

Full paper

A Utility-Based Approach to Some Information Measures

Craig Friedman ^{*}, Jinggang Huang [°] and Sven Sandow [†]

Standard & Poor's, 55 Water Street, 46th Floor, New York, NY 10041, USA

E-Mail: ^{*} craig_friedman@sandp.com, [°] james_huang@sandp.com, [†] ssandow@hotmail.com

Received: 23 May 2006 / Accepted: 5 January 2007 / Published: 20 January 2007

Abstract: We review a decision theoretic, i.e., utility-based, motivation for entropy and Kullback-Leibler relative entropy, the natural generalizations that follow, and various properties of these generalized quantities. We then consider these generalized quantities in an easily interpreted special case. We show that the resulting quantities, share many of the properties of entropy and relative entropy, such as the data processing inequality and the second law of thermodynamics. We formulate an important statistical learning problem – probability estimation – in terms of a generalized relative entropy. The solution of this problem reflects general risk preferences via the utility function; moreover, the solution is optimal in a sense of robust absolute performance.

Keywords: Generalized Entropy, Generalized Kullback-Leibler Relative Entropy, Decision Theory, Expected Utility, Horse Race, Tsallis Entropy, Statistical Learning, Probability Estimation, Risk Neutral Pricing Measure

1 Introduction

It is well known that there are a number of ways to motivate the fundamental quantities of information theory, for example:

- Entropy can be defined as essentially the only quantity that is consistent with a plausible set of information axioms (the approach taken by Shannon (1948) for a communication system—see also Csiszár and Körner (1997) for additional axiomatic approaches).
- The definition of entropy is related to the definition of entropy from thermodynamics (see, for example, Brillouin (1962) and Jaynes (1957)).

There is a substantial literature on various generalizations of entropy, for example, the Tsallis entropy (with an extensive bibliography in Group of Statistical Physics (2006)), which was introduced by Cressie and Read (1984) and Cressie and Read (1989) and used for statistical decisions. There is a literature on income inequality which generalizes entropy (see, for example, Cowell (1981)). A generalization of relative entropy, relative Tsallis entropy (alpha-divergence), was introduced by Liese and Vajda (1987). Österreicher and Vajda (May, 1993), Stummer (1999), Stummer (2001), and Stummer (2004) discussed Tsallis entropy and optimal Bayesian decisions. A number of these generalization are closely related to the material in this paper. In this paper, we discuss entropy and Kullback-Leibler relative entropy from a decision theoretic point of view; in particular, we review utility-based motivations and generalizations of these information theoretic quantities, based on fundamental principles from decision theory applied in a particular "market" context. As we shall see, some generalizations share important properties with entropy and relative entropy. We then formulate a statistical learning problem (estimation of probability measures) in terms of these generalized quantities and discuss optimality properties of the solution.

The connection between information theoretic quantities and gambling is well known. Cover and Thomas (1991) (see Theorem 6.1.2), for example, show that in the horse race setting, an expected wealth growth rate maximizing investor, who correctly understands the probabilities, has expected wealth growth rate equal to that of a clairvoyant (someone who wins every bet) minus the entropy of the horse race. Thus, the entropy can be interpreted as the difference between the wealth growth rates attained by a clairvoyant and an expected wealth growth rate maximizing investor who allocates according to the true measure p . It follows from Cover and Thomas (1991), Theorem 6.1.2 and Equation (6.13) that an investor who allocates according to the measure q rather than p , suffers a loss in expected wealth growth rate of Kullback-Leibler relative entropy $D(p||q)$.

Investors who maximize expected wealth growth rates can be viewed as expected logarithmic utility maximizers; thus, the entropy can be interpreted as the difference in expected logarithmic utility attained by a clairvoyant and expected logarithmic utility maximizing investor who allocates according to the true measure p . By similar reasoning, an investor who allocates according to the measure q rather than p , suffers a loss in expected logarithmic utility of $D(p||q)$. Friedman and Sandow (2003b) extend these ideas and define a generalized entropy and a generalized Kullback-Leibler relative entropy for investors with general utility functions.

In Section 2, we review the generalizations and various properties of the generalized quantities introduced in Friedman and Sandow (2003b). In Section 3 we consider these quantities in a new and natural special case. We show that the resulting quantities, which are still more general than entropy and Kullback-Leibler relative entropy, share many of the properties of entropy and relative entropy, for example, the data processing inequality and the second law of thermodynamics. We note that one of the most popular families of utility functions from finance leads to a monotone transformation of the Tsallis entropy. In Section 4 we formulate a decision-theoretic approach to the problem of learning a probability measure from data. This approach, formulated in terms of the quantities described in Section 3, generalizes the minimum relative entropy approach to estimating probabilistic models. This more general formulation specifically incorporates the risk preferences of an investor who would use the probability measure that is to be learned in order to make decisions in a particular market setting. Moreover, we show that the probability measure that is “learned” is robust in the sense that an investor allocating according to the “learned” measure maximizes his expected utility in the most adverse environment consistent with certain data consistency constraints.

2 (U, \mathcal{O}) -Entropy

According to utility theory (see, for example, Theorem 3, p. 31 of Ingersoll (1987)), a rational investor acts to maximize his expected utility based on the model he believes. Friedman and Sandow (2003b) use this fact to construct general information theoretic quantities that depend on the probability measures associated with potential states of a horse race (defined precisely below). In this section, we review the horse race setting, notions from utility theory and definitions and some of the basic properties of the generalized entropy ((U, \mathcal{O}) -entropy) and the generalized relative entropy (relative (U, \mathcal{O}) -entropy) from Friedman and Sandow (2003b).

2.1 Probabilistic Model and Horse Race

We consider the random variable X with values, x , in a finite set \mathcal{X} .¹ Let $q(x)$ denote $\text{prob}\{X = x\}$ under the probability measure q . We adopt the viewpoint of an investor who uses the model to place bets on a horse race, which we define as follows:²

Definition 1 *A horse race is a market characterized by the discrete random variable X with possible states $x \in \mathcal{X}$, where an investor can place a bet that $X = x$, which pays the odds ratio $\mathcal{O}(x) > 0$ for each dollar wagered if $X = x$, and 0, otherwise.*

We assume that our investor allocates $b(x)$ to the event $X = x$, where

$$\sum_x b(x) = 1. \quad (1)$$

¹We have chosen this setting for the sake of simplicity. One can generalize the ideas in this paper to continuous random variables and to conditional probability models; for details, see Friedman and Sandow (2003b).

²See, for example, Cover and Thomas (1991), Chapter 6.

We note that we have not required $b(x)$ to be nonnegative. For some utility functions, the investor may choose to “short” a particular horse. We also note that we can ease the constraint (1) so as to allow our investor to leverage or withhold cash. This will not lead to any significant change in the final results of this paper (see Friedman and Sandow (2004)).

The investor’s wealth after the bet is

$$W = b(\tilde{x})\mathcal{O}(\tilde{x}), \tag{2}$$

where \tilde{x} denotes the winning horse.

We note that an investor who allocates \$1 of capital, investing $\frac{B}{\mathcal{O}(x)}$ to each state x , where

$$B = \frac{1}{\sum_x \frac{1}{\mathcal{O}(x)}}, \tag{3}$$

receives the payoff B with certainty. This motivates the following definition:

Definition 2 *The riskless bank account payoff, B , is given by*

$$B = \frac{1}{\sum_x \frac{1}{\mathcal{O}(x)}}. \tag{4}$$

In a horse race with superfair odds,³ $B > 1$; in a horse race with subfair odds (there is a “track take”), $B < 1$. In financial markets, with positive “interest rates,” $B > 1$. If the “interest rate” is zero, $B = 1$.

2.2 Utility and Optimal Betting Weights

In order to quantify the benefits, to an investor, of the information in the model q , we consider the investor’s utility function, U .

