Duality Between Maximization of Expected Utility and Minimization of Relative Entropy When Probabilities are Imprecise

Robert F. Nau
The Fuqua School of Business
Duke University
Durham, NC, USA
robert.nau@duke.edu

Victor Richmond R. Jose
The McDonough School of Business
Georgetown University
Washington, DC, USA

Robert L. Winkler
The Fuqua School of Business
Duke University
Durham, NC, USA

Abstract

In this paper we model the problem faced by a risk-averse decision maker with a precise subjective probability distribution who bets against a risk-neutral opponent or invests in a financial market where the beliefs of the opponent or the representative agent in the market are described by a convex set of imprecise probabilities. The problem of finding the portfolio of bets or investments that maximizes the decision maker's expected utility is shown to be the dual of the problem of finding the distribution within the set that minimizes a measure of divergence, i.e., relative entropy, with respect to the decision maker's distribution. In particular, when the decision maker's utility function is drawn from the commonly used exponential/logarithmic/power family, the solutions of two generic utility maximization problems are shown to correspond exactly to the minimization of divergences drawn from two commonly-used parametric families that both generalize the Kullback-Leibler divergence. We also introduce a new parameterization of the exponential/logarithmic/power utility functions that allows the power parameter to vary continuously over all real numbers and which is a natural and convenient parameterization for modeling utility gains relative to a non-zero status quo wealth position.

Keywords. decision theory, decision analysis, relative entropy, utility theory, imprecise probabilities, portfolio optimization

1 Introduction

There are many situations in which it is of interest to measure the distance between two probability distributions – say, **p** and **q** – but the appropriate metric may depend on the field of application. In statistics the relevant metric might be the loss that results from basing an inference or decision on **q** when the true distribution is **p**. In information processing the metric might be the channel capacity that is wasted by using an encoding scheme based on **q** when **p** is the true distribution of a stream of independent signals to be transmitted. In decision analysis the metric might be the value of information that results in the updating of a prior subjective probability distribution **q** to a posterior distribution **p** prior to making a choice. In probability forecasting the metric might be a scoring rule that is used to provide an incentive for a forecaster to report **p** rather than **q** as her prediction if she believes **p** is correct. In finance the metric might be the gain in expected utility that can be achieved by an investor in a market under uncertainty when her personal distribution for future asset prices is **p** and she has the opportunity to trade with a "representative agent" whose probability distribution is **q**. If one of the distributions – say, **q** – is imprecise, then the quantity of interest to be measured may be the distance from **p** to the nearest or farthest of the possible values of **q**.

In this paper[1] we consider the problem of measuring the distance between probability distributions in the case where one is imprecise, and we focus especially on the case of expected-utility gains in a financial market, although we also discuss how all of the applications mentioned above are linked to each other by duality relationships in which an information-theoretic measure of distance – known as a *relative entropy* or *divergence* – can be identified with a loss function or a utility function in a decision or inference problem. The best-known relative entropy measure is the *Kullback-Leibler divergence*, but it has a number of generalizations. We show that two well-known parametric families of generalized divergence, namely the *power* and *pseudospherical* families, have a one-to-one correspondence with the two most commonly used parametric families of scoring rules, and they also have a one-to-one correspondence with the solutions of two canonical investment problems involving the most commonly

[1] This paper is adapted from Jose et al. 2008 with some new material. An earlier, incomplete version, Nau et al. 2007, was presented at ISIPTA '07.

used parametric family of utility functions, namely the *generalized power family* that includes the exponential and logarithmic utility functions as special cases. We also introduce a new parameterization of this family of utility functions that allows the power parameter to vary continuously over all real numbers and which is the most natural and convenient parameterization for modeling utility gains relative to a nonzero status quo wealth position. This parameterization turns out to have the property that it yields an exact agreement between the utility scale and the scales that are conventionally used for the generalized divergences.

Imprecise probabilities naturally arise in the analysis of financial markets under uncertainty wherever those markets are incomplete, which is to say, virtually everywhere. A market is incomplete if some assets have distinct bid and ask prices (or are not priced at all) because of caution or lack of information on the part of buyers and sellers and/or because of transaction costs. The simplest case of a market under uncertainty is one in which assets are purchased at time 0 and sold at time 1, and the uncertainty about asset prices at time 1 is modeled by a finite set of states. Any financial asset in such a market can be constructed from a portfolio of "Arrow securities," where an Arrow security is an asset whose payoff is \$1 in a given state and zero otherwise. The bid and ask prices for a state-i Arrow security can viewed as lower and upper probabilities assigned to state i by the representative agent. Bid and ask prices for more complex assets (which may yield arbitrary payoffs in different states) establish other linear inequality constraints on the probability distribution of the representative agent, so that in general the imprecise beliefs of the representative agent are described by a convex polytope of distributions that is the intersection of all the constraints. This set is non-empty if and only if there are no arbitrage opportunities in the market, a result that is known as the "fundamental theorem of asset pricing" but which was introduced much earlier by de Finetti as the "fundamental theorem of subjective probability." The problem we consider is that of an investor whose (precise) subjective probability distribution is $\mathbf{p}$ and who invests optimally in a market where the imprecise probabilities of the representative agent are described by a convex set Q that is disjoint from $\mathbf{p}$.

2 Generalized measures of entropy and divergence

The *entropy* of a probability distribution, as defined by Shannon (1948), is a measure of the amount of information conveyed by the observation of an event drawn from that distribution. Shannon proved that under the most efficient encoding scheme the average number of bits (binary digits) needed to report the occurrence of an event whose relative frequency is p is proportional to $\ln(1/p) = -\ln(p)$, so the expected number of bits per event to encode events drawn from a distribution $\mathbf{p}$ is proportional to $H(\mathbf{p}) \equiv -\sum_i p_i \ln(p_i)$.[2] This quantity is known as the entropy of the distribution $\mathbf{p}$, because up to a multiplicative constant (namely Boltzmann's constant) it coincides exactly with the definition of the Gibbs entropy of a physical system whose distribution of internal states is $\mathbf{p}$, which in turn is the microscopic interpretation of the macroscopic concept of entropy from classical thermodynamics. If an engineer who had optimized the encoding scheme on the assumption that the distribution was $\mathbf{q}$ subsequently learns or decides (via Bayesian updating or some other method of discovery) that it is actually some other distribution $\mathbf{p}$, then the encoding scheme based on $\mathbf{q}$ is revealed to be suboptimal, and $H(\mathbf{q})$ underestimates the average number of bits per event that are actually being transmitted. A practical measure of the amount of information gained in updating $\mathbf{q}$ to $\mathbf{p}$ is the reduction in the expected number of bits needed to encode an event by re-optimizing for the distribution now believed to be correct, which is known as the Kullback-Leibler (KL) divergence of $\mathbf{p}$ with respect to $\mathbf{q}$:

$$D_{KL}(\mathbf{p}\|\mathbf{q}) \equiv \sum_i p_i(\ln(1/q_i) - \ln(1/p_i)) = \mathbf{E_p}[\ln(\mathbf{p}/\mathbf{q})]. \quad (1)$$

The KL divergence has several very convenient and appealing properties that are often cited as reasons for adopting it as a universal measure of information gain. First, it is naturally *additive* with respect to independent experiments. Suppose that A and B are statistically independent partitions of the state space whose prior marginal probability distributions are $\mathbf{q}_A$ and $\mathbf{q}_B$, so that their prior joint distribution is $\mathbf{q}_A \times \mathbf{q}_B$. Now suppose that independent experiments are performed, which result in the updating of $\mathbf{q}_A$ and $\mathbf{q}_B$ to $\mathbf{p}_A$ and $\mathbf{p}_B$, respectively, so that the posterior joint distribution is $\mathbf{p}_A \times \mathbf{p}_B$. Then the total information gain of the two experiments is the sum of their separate KL divergences:

$$D_{KL}(\mathbf{p}_A \times \mathbf{p}_B\|\mathbf{q}_A \times \mathbf{q}_B) = D_{KL}(\mathbf{p}_A\|\mathbf{q}_A) + D_{KL}(\mathbf{p}_B\|\mathbf{q}_B). \quad (2)$$

