\sum_{i=1}^n p_i x_i = w$ , we get

$$\frac{U(x)}{\lambda^{1/\eta}} = w \left( \sum_{i=1}^n \pi_i^{1/\eta} p_i^{1-1/\eta} \right)^{-1},$$

which yields a unique candidate maximizer given by

$$x_\omega = w \left( \sum_{i=1}^n \pi_i^{1/\eta} p_i^{1-1/\eta} \right)^{-1} \left( \frac{\pi_\omega}{p_\omega} \right)^{1/\eta}.$$

These necessary conditions are also sufficient, and if we plug in the expression for the maximizer into  $U(x)$  we get, after simplification, the indirect utility

$$G(p, w) = w \left( \sum_{i=1}^n \pi_i^{1/\eta} p_i^{1-1/\eta} \right)^{\frac{\eta}{1-\eta}}.$$

By Theorem 1, taking the logarithm of  $G(p, V(p))$ , the expected utility maximizer with CRRA utility chooses to announce the probability distribution  $p^*$  that solves

$$\min_{p \in \Delta(\Omega)} \left\{ \log V(p) + \frac{\eta}{1-\eta} \log \sum_{i=1}^n \pi_i^{1/\eta} p_i^{1-1/\eta} \right\}.$$

As for the special case  $\eta = 1$ , which corresponds to CRRA utility  $u(x) = \log x$ , we get by similar computations the indirect utility



$$G(p, w) = w \prod_{i=1}^n (\pi_i p_i)^{\pi_i},$$

the individual's optimal announcement is then the probability distribution  $p^*$  that solves the minimization problem

$$\min_{p \in \Delta(\Omega)} \left\{ \log V(p) + \sum_{i=1}^n \pi_i \log(\pi_i p_i) \right\}.$$

**Example 6** (*Ambiguity aversion with CARA utility*). One can also combine CARA utility with ambiguity-averse, multiple-prior models. Let us consider a risk-averse, ambiguity-averse individual, who has utility for money  $u(x) = -\exp(-\alpha x)$ , a set of priors  $\mathcal{P} \subseteq \Delta(\Omega)$ , and seeks to maximize the worst-case expected utility over this set of priors. Following Examples 2 and 4, it is straightforward to establish that such an individual can be represented with an indirect utility function of the same form as in Example 1 with  $c(p) = \inf_{q \in \mathcal{P}} \alpha^{-1} \text{KL}(p \parallel q)$ , and so the individual announces the probability distribution  $p^*$  that solves

$$\min_{p \in \Delta(\Omega)} \left\{ \alpha V(p) + \inf_{q \in \mathcal{P}} \text{KL}(p \parallel q) \right\}.$$

**Example 7** (*Gorman polar form*). Gorman (1961) provides necessary and sufficient conditions for utility functions to have Engel curves which are straight lines.<sup>13</sup> We focus on the case in which the Engel curves are parallel across different prices, since our consumption space is unbounded both above and below. Gorman originally proposed the family in order to meaningfully talk about “representative consumers,” though they also have a natural interpretation of utility functions for which wealth effects are absent with respect to some “numeraire” bundle. That is, we want there to be a numeraire bundle  $\beta \in \mathbb{R}^\Omega$  for which for all  $t \in \mathbb{R}$  and all  $x \in \mathbb{R}^\Omega$ ,  $U(x + \beta t) = U(x) + t$ . Gorman calls his utility functions *polar form*, since they are defined in terms of indirect utility. Therefore these form a natural class where the duality result is useful.

Formally, let us take  $\beta \in \mathbb{R}^\Omega$ , where  $\beta \geq 0$  and  $\beta \neq 0$ . Let  $\beta_+ = \{x \in \mathbb{R}_+^\Omega : \beta \cdot x = 1\}$ . Let  $c : \beta_+ \rightarrow \mathbb{R} \cup \{+\infty\}$  be a proper, lower semicontinuous and convex function. Then define

$$G(p, w) = \begin{cases} \frac{w}{\beta \cdot p} + c\left(\frac{p}{\beta \cdot p}\right) & \text{if } \beta \cdot p \neq 0 \\ +\infty & \text{otherwise.} \end{cases}$$

This equation specifies an indirect utility in Gorman polar form.

**Example 8** (*Complete ignorance*). Consider an individual who simply wants to maximize the worst-case payoff, i.e., an individual with utility  $U(x) = \min_{\omega \in \Omega} x_\omega$ . Such an individual is often described as behaving under “complete ignorance” (see, for example, Arrow and Hurwicz, 1972). As discussed in the preceding section, it is straightforward to show that such an individual has indirect utility  $G(p, w) = w$ , and so the individual announces the probability distribution that solves

$$\min_{p \in \Delta(\Omega)} V(p).$$

#### 4. Applications

In this section we illustrate the use of Theorem 1 by several examples of applications. First, we investigate the behavior of individuals who are expected utility maximizers, with CARA and then CRRA utility for money, when these individuals are rewarded according to a proper scoring rule. For such individuals, we provide tight bounds on the misreports for general scoring rules. For the special case of the logarithmic scoring rule (for CARA) and the spherical scoring rule (for CRRA), we also derive the optimal report explicitly as a function of the true belief. Second, we ask what scoring rules should be used to minimize the degree of misreporting of expected CARA/CRRA utility maximizers, subject to a normalizing condition on the scoring rules. Third, still for individuals with these preferences, we show that randomizing over two well-chosen quadratic scoring rules enables an experimenter to elicit both the Arrow-Pratt measure of absolute or relative risk aversion, and the true, unbiased subjective probability of the individuals. Fourth, for individuals whose preferences belong to the more general class of variational preferences introduced in Example 1, we show that providing a random state-contingent payoff in addition to the (nonrandom) reward of a proper scoring rule enables the experimenter to fully identify the individual's preferences (which include, of course, any subjective belief). Hence, we can fairly easily extract a substantial amount of information beyond probabilistic assessments, including risk and ambiguity attitudes, making it possible to elicit unbiased beliefs but also more complex forms of beliefs. For instance, this sort of elicitation method can be useful when individuals

<sup>13</sup> An *Engel curve* is the set of Walrasian demands for a fixed price as wealth varies.

are averse to ambiguity, in which case the belief is captured by a set of possible probability assessments. Relatedly, we show that it is possible to identify the scoring rule that is applied if the individuals are completely ignorant and maximize the worst-case payoff, as in Example 8. Finally, we provide a simple example of comparative statics, showing how reported probabilities vary with respect to uncertainty aversion for the class of variational preferences. Several examples make use of the following lemma, proved in Appendix B.

**Lemma 2.** *If the individual's utility function  $U$  is strictly increasing, i.e.,  $x \geq y$  and  $x \neq y$  implies  $U(x) > U(y)$ , then the optimal announcement for any continuous strictly proper scoring rule has full support.*

4.1. On the behavior of expected CARA utility maximizers

In this subsection, the individual is a subjective expected utility maximizer with CARA utility for money. Her subjective probability distribution (belief)  $\pi$  has full support, and her coefficient of absolute risk aversion is  $\alpha$ . Following Example 4 and Lemma 2, when this individual is rewarded according to a continuous strictly proper scoring rule  $f$ , she announces the unique full-support probability distribution  $p^*$  that solves

$$\min_{p \in \Delta(\Omega)} \left\{ \alpha V(p) + \sum_{i=1}^n p_i \log \frac{p_i}{\pi_i} \right\}, \tag{1}$$

where  $V$  continues to denote the scoring rule's value function, defined on the simplex  $\Delta(\Omega)$ .

The optimization problem of Equation (1) captures the trade-off that the individual faces when deciding on her report. The term  $\sum_{i=1}^n p_i \log \frac{p_i}{\pi_i}$ , which is the Kullback-Leibler divergence from the subjective probability distribution to the reported probability distribution, is minimized exactly when  $p = \pi$ . The term  $V(p)$  is minimized exactly when  $p$  is the report that yields the riskless payoffs (see Lemma 1 above). The coefficient of risk aversion  $\alpha$  indicates how much weight is put on these two conflicting minimization problems. The larger the risk aversion, the more the individual prefers the riskless payoffs, and the greater the degree of misreporting and the less informative the announcement.

Equation (1) makes it easy to relate the reported probability distribution to the (true) subjective probability distribution. For  $q \in \mathbb{R}_+^n$ , let us define the extended value function  $W$  as  $W(q) = \sup_{p \in \Delta(\Omega)} f(p) \cdot q$ . Observe that  $W(q) = V(q)$  for  $q \in \Delta(\Omega)$ , that  $W$  continues to be a (real-valued) convex function, and that, because  $f$  is continuous,  $W$  is differentiable.<sup>14</sup> Extending the value function to the entire nonnegative orthant  $\mathbb{R}_+^n$  enables us to compute partial derivatives and solve the above optimization problem via elementary first-order conditions.