Assumption 1 *Our investor subscribes to the axioms of utility theory.⁴ The investor has a utility function, $U : (W_b, W_s) \rightarrow R$, that*

- (i) *is strictly concave,*
- (ii) *is twice differentiable,*
- (iii) *is strictly monotone increasing,*
- (iv) *has the property $(0, \infty) \subseteq \text{range}(U')$, i.e., there exists a ‘blow-up point’, W_b , with*

$$\lim_{W \rightarrow W_b^+} U'(W) = \infty \tag{5}$$

and a ‘saturation point’, W_s , with

$$\lim_{W \rightarrow W_s^-} U'(W) = 0, \tag{6}$$

and

³See, for example, Cover and Thomas (1991).

⁴Such an investor maximizes his expected utility with respect to the probability measure that he believes (see, for example, Luenberger (1998)).

(v) is compatible with the market in the sense that $W_b < B < W_s$.

One can see easily that condition (v) just means that the wealth associated with the bank account lies in the domain of U .

We note that many popular utility functions (for example, the logarithmic, exponential and power utilities (see Luenberger (1998)) are consistent with these conditions.

According to Utility Theory, a rational investor who believes the measure q allocates so as to maximize the expectation, under q , of his utility function (as applied to his post-bet wealth). In conjunction with (2), this means that this investor allocates with

$$b^*(q) = \arg \max_{b \in \mathcal{B}_X} E_q[U(b, \mathcal{O})], \tag{7}$$

where

$$E_q[U(b, \mathcal{O})] = \sum_{x \in \mathcal{X}} q(x)U(b(x)\mathcal{O}(x)), \tag{8}$$

and

$$\mathcal{B}_X = \{b : \sum_{x \in \mathcal{X}} b(x) = 1\}. \tag{9}$$

The following lemma states how the optimal betting weights can be computed:

Lemma 1 *An expected utility maximizing investor who believes the probability measure q and invests in a horse race with odds ratios \mathcal{O} will allocate*

$$b^*(q)(x) = \frac{1}{\mathcal{O}(x)}(U')^{-1} \left(\frac{\lambda}{q(x)\mathcal{O}(x)} \right), \tag{10}$$

where $(U')^{-1}$ denotes the inverse of the function U' , and λ is the solution of the following equation:

$$\sum_x \frac{1}{\mathcal{O}(x)}(U')^{-1} \left(\frac{\lambda}{q(x)\mathcal{O}(x)} \right) = 1. \tag{11}$$

Under the assumptions of this paper, the solution to (11) exists, is unique, and the optimal allocation, $b^*(q)$, depends continuously on q .

Proof: See, for example, Friedman and Sandow (2003b), Theorem 1.

2.3 Generalization of Entropy and Relative Entropy

We define a utility-based generalization of Kullback Leibler relative entropy via:

Definition 3 *The relative (U, \mathcal{O}) -entropy from the probability measure p to the probability measure q is given by:*

$$D_{U,\mathcal{O}}(p||q) = \sum_x p(x)U(b_p^*(x)\mathcal{O}(x)) - \sum_x p(x)U(b_q^*(x)\mathcal{O}(x)). \tag{12}$$

Interpretation: relative (U, \mathcal{O}) -entropy is the difference between

- (i) the expected (under the measure p) utility of the payoffs if we allocate optimally according to p , and,
- (ii) the expected (under the measure p) utility of the payoffs if we allocate optimally according to the misspecified model, q .

Definition 4 The (U, \mathcal{O}) -entropy of the probability measure p is given by:

$$H_{U,\mathcal{O}}(p) = \sum_x p(x)U(\mathcal{O}(x)) - \sum_x p(x)U(b_p^*(x)\mathcal{O}(x)). \tag{13}$$

Interpretation: (U, \mathcal{O}) -entropy is the difference between the expected utility of the payoffs for a clairvoyant who wins every race and the expected utility under the optimal allocation.

$H_{U,\mathcal{O}}(p)$ and $D_{U,\mathcal{O}}(p||q)$ have important properties summarized in the following theorem.

Known Result 1 $D_{U,\mathcal{O}}(p||q)$ and $H_{U,\mathcal{O}}(p)$ have the following properties

- (i) $D_{U,\mathcal{O}}(p||q) \geq 0$ with equality if and only if $p = q$,
- (ii) $D_{U,\mathcal{O}}(p||q)$ is a strictly convex function of p ,
- (iii) $H_{U,\mathcal{O}}(p) \geq 0$, and
- (iv) $H_{U,\mathcal{O}}(p)$ is a strictly concave function of p .

Proof: See Theorem 2 in Friedman and Sandow (2003b).

Note that for $U(W) = \log(W)$, we recover the entropy and Kullback-Leibler relative entropy under any system of odds ratios.

We note that, solving (12) for $E_p[U(b^*(q), \mathcal{O})]$ and substituting from (13), one can obtain the following decomposition for the expected utility, under the measure p , for an investor who allocates under the (misspecified measure) q :

$$E_p[U(b^*(q), \mathcal{O})] = E_p[U(\mathcal{O})] - H_{U,\mathcal{O}}(p) - D_{U,\mathcal{O}}(p||q). \tag{14}$$

On the right hand side of (14), the first term is the expected utility of an investor who wins every bet (a clairvoyant), the second term is the expected utility loss the investor incurs because the world (characterized by p) is uncertain, and the third term is the expected utility loss related to the investor’s model misspecification.

2.4 Connection with Kullback-Leibler Relative Entropy

It so happens that relative (U, \mathcal{O}) -entropy essentially reduces to Kullback-Leibler relative entropy for a logarithmic family of utilities:

Known Result 2 *The relative (U, \mathcal{O}) -entropy, $D_{U, \mathcal{O}}(p||q)$, is independent of the odds ratios, \mathcal{O} for any candidate model p and prior measure, q , if and only if the utility function, U , is a member of the logarithmic family*

$$U(W) = \gamma_1 \log(W - \gamma) + \gamma_2 \quad , \tag{15}$$

where $\gamma_1 > 0$, γ_2 and $\gamma < B$ are constants. In this case,

$$D_{U, \mathcal{O}}(p||q) = \gamma_1 E_p \left[\log \left(\frac{p}{q} \right) \right] = \gamma_1 D(p||q) \quad . \tag{16}$$

Proof: See Friedman and Sandow (2003a), Theorem 3. It follows that relative (U, \mathcal{O}) -entropy reduces, up to a multiplicative constant, to Kullback-Leibler relative entropy, if and only if the utility is a member of the logarithmic family (16).

3 U -Entropy

We have seen that $D_{U, \mathcal{O}}(p||q) \geq 0$ represents the expected gain in utility from allocating under the true measure p , rather than allocating according to the misspecified measure q under market odds \mathcal{O} for an investor with utility U . One of the main goals of this paper is to explore what happens when we set $q(x)$ equal to the homogeneous expected return measure,

$$q(x) = \frac{B}{\mathcal{O}(x)} \tag{17}$$

Readers familiar with finance will recognize this as the risk neutral pricing measure⁵ generated by the odds ratios. In finance, a risk neutral pricing measure is a measure under which the price of any perfectly hedgeable contingent claim is equal to the discounted (by the bank account) expectation of the payoff of the contingent claim. In the horse race setting, there is only one such measure, given by (17). We shall see in Section 4 that there are compelling reasons to consider this case. In this case,

$$\mathcal{O}(x) = \frac{B}{q(x)} \tag{18}$$

and

$$D_{U, \frac{B}{q}}(p||q) = \sum_x p(x) U \left(B \frac{b_p^*(x)}{q(x)} \right) - \sum_x p(x) U(B),$$

so

$$D_{U, \frac{B}{q}}(p||q) = \sum_x p(x) U \left(B \frac{b_p^*(x)}{q(x)} \right) - U(B) \tag{19}$$

⁵See, for example Duffie (1996). We note that the risk neutral pricing measure generated by the odds ratios need not coincide with any “real world” measure.