Second, and even stronger, the KL divergence has the property of *recursivity* with respect to the splitting of events. Suppose that information is transmitted

[2] Throughout the paper, upper-case functions such as $H(\mathbf{p})$, $D_{KL}(\mathbf{p}\|\mathbf{q})$, $S(\mathbf{r}, \mathbf{p})$, etc., are scalar-valued functions of vector arguments, whereas lower-case functions such as $f(\mathbf{x})$, $\ln(\mathbf{x})$, $u(\mathbf{x})$, etc., are vector-valued functions in which a univariate function is applied elementwise to a vector argument, i.e., $f(\mathbf{x}) = (f(x_1), ..., f(x_n))$.

in a 2-step process, in which two out of n possible states - say, states 1 and 2 - are not distinguished on the first step. If the realized state is neither 1 or 2, the process stops there, but otherwise a second signal is sent to report which of those two has occurred. The probabilities of states 1 and 2 are aggregated in the first step, so the information gain on that step is $D_{KL}(p_1+p_2, p_3, ..., p_n \| q_1+q_2, q_3, ..., q_n)$. On the second step, which occurs with probability (p_1+p_2), the additional gain is $D_{KL}\left(\frac{p_1}{p_1+p_2}, \frac{p_2}{p_1+p_2} \| \frac{q_1}{q_1+q_2}, \frac{q_2}{q_1+q_2}\right)$. The recursivity property of the KL divergence requires the expected total information gain of the two-step process to be the same as that of a one-step process:

$$D_{KL}(\mathbf{p}\|\mathbf{q}) = D_{KL}(p_1+p_2, p_3, ..., p_n \| q_1+q_2, q_3, ..., q_n)$$
$$+(p_1+p_2)D_{KL}\left(\frac{p_1}{p_1+p_2}, \frac{p_2}{p_1+p_2} \| \frac{q_1}{q_1+q_2}, \frac{q_2}{q_1+q_2}\right). \quad (3)$$

The KL divergence is the only distance measure that satisfies both additivity and recursivity, hence it is the divergence that is naturally obtained if those properties are embraced as axioms that an information measure should satisfy. However, in situations other than signal transmission, where the objective may be something other than economizing on bandwidth, these axioms may be unduly restrictive. In applications involving imprecise probabilities, it may be of interest to find the member of a convex set of distributions that is nearest to or farthest from some reference distribution, and desiderata of a distance measure may depend on the inference or decision problem to be solved.

One measure of distance between probability distributions that generalizes the Kullback-Leibler divergence is known as a *Brègman divergence* (Brègman 1967). Any strictly convex function F defines a Brègman divergence $B_F(\mathbf{p}\|\mathbf{r})$ as follows:

$$B_F(\mathbf{p}\|\mathbf{r}) \equiv F(\mathbf{p}) - F(\mathbf{r}) - \nabla F(\mathbf{r}) \cdot (\mathbf{p}-\mathbf{r}). \quad (4)$$

The decision-theoretic significance of a Brègman divergence is that it uniquely determines a *strictly proper scoring rule*, which is a reward function for truthfully eliciting subjective probabilities. As noted by McCarthy (1956) and further elaborated by Hendrickson and Buehler (1971) and Savage (1971), any strictly convex function F can be used to generate a strictly proper scoring rule S as follows:

$$S(\mathbf{r},\mathbf{p}) \equiv F(\mathbf{r}) + \nabla F(\mathbf{r}) \cdot (\mathbf{p}-\mathbf{r}), \quad (5)$$

where $\nabla F(\mathbf{r})$ denotes the gradient (or more generally a subgradient) of F evaluated at $\mathbf{r}$, and conversely F can be recovered from S according to $F(\mathbf{p}) = S(\mathbf{p},\mathbf{p})$. The function $S(\mathbf{r},\mathbf{p})$ is used to "score" a probability forecast in the following way. A forecaster who reports $\mathbf{r}$ to be her probability distribution over the states is given a reward equal to $S(\mathbf{r},\mathbf{e}_i)$ if state i occurs, where $\mathbf{e}_i$ denotes the probability distribution that assigns probability 1 to state i and zero to all other states, i.e., the indicator vector for state i. Because S is linear in $\mathbf{p}$, we have $S(\mathbf{r},\mathbf{p}) = \sum_i p_i S(\mathbf{r},\mathbf{e}_i)$, so the function $S(\mathbf{r},\mathbf{p})$ represents the forecaster's *expected* score if her distribution is $\mathbf{p}$ and she reports distribution $\mathbf{r}$. If $F(\mathbf{p})$ is strictly convex, it follows from the subgradient inequality that $S(\mathbf{r},\mathbf{p})$ is uniquely maximized when $\mathbf{r}=\mathbf{p}$, i.e., when the forecaster honestly reports her probability distribution, which is the defining property of a strictly proper scoring rule.

By construction, the function $F(\mathbf{p}) - S(\mathbf{r},\mathbf{p})$, which represents the forecaster's expected *loss* for reporting $\mathbf{r}$ when her distribution is $\mathbf{p}$, is the Brègman divergence $B_F(\mathbf{p}\|\mathbf{r})$. A Brègman divergence is therefore a decision-theoretic measure of the "information deficit" that is faced by a decision maker who acts on the basis of the distribution $\mathbf{r}$ when the distribution is $\mathbf{p}$. In this capacity, Brègman divergences (and their corresponding strictly proper scoring rules) provide a potentially rich class of loss functions that can be used for robust Bayesian inference, as discussed by Grünwald and Dawid (2004), Dawid (2006), and Gneiting and Raftery (2007). A problem of this kind can be framed as a game against nature in which nature chooses a distribution $\mathbf{p}$ from some convex set $\mathcal{P}$, such as the set of distributions satisfying a mean value constraint. The robust Bayes problem for the decision maker is to determine the distribution $\mathbf{r}$ that minimizes her maximum expected loss over all $\mathbf{p} \in \mathcal{P}$, where the expected loss (in our terms) is the negative expected score $-S(\mathbf{r},\mathbf{p})$. Grünwald and Dawid show that the optimal-expected-loss function, $-F(\mathbf{p})$, is interpretable as a generalized entropy, and minimizing the maximum expected loss is equivalent to maximizing this entropy on the set $\mathcal{P}$. The distribution $\mathbf{r}$ that solves this problem is the one that minimizes $B_F(\mathbf{p}\|\mathbf{r})$ with respect to an uninformative reference distribution $\mathbf{p}_0$ at which the entropy $-F(\mathbf{p})$ is maximized.

3 The pseudospherical and power divergences

In this paper, we will consider a different kind of game and a correspondingly different decision-theoretic measure of information, namely, we will suppose that a risk-averse decision maker with personal probability distribution $\mathbf{p}$ has the opportunity to bet against a non-strategic less-well-informed opponent whose distribution $\mathbf{q}$ is known to lie in some set Q that is disjoint from $\mathbf{p}$, which enables the decision maker to place bets that are profitable in the sense of increasing her expected utility relative to the status quo.

The "information surplus" enjoyed by this decision maker will be shown to be measured by the minimum of a generalized divergence between $\mathbf{p}$ and all $\mathbf{q} \in Q$, but it is generally not a Brègman divergence. The solution of this problem gives rise to families of "weighted" strictly proper scoring rules, in which $\mathbf{q}$ plays the role of a baseline distribution with respect to which the value of the forecaster's information is measured, and they generalize the well-known quadratic, logarithmic, and pseudospherical scoring rules – details are given in Jose et al. (2008).