Specifically, the individual's optimization problem reduces to choosing the vector  $p^* \in \mathbb{R}_+^n$  that solves

$$\min_{p \in \mathbb{R}_+^n} \left\{ \alpha W(p) + \sum_{i=1}^n p_i \log \frac{p_i}{\pi_i} \right\},$$

subject to the constraint  $p_1 + \dots + p_n = 1$ . Noting that the function to minimize is strictly convex on the simplex of probability distributions, and that the optimal announcement has full support, the first-order conditions are necessary and sufficient, and yield that for every  $i \neq j$ ,

$$\alpha \frac{\partial W(p^*)}{\partial p_i} + \log \frac{p_i^*}{\pi_i} = \alpha \frac{\partial W(p^*)}{\partial p_j} + \log \frac{p_j^*}{\pi_j},$$

with

$$\frac{\partial W(q)}{\partial q_i} = f(q)(i).$$

Hence, the first-order conditions give, for every  $i \neq j$ , the equality

$$\alpha f(p^*)(i) + \log \frac{p_i^*}{\pi_i} = \alpha f(p^*)(j) + \log \frac{p_j^*}{\pi_j},$$

so that there exists  $K > 0$  such that for all  $\omega$ ,

$$\pi_\omega = p_\omega^* e^{\alpha f(p^*)(\omega)} K$$

and using that  $\pi_1 + \dots + \pi_n = 1$ , we have

<sup>14</sup> This result is fairly standard in the literature, it follows from the combination of Corollary 23.5.1 with Theorem 25.1 of Rockafellar (1970).



$$K = \frac{1}{p_1^* e^{\alpha f(p^*)(1)} + \dots + p_n^* e^{\alpha f(p^*)(n)}}.$$

In summary, we have the following result.

**Proposition 1.** *Suppose the individual is a subjective expected utility maximizer with constant coefficient of absolute risk aversion  $\alpha$ . If she reports optimally the probability distribution  $p^*$ , then for every state  $\omega \in \{1, \dots, n\}$ , the individual's subjective probability of state  $\omega$ , is given by:*

$$\pi_\omega = \frac{p_\omega^* e^{\alpha f(p^*)(\omega)}}{p_1^* e^{\alpha f(p^*)(1)} + \dots + p_n^* e^{\alpha f(p^*)(n)}}, \quad (2)$$

where  $f$  is the continuous strictly proper scoring rule according to which the individual is rewarded.

Given a continuous strictly proper scoring rule, Proposition 1 makes it simple to infer the subjective probabilities over states for an individual with CARA utility if the coefficient of risk aversion is known or can be estimated.

In general, no closed-form solution enables us to write the optimal announcement  $p^*$  directly as a function of the belief  $\pi$ , this value must be computed numerically. One exception is the logarithmic scoring rule:<sup>15</sup> if  $f(p)(\omega) = K + C \log p_\omega$ , for  $K \in \mathbb{R}$  and  $C > 0$ , it is straightforward to invert Equation (2) to get

$$p_\omega^* = \frac{\pi_\omega^{1/(1+C\alpha)}}{\pi_1^{1/(1+C\alpha)} + \dots + \pi_n^{1/(1+C\alpha)}}. \quad (3)$$

In particular, if  $\alpha \rightarrow 0$ , i.e., if the individual becomes risk neutral, then  $p^*$  converges to the subjective belief  $\pi$ , while if, on the opposite,  $\alpha \rightarrow +\infty$ , i.e., if the individual becomes highly risk averse, then  $p^*$  becomes fully uninformative, converging to the uniform probability distribution—the only announcement that guarantees riskless payoffs. Similarly, for an individual with a given risk aversion, as  $C \rightarrow 0$ , i.e., as the stakes vanish, the individual tends to report her true belief, and as  $C \rightarrow +\infty$ , i.e., as the stakes become infinitely large, the reported probabilities become totally uninformative.

Proposition 1 can also be used to provide bounds on misreports, noting that

$$\frac{p_\omega^*}{\pi_\omega} = \sum_{i=1}^n p_i^* e^{\alpha [f(p^*)(i) - f(p^*)(\omega)]}.$$

The following corollary is immediate.

**Corollary 2.** *Given a continuous strictly proper scoring rule  $f$ , let the worst-case difference of scores across states be  $m = \max_{p,i,j} |f(p)(i) - f(p)(j)|$ . Under the assumptions of Proposition 1, for all states  $\omega$ , the reported probability  $p_\omega^*$  satisfies*

$$\pi_\omega e^{-\alpha m} \leq p_\omega^* \leq \pi_\omega e^{\alpha m}.$$

Corollary 2 implies that one can uniformly bound misreports, and in particular, as  $m \rightarrow 0$ , the worst-case difference between announced probabilities and true beliefs vanishes.<sup>16</sup>

As a concrete example, if a subject who participates in an experiment is assumed to conform to the expected utility model and to have an approximately constant Arrow-Pratt coefficient of absolute risk aversion for the range of payments considered in the experiment (say, in US dollars), for instance 0.01—a conservative estimate, most studies estimate the coefficient of absolute risk aversion to be of the order of  $10^{-3}$  and below—and if the experimenter employs a quadratic scoring rule that pays off between \$0 and \$10, then the misreporting of probabilities can only be up to the order of 10% of the true subjective probability.

#### 4.2. On the behavior of expected CRRA utility maximizers

The individual is now a subjective expected utility maximizer with CRRA utility for money. The coefficient of relative risk aversion is  $\eta \geq 0$ , i.e., the utility for money is  $u(x) = x^{1-\eta}/(1-\eta)$  if  $\eta \neq 1$ , and  $u(x) = \log x$  if  $\eta = 1$ .<sup>17</sup> The individual has a full-support subjective probability distribution  $\pi$ . The goal of this subsection is to relate the optimal reports of probability

<sup>15</sup> As stated in Section 2, although this scoring rule does not always take finite values, Theorem 1 continues to apply. Note that payoffs here are guaranteed to be finite because the reported probabilities are positive for all states.

<sup>16</sup> For a given scoring rule, these bounds are, in general, not tight. The same remark holds for Corollary 3 below.

<sup>17</sup> This formulation of the utility for money implicitly assumes no initial wealth. It is without loss of generality: any initial wealth can be included as part of the scoring rule payments. An individual with initial wealth  $w$  paid according to the scoring rule  $f$  behaves the same way as an individual with no initial wealth paid according to the scoring rule  $w + f$ .

assessments to the true beliefs, as we just did for the case of CARA utilities. We derive the result for the case  $\eta \neq 1$ . Following similar steps, it is easily seen that the results continue to hold for  $\eta = 1$ .

By Example 5 and Lemma 2, when rewarded by a continuous strictly proper scoring rule  $f$  which always outputs positive values, the individual announces the probability distribution  $p^*$  that solves

$$\min_{p \in \Delta(\Omega)} \left\{ \log V(p) + \frac{\eta}{1-\eta} \log \sum_{i=1}^n \pi_i^{1/\eta} p_i^{1-1/\eta} \right\},$$

with  $V$  the value function associated to  $f$ . As in the previous subsection, this optimization problem captures the tradeoff between minimizing the risk of the reward, obtained for a report  $p$  that minimizes the first term  $\log V(p)$ , and minimizing the “distance” between the reported belief and the true belief, captured by the second term. Indeed, it is straightforward to show that the second term is minimized precisely when  $p = \pi$ .

We continue to define the extended value function  $W$  on  $\mathbb{R}_+^n$  as  $W(q) = \sup_{p \in \Delta(\Omega)} f(p) \cdot q$ . The optimization problem now reduces to choosing the vector  $p^* \in \mathbb{R}_+^n$  that solves

$$\min_{p \in \mathbb{R}_+^n} \left\{ \log W(p) + \frac{\eta}{1-\eta} \log \sum_{i=1}^n \pi_i^{1/\eta} p_i^{1-1/\eta} \right\},$$

subject to the constraint  $p_1 + \dots + p_n = 1$ . As it turns out, the constraint is not binding in this case. The first-order conditions are necessary and sufficient and give, for every state  $\omega$ , the equality

$$\frac{f(p^*)(\omega)}{W(p^*)} = \frac{\pi_\omega^{1/\eta} (p_\omega^*)^{-1/\eta}}{\sum_{i=1}^n \pi_i^{1/\eta} (p_i^*)^{1-1/\eta}}.$$

This equality gives us the ratios

$$\frac{\pi_i}{\pi_j} = \frac{(f(p^*)(i))^\eta p_i^*}{(f(p^*)(j))^\eta p_j^*},$$

for  $i \neq j$ , which, in turn, yield the expression of the individual’s subjective probabilities as a function of the reported probabilities:

$$\pi_\omega = \frac{(f(p^*)(\omega))^\eta p_\omega^*}{(f(p^*)(1))^\eta p_1^* + \dots + (f(p^*)(n))^\eta p_n^*}.$$

This result is formally stated as follows.

**Proposition 2.** *Suppose the individual is a subjective expected utility maximizer with constant coefficient of relative risk aversion  $\eta$ . If she reports optimally the probability distribution  $p^*$ , then for every state  $\omega \in \{1, \dots, n\}$ , the individual’s subjective probability of state  $\omega$  is given by*

$$\pi_\omega = \frac{(f(p^*)(\omega))^\eta p_\omega^*}{(f(p^*)(1))^\eta p_1^* + \dots + (f(p^*)(n))^\eta p_n^*}, \tag{4}$$

where  $f$  is the continuous strictly proper scoring rule outputting positive values according to which the individual is rewarded.