In this special case, $D_{U, \frac{B}{q}}(p||q)$ can be interpreted as the excess performance from allocating according to the real world measure p over the (risk-neutral) measure q derived from the odds ratios.⁶

Likewise, under (18), we have

$$\begin{aligned} H_{U, \frac{B}{q}}(p) &= \sum_x p(x)U\left(\frac{B}{q(x)}\right) - \sum_x p(x)U\left(B\frac{b_p^*(x)}{q(x)}\right) \\ &= \sum_x p(x)U\left(\frac{B}{q(x)}\right) - D_{U, \frac{B}{q}}(p||q) - U(B). \end{aligned}$$

In the special case where q is the uniform distribution, $q(x) = \frac{1}{|\mathcal{X}|}$, which we denote by $\frac{1}{|\mathcal{X}|}$, we obtain

$$H_{U, |\mathcal{X}|B}(p) = \sum_x p(x)U(|\mathcal{X}|B) - D_{U, B|\mathcal{X}|}\left(p||\frac{1}{|\mathcal{X}|}\right) - U(B). \tag{20}$$

For simplicity, we assume that $B = 1$ and $U(B) = 0$ from now on.⁷

3.1 Definitions

Motivated by (19), we make the following definition:

Definition 5 *The U -relative entropy from the probability measure p to the probability measure q is given by*

$$D_U(p||q) = \sup_{b \in \mathcal{B}_x} \sum_x p(x)U\left(\frac{b(x)}{q(x)}\right). \tag{21}$$

Motivated by (20), we define the U -entropy:⁸

Definition 6 *The U -entropy of the probability measure p is given by*

$$H_U(p) = U(|\mathcal{X}|) - D_U\left(p||\frac{1}{|\mathcal{X}|}\right). \tag{22}$$

Of course, these specializations of (U, \mathcal{O}) -entropy and relative (U, \mathcal{O}) -entropy inherit all of the properties stated in Section 2. It is easy to show that for logarithmic utilities, they reduce to entropy and Kullback-Leibler relative entropy, respectively, so these quantities are generalizations of entropy and Kullback-Leibler relative entropy, respectively. Next, we generalize the above definitions to conditional probability measures. To keep our notation simple, we use $p(x)$ and $p(y)$ to represent the probability distributions for the random variables X and Y and use both $H_U(p)$ and $H_U(X)$ to represent the U -entropy for the random variable X which has probability measure p .

Let

$$\mathcal{B}_y = \{b : \sum_{y \in \mathcal{Y}} b(y) = 1\}. \tag{23}$$

⁶Stutzer (1995), provides a similar interpretation in a slightly different setting.

⁷It is straightforward to develop the material below under more general assumptions.

⁸We note that this definition of U -entropy is quite similar to, but not the same as, the definition of u -entropy in Slomczyński and Zastawniak (2004).

Definition 7 The conditional relative U -entropy from the probability measure p to the probability measure q is given by

$$\begin{aligned} D_U(p(y|x)||q(y|x)) &= \sum_x p(x) D_U(p(y|X=x)||q(y|X=x)) \\ &= \sum_x p(x) \sup_{b(\cdot|x) \in \mathcal{B}_y} \sum_y p(y|x) U\left(\frac{b(y|x)}{q(y|x)}\right) \end{aligned}$$

Definition 8 The conditional U -entropy from the probability measure p to the probability measure q is given by

$$\begin{aligned} H_U(Y|X) &= \sum_{x \in \mathcal{X}} p(x) H_U(Y|X=x) \\ &= U(|\mathcal{Y}|) - \sum_{x \in \mathcal{X}} p(x) \sup_{b(\cdot|x) \in \mathcal{B}_y} \sum_y p(y|x) U(|\mathcal{Y}|b(y|x)) \end{aligned}$$

We note that we could have stated the preceding definitions as special cases of conditional (U, \mathcal{O}) -entropy and conditional relative (U, \mathcal{O}) -entropy as in Friedman and Sandow (2003b). The latter quantities can be interpreted in the context of a conditional horse race. The following definition generalizes mutual information.

Definition 9 The mutual U -information between the probability measures p and q is defined as

$$I_U(X; Y) = D_U(p(x, y)||p(x)p(y)) \tag{24}$$

3.2 Properties of U -Entropy and Relative U -Entropy

As a special case of Known Result 1, we have

Corollary 1 The generalized relative entropy, $D_U(p||q)$, and the generalized entropy, $H_U(p)$, have the following properties

- (i) $D_U(p||q) \geq 0$ with equality if and only if $p = q$,
- (ii) $D_U(p||q)$ is a strictly convex function of p ,
- (iii) $H_U(p) \geq 0$, and
- (iv) $H_U(p)$ is a strictly concave function of p .

We now establish that many properties that hold for the classical quantities of information theory also hold in this more general setting.

Theorem 1 $H_U(Z|X, Y) \leq H_U(Z|X)$

Proof. By definition,

$$H_U(Z|X, Y) = \sum_{x,y} p(x, y) H_U(Z|X = x, Y = y) \tag{25}$$

$$= \sum_{x,y} p(x, y) \left(U(|\mathcal{Z}|) - D_U \left(p(z|x, y) \parallel \frac{1}{|\mathcal{Z}|} \right) \right)$$

$$= U(|\mathcal{Z}|) - \sum_x p(x) \sum_y p(y|x) D_U \left(p(z|x, y) \parallel \frac{1}{|\mathcal{Z}|} \right) \tag{26}$$

$$\leq U(|\mathcal{Z}|) - \sum_x p(x) D_U \left(\sum_y p(y|x) p(z|x, y) \parallel \frac{1}{|\mathcal{Z}|} \right) \tag{27}$$

$$\text{(by convexity of } D_U(\cdot \parallel \frac{1}{|\mathcal{Z}|}) \text{)} \tag{28}$$

$$= U(|\mathcal{Z}|) - \sum_x p(x) D_U \left(p(z|x) \parallel \frac{1}{|\mathcal{Z}|} \right) \tag{29}$$

$$= H_U(Z|X) \quad \square \tag{30}$$

Corollary 2 $H_U(Y|X) \leq H_U(Y)$ (*Conditioning Reduces Entropy*)

Corollary 3 (*Shuffles increase entropy*). If T is a random shuffle (permutation) of a deck of cards and X is the initial position of the cards in the deck and if the choice of shuffle T is independent of X , then

$$H_U(X) \leq H_U(TX) \tag{31}$$

where TX is the permutation of the deck induced by the shuffle T .

Proof. We follow the proof sketched on Page 48 of Cover and Thomas (1991). First, notice that for any fixed permutation $T = t$, from the definition of H_U , we have

$$H_U(tX) = H_U(X) \tag{32}$$

(since we have only reordered the states, but we have not changed the probabilities associated with the states). So

$$H_U(TX) \geq H_U(TX|T)$$

$$\text{(by Corollary 2)}$$

$$= \sum_t p(t) H_U(TX|T = t)$$

$$= \sum_t p(t) H_U(tX)$$

$$\text{(by (32))}$$

$$= H_U(X) \quad \square$$

Theorem 2 *If the random variables X, Y, Z form a Markov chain $X \rightarrow Y \rightarrow Z$, then $H_U(Z|X, Y) = H_U(Z|Y)$*

Proof.

$$\begin{aligned} H_U(Z|X, Y) &= \sum_{x,y} p(x, y) \left(U(|\mathcal{Z}|) - D_U \left(p(z|x, y) \parallel \frac{1}{|\mathcal{Z}|} \right) \right) \\ &= U(|\mathcal{Z}|) - \sum_{x,y} p(x, y) D_U \left(p(z|y) \parallel \frac{1}{|\mathcal{Z}|} \right) \\ &\quad \text{(by the Markov property)} \\ &= U(|\mathcal{Z}|) - \sum_y p(y) D_U \left(p(z|y) \parallel \frac{1}{|\mathcal{Z}|} \right) \\ &= H_U(Z|Y) \quad \square \end{aligned}$$