There are various functional forms that could be used to define a divergence of $\mathbf{p}$ with respect to $\mathbf{q}$, and the one we that we find most compelling, for both practical and theoretical reasons, is that for a given p_i the divergence should depend on q_i only through the ratio p_i/q_i, which is the marginal value of a bet on state i: a \$1 bet yields a payoff of $\$1/q_i$ when that state occurs and zero otherwise, because this is a fair payoff from the perspective of the opponent, and its expected value for the decision maker is $\$p_i/q_i$. More generally, whenever low-probability states are explicitly distinguished in the setup of a decision model, it is usually because they have large consequences, in which case relative rather than absolute errors in probability estimation are what matter. Another rationale is illustrated by the following example: suppose that the state space consists of 4 states formed by the Cartesian product of two binary events E and F, and suppose it happens that the decision maker and her opponent both agree on the probability of F and they also agree that E and F are statistically independent. Then it seems reasonable that the marginal value of a bet on any state should depend only on the extent of disagreement about the probability of E, and this requires it to depend only on the ratio of the two agents' probabilities for that state, which divides out the common probability of F.

The measurement of distance between two probability distributions in terms of ratios has a long history in statistics and information theory, and it is the basis of another kind of generalized divergence known as an *f-divergence* (Csiszár 1967). If f is a strictly convex function, the corresponding f-divergence is defined as

$$D_f(\mathbf{p}\|\mathbf{q}) \equiv E_{\mathbf{p}}[f(\mathbf{p}/\mathbf{q})]. \tag{6}$$

Divergences of this general form have been widely used in statistics for many years as (seemingly) utility-free measures of the value of the information - e.g., Goel (1983) uses f-divergence to define a "conditional amount of sample information" for measuring prior-to-posterior information gains in Bayesian hierarchical models. More recently it has been recognized that f-divergences are interpretable as measures of expected utility gains that are available to decision makers who have opportunities to bet against less-well-informed opponents or to invest in financial markets, as will be more fully discussed in later sections of this paper.

As noted above, the KL divergence is the only distance measure that satisfies the axioms of both additivity and recursivity. However, it has been discovered that weakenings of these axioms lead to several interesting parametric families of f-divergences (or transformations thereof) which have their own merits and their own applications. Havrda and Chavrát (1967) defined a quantity that they called the *directed divergence of order β between* $\mathbf{p}$ *and* $\mathbf{q}$, and variants of this divergence, which are equivalent up to a scale factor, were discussed by Rathie and Kannappan (1972), Cressie and Read (1984), and Haussler and Opper (1997). Cressie and Read referred to this quantity as the *power divergence*, and that term will be adopted here. The power divergence (as originally introduced by Havrda and Chavrát) is defined for all $\beta \in \mathbb{R}$ by:

$$D^{\mathbf{P}}_{\beta}(\mathbf{p}\|\mathbf{q}) \equiv \frac{E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}]-1}{\beta(\beta-1)}, \tag{7}$$

which is an f-divergence based on the normalized power function $f_\beta(x) = (x^{\beta-1}-1)/(\beta(\beta-1))$.[3] The cases of $\beta = -1$, 0, $\frac{1}{2}$, 1, and 2 are of special interest. At $\beta = 1$, the power divergence between $\mathbf{p}$ and $\mathbf{q}$ is equal to the KL divergence $D_{KL}(\mathbf{p}\|\mathbf{q})$, and at $\beta = 0$ it is the reverse KL divergence $D_{KL}(\mathbf{q}\|\mathbf{p})$. In fact, $D^{\mathbf{P}}_{\beta}(\mathbf{p}\|\mathbf{q})$ is antisymmetric around $\beta = \frac{1}{2}$ in the sense that $D^{\mathbf{P}}_{\beta}(\mathbf{p}\|\mathbf{q}) = D^{\mathbf{P}}_{1-\beta}(\mathbf{q}\|\mathbf{p})$, i.e., the reverse divergence is obtained by replacing β with $1-\beta$ for any value of β. The case $\beta = \frac{1}{2}$ has perfect symmetry, i.e., $D^{\mathbf{P}}_{1/2}(\mathbf{p}\|\mathbf{q}) = D^{\mathbf{P}}_{1/2}(\mathbf{q}\|\mathbf{p})$, and it reduces to

$$D^{\mathbf{P}}_{1/2}(\mathbf{p}\|\mathbf{q}) = 4\left(1 - \sum_{j=1}^{n} \sqrt{p_j q_j}\right), \tag{8}$$

which is proportional to the *squared Hellinger distance* between $\mathbf{p}$ and $\mathbf{q}$, as noted by Haussler and Opper (1997). The Hellinger distance $D_H(\mathbf{p}\|\mathbf{q})$ is widely used in statistics and is defined by

$$D_H(\mathbf{p}\|\mathbf{q}) \equiv \left(\sum_{j=1}^{n} \left(\sqrt{p_j} - \sqrt{q_j}\right)^2\right)^{1/2}, \tag{9}$$

whence

$$D^{P}_{1/2}(\mathbf{p}\|\mathbf{q}) = 2D_H(\mathbf{p}\|\mathbf{q})^2. \tag{10}$$

[3] $f_\beta(x)$ converges to $\ln(x)$ as $\beta \to 1$, but it goes to $\pm\infty$ as β approaches zero from above or below. Nevertheless, (7) is a continuous function of β at $\beta = 0$ by virtue of the special nature of the argument of f_β and its behavior inside the expectation: the individual terms go to $\pm\infty$, but their expectation converges. Note also that f_β is antisymmetric around $\beta = 1/2$ in the following way $f_\beta(x^\beta) = f_{1-\beta}(x^{1-\beta})$, which parallels a similar property of the divergences and utility functions discussed here.

At $\beta = 2$ the power divergence reduces to (a multiple of) another well-known divergence, the *Chi-square divergence* (Pearson 1900):

$$D_2^{\mathbf{P}}(\mathbf{p}\|\mathbf{q}) = \frac{1}{2}\left(E_{\mathbf{p}}[\mathbf{p}/\mathbf{q}] - 1\right) = \frac{1}{2}\chi^2(\mathbf{p}\|\mathbf{q}), \qquad (11)$$

while at $\beta = -1$ it is the reverse Chi-square divergence $\frac{1}{2}\chi^2(\mathbf{q}\|\mathbf{p})$.

The power divergence is generally neither additive nor recursive, but it satisfies two slightly weaker properties for all values of β. First, it satisfies the following *pseudoadditivity* property with respect to independent partitions A and B:

$$\begin{aligned} D_\beta^{\mathbf{P}}(\mathbf{p}_A \times \mathbf{p}_B\|\mathbf{q}_A \times \mathbf{q}_B) &= D_\beta^{\mathbf{P}}(\mathbf{p}_A\|\mathbf{q}_A) + D_\beta^{\mathbf{P}}(\mathbf{p}_B\|\mathbf{q}_B) \\ &+\beta(\beta-1)D_\beta^{\mathbf{P}}(\mathbf{p}_A\|\mathbf{q}_A)D_\beta^{\mathbf{P}}(\mathbf{p}_B\|\mathbf{q}_B). \qquad (12)\end{aligned}$$

Second, it satisfies the following *pseudorecursivity* property with respect to the splitting of events (Rathie and Kannappan 1972, Cressie and Read 1984):

$$\begin{aligned} D_\beta^{\mathbf{P}}(\mathbf{p}\|\mathbf{q}) &= D_\beta^{\mathbf{P}}(p_1+p_2, p_3, ..., p_n\|q_1+q_2, q_3, ..., q_n) \\ &+ (p_1+p_2)\left(\frac{p_1+p_2}{q_1+q_2}\right)^{\beta-1} \qquad (13) \\ &\times D_\beta^{\mathbf{P}}\left(\frac{p_1}{p_1+p_2}, \frac{p_2}{p_1+p_2}\|\frac{q_1}{q_1+q_2}, \frac{q_2}{q_1+q_2}\right). \end{aligned}$$