It is easily verified that as  $\eta \rightarrow 0$ , i.e., as the individual becomes increasingly risk neutral, the reported probabilities becomes confounded with the true subjective probabilities.

For the special case of the spherical scoring rule, it is possible to invert Equation (4) and write the optimal announcement  $p^*$  explicitly as a function of the belief  $\pi$ . Strictly speaking, the spherical scoring rule does not take positive values everywhere, but it is straightforward to see that an individual with CRRA preferences and a full-support subjective probability distribution would always maximize her expected utility by choosing an announcement that leads to positive payoffs. Recall that the spherical scoring rule is written

$$f(p)(\omega) = C \cdot \frac{p_\omega}{\sqrt{p_1^2 + \dots + p_n^2}},$$

for an arbitrary positive constant  $C$  that captures the magnitude of the rewards. In this case, we get

$$p_\omega^* = \frac{\pi_\omega^{1/(1+\eta)}}{\pi_1^{1/(1+\eta)} + \dots + \pi_n^{1/(1+\eta)}}. \tag{5}$$

Note the similarity with Equation (3) discussed in the previous subsection, in spite of the fact that the spherical scoring rule used here is very different from the logarithmic scoring rule used for individuals with CARA utilities. It is also worth observing that, as opposed to the case of CARA utilities, the reported probabilities do not depend on the amplitude of the payments. This is because the marginal utility at low wealth levels grows unbounded.

To provide bounds on misreports for general scoring rules, we note that according to Proposition 2, we have

$$\frac{p_\omega^*}{\pi_\omega} = \sum_{i=1}^n p_i^* \left( \frac{f(p^*(i))}{f(p^*(\omega))} \right)^\eta.$$

We then get the following corollary.

**Corollary 3.** *Given a continuous strictly proper scoring rule  $f$ , let the largest ratio of scores across states be  $m = \sup_{p,i,j} |f(p)(i)/f(p)(j)|$ . If  $m < \infty$ , then under the assumptions of Proposition 1, for all states  $\omega$ , the reported probability  $p_\omega^*$  satisfies*

$$\pi_\omega \frac{1}{m^\eta} \leq p_\omega^* \leq \pi_\omega m^\eta.$$

Note the subtle difference with Corollary 2: with CARA utilities, one increases accuracy by decreasing the difference between large and small scores, while with CRRA utilities, one increases accuracy by decreasing their ratios.

#### 4.3. Minimizing misreporting for expected CARA/CRRA utility maximizers

The subsections above tell us the extent to which risk-averse individuals can misreport when rewarded according to a strictly proper scoring rule. In this subsection, we ask what scoring rules minimize this misreporting behavior, subject to a normalization condition. To simplify matters, we consider the special case of a state that captures the occurrence or non-occurrence of a random event, that is, with a slight abuse of notation,  $\Omega = \{0, 1\}$ , and  $\omega \in \Omega$  denotes the outcome of the event, with  $\omega = 1$  if the event occurs and  $\omega = 0$  otherwise. We assume that the individual conforms to the expected utility model with CARA or CRRA utility for money. We start with and mostly focus on CARA utility; the results remain essentially the same for CRRA utility and so this case is only briefly discussed at the end of the subsection.

So let us assume the individual believes that the event occurs with probability  $\pi \in (0, 1)$ , and assume her Arrow-Pratt coefficient of absolute risk aversion is  $\alpha$ . Note that, for the special case of eliciting event probabilities, we find it convenient to abuse notation and use the symbols  $\pi, p, q$  to denote event probabilities as opposed to probability distributions in the simplex  $\Delta(\Omega)$ .

Suppose the individual is rewarded according to a differentiable strictly proper scoring rule  $f$ . For an individual who reports  $p$ , the value

$$D(\pi, p) \equiv \max \left\{ \frac{\pi/(1-\pi)}{p/(1-p)}, \frac{p/(1-p)}{\pi/(1-\pi)} \right\},$$

which takes the maximum of the two odds ratios, provides a convenient measure of the amount of misreporting. Clearly,  $D(\pi, p) = 1$  is minimized for  $p = \pi$  (case of truthful report), and  $D(\pi, p) > 1$  if  $p \neq \pi$  (case of misreport). As  $p$  deviates further away from  $\pi$ ,  $D(\pi, p)$  grows arbitrarily large.

If we seek to minimize the degree of misreporting, then of course following Corollary 2 we can choose a scoring rule that delivers near-constant payoffs across reports and states. However, if this choice limits the degree of misreporting in risk-averse individuals, it also provides little incentive to report close to one's true belief, even for risk-neutral individuals.

To circumvent this issue, we impose a lower bound on the penalty that a risk-neutral individual gets when she deviates from the truth, risk neutrality being a natural benchmark. With a slight abuse of notation, let  $f(p + \varepsilon, p)$  be the expected score obtained if the event occurs with probability  $p$  while the individual reports probability  $p + \varepsilon$ . Because, by definition of a strictly proper scoring rule, the individual maximizes the expected reward when  $\varepsilon = 0$ , the term  $f(p + \varepsilon, p)$  is of second order in  $\varepsilon$ . Formally, we can write  $f(p)(\omega) = V(p) + (\omega - p)V'(p)$ , where  $V$  is the value function of  $f$  (see, for example, Gneiting and Raftery, 2007). Since  $f$  is differentiable, its value function  $V$  is twice differentiable. Let  $p \in (0, 1)$ . Applying Taylor's theorem to  $V$  and then  $V'$  at  $p$  yield

$$V(p + \varepsilon) = V(p) + \varepsilon V'(p) + \frac{1}{2} \varepsilon^2 V''(p) + o(\varepsilon^2),$$

$$V'(p + \varepsilon) = V'(p) + \varepsilon V''(p) + o(\varepsilon),$$

so that

$$f(p + \varepsilon, p) = f(p, p) - \frac{1}{2} \varepsilon^2 V''(p) + o(\varepsilon^2),$$

and thus  $f(p + \varepsilon, p)$  is of second order in  $\varepsilon$ , since  $V'' > 0$  by strict convexity of  $V$ .

The coefficient of this second-order term determines the extent to which the scoring rule deters deviations from the truth. We therefore consider the following condition:

$$\lim_{\varepsilon \rightarrow 0} \frac{f(p, p) - f(p + \varepsilon, p)}{\varepsilon^2} \geq L \quad (6)$$

where  $L > 0$  is a fixed parameter. The greater  $L$ , the more the scoring rule deters deviations for a risk-neutral individual. Let  $\mathcal{C}(L)$  be the set of differentiable strictly proper scoring rules  $f$  that satisfy Equation (6).

We can now state our goal formally. Fixing the coefficient of risk aversion  $\alpha$ , for a scoring rule  $f$ , let  $r_f(\pi)$  denote the optimal announcement as a function of the true belief  $\pi$ . For given values of  $L$  and  $\alpha$ , we ask what scoring rules  $f \in \mathcal{C}(L)$  minimize the amount of misreporting in the worst case.

**Proposition 3.** *For any parameter  $L > 0$  and any coefficient of absolute risk aversion  $\alpha > 0$ , a scoring rule  $f \in \mathcal{C}(L)$  minimizes the worst-case degree of misreporting, captured by  $\sup_{\pi \in (0,1)} D(\pi, r_f(\pi))$ , if and only if  $f(p)(\omega) = K - L(\omega - p)^2$ , for  $K$  an arbitrary constant.*

Hence, Proposition 3 states that the scoring rules that minimize the degree of misreporting subject to the incentives normalization condition are the quadratic scoring rules. Moreover, for such quadratic scoring rules, Equation (6) holds for all  $\varepsilon$ , not only in the limit case, and the inequality is tight:  $\forall \varepsilon, f(p, p) - f(p + \varepsilon, p) = L\varepsilon^2$ . The proof of Proposition 3 is in Appendix B.

If instead of having CARA utility for money, the individual has CRRA utility for money, then the problem has a very similar structure and, not surprisingly, a similar solution: fixing the maximum reward  $R_{\text{MAX}}$ , the scoring rule that minimizes the worst-case degree of misreporting is the quadratic scoring rule  $R_{\text{MAX}} - L(\omega - p)^2$ . Without fixing the maximum reward, no scoring rule minimizes misreporting in the sense that, by using the quadratic scoring rule  $K - L(\omega - p)^2$  with  $K$  arbitrarily large, the reports become arbitrarily close to the true beliefs by Corollary 3. We refer the reader to the proof of Proposition 3 for further details.

#### 4.4. Eliciting both subjective probabilities and coefficients of risk aversion

Suppose an experimenter wants to elicit the subjective probability of a random event from a subject assumed to behave as an expected utility maximizer with CARA or CRRA utility for money. We continue to use the notation and state interpretation of the previous subsection. If the experimenter knows the coefficient of risk aversion, then as long as the scoring rule is continuous and strictly proper, he can infer the subjective probability from the reported probability following Propositions 1 and 2. If the experimenter does not know the coefficient of risk aversion, he can attempt to elicit it via a separate experiment. However, the application of separate elicitation schemes on the same subject is problematic in models of choice under uncertainty, where hedging behavior plays a major role and leads to distortions (see Azrieli et al., 2018 for details). To circumvent the issue, the experimenter must elicit the information all at once, and so, in particular, obtain a truthful report of the subject's belief in spite of her unknown aversion to risk.