Corollary 4 *If the random variables X, Y, Z form a Markov chain $X \rightarrow Y \rightarrow Z$, then $H_U(Z|X) \geq H_U(Z|Y)$ and $H_U(X|Z) \geq H_U(X|Y)$*

Proof. We prove the first inequality by noting that by Corollary 2,

$$H_U(Z|X) \geq H_U(Z|X, Y), \quad (33)$$

and by Theorem 2,

$$H_U(Z|X, Y) = H_U(Z|Y). \quad (34)$$

The second inequality follows from the fact that $X \rightarrow Y \rightarrow Z$ is equivalent to $Z \rightarrow Y \rightarrow X$. \square

Corollary 5 *The conditional entropy $H_U(X_n|X_1)$ increases with n for a stationary Markov process.*

Proof. The proof is essentially the same as that given on 36 of Cover and Thomas (1991). \square

In general, the chain rule for relative entropy does not hold, i.e.

$$D_U(p(x, y) \parallel q(x, y)) \neq D_U(p(x) \parallel q(x)) + D_U(p(y|x) \parallel q(y|x)) \quad (35)$$

However, we do have the following two inequalities.

Theorem 3

(i) $D_U(p(x) \parallel q(x)) \leq D_U(p(x, y) \parallel q(x, y))$, and

(ii) $D_U(p(y|x) \parallel q(y|x)) \leq D_U(p(x, y) \parallel q(x, y))$.

Proof. By definition,

$$D_U(p(x, y) \| q(x, y)) = \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}}} \sum_{x, y} p(x, y) U \left(\frac{b(x, y)}{q(x, y)} \right) \tag{36}$$

$$\geq \sup_{b \in \mathcal{B}_{\mathcal{X}}} \sum_{x, y} p(x, y) U \left(\frac{b(x)q(y|x)}{q(x)q(y|x)} \right) \tag{37}$$

$$\text{(since } b(x)q(y|x) \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}} \text{)} \tag{38}$$

$$= \sup_{b \in \mathcal{B}_{\mathcal{X}}} \sum_x p(x) \sum_y p(y|x) U \left(\frac{b(x)}{q(x)} \right) \tag{39}$$

$$= D_U(p(x) \| q(x)) \tag{40}$$

and

$$D_U(p(x, y) \| q(x, y)) = \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}}} \sum_{x, y} p(x, y) U \left(\frac{b(x, y)}{q(x, y)} \right) \tag{41}$$

$$= \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}}} \sum_x p(x) \sum_y p(y|x) U \left(\frac{b(x, y)}{q(x, y)} \right) \tag{42}$$

$$\geq \sum_x p(x) \sup_{b(\cdot|x) \in \mathcal{B}_{\mathcal{Y}}} \sum_y p(y|x) U \left(\frac{b(y|x)q(x)}{q(y|x)q(x)} \right) \tag{43}$$

$$\text{(since } \tilde{b}(x, y) = b(y|x)q(x) \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}} \text{)} \tag{44}$$

$$= D_U(p(y|x) \| q(y|x)) \tag{45}$$

□

Theorem 4 (More refined horse races have higher relative U -entropies) Let $T : \mathcal{X} \rightarrow \mathcal{Y}$ be any onto function and $p_T(y), q_T(y)$ be the induced probabilities on \mathcal{Y} , i.e.,

$$p_T(y) = \sum_{\{x:T(x)=y\}} p(x). \tag{46}$$

Then

$$D_U(p_T(y) \| q_T(y)) \leq D_U(p(x) \| q(x)) \tag{47}$$

Proof.

$$D_U(p_T(y) \| q_T(y)) = \sup_{b \in \mathcal{B}_{\mathcal{Y}}} \sum_y p_T(y) U \left(\frac{b(y)}{q_T(y)} \right) \tag{48}$$

$$= \sup_{b \in \mathcal{B}_{\mathcal{Y}}} \sum_y \left(\sum_{\{x:T(x)=y\}} p(x) \right) U \left(\frac{b(y)}{\sum_{\{x:T(x)=y\}} q(x)} \right) \tag{49}$$

Define $\tilde{b}(x) = \frac{q(x)b(T(x))}{\sum_{\{x':T(x')=T(x)\}} q(x')}$ Then $\tilde{b}(x) \in \mathcal{B}_X$ and

$$\sum_y \left(\sum_{\{x:T(x)=y\}} p(x) \right) U \left(\frac{b(y)}{\sum_{\{x:T(x)=y\}} q(x)} \right) = \sum_x p(x) U \left(\frac{\tilde{b}(x)}{q(x)} \right) \tag{50}$$

$$\leq D_U(p(x)||q(x)) \tag{51}$$

□

Theorem 4 can also be regarded as a special case of Theorem 3, part (i).

Theorem 5 (second law of thermodynamics) Let μ_n and μ'_n be the probability distributions (at time n) of two different Markov chains arising from the same transition mechanism but from different starting distributions. Then $D_U(\mu_n||\mu'_n)$ decreases with n .

Proof. Let the corresponding joint mass function be denoted by $p(x_n, x_{n+1})$ and $q(x_n, x_{n+1})$ and let $r(\cdot|\cdot)$ denote the probability transition function for the Markov chain. Then $p(x_n, x_{n+1}) = p(x_n)r(x_{n+1}|x_n)$ and $q(x_n, x_{n+1}) = q(x_n)r(x_{n+1}|x_n)$.

$$\begin{aligned} D_U(p(x_n, x_{n+1})||q(x_n, x_{n+1})) &= \sup_{b \in \mathcal{B}_X \times \mathcal{X}} \sum_{x_n} \sum_{x_{n+1}} p(x_n, x_{n+1}) U \left(\frac{b(x_n, x_{n+1})}{q(x_n, x_{n+1})} \right) \\ &= \sup_{b \in \mathcal{B}_X \times \mathcal{X}} \sum_{x_n} p(x_n) \sum_{x_{n+1}} r(x_{n+1}|x_n) U \left(\frac{b(x_n, x_{n+1})}{q(x_n, x_{n+1})} \right) \\ &\leq \sup_{b \in \mathcal{B}_X \times \mathcal{X}} \sum_{x_n} p(x_n) U \left(\sum_{x_{n+1}} r(x_{n+1}|x_n) \frac{b(x_n, x_{n+1})}{q(x_n)r(x_{n+1}|x_n)} \right) \\ &\quad \text{(Here, we have used Jensen's inequality} \\ &\quad \text{and the fact that } U \text{ is concave.)} \\ &= \sup_{b \in \mathcal{B}_X \times \mathcal{X}} \sum_{x_n} p(x_n) U \left(\frac{\sum_{x_{n+1}} b(x_n, x_{n+1})}{q(x_n)} \right) \\ &\leq \sup_{b \in \mathcal{B}_X} \sum_{x_n} p(x_n) U \left(\frac{b(x_n)}{q(x_n)} \right) \\ &\quad \text{(since } b(x_n) = \sum_{x_{n+1}} b(x_n, x_{n+1}) \in \mathcal{B}_X) \\ &= D_U(p(x_n)||q(x_n)) \end{aligned}$$

By Theorem 3,

$$D_U(p(x_n, x_{n+1})||q(x_n, x_{n+1})) \geq D_U(p(x_{n+1})||q(x_{n+1})), \tag{52}$$

hence, $D_U(p(x_n)||q(x_n)) \geq D_U(p(x_n, x_{n+1})||q(x_n, x_{n+1})) \geq D_U(p(x_{n+1})||q(x_{n+1}))$. □

Corollary 6 *The relative U-entropy $D_U(\mu_n||\mu)$ between a distribution μ_n on the states at time n and a stationary distribution μ decreases with n .*

Proof: This is a special case of Theorem 5, where $\mu'_n = \mu$ for any n . \square

Corollary 7 *If the uniform distribution $\frac{1}{|\mathcal{X}|}$ is a stationary distribution, then entropy $H_U(\mu_n)$ increases with n .*

Theorem 6

(i) $I_U(X;Y) = I_U(Y;X)$, and

(ii) $I_U(X;Y) \leq I_U(X;Y,Z)$.