Pseudoadditivity reduces to additivity in both of the special cases $\beta = 0$ and $\beta = 1$ (both the KL divergence and the reverse KL divergence are additive), while pseudorecursivity reduces to recursivity only in the special case $\beta = 1$. Also note that for $\beta \in (0, 1)$ the power divergence is *subadditive*, i.e., $D_\beta^{\mathbf{P}}(\mathbf{p}_A \times \mathbf{p}_B\|\mathbf{q}_A \times \mathbf{q}_B) \le D_\beta^{\mathbf{P}}(\mathbf{p}_A\|\mathbf{q}_A) + D_\beta^{\mathbf{P}}(\mathbf{p}_B\|\mathbf{q}_B)$, while for $\beta < 0$ or $\beta > 1$ it is *superadditive*, i.e., $D_\beta^{\mathbf{P}}(\mathbf{p}_A \times \mathbf{p}_B\|\mathbf{q}_A \times \mathbf{q}_B) \ge D_\beta^{\mathbf{P}}(\mathbf{p}_A\|\mathbf{q}_A) + D_\beta^{\mathbf{P}}(\mathbf{p}_B\|\mathbf{q}_B)$.

A different form of generalized entropy was introduced by Arimoto (1971) and further elaborated by Sharma and Mittal (1975), Boekee and Van der Lubbe (1980) and Lavenda and Dunning-Davies (2003). Arimoto's generalized entropy of order β is defined for $\beta > 0$ as follows:

$$\frac{\beta}{\beta-1}\left(E_{\mathbf{p}}[\mathbf{p}^{\beta-1}]^{1/\beta} - 1\right). \qquad (14)$$

(Here β corresponds to the term $1/\beta$ in Arimoto's original presentation and to the term R in Boekee and Van der Lubbe's presentation.) The factor of β in the numerator plays no essential role when β is restricted to be positive, and without it the measure is actually valid for all real β and closely related to the pseudospherical scoring rule (Jose et al. 2008). The corresponding relative entropy measure, which we will henceforth call the *pseudospherical divergence of order* β between $\mathbf{p}$ and $\mathbf{q}$, is obtained by introducing a reference distribution $\mathbf{q}$ and dividing out the unnecessary factor of β,

	$D_\beta^{\mathbf{P}}(\mathbf{p}\|\mathbf{q})$	$D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{q})$
$\beta = -1$	$\frac{1}{2}\chi^2(\mathbf{q}\|\mathbf{p})$	$\frac{1}{2}\left(1 - (\chi^2(\mathbf{q}\|\mathbf{p}) + 1)^{-1}\right)$
$\beta = 0$	$D_{KL}(\mathbf{q}\|\mathbf{p})$	$1 - \exp(-D_{KL}(\mathbf{q}\|\mathbf{p}))$
$\beta = \frac{1}{2}$	$2D_H(\mathbf{p}\|\mathbf{q})^2$ $= 2D_H(\mathbf{q}\|\mathbf{p})^2$	$2\left(1 - \left(1 - \frac{1}{2}D_H(\mathbf{p}\|\mathbf{q})^2\right)^2\right)$
$\beta = 1$	$D_{KL}(\mathbf{p}\|\mathbf{q})$	$D_{KL}(\mathbf{p}\|\mathbf{q})$
$\beta = 2$	$\frac{1}{2}\chi^2(\mathbf{p}\|\mathbf{q})$	$\sqrt{\chi^2(\mathbf{p}\|\mathbf{q}) + 1} - 1$

Table 1: Special cases of power and pseudospherical divergences

$$D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{q}) \equiv \frac{\left(E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}]\right)^{1/\beta} - 1}{\beta - 1}. \qquad (15)$$

This is a nonlinear transformation of the power divergence, hence it can also be expressed as a function of other well-known divergences for special cases of β, as summarized in Table 1, which highlights the antisymmetry of the power divergence around $\beta = \frac{1}{2}$.

Like the power divergence, the pseudospherical divergence satisfies a pseudoadditivity property:

$$\begin{aligned} D_\beta^{\mathbf{S}}(\mathbf{p}_A \times \mathbf{p}_B\|\mathbf{q}_A \times \mathbf{q}_B) &= D_\beta^{\mathbf{S}}(\mathbf{p}_A\|\mathbf{q}_A) + D_\beta^{\mathbf{S}}(\mathbf{p}_B\|\ \mathbf{q}_B) \\ &+(\beta-1)D_\beta^{\mathbf{S}}(\mathbf{p}_A\|\mathbf{q}_A)D_\beta^{\mathbf{S}}(\mathbf{p}_B\|\mathbf{q}_B). \qquad (16)\end{aligned}$$

The coefficient of the cross-term in this case is $\beta - 1$, not $\beta(\beta - 1)$, and hence $D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{q})$ is subadditive for $\beta < 1$ and superadditive for $\beta > 1$. However, the pseudospherical divergence is generally not pseudorecursive, and it is not an f-divergence, although it is monotonically related to one.

4 The family of normalized linear-risk-tolerance utility functions

In the optimization problems to be discussed in the following section of the paper, the decision maker's utility function will be assumed to be drawn from the most commonly used parametric family of utility functions, namely the generalized power family that includes the exponential and logarithmic utility functions as limiting cases. The utility functions from this family will be parameterized here as:

$$\begin{aligned} u_\beta(x) &\equiv \frac{1}{\beta-1}((1+\beta x)^{(\beta-1)/\beta} - 1) \text{ if } \beta x > -1 \\ u_\beta(x) &\equiv -\infty \text{ otherwise,} \end{aligned}$$

for all $\beta \in \mathbb{R}$. This parameterization, which was introduced by Jose et al. (2008), has two key properties. First, $u_\beta(0) = 0$ and $u'_\beta(0) = 1$, so that for every β the graph of u_β passes through the origin and has a slope of unity there. Second, the corresponding *risk tolerance function* $\tau_\beta(x)$, which is the reciprocal of the Pratt-Arrow risk aversion measure, is a linear function of wealth with slope equal to β and intercept

$\beta=-1$	quadratic utility	$u_{-1}(x)=-\frac{1}{2}((1-x)^2-1)$
$\beta=0$	exponential utility	$u_0(x)=1-\exp(-x)$
$\beta=\frac{1}{2}$	reciprocal utility	$u_{1/2}(x)=2\left(1-\frac{1}{1+x/2}\right)$
$\beta=1$	logarithmic utility	$u_1(x)=\ln(1+x)$
$\beta=2$	square-root utility	$u_2(x)=\sqrt{1+2x}-1$

Table 2: Examples of normalized linear-risk-tolerance utility functions

equal to 1: $\tau_\beta(x) \equiv -u'_\beta(x)/u''_\beta(x) = 1+\beta x$.[4] Thus, risk tolerance as well as marginal utility is normalized to a value of 1 at $x=0$. This amounts to choosing the unit of money to be the status quo risk tolerance (which is without loss of generality when there is a single risk-averse agent) and then choosing the unit of utility to be the status quo marginal utility of money (which is also without loss of generality and which yields money-utile parity at the status quo). Henceforth we will refer to u_β as a *normalized linear-risk-tolerance* (normalized LRT) utility function. The advantages of this normalization are that (a) it is a natural one for modeling utility gains and losses relative to the status quo rather than relative to some hypothetical zero-point of wealth at which utility goes to minus-infinity, and (b) for fixed $x, u_\beta(x)$ is a continuous function of β on the entire real line, so that it sweeps out the widest possible spectrum of local risk attitudes. (Utility functions with the property of linear risk tolerance but without this useful normalization are known as hyperbolic-absolute-risk-aversion (HARA) utility functions in the literature of financial economics, and they typically use different parameterizations for different ranges of the power parameter.) Some important special cases of $u_\beta(x)$ are given in Table 2.