Propositions 1 and 2 can be applied to show that randomizing over two different quadratic scoring rules enables the experimenter to elicit both the subjective probability and the coefficient of absolute (if CARA) or relative (if CRRA) risk aversion simultaneously in a very simple fashion. For example, let us consider the following protocol:

**First stage:** The subject reports her coefficient of (absolute or relative) risk aversion and her assessed probability that the event occurs.

**Second stage:** The experimenter draws a fair coin. If heads, the experimenter selects the quadratic scoring rule  $f_H(p)(\omega) \equiv 1 - \frac{1}{3}(p - \omega)^2$ . If tails, the experimenter selects the quadratic scoring rule  $f_T(p)(\omega) \equiv 1 - \frac{2}{3}(p - \omega)^2$ .

**Third stage:** The experimenter computes what would be the optimal announcement of the subject for the selected scoring rule—accounting for the subject's risk aversion—and pays the subject accordingly.

**Proposition 4.** *In the protocol just described, the subject reports her coefficient of risk aversion and subjective probability truthfully as a strict best response.*

The proof of Proposition 4 is in Appendix B. Randomization is useful in this experiment because it induces the subject, who, in the first stage, does not know which scoring rule will eventually be applied, to communicate information so as to optimize simultaneously on both scoring rules. And, as it turns out, Propositions 1 and 2 imply that the only way for the subject to optimize simultaneously on both scoring rules is to be truthful: misreporting risk aversion or subjective probability or both generates a strictly suboptimal utility in at least one of the two scoring rules.

Two points deserve mention.

First, the randomization in this protocol is very different from the idea that consists in paying in probability currency or lottery tickets, described for example in Savage (1971). With probability currencies, there are only two possible fixed payoffs and the probability of getting the higher payoff is a function of the report and event outcome. In effect, working

with two fixed payoffs “linearizes” the preferences of the subject, who then acts as if she were neutral to risk. Thus, paying in probability currency does not allow to infer the subject’s risk aversion. Instead, in this protocol, the payoffs are not fixed, they continue to be determined by quadratic scoring rules, while the randomization remains independent of the subject’s announcements and the event outcome.

Second, this protocol is a direct mechanism. Of course, in practice, subjects may not be familiar with measures of risk aversion and so may find themselves unable to provide a meaningful announcement for  $\alpha$ . To remedy this problem, one may consider an indirect mechanism instead, in which the subject is asked to compare a range of lotteries, and from which the coefficient of risk aversion can be uniquely determined in a revealed-preference fashion. Propositions 1 and 2 imply that the subject is strictly penalized if she misreports her true subjective belief, in spite of her aversion to risk.

#### 4.5. Identification of utilities and scoring rules

In this subsection, we return to the general case  $\Omega = \{1, \dots, n\}$ , and apply the dual characterization to solve two problems of identification.

In the first problem of identification, we ask how to identify the utility function  $U$  of an individual when  $U$  is a translation-invariant utility (see Example 1 for a definition), i.e., when the individual’s preferences are variational.

**Proposition 5.** *Let  $f$  be a continuous strictly proper scoring rule. Suppose that two individuals have translation-invariant utility functions. If, for all side state-contingent payoffs  $y \in \mathbb{R}^n$ , the respective optimal announcements of the two individuals rewarded according to  $f + y$  are the same, then the two individuals have the same utility function up to an additive constant.*

Proposition 5 states that, no matter the baseline scoring rule being used, by varying the side state-contingent payoffs one can fully identify the individual’s preferences. The proof is in Appendix B. To understand the main idea, let us consider the special case of two individuals  $k = 1, 2$ , whose indirect utility functions are given by  $G_k(p, w) = w + c_k(p)$ , with  $c_k$  convex and real-valued. Fix a continuous strictly proper scoring rule  $f$  with value function  $V$ , and assume that  $c_1, c_2, V$  are extended as convex functions on the nonnegative orthant  $\mathbb{R}_+^n$  and are differentiable on this domain. Under the conditions of Proposition 5, and applying Theorem 1, for any  $y \in \mathbb{R}_+^n$ , both individuals choose to announce the vector of probabilities  $p^*$  that solves, for  $k = 1, 2$ ,

$$\min_{p \in \mathbb{R}_+^n} \{c_k(p) + V(p) + y \cdot p\},$$

subject to the condition  $p_1 + \dots + p_n = 1$ . If  $p_i^* > 0$  for all  $i$ , the first-order conditions that determine  $p^*$  are that for every  $i \neq j$ , and for  $k = 1, 2$ ,

$$\frac{\partial c_k(p^*)}{\partial p_i} + \frac{\partial V(p^*)}{\partial p_i} + y_i = \frac{\partial c_k(p^*)}{\partial p_j} + \frac{\partial V(p^*)}{\partial p_j} + y_j.$$

Hence, we have

$$\left[ \frac{\partial c_k(p^*)}{\partial p_i} - \frac{\partial c_k(p^*)}{\partial p_j} \right] = y_j - y_i + \left[ \frac{\partial V(p^*)}{\partial p_j} - \frac{\partial V(p^*)}{\partial p_i} \right].$$

Roughly, this equality implies that, as  $y$  varies,  $p^*$  covers the entire simplex. Then, we also have

$$\frac{\partial c_1(p^*)}{\partial p_i} - \frac{\partial c_1(p^*)}{\partial p_j} = \frac{\partial c_2(p^*)}{\partial p_i} - \frac{\partial c_2(p^*)}{\partial p_j},$$

which implies that, at every possible  $p^*$ ,  $c_1$  and  $c_2$  share the same subgradients, which in turn implies that  $c_1$  and  $c_2$  are equal up to an additive constant, and so that the individuals have the same utility function, up to a constant.

The main application of Proposition 5 is in preference elicitation: an experimenter can elicit precisely, as a strict best response, the utility function of subjects who have translation-invariant utilities. This utility function includes all information about the subject’s beliefs, whether simple as with expected utility preferences or more complex as with ambiguity averse preferences, along with her risk and ambiguity attitudes. To do so, the experimenter can utilize a mechanism of the sort described in the preceding subsection. Letting  $s$  denote the quadratic scoring rule, for example, the experimenter could operate as follows. First, he asks the subject to communicate her entire utility function—either directly, or indirectly via questionnaires and a sequence of binary choices over lotteries. No payments are made at this stage. Second, the experimenter draws  $y$  at random, with a full-support probability measure over  $\mathbb{R}^n$ .<sup>18</sup> Third, he rewards the subject with the

<sup>18</sup> Of course, in practice, the support will be bounded. For example, one can randomize uniformly over a range of possible payments. Having bounded support does not induce any sort of misreporting, but may possibly make “reporting the truth” a non-unique best response. The larger the support, the closer we get to the unique best response.

scoring rule  $f = s + y$ , i.e., the baseline quadratic scoring rule plus the random side payments, using as input the optimal announcement the subject would have made had she confronted this scoring rule directly.

We now turn to the second problem of identification. In this problem, the preferences are fixed: we consider an individual who behaves under “complete ignorance” as in Example 8, whose utility function is

$$U(x) = \min_{\omega \in \Omega} x_{\omega}.$$

We refer to these preferences as *complete maxmin* preferences (to distinguish with the preferences of maxmin decision makers over nondegenerate set of priors, as in Example 2).

**Proposition 6.** *Let  $f$  and  $g$  be continuous strictly proper scoring rules. If, for all side state-contingent payoffs  $y \in \mathbb{R}^n$ , the optimal announcement of an individual with complete maxmin preferences is the same under both  $f + y$  and  $g + y$ , then  $f$  and  $g$  are identical up to an additive constant.*

Therefore, it is possible to fully identify the scoring rule being applied, up to an additive constant, by varying side payments while observing the probabilities announced by the individual. The proof is in Appendix B. Note that, although we find it convenient to have unrestricted side payments, we can also restrict them to take values in a much smaller set, such as the set for which  $\sum_{i=1}^n y_i = 0$ .

Though Proposition 6 is stated in terms of proper scoring rules, it also has an interpretation in terms of Robinson Crusoe economies. Namely, suppose two given convex and closed technologies with free disposability, and an individual with Leontief preferences. If the equilibrium prices under the two technologies are the same when varying endowments, then the two technologies must coincide up to some constant.

Of course, there would be no hope of establishing such a result under the hypothesis that the individual is risk neutral rather than maxmin, since the risk-neutral individual always announces her true belief, irrespective of which scoring rule is applied and independently of any side payments. More generally, the probability distributions announced by individuals who confront a given scoring rule reflect information regarding her preferences and subjective beliefs, and also information regarding the scoring rule. At one end of the spectrum, an individual who forms a subjective assessment of the state probabilities and is risk neutral makes reports that fully identify her beliefs, but conveys no information on the scoring rule being applied. At the other end, an individual who behaves under complete ignorance makes reports that reflect no particular probabilistic beliefs, but allows to fully identify the scoring rule being applied (up to an additive constant).