Proof. (i) is obvious. By definition,

$$I_U(X;Y) = \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}}} \sum_{x,y} p(x,y)U\left(\frac{b(x,y)}{p(x)p(y)}\right) \tag{53}$$

$$= \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}}} \sum_{x,y,z} p(x,y,z)U\left(\frac{b(x,y)p(z|y)}{p(x)p(y,z)}\right) \tag{54}$$

$$\leq \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}}} \sum_{x,y,z} p(x,y,z)U\left(\frac{b(x,y,z)}{p(x)p(y,z)}\right) \tag{55}$$

$$= I_U(X;Y,Z) \tag{56}$$

\square

Theorem 7 (Data processing inequality): *If the random variables X,Y,Z form a Markov chain $X \rightarrow Y \rightarrow Z$, then $I_U(X;Y) \geq I_U(X;Z)$.*

Proof. We will first show that $I_U(X;Y) = I_U(X;Y,Z)$. From the previous result,

$$I_U(X;Y) \leq I_U(X;Y,Z) \tag{57}$$

On the other hand,

$$I_U(X;Y,Z) = \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}}} \sum_{x,y,z} p(x,y,z)U\left(\frac{b(x,y,z)}{p(x)p(y,z)}\right) \tag{58}$$

$$= \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}}} \sum_{x,y} p(x,y) \sum_z p(z|x,y)U\left(\frac{b(x,y,z)}{p(x)p(y,z)}\right) \tag{59}$$

$$\leq \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}}} \sum_{x,y} p(x,y)U\left(\sum_z p(z|x,y) \frac{b(x,y,z)}{p(x)p(y,z)}\right) \tag{60}$$

$$= \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}}} \sum_{x,y} p(x,y)U\left(\frac{\sum_z b(x,y,z)}{p(x)p(y)}\right) \tag{61}$$

$$\text{(since } p(z|x,y) = p(z|y)\text{, by the Markov property)} \tag{62}$$

$$\leq \sup_{b \in \mathcal{B}_{\mathcal{X} \times \mathcal{Y}}} \sum_{x,y} p(x,y)U\left(\frac{b(x,y)}{p(x)p(y)}\right) \tag{63}$$

$$= I_U(X;Y) \tag{64}$$

so $I_U(X; Y) = I_U(X; Y, Z) \geq I_U(X; Z)$. \square

Corollary 8 *If $Z = g(Y)$, we have $I_U(X; Y) \geq I_U(X; g(Y))$.*

3.3 Power Utility

Consider the power utility, given by

$$U_\kappa(W) = \frac{W^{1-\kappa} - 1}{1 - \kappa}, \kappa \geq 0, \kappa \neq 1. \tag{65}$$

We note that for each fixed wealth level, taking limits as $\kappa \rightarrow 1$, the power utility approaches the logarithmic utility.

In this section, we compute the power utility's U -entropy and relative U -entropy, explore limiting behavior and make comparisons with Tsallis entropy. We note that the power utility is commonly used (see, for example, Morningstar (2002)), since it has desirable optimality properties (see Stutzer (2003)) and the constant relative risk aversion property. The latter property follows from the fact that, for the power utility, the relative risk aversion coefficient (see, for example, Luenberger (1998)) is

$$R_\kappa(W) = -W \frac{U''_\kappa(W)}{U'_\kappa(W)} = \kappa. \tag{66}$$

From Friedman and Sandow (2003b), with $\mathcal{O}(x) = \frac{1}{q(x)}$, we have

$$b^*(p)(x) = q(x) \left(\frac{p(x)}{\lambda q(x)} \right)^{\frac{1}{\kappa}} \tag{67}$$

where

$$\lambda = \left(\sum_x q(x) \left(\frac{p(x)}{q(x)} \right)^{\frac{1}{\kappa}} \right)^\kappa \tag{68}$$

After some algebra, we obtain

$$D_{U_\kappa}(p||q) = \frac{\left(\sum_x p(x)^{\frac{1}{\kappa}} q(x)^{1-\frac{1}{\kappa}} \right)^\kappa}{1 - \kappa} - \frac{1}{1 - \kappa}. \tag{69}$$

3.3.1 U -Entropy for Large Relative Risk Aversion

Let us consider the limit of infinite risk aversion, which for the power utility corresponds to $\kappa \rightarrow \infty$, as can be seen from (66). It follows from (67) and (68), that a power utility investor with infinite risk aversion invests all his money in the bank account, i.e., allocates according to

$$b^*(p)(x) = q(x) = \frac{B}{\mathcal{O}(x)} = \frac{1}{\mathcal{O}(x)}, \tag{70}$$

no matter what his belief-measure, p , is, so that his after-bet wealth is $B = 1$ with certainty. Such an investor makes no use of the information provided by the model, so no model can outperform another and, therefore,

$$\lim_{\kappa \rightarrow \infty} D_{U_\kappa}(p||q) = 0, \tag{71}$$

for any measures p, q and any utility function U .

What happens if the relative risk aversion of a power-utility investor, i.e., κ , is large but finite? To answer this question, we expand (69) as follows.

$$\begin{aligned} D_{U_\kappa}(p||q) &= \frac{1}{1-\kappa} \left[\sum_x \left(\frac{p(x)}{q(x)} \right)^{\frac{1}{\kappa}} q(x) \right]^\kappa - \frac{1}{1-\kappa} \\ &= \frac{1}{1-\kappa} \left[\sum_x \left(1 + \frac{1}{\kappa} \log \frac{p(x)}{q(x)} + \frac{1}{2\kappa^2} \left(\log \frac{p(x)}{q(x)} \right)^2 + O(\kappa^{-3}) \right) q(x) \right]^\kappa \\ &\quad - \frac{1}{1-\kappa} \quad \left(\text{since } z^{\frac{1}{\kappa}} = 1 + \frac{1}{\kappa} \log(z) + \frac{1}{2\kappa^2} (\log z)^2 + O(\kappa^{-3}) \right) \\ &= \frac{1}{1-\kappa} \left[1 + \frac{1}{\kappa} \sum_x q(x) \log \frac{p(x)}{q(x)} + \frac{1}{2\kappa^2} \sum_x q(x) \left(\log \frac{p(x)}{q(x)} \right)^2 + O(\kappa^{-3}) \right]^\kappa \\ &\quad - \frac{1}{1-\kappa} \\ &= \frac{1 - e^{\sum_x q(x) \log \frac{p(x)}{q(x)}}}{\kappa} + \frac{1 - e^{\sum_x q(x) \log \frac{p(x)}{q(x)}} \left[1 + \frac{1}{2} \text{var}_q \left(\log \frac{p(x)}{q(x)} \right) \right]}{\kappa^2} + O(\kappa^{-3}), \end{aligned}$$

where the last equality follows from the fact that

$$\frac{1}{1-\kappa} \left(1 + \frac{a}{\kappa} + \frac{b}{2\kappa^2} + \frac{c}{\kappa^3} \right)^\kappa - \frac{1}{1-\kappa} = \frac{1 - e^a}{\kappa} + \frac{1 - e^a \left(1 + \frac{b-a^2}{2} \right)}{\kappa^2} + O(\kappa^{-3}).$$

So

$$D_{U_\kappa}(p||q) = \frac{1 - e^{-D(q||p)}}{\kappa} + \frac{1 - e^{-D(q||p)} \left[1 + \frac{1}{2} \text{var}_q \left(\log \frac{p(x)}{q(x)} \right) \right]}{\kappa^2} + O(\kappa^{-3}). \tag{72}$$

Thus, we have related, for large κ , $D_{U_\kappa}(p||q)$ to Kullback-Leibler relative entropy.