The utility functions $\{u_\beta\}$ exhibit their own form of anti-symmetry around $\beta=\frac{1}{2}$, namely that $u_{1-\beta}(x) = -u_\beta(-x)$, or equivalently $u_\beta(-u_{1-\beta}(-x)) = x$. In other words, the graph of $u_{1-\beta}(x)$ is obtained from the graph of $u_\beta(x)$ by reflecting it around the line $y=-x$. The power (exponent) in u_β is the term $(\beta-1)/\beta$, which has the property that $((\beta-1)/\beta)^{-1} = ((1-\beta)-1)/(1-\beta)$, so that swapping β for $1-\beta$ results in another power utility function whose power is the reciprocal of the original. Thus, under this parameterization, the reciprocal utility function ($\beta=\frac{1}{2}$) is its own reflection around the line $y=-x$, the exponential and log utility functions ($\beta=0$ and $\beta=1$) are reflections of each other, as illustrated in Figure 1, and the power utility function with exponent δ is the reflection of the power utility function with exponent $1/\delta$ for any positive or negative δ other than 0 or 1.

[4] The decision maker's risk tolerance is the parameter that determines the mean-variance tradeoffs she is willing to make on the margin. To a second-order approximation, the amount that she is willing to pay for a risky asset whose payoff distribution has mean μ and variance σ^2 is equal to $\mu-\sigma^2/2\tau$, where τ is her risk tolerance. In other words, her *risk premium* for such an asset, which is the amount by which she devalues it relative to its expected value, is $\sigma^2/2\tau$. In general a decision maker's risk tolerance may be expected to change as her wealth changes, and with this utility function her risk tolerance is a linear function of wealth with slope coefficient β.

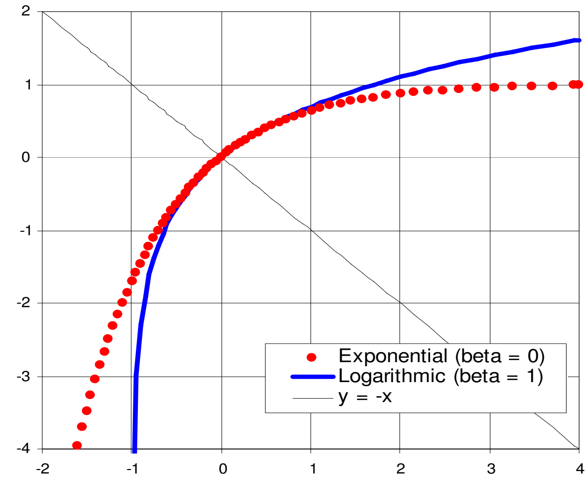

Figure 1: Reflection property of normalized LRT utility functions around $y=-x$

5 Duality between maximization of expected utility and minimization of relative entropy in incomplete markets

We now consider two generic optimization problems in which a risk averse decision maker with probability distribution $\mathbf{p}$ invests in an incomplete financial market where bid-ask spreads in asset prices are determined by a convex set Q of imprecise probabilities representing the beliefs of a risk-neutral representative agent, as noted in the introduction. The problem of expected-utility maximization in incomplete markets has been widely studied in the mathematical finance literature in recent years, and it has been shown that there is a duality relationship between maximization of expected utility and minimization of an appropriate divergence (e.g., Frittelli 2000, Rouge and El Karoui 2000, Goll and Rüschendorf 2001, Delbaen et al. 2002, Slomczyński and Zastawniak 2004, Ilhan et al. 2004, Samperi 2005). Most of this literature has focused on the case of exponential utility, for which the dual problem is the minimization of the reverse KL divergence $D_{KL}(\mathbf{q}\|\mathbf{p})$, as well as on issues that arise in multi-period or continuous-time markets. In this section we will show that in a single-period or two-period market, there is a duality relation between the pseudospherical or power divergence and the solution of an expected-utility-maximization problem in which the utility function is drawn from the normalized linear-risk-tolerance family.

Let $\mathbf{x} \in \mathbb{R}^n$ denote the vector of monetary payoffs to the decision maker, and let $u_\beta(\mathbf{x}) \equiv (u_\beta(x_1), ..., u_\beta(x_n))$ denote the vector of utilities that the function u_β yields when applied to $\mathbf{x}$. An incomplete, single-period market can either be parameterized in terms of an $m \times n$ matrix $\mathbf{A}$ whose rows are the (net) payoff vectors of available assets, i.e., $\mathbf{A} = \{a_{ij}\}$ where a_{ij} is the net payoff to the decision maker of one unit of the i^{th}asset in state j, or else in terms of a $k \times n$ matrix $\mathbf{Q}$ whose rows are risk neutral probability distributions that support the asset prices, i.e., $\mathbf{Q} = \{q_{ij}\}$ where q_{ij} is the probability of state j under the i^{th} risk neutral distribution. The rows of $\mathbf{Q}$ are the extremal risk-neutral probability distributions assigning non-positive expectation to all the rows of $\mathbf{A}$, i.e., the rows of $-\mathbf{Q}$ are the dual cone of the rows of $\mathbf{A}$. The parameterization in terms of $\mathbf{Q}$ will be adopted here. Let $\mathbf{x}$ denote an arbitrary n-vector of monetary payoffs to the decision maker (an element of $\mathbb{R}^n$), and let $\mathbf{z}$ denote an arbitrary k-vector of non-negative weights summing to one (an element of $\triangle^k$, the unit simplex in $\mathbb{R}^k$). As before, let $\mathbf{p}$ denote the decision maker's subjective probability distribution, and henceforth let $\mathbf{q}$ denote one of many possible probability distributions attributable to a risk-neutral trading opponent: the representative agent.

In the first generic decision problem (**"S"**), there is a single time period in which consumption occurs, the decision maker has a single-attribute LRT utility function $u_\beta(x)$, and her objective is to find the payoff vector $\mathbf{x}$ that maximizes her subjective expected utility subject to the self-financing constraint $E_\mathbf{q}[\mathbf{x}] \leq 0$. The decision maker's optimal expected utility, denoted $U^\mathbf{S}(\mathbf{p}\|\mathbf{q})$, is determined by solving:

Primal Problem S:

$$U^\mathbf{S}_\beta(\mathbf{p}\|\mathbf{Q}) \equiv \max_{\mathbf{x}\in\mathbb{R}^n} E_\mathbf{p}[u_\beta(\mathbf{x})] \text{ subject to } \mathbf{Qx} \leq \mathbf{0}$$

Note that $-\mathbf{Qx}$ is the k-vector of the opponent's expected values for payoff vector $\mathbf{x}$ under all the extremal risk neutral distributions, hence the condition $\mathbf{Qx} \leq \mathbf{0}$ means that $\mathbf{x}$ yields non-negative expected value to the opponent under all those distributions.