Proposition 6 is of theoretical nature, but it also means that, in principle, an experimenter could test if the subjects report meaningful probabilities, or if, on the opposite, the subjects' reporting behavior is dominated by risk or ambiguity aversion. For example, suppose an experimenter works with a large number of subjects, and requests from each subject, individually, a subjective probability distribution over relevant states. Beliefs may be similar across subjects, or may differ (one may think, for example, of an experiment that tests conformity). The experimenter is unsure about whether the subjects are probabilistically sophisticated and close to being risk neutral, in which case reports are meaningful, or if the subjects are closer to behave under complete ignorance, in which case reports are meaningless. If the experimenter uses a single scoring rule for all subjects, the experimenter cannot distinguish between the risk-neutral probabilistically-sophisticated subjects who form a common belief and the subjects who behave under complete ignorance. Proposition 6 implies that, by applying different scoring rules on different groups of subjects, the experimenter can tell if the subjects' reports reflect meaningful information on their beliefs.

#### 4.6. Comparative statics with respect to uncertainty aversion

This last example illustrates the use of the dual characterization to obtain comparative statics. We focus on the special case of binary states that describe the outcome of a random event; as before,  $\Omega = \{0, 1\}$ ,  $\omega = 1$  if the event occurs, and  $\omega = 0$  otherwise. We continue to abuse notation and use  $p$  to represent event probabilities.

Following the definition of Yaari (1969), we say that an individual with utility function  $U_1$  is more *uncertainty averse* than an individual with utility function  $U_2$  if, for every state-independent payoff  $(K, K) \in \mathbb{R}^2$ , and every state-contingent payoff  $(x_0, x_1) \in \mathbb{R}^2$ ,  $U_1(x_0, x_1) \geq U_1(K, K)$  implies  $U_2(x_0, x_1) \geq U_2(K, K)$ ; that is, if the individual with  $U_1$  prefers a given risky payoff to a given safe payoff, then so does the individual with  $U_2$ .

Let us fix two utility functions  $U_k$ ,  $k = 1, 2$ , in the class of variational preferences, with  $U_1$  more uncertainty averse than  $U_2$ . The indirect utility for  $U_k$  is

$$G_k(p, w) = w + c_k(p)$$

for  $c_k$  a proper convex lower semicontinuous function (see Example 1). Let us impose the normalization  $\min_{p \in [0, 1]} c_k(p) = 0$ . This normalization is without loss of generality since utility-based preferences are defined up to an additive constant. Then, the requirement that  $U_1$  is more uncertainty averse than  $U_2$  is equivalent to the requirement that for every  $p \in [0, 1]$ ,  $c_2(p) \geq c_1(p)$ . Intuitively,  $c_k(p)$  captures the “degree of uncertainty aversion at event probability  $p$ .” Assume, for simplicity, that  $c_k$ ,  $k = 1, 2$ , is real-valued, strictly convex, and twice continuously differentiable. Together with the normalization above,



this assumption implies existence of a unique probability  $\pi \in [0, 1]$  such that  $c_1(\pi) = c_2(\pi) = 0$ . Let  $f$  be a continuously differentiable scoring rule. This differentiability implies that the value function  $V$  is twice continuously differentiable (which, as discussed before, follows from the subgradient representation of scoring rules).

Consider an individual with variational preferences described by the indirect utility  $G(p, w) = w + \lambda c_1(p) + (1 - \lambda)c_2(p)$ , with  $\lambda \in (0, 1)$ . The parameter  $\lambda$  captures the extent to which the individual is uncertainty averse: as  $\lambda$  increases, the individual becomes more uncertainty averse. The probability  $\pi$  can be interpreted as the belief of the individual regarding the event. In particular, for the case of a risk-neutral individual or an expected utility maximizer,  $\pi$  is exactly the subjective probability.

By Proposition 2, when rewarded according to  $f$ , the individual chooses to announce the probability  $p^*$  that solves

$$\min_{p \in [0, 1]} \{V(p) + \lambda c_1(p) + (1 - \lambda)c_2(p)\}.$$

Suppose  $p^* \in (0, 1)$  and  $p^* \neq \pi$ . Then,  $p^*$  is determined by the first-order condition

$$V'(p^*) + \lambda c_1'(p^*) + (1 - \lambda)c_2'(p^*) = 0.$$

To understand how the optimal probability announcement  $p^*$  varies as the individual becomes increasingly uncertainty averse, we apply the implicit function theorem and get

$$\frac{dp^*}{d\lambda} = \frac{c_2'(p^*) - c_1'(p^*)}{V''(p^*) + \lambda c_1''(p^*) + (1 - \lambda)c_2''(p^*)}.$$

The denominator is positive by strict convexity of  $V + \lambda c_1 + (1 - \lambda)c_2$ , so we have that  $dp^*/d\lambda$  has the sign of  $c_2'(p^*) - c_1'(p^*)$ . Therefore, the fact that the individual becomes more uncertainty averse is not enough to tell if the individual's announcement moves closer to or further away from the belief  $\pi$ .

Observe that, if  $p^* > \pi$ , then  $c_1'(p^*) > 0$ , so that  $c_2'(p^*) - c_1'(p^*) > 0$  exactly when  $c_2'(p^*)/c_1'(p^*) > 1$ , and that if  $p^* < \pi$ , then  $c_1'(p^*) < 0$ , so that  $c_2'(p^*) - c_1'(p^*) < 0$  exactly when  $c_2'(p^*)/c_1'(p^*) > 1$ . In either case, the individual moves away from her true belief exactly when  $c_2'(p^*)/c_1'(p^*) > 1$ .

We conclude that the relative rate of change of uncertainty aversion of the two preferences  $U_1$  and  $U_2$  at the current optimal announcement  $p^*$ , as measured by the ratio

$$\frac{c_2'(p^*)}{c_1'(p^*)},$$

is the appropriate statistic by which to judge whether a change in uncertainty aversion leads to a more accurate announcement. If  $c_2'(p^*)/c_1'(p^*) > 1$ , then as the individual becomes more uncertainty averse, her announcement moves away from her belief  $\pi$ . If, instead,  $c_2'(p^*)/c_1'(p^*) < 1$ , then as the individual becomes more uncertainty averse, her announcement moves closer to her belief  $\pi$ .

## 5. Conclusion and related literature

The literature on scoring rules is vast. The first characterization of proper scoring rules, based on the subgradients of convex functions, goes back to McCarthy (1956); see also the generalizations of Savage (1971) and Fang et al. (2010). Gneiting and Raftery (2007) provide a recent survey of the literature.

Winkler and Murphy (1970) are the first to depart from the risk neutrality assumption, by studying how the curvature of utility for money affects behavior under the quadratic scoring rule for expected utility maximizers. Continuing to assume the quadratic scoring rule and considering binary events, Kadane and Winkler (1988) calculate an explicit formula that expected utility maximizers would use when they have nontrivial risk attitudes. They show that risk averse individuals tend to report more uniform probabilities. Armantier and Treich (2013) extend this result to a larger class of scoring rules, and, in particular, study how the bias towards uniform probability assessments vary with risk aversion. Offerman et al. (2009) derive a related result in a more general decision-theoretic model. Grünwald and Dawid (2004) and Chambers (2008) uncover optimal behavior in the context of risk-neutral individuals with multiple priors. Bickel (2007) establishes properties of individuals with CARA-style expected utility preferences: for example, he shows that one can add a constant payoff to each action in the profile of a scoring rule without changing behavior. He attributes this to what he calls the “delta” property; something economists would call translation invariance or quasilinearity, and is characterized by the variational model. This result is implied by Theorem 1. Jose et al. (2008) discusses a duality related to that of Grünwald and Dawid (2004), but with a different aim: they seek to understand how a risk-averse expected utility maximizer will “bet” against a given set of priors.

We show that the quadratic scoring rule is optimal in the sense that it minimizes deviations from truth-telling for both CARA and CRRA utilities for money. The quadratic scoring rule has other desirable properties. For example, Selten (1998) shows that it is the unique scoring rule that is proper, has symmetric losses (meaning  $V(p) - p \cdot f(q) = V(q) - q \cdot f(p)$  for all  $p$  and  $q$ ), is invariant to adding zero-probability events, and is invariant to permutations of the labeling of events.

Assuming a rational inattention model, Tsakas (2018) proves that, for a certain family of inattention costs, the quadratic scoring rule virtually eliminates incentives to collect new information about the event in question.

A different stream of literature, pioneered by Maccheroni et al. (2006) and Cerreia-Vioglio et al. (2011b), exploits the duality between indirect and direct utility in order to study properties of uncertainty aversion. These works use the richer structure of Anscombe and Aumann (1963) acts, which allows for the separate study of uncertainty and risk. The duality investigated in Cerreia-Vioglio et al. (2011b) and discussed in detail in Cerreia-Vioglio et al. (2011a) corresponds to the one being used in this paper.<sup>19</sup> Although these works are concerned with uncertainty aversion, many of the results and characterizations continue to apply to the setup of this paper. Further, because of risk attitudes, it is often advocated that individuals be paid in probability currency, instead of monetary terms. It relies on the idea that, over purely risky prospects, individuals will likely conform to expected utility behavior (indeed this is the framework upon which the analysis of Anscombe and Aumann (1963) is built). In such a framework, the model of Cerreia-Vioglio et al. (2011b) is the appropriate one for studying elicitation questions. The natural counterparts of the following results hold as stated, we focus here on monetary payoffs for simplicity. Similarly, one could restrict the domain of scoring rules to take only nonnegative values, and describe “homogeneous” utility indexes  $U$ ; these would correspond to indirect utility functions  $G$  for which  $G(p, w) = wG(p, 1)$ . Finally, comparative statics on the indirect utility function in terms of the risk aversion relation of Yaari (1969) exists in Cerreia-Vioglio et al. (2011b) in the form of comparative statics on uncertainty aversion.