3.3.2 Relation with Tsallis Entropy

Another generalization of entropy, which recently has been a topic of increasing interest in physics, is the Tsallis entropy (see Tsallis (1988)). The Tsallis entropy provides a theoretical framework for a nonextensive thermodynamics, i.e., for a thermodynamics for which the entropy of a system made up of two independent subsystems is not simply the sum of the entropies of the subsystems. Examples for physical application of such a theory might be astronomical self-gravitating systems (see, for example, Plastino and Plastino (1999)).

The Tsallis entropy can be generalized to the Tsallis relative entropy (see, for example, Tsallis and Brigatti (2004)) as follows:

$$D_\alpha^T(p||q) = \frac{\sum_x p^\alpha(x) q^{1-\alpha}(x) - 1}{\alpha - 1}, \alpha \in R. \tag{73}$$

This entropy recovers the standard Boltzmann-Gibbs entropy $\alpha \rightarrow 1$. It turns out that there is a simple monotonic relationship between our relative U -entropy and the relative Tsallis entropies:

$$D_{U_\kappa}(p||q) = \frac{\left(D_{\frac{1}{\kappa}}^T(p||q) \left(\frac{1}{\kappa} - 1\right) + 1\right)^\kappa - 1}{1 - \kappa}. \quad (74)$$

Also note that

$$H_{U_\kappa}(X) = U_\kappa(|\mathcal{X}|) - D_{U_\kappa}\left(p||\frac{1}{|\mathcal{X}|}\right) = \frac{|\mathcal{X}|^{1-\kappa} \left[1 - \left(\sum_x p(x)^{\frac{1}{\kappa}}\right)^\kappa\right]}{1 - \kappa}. \quad (75)$$

Comparing this to the Tsallis entropy:

$$H_\alpha^T(X) = \frac{\sum_x p^\alpha(x) - 1}{1 - \alpha}; \quad (76)$$

we have the following simple monotonic relationship between the two entropies

$$H_{U_\kappa}(X) = \frac{\left[1 - \left(H_{\frac{1}{\kappa}}^T(X) \left(1 - \frac{1}{\kappa}\right) + 1\right)^\kappa\right] |\mathcal{X}|^{1-\kappa}}{1 - \kappa}. \quad (77)$$

4 Application: Probability Estimation via Relative U -Entropy Minimization

Information theoretic ideas such as the maximum entropy principle (see, for example, Jaynes (1957) and Golan et al. (1996)) or the closely related minimum relative entropy (MRE) principle (see, for example Lebanon and Lafferty (2001)) have assumed major roles in statistical learning.

We discuss these principles from the point of view of a model user who would use a probabilistic model to make bets in the horse race setting of Section 2.⁹ As we shall see, probability models estimated by these principles are robust in the sense that they maximize “worst case” model performance, in terms of expected utility, for a horse race investor with the three parameter logarithmic utility (15) in a particular market (horse race).

As such, probability models estimated by MRE methods are special in the sense that they are tailored to the risk preferences of logarithmic-family-utility investors. However, not all risk preferences can be expressed by utility functions in the logarithmic family. In the financial community, a substantial percentage, if not a majority, of practitioners implement utility functions outside the logarithmic family (15) (see, for example, Morningstar (2002)). This is not surprising – other utility functions may more accurately reflect the risk preferences of certain investors, or possess important defining properties or optimality properties.

In order to tailor the statistical learning problem to the specific utility function of the investor, one can replace the usual MRE problem with a generalization, based on relative U -entropy, as indicated below. Before stating results for minimum relative U -entropy (MRUE) probability estimation, we review results for MRE probability estimation.

⁹It is possible to consider more general settings, such as incomplete markets, but a number of results depend in a fundamental way on the horse race setting.

4.1 MRE and Dual Formulation

The usual MRE formulation (see, for example, Lebanon and Lafferty (2001) and Kitamura and Stutzer (1997)) is given by:¹⁰

Problem 1 Find

$$p^* = \arg \min_{p \in Q} D(p||p^0) \quad (78)$$

$$s.t. E_p[f^j] - E_{\tilde{p}}[f^j] = 0, j = 1, \dots, J, \quad (79)$$

where each f^j represents a feature (i.e., a function that maps \mathcal{X} to \mathbf{R}), p^0 is a probability measure, \tilde{p} represents the empirical measure, and Q is the set of all probability measures.

The feature constraints (79) can be viewed as a mechanism to enforce "consistency" between the model and the data. Typically, p^0 is interpreted as representing prior beliefs, before modification in light of model training data. Below, we discuss what happens when p^0 is a benchmark measure or the pricing measure generated by the odds ratios. The solution to Problem 1 will be the measure satisfying the feature constraints that is closest to p^0 (in the sense of relative entropy).

In many settings, the above (primal) problem is not as easy to solve numerically as the following dual problem:¹¹

Problem 2 Find

$$\beta^* = \arg \max_{\beta \in \mathbf{R}^J} L(\hat{p}_\beta), \quad (80)$$

$$where L(p) = \sum_x \tilde{p}(x) \log(p(x)) \quad (81)$$

is the log-likelihood function,

$$\hat{p}_\beta(x) = \frac{1}{Z} p^0(x) e^{\beta^T f(x)} \quad (82)$$

$$and Z = \sum_x p^0(x) e^{\beta^T f(x)}, \quad (83)$$

and $\tilde{p}(x)$ denotes the empirical probability that $X = x$.

We note that the primal problem (Problem 1) and the dual problem (Problem 2) have the same solution, in the sense that $p^* = \hat{p}_{\beta^*}$.

¹⁰MRE problems can be stated for *conditional* probability estimation and with regularization. We keep the context and notation as simple as possible by confining our discussion to unconditional estimation without regularization. Extensions are straightforward.

¹¹See, for example, Lebanon and Lafferty (2001).

4.2 Robustness of MRE when the Prior Measure is the Odds Ratio Pricing Measure

In a certain circumstance, the MRE measure (i.e., the solution of Problem 1) has a desirable robustness property, which we discuss in this section. First we state the following minimax result from Grünwald and Dawid (2004).

Known Result 3 *If there exists a probability measure $\pi \in K$, with $\pi_x > 0, \forall x \in \mathcal{X}$, then*

$$\inf_{p' \in K} (-H(p')) = \inf_{p \in K} \sup_{p' \in Q} E_p [\log p'] = \sup_{p' \in Q} \inf_{p \in K} E_p [\log p'] , \tag{84}$$

where Q is the set of all possible probability measures, K is a closed convex set, such as the set of all $p \in Q$ that satisfy the feature constraints, (79). Moreover, a saddle point is attained.

Replacing $\log p'$ with $\log p' - \log p^0$ in (84) doesn't change the finiteness or the convexity properties of the function under the infsup or the supinf. Since the proof of Known Result 3 relies only on those properties of this function, we have

$$\inf_{p \in K} \sup_{p' \in Q} E_p [\log p' - \log p^0] = \sup_{p' \in Q} \inf_{p \in K} E_p [\log p' - \log p^0] , \tag{85}$$

and a saddle point is attained here as well. Further using the fact that

$$D(p||p^0) = \sup_{p' \in Q} E_p [\log p' - \log p^0] , \tag{86}$$

which follows from the information inequality (see, for example, Cover and Thomas (1991)), we obtain

$$\begin{aligned} \inf_{p \in K} D(p||p^0) &= \sup_{p' \in Q} \inf_{p \in K} E_p [\log p' - \log p^0] \\ &= \sup_{p \in Q} \inf_{p' \in K} E_{p'} [\log p - \log p^0] . \end{aligned} \tag{87}$$

Using further $b^*(p) = p$,¹² for $U(W) = \log W$, we see that (87) leads to the following corollary:

Corollary 9 *If $U(W) = \log W$, and*

$$p^0(x) = \frac{B}{\mathcal{O}(x)} = \frac{1}{\mathcal{O}(x)}, \forall x \in \mathcal{X}, \tag{88}$$

then

$$p^* = \arg \min_{p \in K} D(p||p^0) = \arg \max_{p \in Q} \min_{p' \in K} E_{p'} [U(b^*(p), \mathcal{O})], \tag{89}$$

where Q is the set of all possible probability measures and K is the set of all $p \in Q$ that satisfy the feature constraints, (79).¹³

¹²See, for example, Cover and Thomas (1991), Theorem 6.1.2.