In the second problem (**"P"**), there are two periods in which consumption occurs and the decision maker with probability distribution $\mathbf{p}$ has a quasilinear utility function $u_\beta(a, b) = a + u_\beta(b)$ where a is money consumed at time 0 and b is money consumed at time 1. Under the normalized LRT family of utility functions, the marginal rate of substitution between time-0 consumption and time-1 consumption is equal to unity at $x = 0$ in this problem, as though in the status quo the decision maker is indifferent between consuming the next dollar at time 0 or time 1. The decision maker's objective is to choose a vector $\mathbf{x}$ of time-1 payoffs to be purchased from time-0 funds at market prices so as to maximize the total expected utility of consumption in both periods. The time-0 cost of purchasing $\mathbf{x}$ is $E_\mathbf{q}[\mathbf{x}]$, so the optimal expected utility, denoted $U^\mathbf{P}(\mathbf{p}\|\mathbf{q})$, is the solution of:

Primal Problem P:

$$U^\mathbf{P}_\beta(\mathbf{p}\|\mathbf{Q}) \equiv \max_{y\in\mathbb{R},\mathbf{x}\in\mathbb{R}^n} E_\mathbf{p}[u_\beta(\mathbf{x})] - y \text{ subject to } \mathbf{Qx} \leq y\mathbf{1}$$

Henceforth, let $\mathbf{x}^\mathbf{S}_\beta(\mathbf{p}\|\mathbf{q})$ and $\mathbf{x}^\mathbf{P}_\beta(\mathbf{p}\|\mathbf{q})$ denote the solutions of Problems $\mathbf{S}$ and $\mathbf{P}$, with i^{th} elements $x^\mathbf{S}_{\beta,i}(\mathbf{p}\|\mathbf{q})$ and $x^\mathbf{P}_{\beta,i}(\mathbf{p}\|\mathbf{q})$, respectively. Let $\mathbf{z} \in \triangle^k$ denote a vector of weights, so that $\mathbf{z}^T\mathbf{Q}$ is a mixture of the rows of $\mathbf{Q}$, which is an element of the convex polytope Q of risk neutral distributions. Our main result is that the utility gains to the decision maker under problems $\mathbf{S}$ and $\mathbf{P}$ are, respectively, the minima of the pseudospherical and power divergences between $\mathbf{p}$ and all $\mathbf{q} \in Q$ for the same β.

THEOREM (Jose et al. 2008):

(a) In an incomplete, single-period market, maximization of expected linear-risk-tolerance utility with risk tolerance coefficient β (Primal Problem $\mathbf{S}$) is dual to minimization of the pseudospherical divergence of order β between the decision maker's subjective distribution $\mathbf{p}$ and a risk neutral distribution $\mathbf{q}$ consistent with asset prices. That is, the corresponding dual problem is:

$$\textbf{Dual Problem S: } D^\mathbf{S}_\beta(\mathbf{p}\|\mathbf{Q}) \equiv \min_{\mathbf{z}\in\triangle^k} D^\mathbf{S}_\beta(\mathbf{p}\|\mathbf{z}^T\mathbf{Q}).$$

Their optimal objective values are the same and the optimal values of the decision variables in one problem are equal to the normalized optimal values of the Lagrange multipliers in the other.

(b) In an incomplete, two-period market, maximization of expected quasilinear linear-risk-tolerance utility with second-period risk tolerance coefficient β (Primal Problem $\mathbf{P}$) is equivalent to minimization of the power divergence of order β between the decision maker's subjective distribution $\mathbf{p}$ and a risk neutral distribution $\mathbf{q}$ consistent with asset prices (Dual Problem $\mathbf{P}$). Their optimal objective values are the same and the optimal values of the decision variables in one problem are equal to the normalized optimal values of the Lagrange multipliers in the other. That is, the corresponding dual problem is:

$$\textbf{Dual Problem P:} D_\beta^{\mathbf{P}}(\mathbf{p}\|\mathbf{Q}) \equiv \min_{\mathbf{z}\in\triangle^k} D_\beta^{\mathbf{P}}(\mathbf{p}\|\mathbf{z}^T\mathbf{Q}).$$

Proof: For part (a), Lagrangian relaxation is applicable because the primal problem has a strictly concave, continuously differentiable objective function and linear constraints. Let $\boldsymbol{\lambda}$ denote the vector of Lagrange multipliers associated with the constraints $\mathbf{Qx} \leq \mathbf{0}$. The Lagrangian relaxation of Primal Problem $\mathbf{S}$ is then $\min_{\boldsymbol{\lambda}\in\mathbb{R}^{k+}} L(\boldsymbol{\lambda})$ where

$$L(\boldsymbol{\lambda}) = \max_{\mathbf{x}\in\mathbb{R}^n} E_{\mathbf{p}}[u_\beta(\mathbf{x})] - \boldsymbol{\lambda}^T\mathbf{Qx}. \tag{17}$$

The Lagrangian $L(\boldsymbol{\lambda})$ is an unconstrained maximum of a continuously differentiable concave function, so it can be solved for $\mathbf{x}$ in terms of $\boldsymbol{\lambda}$ by setting $\nabla(E_{\mathbf{p}}[u_\beta(x)] - \boldsymbol{\lambda}^T\mathbf{Q}x) = \mathbf{0}$, which yields

$$\mathbf{x} = \frac{1}{\beta}\left(\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^T\mathbf{Q}}\right)^\beta - 1\right), \tag{18}$$

whence

$$\begin{aligned} L(\boldsymbol{\lambda}) &= E_{\mathbf{p}}\left[\frac{1}{\beta-1}\left[\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^T\mathbf{Q}}\right)^{\beta-1} - 1\right]\right] && (19)\\ &\quad -\boldsymbol{\lambda}^T\mathbf{Q}\left[\frac{1}{\beta}\left[\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^T\mathbf{Q}}\right)^{\beta} - 1\right]\right] \\ &= \frac{1}{\beta-1}\left[E_{\mathbf{p}}\left[\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^T\mathbf{Q}}\right)^{\beta-1}\right] - 1\right] && (20)\\ &\quad -\frac{1}{\beta}\left[E_{\mathbf{p}}\left[\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^T\mathbf{Q}}\right)^{\beta-1}\right] - \mathbf{1}^T(\boldsymbol{\lambda}^T\mathbf{Q})\right]. && (21)\end{aligned}$$

In the optimal solution $\boldsymbol{\lambda}^*$, where the constraints are satisfied, the second term will be zero, which implies

$$\mathbf{1}^T(\boldsymbol{\lambda}^{*T}\mathbf{Q}) = E_{\mathbf{p}}\left[\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^{*T}\mathbf{Q}}\right)^{\beta-1}\right] \tag{22}$$

and consequently

$$L(\boldsymbol{\lambda}^*) = \frac{1}{\beta-1}\left(E_{\mathbf{p}}\left[\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^{*T}\mathbf{Q}}\right)^{\beta-1}\right] - 1\right). \tag{23}$$

Now let $\mathbf{z}^* = \boldsymbol{\lambda}^*/\mathbf{1}^T\boldsymbol{\lambda}^*$ be the probability distribution that is obtained by normalization of the optimal Lagrange multipliers $\boldsymbol{\lambda}^*$. Then it follows from (21) that:

$$\mathbf{z}^{*T}\mathbf{Q} = \frac{\boldsymbol{\lambda}^{*T}\mathbf{Q}}{E_{\mathbf{p}}[(\mathbf{p}/\boldsymbol{\lambda}^{*T}\mathbf{Q})^{\beta-1}]}. \tag{24}$$

The pseudospherical divergence between $\mathbf{p}$ and $\mathbf{z}^{*T}\mathbf{Q}$ can therefore be expressed in terms of $\boldsymbol{\lambda}^*$ as:

$$\begin{aligned} &D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{z}^{*T}\mathbf{Q}) \\ &= \frac{(E_{\mathbf{p}}[(\mathbf{p}/\mathbf{z}^{*T}\mathbf{Q})^{\beta-1}])^{1/\beta} - 1}{\beta-1} \\ &= \frac{(E_{\mathbf{p}}[(E_{\mathbf{p}}[(\mathbf{p}/\boldsymbol{\lambda}^{*T}\mathbf{Q})^{\beta-1}](\mathbf{p}/\boldsymbol{\lambda}^{*T}\mathbf{Q}))^{\beta-1}])^{1/\beta} - 1}{\beta-1} \\ &= \frac{(E_{\mathbf{p}}[(\mathbf{p}/\boldsymbol{\lambda}^{*T}\mathbf{Q})^{\beta-1}])^{1-1/\beta}(E_{\mathbf{p}}[(\mathbf{p}/\boldsymbol{\lambda}^{*T}\mathbf{Q})^{\beta-1}])^{1/\beta} - 1}{\beta-1} \\ &= \frac{1}{\beta-1}\left(E_{\mathbf{p}}\left[\left(\frac{\mathbf{p}}{\boldsymbol{\lambda}^{*T}\mathbf{Q}}\right)^{\beta-1}\right] - 1\right) \\ &= L(\boldsymbol{\lambda}^*), \end{aligned} \tag{25}$$