This paper adds to the literature by considering a general specification of utility together with general proper scoring rules. Previous works assume that preference is expected utility, or consider the special case of risk neutrality with ambiguity aversion, or focus on the quadratic scoring rule. In contrast, our preferences do not need to reference any concept of likelihood or state-contingent utility payoffs whatsoever, and our results do not require particular restrictions on the scoring rule being used, except for its continuity. The duality approach we pursue allows a sharp and simple characterization of individual behavior, which allows for a range of applications. We leave the question of infinite state spaces, and the study of elicitation mechanisms utilizing objective randomization devices, to future research.

## Appendix A. Proofs of Section 2

### A.1. Proof of Theorem 1

As a first step, consider the set  $K = \text{co}(f(\Delta(\Omega)))$ , the convex hull of the image of  $f$ . Observe that  $K$  is itself compact, since  $\mathbb{R}^\Omega$  is finite-dimensional (see Corollary 5.18 of Aliprantis and Border, 1999) and  $f(\Delta(\Omega))$  is compact due to continuity of  $f$ . We will show that there is a unique maximizer of  $U$  across the set  $K$ , and that this maximizer coincides with  $\arg \max_{p \in \Delta(\Omega)} U(f(p))$ .

Let  $x^* \in \arg \max_{x \in K} U(x)$  (such a maximizer is guaranteed to exist due to continuity of  $U$  and compactness of  $K$ ). Let  $Y = \{y : U(y) > U(x^*)\}$ . This set is open (by continuity of  $U$ ) and convex (by quasiconcavity of  $U$ ). By Theorem 5.50 of Aliprantis and Border (1999), as  $K \cap Y = \emptyset$ , the sets  $K$  and  $Y$  can be separated by a hyperplane. This hyperplane can be normalized to have direction in  $\Delta(\Omega)$ , by the fact that  $U$  is increasing. Let us call this direction  $p^*$ . Observe that the hyperplane of direction  $p^*$  passes through  $x^*$ , as for any  $\epsilon > 0$ ,  $x^* + \epsilon(1, \dots, 1) \in Y$ . Hence, we conclude that for all  $x \in K$ ,  $p^* \cdot x \leq p^* \cdot x^*$ ; i.e.,  $x^*$  maximizes  $p^* \cdot x$  subject to  $x \in K$ . Clearly  $f(p^*) \in K$  satisfies this inequality. We claim that it is the unique such element of  $K$ . So, by contradiction, let  $\hat{x} \in K$ , where  $\hat{x} \neq f(p^*)$ . Then there are  $p_1, \dots, p_n \in \Delta(\Omega)$ , not all equal to  $p^*$  and  $\lambda_i \geq 0$ ,  $\sum_{i=1}^n \lambda_i = 1$  for which  $\hat{x} = \sum_{i=1}^n \lambda_i f(p_i)$ . But by strict properness, we then obtain  $p^* \cdot \hat{x} < p^* \cdot f(p^*)$ , contradicting the fact that  $\hat{x}$  maximizes  $p^* \cdot x$  subject to  $x \in K$ . So  $x^* = f(p^*)$ . Further,  $f(p^*)$  is the unique maximizer of  $U$  in  $K$ . Indeed, observe that by continuity and monotonicity, the closure of  $Y$  is  $\{y : U(y) \geq U(f(p^*))\}$ . Hence if  $x' \in \arg \max_{x \in K} U(x)$ , then  $p^* \cdot x' \geq p^* \cdot f(p^*)$ , which we have shown to be impossible.

Since  $p^*$  separates  $K$  and  $Y$ , and again by continuity and monotonicity  $U$ , we have that  $U(y) \geq U(f(p^*))$  (i.e.,  $y$  in the closure of  $Y$ ) implies  $p^* \cdot y \geq p^* \cdot f(p^*)$ . We claim that this fact, in turn, implies  $U(f(p^*)) = G(p^*, p^* \cdot f(p^*))$ . To see this, suppose by means of contradiction that there is  $y$  for which  $p^* \cdot y \leq p^* \cdot f(p^*)$  and  $U(y) > U(f(p^*))$ . By continuity of  $U$ , we conclude that there is  $y^*$  for which  $p^* \cdot y^* < p^* \cdot f(p^*)$  and  $U(y^*) > U(f(p^*))$ , contradicting the fact that  $U(y^*) \geq U(f(p^*))$  implies  $p^* \cdot y^* \geq p^* \cdot U(f(p^*))$ . So,  $U(f(p^*)) = G(p^*, p^* \cdot f(p^*))$ .

Finally, let  $\hat{p} \neq p^*$ . Then  $\hat{p} \cdot f(p^*) < \hat{p} \cdot f(\hat{p})$ , by strict properness. Therefore, there is  $\epsilon > 0$  for which  $\hat{p} \cdot (f(p^*) + \epsilon(1, \dots, 1)) < \hat{p} \cdot f(\hat{p})$ , and since  $U$  is monotonic, we then conclude that  $G(\hat{p}, \hat{p} \cdot f(\hat{p})) > U(f(p^*))$ .

Overall, we get that  $p^*$  is the unique minimizer of the function  $p \mapsto G(p, p \cdot f(p))$ .

### A.2. Proof of Lemma 1

First, suppose that  $p^* \in \Delta(\Omega)$  has the property that for all  $\omega, \omega' \in \Omega$ ,  $f(p^*)(\omega) = f(p^*)(\omega') = V(p^*)$ . Then, by definition, for any  $p \neq p^*$ ,  $V(p^*) = p \cdot f(p^*) \leq p \cdot f(p) = V(p)$ .

Conversely, suppose that  $p^* \in \Delta_{++}(\Omega)$  (meaning that  $p^*$  has full support) and  $p^* \in \arg \min_{p \in \Delta(\Omega)} V(p)$ . Observe that by properness, for any  $p$ ,  $0 \leq V(p) - V(p^*) \leq (p - p^*) \cdot f(p)$ , so  $(p - p^*) \cdot f(p) \geq 0$ . Pick any two states,  $\omega', \omega'' \in \Omega$ ,

<sup>19</sup> Specifically, they apply the duality to von Neumann-Morgenstern utilities under more general framework.

$\omega' \neq \omega''$ . Because  $p^* \in \Delta_{++}(\Omega)$ , there exists  $\epsilon > 0$  small enough so that  $p_\epsilon$  as defined by the following is an element of  $\Delta(\Omega)$ :  $p_\epsilon(\omega') = p^*(\omega') + \epsilon$ ,  $p_\epsilon(\omega'') = p^*(\omega'') - \epsilon$ , and finally  $p_\epsilon(\omega) = p^*(\omega)$  for all  $\omega \notin \{\omega', \omega''\}$ . Observe then that  $\epsilon(f(p_\epsilon)(\omega') - f(p_\epsilon)(\omega'')) \geq 0$ . Since  $\epsilon > 0$  was arbitrary, we conclude by continuity of  $f$  that  $f(p^*)(\omega') \geq f(p^*)(\omega'')$ . The result of the lemma follows from the observation that  $\omega', \omega''$  were arbitrary.

## Appendix B. Proofs of Section 4

### B.1. Proof of Lemma 2

Let  $\mathbf{1}_\omega$  be the vector of  $\mathbb{R}^n$  whose elements are all zero except for the element  $\omega$ , that takes value 1. The proof of Theorem 1 shows that if  $p^*$  is the optimal announcement, then  $U(y) \geq U(f(p^*))$  implies  $p^* \cdot y \geq p^* \cdot f(p^*)$ . Suppose by means of contradiction that there is  $\omega$  for which  $p_\omega^* = 0$ . Observe by strict monotonicity of  $U$  that  $U(f(p^*) + \mathbf{1}_\omega) > U(f(p^*))$ . So there exists  $\epsilon > 0$  for which  $U(f(p^*) + \mathbf{1}_\omega - \epsilon(1, \dots, 1)) > U(f(p^*))$ . But since  $p_\omega^* = 0$ ,  $p^* \cdot (f(p^*) + \mathbf{1}_\omega - \epsilon(1, \dots, 1)) < p^* \cdot f(p^*)$ , a contradiction.