¹³For ease of exposition, we have proved this result for $U(\cdot) = \log(\cdot)$. The same result holds for utilities in the generalized logarithmic family (15).

Thus, if p^0 is given by pricing measure generated by the odds ratios, then the MRE measure, p^* , is robust in the following sense: by choosing p^* a rational (expected-utility optimizing) investor maximizes his (model-based optimal) expected utility in the most adverse environment consistent with the feature constraints.

Here, p^0 has a specific and non-traditional interpretation. It no longer represents prior beliefs about the real world measure that we estimate with p^* . Rather, it represents the risk neutral pricing measure consistent with the odds ratios. If a logarithmic utility investor wants the measure that he estimates to have the optimality property (89), he is not free to choose p^0 to represent his prior beliefs, in general. However, if he sets p^0 to the risk neutral pricing measure determined by the odds ratios, then the optimality property (89) will hold. In this case, p^* , can be thought of as a function of the odds ratios, rather than an update of prior beliefs.

It is this property, (89), that we seek to generalize to accommodate risk preferences more general than the ones compatible with a logarithmic utility function.

4.3 General Risk Preferences and Robust Relative Performance

In this section, we state a known result that we will use below to generalize the robustness property (89).

The MRE problem is generalized in Friedman and Sandow (2003a) to a minimum relative (U, \mathcal{O}) entropy problem. The solution to this generalized problem is robust in the following sense:

Known Result 4

$$\arg \min_{p \in K} D_{U, \mathcal{O}}(p || p^0) = \arg \max_{p \in Q} \min_{p' \in K} \{ E_{p'} [U(b^*(p), \mathcal{O})] - E_{p'} [U(b^*(p^0), \mathcal{O})] \} ,$$

where Q is the set of all possible probability measures, K is the set of all $p \in Q$ that satisfy the feature constraints, (79).

This is only a slight modification of a result from Grünwald and Dawid (2002) and is based on the logic from Topsøe (1979). A proof can be found in Friedman and Sandow (2003a).

Known Result 4 states the following: minimizing $D_{U, \mathcal{O}}(p || p^0)$ with respect to $p \in K$ is equivalent to searching for the measure $p^* \in Q$ that maximizes the worst-case (with respect to the potential true measures, $p' \in K$) relative model performance (over the benchmark model, p^0), in the sense of expected utility. The optimal model, p^* , is robust in the sense that for any other model, p , the worst (over potential true measures $p' \in K$) relative performance is even worse than the worst-case relative performance under p^* . We do not know the true measure; an investor who makes allocation decisions based on p^* is prepared for the worst that “nature” can offer.

4.4 Robust Absolute Performance and MRUE

We now address the following question: Is it possible to formulate a generalized MRE problem for which the solution is robust in the absolute sense of (89), as opposed to the relative sense in Known Result 4? This is indeed possible, as the following Corollary, which follows directly from Known Result 4, indicates:

Corollary 10 *If*

$$p^0(x) = \frac{B}{\mathcal{O}(x)} = \frac{1}{\mathcal{O}(x)}, \forall x \in \mathcal{X}, \tag{90}$$

then

$$p^* = \arg \min_{p \in K} D_U(p||p^0) = \arg \max_{p \in Q} \min_{p' \in K} E_{p'}[U(b^*(p), \mathcal{O})], \tag{91}$$

where Q is the set of all possible probability measures and K is the set of all $p \in Q$ that satisfy the feature constraints, (79).

This Corollary states that, by choosing p^* , a rational (expected-utility optimizing) investor maximizes his (model-based optimal) expected utility in the most adverse environment consistent with the feature constraints. Therefore, an investor with a general utility function, who bets on this horse race and wants to maximize his worst case (in the sense of (91)) expected utility can set p^0 to the risk neutral pricing measure determined by the odds ratios and solve the following problem:¹⁴

Problem 3 (MRUE Problem)

Find

$$p^* = \arg \min_{p \in Q} D_U(p||p^0) \tag{92}$$

$$s.t. E_p[f^j] - E_{\tilde{p}}[f^j] = 0, j = 1, \dots, J, \tag{93}$$

where each f^j represents a feature (i.e., a function that maps \mathcal{X} to \mathbf{R}), p^0 represents the prior measure, \tilde{p} represents the empirical measure, and Q is the set of all probability measures.

As before, p^0 , here, has a specific and non-traditional interpretation. It no longer represents prior beliefs about the real world measure that we estimate with p^* . Rather, it represents the risk neutral pricing measure consistent with the odds ratios. If an investor wants the measure that he estimates to have the optimality property (91), he is not free to choose p^0 to represent his prior beliefs, in general. To attain the optimality property (91), he can set p^0 to represent the risk neutral pricing measure determined by the odds ratios. The MRUE measure, p^* , is a function of the odds ratios (which are incorporated in the measure p^0).

In order to obtain the dual to Problem 3, we specialize, in Problem 3 of Friedman and Sandow (2003a), relative (U, \mathcal{O}) -entropy to relative U -entropy by setting the “prior” measure to the risk neutral pricing measure generated by the odds ratios, $p^0(x) = \frac{1}{\mathcal{O}(x)}$. We obtain

Problem 4 (Dual of MRUE Problem)

¹⁴As noted above, we keep the context and notation as simple as possible by confining our discussion to unconditional estimation without regularization. Extensions are straightforward.

Find

$$\beta^* = \arg \max_{\beta} \sum_x \tilde{p}(x) U(b^*(\hat{p}_{\beta})(x) \mathcal{O}(x)) , \tag{94}$$

$$\text{where } b^*(\hat{p}_{\beta})(x) = \frac{1}{\mathcal{O}(x)} (U')^{-1} \left(\frac{\lambda^*}{\hat{p}_{\beta}(x) \mathcal{O}(x)} \right) \tag{95}$$

$$\text{and } \hat{p}_{\beta}(x) = \frac{\lambda^*}{\mathcal{O}(x) U' (U^{-1}(\beta^T f(x) - \mu^*))} , \tag{96}$$

$$\text{where } \mu^* \text{ solves } 1 = \sum_x \frac{U^{-1}(\beta^T f(x) - \mu^*)}{\mathcal{O}(x)} , \tag{97}$$

$$\text{and } \lambda^* = \left\{ \sum_x \frac{1}{\mathcal{O}(x) U' (U^{-1}(\beta^T f(x) - \mu^*))} \right\}^{-1} . \tag{98}$$

Problem 4 is easy to interpret.¹⁵ The objective function of Problem 4 is the utility (of the expected utility maximizing investor) averaged over the training sample. Thus, our dual problem is a maximization of the training-sample averaged utility, where the utility function, U , is the utility function on which the U -entropy depends.

We note that the primal problem (Problem 3) and the dual problem (Problem 4) have the same solution, in the sense that $p^* = \hat{p}_{\beta^*}$.

This problem is a J -dimensional (J is the number of features), unconstrained, concave maximization problems (see Boyd and Vandenberghe (2004), p. 159 for the concavity). The primal problem, Problem 3, on the other hand, is an m -dimensional (m is the number of states) convex minimization with linear constraints. The dual problem, Problem 4, may be easier to solve than the primal problem, Problem 3, if $m > J$. For conditional probability models, the dual problem will always be easier to solve than the primal problem.