which is the optimal objective value of the primal problem. Furthermore $\mathbf{z}^* = \boldsymbol{\lambda}^*/\mathbf{1}^T\boldsymbol{\lambda}^*$ must also minimize $D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{z}^T\mathbf{Q})$ over all $\mathbf{z} \in \triangle^k$, because if there were some other $\mathbf{z}^{**} \in \triangle^k$ such that $D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{z}^{**T}\mathbf{Q}) < D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{z}^{*T}\mathbf{Q})$, then it would be possible to find some $\boldsymbol{\lambda}^{**} \in \mathbb{R}^{k+}$ proportional to $\mathbf{z}^{**}$ such that $\mathbf{z}^{**T}\mathbf{Q} = \boldsymbol{\lambda}^{**T}\mathbf{Q}/(\mathbf{E_p}[(\mathbf{p}/(\boldsymbol{\lambda}^{**T}\mathbf{Q}))^{\boldsymbol{\beta}-\mathbf{1}}])$. By construction this $\boldsymbol{\lambda}^{**}$ would satisfy $E_{\mathbf{p}}[(\mathbf{p}/\boldsymbol{\lambda}^{**T}\mathbf{Q})^{\beta-1}] - \mathbf{1}^T(\boldsymbol{\lambda}^{**T}\mathbf{Q}) = 0$, implying $L(\boldsymbol{\lambda}^{**}) = D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{z}^{**T}\mathbf{Q})$, and it would follow that $L(\boldsymbol{\lambda}^{**}) < L(\boldsymbol{\lambda}^*)$, contradicting the assumption that $\boldsymbol{\lambda}^*$ was optimal.

For part (b), the problem of finding the feasible risk neutral distribution that minimizes the power divergence of order β:

$$\min_{\mathbf{z}\in\triangle^k} D_\beta^{\mathbf{P}}(\mathbf{p}\|\mathbf{z}^T\mathbf{Q}), \tag{26}$$

is equivalent to the Lagrangian problem $\min_{\boldsymbol{\lambda}\in\triangle^k} L(\boldsymbol{\lambda})$, where $L(\boldsymbol{\lambda}) = \max_{\mathbf{x}\in\mathbb{R}^n} E_{\mathbf{p}}[u_\beta(x)] - \boldsymbol{\lambda}^T\mathbf{Q}x$ is the same Lagrangian that was used in the proof of part (a) to minimize the pseudospherical divergence, except that here $\boldsymbol{\lambda}$ is constrained to be in the simplex, not just the non-negative orthant ($\boldsymbol{\lambda} \in\triangle^k$ rather than $\boldsymbol{\lambda} \in\mathbb{R}^{k+}$), which requires a Lagrange multiplier for the constraint $\mathbf{1}^T\mathbf{q} = 1$ in addition to the m Lagrange multipliers for the constraints $\mathbf{Aq} \geq \mathbf{0}$. The latter divided by the former are equal to the optimal values of the decision variables in Primal Problem $\mathbf{P}$ multiplied by $-\beta$. The power divergence is minimized by the same risk neutral distribution $\mathbf{q}^* = \mathbf{z}^{*T}\mathbf{Q}$ that minimizes the pseudospherical divergence (for the same $\mathbf{p}$, β and $\mathbf{Q}$), because they are both monotonic functions of $E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}]$. The optimal value of $\boldsymbol{\lambda}$ is a unit vector selecting the largest element of $\mathbf{Qx}$. Let z denote this largest element. Then $\min_{\boldsymbol{\lambda}\in\triangle^k}\max_{\mathbf{x}\in\mathbb{R}^n} E_{\mathbf{p}}[u_\beta(x)] - \boldsymbol{\lambda}^T\mathbf{Q}x$ is equivalent to $\max_{\mathbf{x}\in\mathbb{R}^n} E_{\mathbf{p}}[u_\beta(\mathbf{x})] - z$ subject to $\mathbf{Qx} \leq z\mathbf{1}$.■

The special case $\beta = 1$ corresponds to log utility in the primal problem and KL divergence in the dual problem, while $\beta = 0$ corresponds to exponential utility in the primal problem and reverse KL divergence in the dual problem, and the cases $\beta = 1/2$ and $\beta = 2$ are related to the squared Hellinger distance and the Chi-square divergence as shown in the right-hand column of Table 1. Because the pseudospherical divergence is a monotonic transformation of the power divergence, the distribution $\mathbf{q}$ $(= \mathbf{z}^T\mathbf{Q})$ that solves Dual Problem $\mathbf{S}$ is the same one that solves Dual Problem $\mathbf{P}$, although the objective values and the primal payoff vectors are generally different. The power divergence is always strictly greater than the pseudospherical divergence $(D_\beta^{\mathbf{P}}(\mathbf{p}\|\mathbf{q}) > D_\beta^{\mathbf{S}}(\mathbf{p}\|\mathbf{q}))$ except at $\beta = 1$, as pointed out earlier, but this inequality is further illuminated by a comparison of the corresponding Lagrangian relaxation problems: the minimization of $L(\boldsymbol{\lambda})$ over $\boldsymbol{\lambda} \in \triangle^k$ must yield a result greater than or equal to its minimization over the larger set $\boldsymbol{\lambda} \in \mathbb{R}^{k+}$, whether or not the market is complete.

Versions of the same duality theorem have been discussed in the mathematical finance literature, as noted above, although the full spectrum of LRT utility and its closed-form solution have not previously been characterized. The details of the correspondence between our results and those of Goll and Rüschendorf (2001) are given in Jose et al. (2008).

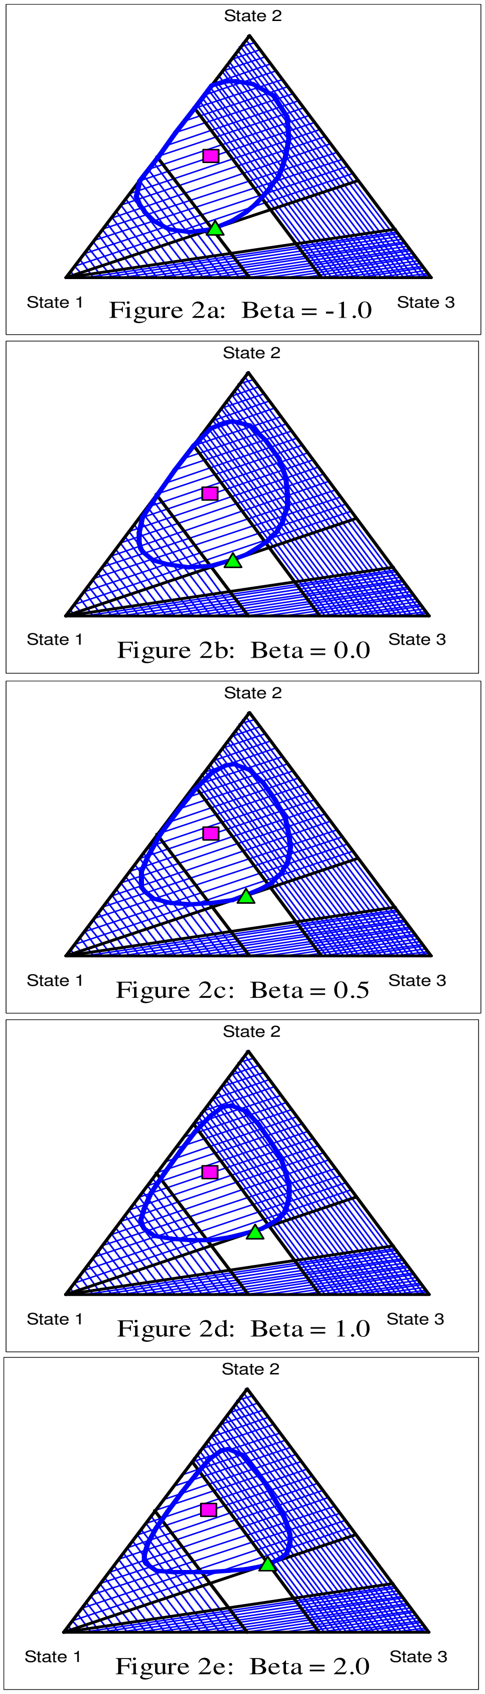

Figures 2a-e: Minimum-divergence solution for the power divergence with $\mathbf{p} = (0.35, 0.5, 0.15)$ for $\beta = -1$, 0, 0.5, 1, and 2.