### B.2. Proof of Proposition 3

A direct application of Proposition 1 yields that the optimal report  $p^*$  and the subjective probability  $\pi$  satisfy the equality

$$\frac{\pi}{1 - \pi} = \frac{p^*}{1 - p^*} e^{\alpha[f(p^*)(1) - f(p^*)(0)]}.$$

That is, the odds-ratio of the subjective belief to the optimal announcement is equal to  $\exp(\alpha[f(p^*)(1) - f(p^*)(0)])$ . In particular, fixing  $f$  and  $\alpha$ , as  $\pi$  varies over the open range  $(0, 1)$ , optimal reports span the range  $(0, 1)$ . Hence, the image of  $r_f$  is  $(0, 1)$ .

Recall that we have  $f(p)(\omega) = V(p) + (\omega - p)V'(p)$ , which owes to the fact that  $f$  captures the subgradients of the value function  $V$  (Gneiting and Raftery, 2007). Taylor's theorem then implies that

$$\lim_{\epsilon \rightarrow 0} \frac{f(p, p) - f(p + \epsilon, p)}{\epsilon^2} = \frac{1}{2} V''(p).$$

Thus, formally, the problem is to find the smooth strictly proper scoring rules  $f$  such that  $D(\pi, r(\pi))$  is minimized subject to the condition that  $V'' \geq 2L$ , where  $V$  is the value function associated to  $f$ . Since  $f(1, 1) - f(1, 0) = V'(1)$  and  $f(0, 0) - f(0, 1) = -V'(0)$ , the problem reduces to minimizing

$$\max_{p \in [0, 1]} |f(p)(1) - f(p)(0)| = \max \{V'(1), -V'(0)\},$$

subject to  $V'' \geq 2L$ , where we used the monotonicity of  $p \mapsto f(p)(\omega)$ . We immediately get  $V'(1) = -V'(0)$  and  $V'' = 2L$ , which implies  $V'(p) = 2L(p - 1/2)$  and  $V(p) = K + L(p^2 - p)$ . Hence, applying the formula  $f(p)(\omega) = V(p) + (\omega - p)V'(p)$ , we get  $f(p)(\omega) = K - L(p - \omega)^2$  where  $K$  is an arbitrary constant.

If the individual is an expected CRRA utility maximizer, then Corollary 3 implies that one can get arbitrarily low distortions in reported assessments by choosing to add to the state-contingent rewards a fixed payment large enough. Hence, let us impose the condition that  $f(p)(\omega) \leq R_{\text{MAX}}$  for an arbitrary  $R_{\text{MAX}} > 0$ , which captures the fact that the payments of the scoring rule are bounded. In addition, we impose positive payoffs to be within the domain of our CRRA preferences as defined in Example 5.

The steps are analogous to those above. A direct application of Proposition 2 yields that the optimal report  $p^*$  and the subjective probability  $\pi$  satisfy the equality

$$\frac{\pi}{1 - \pi} = \frac{p^*}{1 - p^*} \left( \frac{f(p)(1)}{f(p)(0)} \right)^\eta.$$

Hence, finding the smooth scoring rule that minimizes the worst-case degree of misreporting reduces to minimizing

$$\max \left\{ \max_{p \in [0, 1]} \frac{f(p)(1)}{f(p)(0)}, \max_{p \in [0, 1]} \frac{f(p)(0)}{f(p)(1)} \right\} = \max \left\{ \frac{V(0)}{V(0) + V'(0)}, \frac{V(1)}{V(1) - V'(1)} \right\},$$

where we note that  $f(1, 1) = V(1)$ ,  $f(1, 0) = V(1) - V'(1)$ ,  $f(0, 0) = V(0)$ ,  $f(0, 1) = V(0) + V'(0)$ . As for case of CARA utility, the constraint  $V'' \geq 2L$  is binding, which yields as solution the quadratic scoring rule  $f(p)(\omega) = R_{\text{MAX}} - L(p - \omega)^2$  (for  $R_{\text{MAX}} > L$  to ensure positiveness of the payoffs.)

B.3. Proof of Proposition 4

Let us focus on the case of CARA utility for money. The case of CRRA utility for money is proved in a similar fashion.

Of course, because the experimenter acts in the best interest of the subject, making truthful announcements is always a best response. The core of the proof is the argument that the truthful best response is strict.

Because of the randomization that occurs in stage 2, the subject is induced to make reports that provide her with the optimal payments for both scoring rules simultaneously.

Suppose that  $(\alpha, \pi)$  captures the true pair (coefficient of absolute risk aversion, subjective probability), and that the announcement  $(\alpha', \pi')$  is optimal. Let  $p_H$  be the probability the experimenter plugs into scoring rule  $f_H$  when the subject announces  $(\alpha, \pi)$ , and  $p_T$  the probability he plugs into scoring rule  $f_T$ . The fact that payments remain optimal for  $(\alpha', \pi')$  with both  $f_H$  and  $f_T$  implies that the experimenter would continue to plug in  $p_H$  and  $p_T$ , respectively, when the subject announces  $(\alpha', \pi')$ .

For scoring rule  $f = f_H, f_T$ , let  $\Delta f(p)$  denote the difference  $f(p)(1) - f(p)(0)$ . Proposition 1 then implies the following equalities:

$$\frac{\pi}{1 - \pi} = \frac{p_H}{1 - p_H} e^{\alpha \Delta f_H(p_H)}, \tag{7}$$

$$\frac{\pi'}{1 - \pi'} = \frac{p_H}{1 - p_H} e^{\alpha' \Delta f_H(p_H)}, \tag{8}$$

$$\frac{\pi}{1 - \pi} = \frac{p_T}{1 - p_T} e^{\alpha \Delta f_T(p_T)}, \tag{9}$$

$$\frac{\pi'}{1 - \pi'} = \frac{p_T}{1 - p_T} e^{\alpha' \Delta f_T(p_T)}. \tag{10}$$

Dividing (10) by (9) and (8) by (7), we obtain

$$\frac{\pi'(1 - \pi)}{\pi(1 - \pi')} = e^{\Delta f_T(p_T)(\alpha' - \alpha)} = e^{\Delta f_H(p_H)(\alpha' - \alpha)},$$

which implies  $\Delta f_T(p_T) = \Delta f_H(p_H)$ . Let us now divide (9) by (7), we obtain

$$\frac{p_T}{1 - p_T} = \frac{p_H}{1 - p_H},$$

and thus  $p_T = p_H$ . Putting those two facts together, we get  $\alpha = \alpha'$  and  $\pi = \pi'$ .

B.4. Proof of Proposition 5

Consider two individuals indexed  $k = 1, 2$  with utility function  $U_k$  and indirect utility function  $G_k$ . From Example 1, we can write  $G_k(p, w) = w + c_k(p)$  for a proper convex and lower semicontinuous function  $c_k$  on  $\Delta(\Omega)$ . We extend  $c_k$  on  $\mathbb{R}_+^n$  as follows: let  $c_k(q) = +\infty$  if  $q \in \mathbb{R}_+^n$  and  $q \notin \Delta(\Omega)$ . Observe that  $c_k, k = 1, 2$ , continues to be a proper lower semicontinuous convex function.

Let us fix a continuous strictly proper scoring rule  $f$ , with value function  $V$ . We also extend  $V$  on the nonnegative orthant  $\mathbb{R}_+^n$  by defining, for all  $q \in \mathbb{R}_+^n$ ,

$$W(q) = \sup_{p \in \Delta(\Omega)} f(p) \cdot q,$$

as we did in other parts of this paper. Recall that  $W$  is a differentiable real-valued convex function, and that  $W(p) = V(p)$  for  $p \in \Delta(\Omega)$ .

Now, let  $p^*$  be in the (relative) interior of  $\Delta(\Omega)$  and such that  $c_1(p^*) < +\infty$  (recall that  $c_1$  may be infinite-valued, but there exists at least one such  $p^*$ ). Because  $W$  is differentiable, it has a unique subgradient  $z$  at  $p^*$ . Then, let  $z'$  be any subgradient of  $c_1$  at  $p^*$ , and let  $y = -(z + z')$ .

We begin with the observation that, if the first individual is rewarded according to scoring rule  $f$  plus side state-contingent payoff  $y$ , then by Theorem 1, the individual chooses to report  $p^*$ . Indeed, for all  $p \in \Delta(\Omega)$ ,

$$W(p^*) + z \cdot (p - p^*) \leq W(p)$$

and

$$c_1(p^*) + z' \cdot (p - p^*) \leq c_1(p)$$

so

$$c_1(p^*) + W(p^*) + y \cdot p^* \leq c_1(p) + W(p) + y \cdot p,$$

and thus  $p^*$  minimizes  $G_1(p, W(p) + y \cdot p) = G_1(p, V(p) + y \cdot p)$  over  $\Delta(\Omega)$ , where  $p \mapsto V(p) + y \cdot p$  is the value function associated to the combination of scoring rule and outside payoff.

Suppose that the second individual also finds it optimal to announce  $p^*$  when rewarded according to the same combination  $f + y$ . Then,  $c_2(p^*) < \infty$ , and further, Theorem 1 implies that for all  $p \in \Delta(\Omega)$ ,

$$c_2(p^*) + W(p^*) + y \cdot p^* \leq c_2(p) + W(p) + y \cdot p,$$

and hence,

$$W(p^*) - W(p) \leq c_2(p) - c_2(p^*) + y \cdot (p - p^*),$$

inequality that continues to hold for all  $p \in \mathbb{R}_+^n$ , as  $c_2$  is positively infinite outside of  $\Delta(\Omega)$ .