As we have seen above, in cases where odds ratios are available, the MRUE problem yields a solution with the optimality property (91). However, in real statistical learning applications, as mentioned above, it is often the case that odds ratios are not observable. In this case, the builder of a statistical learning model can use assumed odds ratios, on which the model will depend. Given the relation

$$p^0(x) = \frac{B}{\mathcal{O}(x)} = \frac{1}{\mathcal{O}(x)}, \forall x \in \mathcal{X}, \tag{99}$$

as a perhaps more convenient alternative, the model-builder can directly specify a risk neutral pricing measure consistent with the assumed odds ratios. Either way, the model will possess the optimality property (91) under the odds ratios consistent with the assumption. The necessity of providing a risk neutral pricing measure, perhaps, imposes an onus on the MRUE modeler comparable to the onus of finding a prior for the MRE modeler. However, we note that, as for MRE models, the importance of p^0 will diminish as the number of feature constraints grows.

For logarithmic-family utility functions, as a practical matter, even though the MRE and MRUE models require different interpretations and, possibly, different choices for p^0 , the numerical implementations will be the same under the appropriate assumptions.

¹⁵A version that is more easily implemented is given in the Appendix.

5 Conclusions

Motivated by decision theory, we have defined a generalization of entropy (U -entropy) and a generalization of relative entropy (relative U -entropy). We have shown that these generalizations retain a number of properties of entropy and relative entropy that are not retained by the generalizations $D_{U,\mathcal{O}}$ and $H_{U,\mathcal{O}}$ presented in Friedman and Sandow (2003a).

We have used relative U -entropy to formulate an approach to the probability measure estimation problem (Problem 3). Problem 3 generalizes the absolute robustness properties of the solution to the MRE problem, incorporating risk preferences more general than those that can be expressed via the logarithmic family (15).

Depending on the situation, the model builder may prefer one of the methods listed below:

I General Risk Preferences (not expressible by (15)), real or assumed odds ratios available

- (i) *MRUE, Problem 3:* If p^0 is the pricing measure generated by the odds ratios, we get robust absolute performance (in the market described by the odds ratios) in the sense of Corollary 10.
- (ii) *Minimum Relative (U, \mathcal{O})-entropy Problem from Friedman and Sandow (2003a):* If p^0 represents a benchmark model (possibly prior beliefs), we get robust relative outperformance (relative to the benchmark model, in the market described by the odds ratios) in the sense of Known Result 4.

II Special Case: Logarithmic Family Risk Preferences (15), odds ratios need not be available

MRE, Problem 1: If p^0 represents a benchmark model (possibly prior beliefs), we get robust relative outperformance with respect to the benchmark model, under any odds ratios.

6 Appendix

If we specialize relative (U, \mathcal{O})-entropy to relative U -entropy, substituting

$$p^0(x) = \frac{B}{\mathcal{O}(x)} = \frac{1}{\mathcal{O}(x)}, \forall x \in \mathcal{X}, \tag{100}$$

in Problem 4 of Friedman and Sandow (2003a), we obtain a version of the dual problem that is easier to implement, though the objective function is not as easy to interpret as the objective function in Problem 4.

Problem 5 (Easily Implemented Version of Dual Problem)

$$\text{Find } \beta^* = \arg \max_{\beta} \{ \beta^T E_{\bar{p}}[f] - \mu^* \} , \tag{101}$$

$$\tag{102}$$

$$\text{where } \mu^* \text{ solves } 1 = \sum_x p^0(x) U^{-1}(\beta^T f(x) - \mu^*) . \tag{103}$$

Once Problem 5 is solved, the solution to Problem 3 is simply,

$$p^*(x) = \lambda^* \frac{p^0(x)}{U'(U^{-1}(\beta^{*T} f(x) - \mu^*))} \quad (104)$$

where λ^* is the positive constant that makes $\sum_x p^*(x) = 1$.

References

- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge, 2004.
- L. Brillouin. *Science and Information Theory*. Academic Press, New York, 1962.
- T. Cover and J. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.
- F. Cowell. Additivity and the entropy concept: An axiomatic approach to inequality measurement. *Journal of Economic Theory*, 25:131–143, 1981.
- N. Cressie and T. Read. Multinomial goodness of fit tests. *journal of the royal statistical society b*, 46(3):440–464, 1984.
- N. Cressie and T. Read. Pearson's x^2 and the loglikelihood ratio statistic g^2 : a comparative review. *International statistical review*, 57(1):19–43, 1989.
- I. Csiszár and J. Körner. *Information theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, New York, 1997.
- D. Duffie. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, 1996.
- C. Friedman and S. Sandow. Learning probabilistic models: An expected utility maximization approach. *Journal of Machine Learning Research*, 4:291, 2003a.
- C. Friedman and S. Sandow. Model performance measures for expected utility maximizing investors. *International Journal of Theoretical and Applied Finance*, 6(4):355, 2003b.
- C. Friedman and S. Sandow. Model performance measures for leveraged investors. *International Journal of Theoretical and Applied Finance*, 7(5):541, 2004.
- A. Golan, G. Judge, and D. Miller. *Maximum Entropy Econometrics*. Wiley, New York, 1996.
- Group of Statistical Physics. Nonextensive statistical mechanics and thermodynamics: Bibliography. *Working Paper*, <http://tsallis.cat.cbpf.br/TEMUC0.pdf>, 2006.
- P. Grünwald and A. Dawid. Game theory, maximum generalized entropy, minimum discrepancy, robust bayes and pythagoras. *Proceedings ITW*, 2002.
- P. Grünwald and A. Dawid. Game theory, maximum generalized entropy, minimum discrepancy, and robust bayesian decision theory. *Annals of Statistics*, 32(4):1367–1433, 2004.

- J. Ingersoll. *Theory of Financial Decision Making*. Rowman and Littlefield, New York, 1987.
- E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106:620, 1957.
- Y. Kitamura and M. Stutzer. An information-theoretic alternative to generalized method of moments. *Econometrica*, 65(4):861–874, 1997.
- G. Lebanon and J. Lafferty. Boosting and maximum likelihood for exponential models. *Technical Report CMU-CS-01-144 School of Computer Science Carnegie Mellon University*, 2001.
- F. Liese and I. Vajda. *Convex Statistical Distances*. Teubner, Leipzig, 1987.
- D. Luenberger. *Investment Science*. Oxford University Press, New York, 1998.
- Morningstar. The new morningstar ratingTM methodology. <http://www.morningstar.dk/downloads/MRARdefined.pdf>, 2002.
- F. Österreicher and I. Vajda. Statistical information and discrimination. *IEEE transactions on information theory*, 39(3):1036–1039, May, 1993.
- A. Plastino and A. R. Plastino. Tsallis entropy and Jaynes' information theory formalism. *Brazilian Journal of Physics*, 29(1):50, 1999.
- C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27: 379–423 and 623–656, Jul and Oct 1948.
- W. Slomczyński and T. Zastawniak. Utility maximizing entropy and the second law of thermodynamics. *Annals of Probability*, 32(3A):2261, 2004.
- W. Stummer. On a statistical information measure of diffusion processes. *statistics and decisions*, 17:359–376, 1999.
- W. Stummer. On a statistical information measure for a Samuelson-Black-Scholes model. *statistics and decisions*, 19:289–313, 2001.
- W. Stummer. *Exponentials, Diffusions, Finance, Entropy and Information*. Shaker, 2004.
- M. Stutzer. A bayesian approach to diagnosis of asset pricing models. *Journal of Econometrics*, 68:367–397, 1995.
- M. Stutzer. Portfolio choice with endogenous utility: A large deviations approach. *Journal of Econometrics*, 116:365–386, 2003.
- F. Topsøe. Information theoretical optimization techniques. *Kybernetika*, 15(1):8, 1979.
- C. Tsallis. Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52:479, 1988.

C. Tsallis and E. Brigatti. Nonextensive statistical mechanics: A brief introduction. *Continuum Mechechanics and Thermodynammics*, 16:223–235, 2004.

©2007 by MDPI (<http://www.mdpi.org>). Reproduction for noncommercial purposes permitted.