6 Illustration of the geometry of the divergence-minimization problem

To visualize the preceding results, consider a simple example in which there are three states and (only) lower and upper bounds of 0.3 and 0.5 are given for the probability of state 1 and lower and upper bounds of 0.6 and 0.8 are given for the conditional probability of state 3 given not-state-1. The set Q of probability distributions that satisfies these constraints is the unshaded quadrilateral in the lower center of the simplex in Figures 2a-e. Let the reference distribution be $\mathbf{p} = (0.35, 0.5, 0.15)$, which is the square dot in the upper left. Figures 2a-e show the solution of the dual problem of finding the element of Q that minimizes the pseudospherical or power divergence between itself and $\mathbf{p}$ for $\beta = -1$, 0, 0.5, 1, and 2. The triangular dot is the minimum-divergence solution, and the contour (level curve) that passes through it is also shown. In this case, the solution moves from the left to the right of the upper edge of the quadrilateral as β increases from -1 to 2. Also, the contours become more triangular in shape as β increases, flattening more near the edges of the simplex, because as $\mathbf{q}$ approaches an edge of the simplex, q_i goes to zero for some i, and the term $(p_i/q_i)^{\beta-1}$ in the divergence calculation blows up faster for larger values of β as that edge is approached.

7 Discussion

A financial market under uncertainty provides one of the purest and most economically important examples of a situation in which subjective beliefs – in this case those of a risk neutral representative agent with whom individual investors may trade – are represented by imprecise probabilities that are subject to direct measurement. The measurement process, which consists of setting bid and ask prices for portfolios of Arrow securities, is essentially the same operational method of eliciting subjective probabilities that was introduced by de Finetti, and it naturally leads to a representation of beliefs in the form of a convex polytope of probability distributions. In this paper we have considered the decision problem faced by a risk-averse investor in such a market when her risk preferences are represented by a utility function drawn from the generalized power family, which is the family most commonly used in finance theory and applied decision analysis. Under a natural (but novel) parameterization of the generalized power utility function, the investor's optimal expected utility is equal to the minimum of a generalized divergence between her own distribution and the nearest element of the polytope that characterizes the imprecise beliefs of the representative agent, where the generalized divergence is drawn from a parametric family that generalizes the Kullback-Lielber divergence. We have also pointed out connections with recent developments in the use of generalized divergences in robust Bayesian statistics. These results highlight the interconnections among information theory, Bayesian statistics, decision analysis, and finance theory with respect to the program of modeling imprecise probabilities.

References

[1] ARIMOTO, S. 1971. Information-theoretical considerations on estimation problems. *Infor. Contr.* **19**:181-194.

[2] BOEKEE, D. E., J. C. A. VAN DER LUBBE. 1980. The R-norm information measure. *Infor. Contr.* **45**:136-155.

[3] BRÈGMAN, L. 1967. The relaxation method of finding the common points of convex sets and its application to the solution of problems in convex programming. *USSR Comput. Math. and Math. Phys.*, **7**:200-217.

[4] CRESSIE, N., T. R. C. READ. 1984. Multinomial goodness of fit. *J. Royal Stat. Soc. B* **46**(3):440-464.

[5] CSISZÁR, I. 1967 Information type measures of differences of probability distribution and indirect observations. *Studia Math. Hungarica* **2**:299-318

[6] DAWID, A.P. 2006. The geometry of proper scoring rules. Research Report No. 268, Department of Statistical Science, University College, London. (April 2006)

[7] DELBAEN, F., P. GRANDITS, T. RHEINLÄNDER, D. SAMPERI, M. SCHWEIZER, C. STRICKER. 2002. Exponential hedging and entropy penalties. *Math. Finance.* **12:**99-123.

[8] FRITTELLI, M. 2000. The minimal entropy martingale measure and the valuation problem in incomplete markets. *Math. Finance.* **10:**39-52.

[9] GNEITING, T., A. RAFTERY. 2007. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **102**:359-378

[10] GOEL, P. 1983. Information measures and Bayesian hierarchical models. *J. Am. Stat. Assoc.* **78**:408-410.

[11] GOLL, T., L. RÜSCHENDORF. 2001. Minimax and minimal distance martingale measures and their relationship to portfolio optimization. *Finance Stoch.* **5**:557-581

[12] GRÜNWALD, P.D., A.P. DAWID. 2004. Game theory, maximum entropy, minimum discrepancy, and robust Bayesian decision theory. *Ann. Statist.* **32**: 1367-1433.

[13] HAVRDA, J., F. CHAVRÁT. 1967. Quantification method of classification processes: the concept of structural α-entropy. *Kybernetika* **3**:30-35.

[14] HAUSSLER, D., M. OPPER. 1997. Mutual information, metric entropy, and cumulative relative entropy risk. *Ann. Statist.* **25**: 2451-2492.

[15] HENDRICKSON, A. D., R. J. BUEHLER 1971. Proper scores for probability forecasters. *Ann. Math. Stat.* **42**:1916-1921.

[16] ILHAN, A., M. JONSSON, R. SIRCAR. 2004. Portfolio optimization with derivatives and indifference pricing. Working Paper, Princeton University.

[17] JOSE, V.R.R., R. F. NAU, R. L. WINKLER. 2008. Scoring rules, generalized entropy, and utility maximization. *Oper. Res.* **56**:1146-1147

[18] LAVENDA, B. H., J. DUNNING-DAVIES. 2003. Qualms concerning Tsallis's condition of pseudo-additivity as a definition of non-extensivity. http://arxiv.org/abs/cond-mat/0312132

[19] MCCARTHY, J. 1956. Measures of the value of information. *Proc. Nat. Acad. Sci. USA* **42**:654-655.

[20] NAU, R. F., V. R. R. JOSE, R. L. WINKLER. 2007. Scoring Rules, entropy and imprecise probabilities. *Proceedings of the 5th ISIPTA Conference*, 307-315.

[21] PEARSON, K. 1900. On a criterion that a given system of deviations from the probable in the case of correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *London Edinburgh Dublin Philos. Mag. J. Sci. Series 5* **50**:157-175.

[22] RATHIE, P. N., P. KANNAPPAN. 1972. A directed-divergence function of type β. *Infor. Contr.* **20**:38-45.

[23] ROUGE, R., N. EL KAROUI. 2000. Pricing via utility maximization and entropy. *Math. Finance.* **10:**259-276.

[24] SAMPERI, D. 2005. Model selection using entropy and geometry: complements to the six-author paper. Working paper, Decision Synergy.

[25] SAVAGE, L. J. 1971. Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* **66**:783-801.

[26] SELTEN, R. 1998. Axiomatic characterization of the quadratic scoring rule. *Exper. Econ.* **1**:43-62.

[27] SHANNON, C. E. 1948. A mathematical theory of communication. *Bell Sys. Tech. J.* **27**:379-423.

[28] SHARMA, B. D., D. P. MITTAL. 1975. New non-additive measures of entropy for discrete probability distributions *J. Math. Sci.* **10**:28-40.

[29] SLOMCZYŃSKI, W., T ZASTAWNIAK. 2004. Utility maximizing entropy and the second law of thermodynamics. *Ann. Prob.* **32**:2261-2285.