Let the functions  $g, h$  be defined on  $\mathbb{R}_+^n$  as  $g(p) = W(p^*) - W(p)$  and  $h(p) = c_2(p) - c_2(p^*) + y \cdot (p - p^*)$ . We have  $g(p^*) = h(p^*) = 0$ ,  $g \leq h$ ,  $g$  is concave and  $h$  is proper convex and lower semicontinuous. Applying the separating hyperplane theorem to the (interior of) the hypograph of  $g$  and the epigraph of  $h$  (when nonempty), there exists a hyperplane that lies weakly above  $g$  and weakly below  $h$ . In addition, because  $W$  has a unique subgradient  $z$  at  $p^*$ ,  $g$  has a unique supergradient  $-z$  at  $p^*$ , and there is a unique supporting hyperplane of the hypograph of  $g$  at  $p^*$ . Hence,  $-z$  is a subgradient of  $h$  at  $p^*$ , which means that for every  $p \in \mathbb{R}_+^n$ ,

$$h(p^*) - z \cdot (p - p^*) \leq h(p)$$

or equivalently, using that  $y = -(z + z')$ ,

$$c_2(p^*) + z' \cdot (p - p^*) \leq c_2(p).$$

Therefore,  $z'$  is a subgradient of  $c_2$  at  $p^*$ .

To sum up, if the two individuals make identical reports when rewarded under the same combination of scoring rule and side payments, we have that for every  $p^*$  such that  $c_1(p^*) < +\infty$ ,  $c_2(p^*) < +\infty$  and any subgradient of  $c_1$  at  $p^*$  is also a subgradient of  $c_2$  at  $p^*$ —and conversely.

Using that proper convex functions are determined, up to an additive constant, by their subgradient correspondence (see Theorem 24.9 of Rockafellar, 1970), we get that  $c_1$  and  $c_2$  are identical up to an additive constant, which implies that  $U_1$  and  $U_2$  are identical up to an additive constant.

### B.5. Proof of Proposition 6

The indirect utility associated with complete maxmin preferences is  $G(p, w) = w$ . If  $V$  is the value function associated to the scoring rule being applied, and  $y$  is the side state-contingent payoff, then by Theorem 1, the optimal report  $p^*$  solves

$$\min_{p \in \Delta(\Omega)} G(p, V(p) + y \cdot p) = \min_{p \in \Delta(\Omega)} \{V(p) + y \cdot p\}.$$

Hence,

$$V(p) \geq V(p^*) + y \cdot (p^* - p) = V(p^*) - y \cdot (p - p^*),$$

and  $-y$  is a subgradient of  $V$  at  $p^*$ . Conversely, if the inequality  $V(p) \geq V(p^*) - y \cdot (p - p^*)$  is satisfied, then  $p^*$  is the optimal announcement of the individual.

We can extend  $V$  to a function  $V^* : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{+\infty\}$  with output in the extended real line by defining  $V^*(p) = V(p)$  when  $p \in \Delta(\Omega)$  and  $V^*(p) = +\infty$  otherwise. Observe that the subgradient inequalities continue to hold, that is,  $-y$  is a subgradient of  $V^*$  at  $p^* \in \Delta(\Omega)$  if and only if  $p^*$  is the optimal announcement of the individual.

This observation allows us to completely determine the subgradient correspondence  $\partial V^*$  via the optimal announcements of the individual as the side state-contingent payoffs vary. Further,  $V^*$  is determined up to a constant by its subgradient correspondence (see Theorem 24.9 of Rockafellar, 1970).

Finally, since  $V^*$  is determined up to a constant, so is  $V$ . And since the scoring rule being applied is continuous, it is uniquely determined by  $V$ . Adding a constant to  $V$  adds the same (state-independent) constant to  $f$ .

## References

- Aliprantis, C.D., Border, K.C., 1999. *Infinite Dimensional Analysis*, second ed.. Springer.
- Anscombe, F.J., Aumann, R.J., 1963. A definition of subjective probability. *Ann. Math. Stat.* 34 (1), 199–205.
- Armantier, O., Treich, N., 2013. Eliciting beliefs: proper scoring rules, incentives, stakes and hedging. *Eur. Econ. Rev.* 62, 17–40.
- Arrow, K.J., Hurwicz, L., 1972. An optimality criterion for decision-making under ignorance. In: Shackle, G.L.S., Carter, C.F., Ford, J.L. (Eds.), *Uncertainty and Expectations in Economics: Essays in Honour of G.L.S. Shackle*. A.M. Kelley, pp. 1–11.

- Azrieli, Y., Chambers, C.P., Healy, P.J., 2018. Incentives in experiments: a theoretical analysis. *J. Polit. Econ.* 126 (4), 1472–1503.
- Bickel, J.E., 2007. Some comparisons among quadratic, spherical, and logarithmic scoring rules. *Decis. Anal.* 4 (2), 49–65.
- Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., Montrucchio, L., 2011a. Complete monotone quasiconcave duality. *Math. Oper. Res.* 36 (2), 321–339.
- Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., Montrucchio, L., 2011b. Uncertainty averse preferences. *J. Econ. Theory* 146 (4), 1275–1330.
- Chambers, C.P., 2008. Proper scoring rules for general decision models. *Games Econ. Behav.* 63 (1), 32–40.
- Cornes, R., 1992. *Duality and Modern Economics*. Cambridge University Press.
- Fang, F., Stinchcombe, M.B., Whinston, A.B., 2010. Proper scoring rules with arbitrary value functions. *J. Math. Econ.* 46 (6), 1200–1210.
- Gneiting, T., Raftery, A.E., 2007. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* 102 (477), 359–378.
- Good, I.J., 1952. Rational decisions. *J. R. Stat. Soc. B* 14 (1), 107–114.
- Gorman, W.M., 1961. On a class of preference fields. *Metroeconomica* 13 (2), 53–56.
- Grünwald, P.D., Dawid, A.P., 2004. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Ann. Stat.* 32 (4), 1367–1433.
- Jose, V.R.R., Nau, R.F., Winkler, R.L., 2008. Scoring rules, generalized entropy, and utility maximization. *Oper. Res.* 56 (5), 1146–1157.
- Kadane, J.B., Winkler, R.L., 1988. Separating probability elicitation from utilities. *J. Am. Stat. Assoc.* 83 (402), 357–363.
- Konüs, A.A., 1939. The problem of the true index of the cost of living. *Econometrica* 7 (1), 10–29.
- Lau, L.J., 1969. Duality and the structure of utility functions. *J. Econ. Theory* 1 (4), 374–396.
- Maccheroni, F., Marinacci, M., Rustichini, A., 2006. Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74 (6), 1447–1498.
- Machina, M.J., Schmeidler, D., 1992. A more robust definition of subjective probability. *Econometrica* 60 (4), 745–780.
- Mas-Colell, A., Whinston, M.D., Green, J.R., 1995. *Microeconomic Theory*, 1st ed. Oxford University Press.
- McCarthy, J., 1956. Measures of the value of information. *Proc. Natl. Acad. Sci.* 42 (9), 654–655.
- Nyarko, Y., Schotter, A., 2002. An experimental study of belief learning using elicited beliefs. *Econometrica* 70 (3), 971–1005.
- Offerman, T., Sonnemans, J., Van de Kuilen, G., Wakker, P.P., 2009. A truth serum for non-Bayesians: correcting proper scoring rules for risk attitudes. *Rev. Econ. Stud.* 76 (4), 1461–1489.
- Rockafellar, R.T., 1970. *Convex Analysis*. Princeton University Press.
- Roth, A.E., Malouf, M.W., 1979. Game-theoretic models and the role of information in bargaining. *Psychol. Rev.* 86 (6), 574–594.
- Roy, R., 1947. La distribution du revenu entre les divers biens. *Econometrica* 15 (3), 205–225.
- Savage, L.J., 1971. Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* 66 (336), 783–801.
- Selten, R., 1998. Axiomatic characterization of the quadratic scoring rule. *Exp. Econ.* 1 (1), 43–61.
- Selten, R., Sadrieh, A., Abbink, K., 1999. Money does not induce risk neutral behavior, but binary lotteries do even worse. *Theory Decis.* 46 (3), 213–252.
- Shepherd, R.W., 1970. *Theory of Cost and Production Functions*. Princeton University Press.
- Strzalecki, T., 2011. Axiomatic foundations of multiplier preferences. *Econometrica* 79 (1), 47–73.
- Tsakas, E., 2018. *Robust Scoring Rules*. Maastricht University. working paper.
- Ville, J., Newman, P.K., 1951–1952. The existence-conditions of a total utility function. *Rev. Econ. Stud.* 19 (2), 123–128.
- Weymark, J.A., 1980. Duality results in demand theory. *Eur. Econ. Rev.* 14 (3), 377–395.
- Winkler, R.L., Murphy, A.H., 1970. Nonlinear utility and the probability score. *J. Appl. Meteorol.* 9 (1), 143–148.
- Yaari, M.E., 1969. Some remarks on measures of risk aversion and on their uses. *J. Econ. Theory* 1 (3), 315–329.