Duality between maximization of expected utility and minimization of relative entropy when probabilities are imprecise

Robert Nau, Duke University
Victor R. R. Jose, Georgetown University
Robert Winkler, Duke University

Subjective probabilities as prices

- Fundamental theorem of subjective probability (de Finetti):
  - Your prices for gambles are *coherent* (avoid a Dutch book) iff there exists a probability dist'n w.r.t. every gamble's price is its expected value
- Fundamental theorem of asset pricing ?

Subjective probabilities as prices

- Fundamental theorem of subjective probability (de Finetti):
  - Your prices for gambles are *coherent* (avoid a Dutch book) iff there exists a probability dist'n w.r.t. every gamble's price is its expected value
- Fundamental theorem of asset pricing ?
  - Same theorem! Just replace "you" with "market," "gamble" with "asset" and "Dutch book" with "arbitrage."
  - Minor differences: (i) EV may discounted at risk free rate if payoffs are received later, (ii) the probability dist'n does not represent any one person's beliefs but those of a "risk neutral representative agent," and (iii) they may not measure pure belief— could be distorted by risk aversion

Incompleteness → Imprecision

- SP with incomplete preferences: Your *buying and selling prices* for gambles are coherent iff there exists a *convex set* of probability dist'ns w.r.t. every gamble's buying [selling] price is less [greater] than its expected value

- AP with incomplete markets: *Bid-ask spreads* in asset prices are coherent iff etc. etc.
  - *Incomplete financial markets provide the purest and most economically important example of imprecise probabilities with precise lower and upper bounds*

In this paper…

- We consider the decision problem of a risk averse agent with precise probabilities who bets against a risk neutral opponent (or market) with imprecise probabilities

- **Theorem:** the expected utility maximization problem of the agent is the *dual* of the problem of finding the distribution in the opponent's set that minimizes a generalized measure of divergence (relative entropy) w.r.t. his own distribution.

In particular (what’s really new)…

- When the agent’s utility function is a *power utility function* (of which logarithmic and exponential are limiting cases), the divergence in the dual problem is a *power divergence* (of which the KL divergence, Chi-square divergence, and Hellinger distance are special cases).

- The *pseudospherical divergence*, which is a transformation of the power divergence, arises in the dual solution of an alternative utility maximization problem.

- Kullback-Leibler divergence

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

- Power divergence

$$D_\beta^{\mathbf{P}}(\mathbf{p}\|\mathbf{q}) \equiv \frac{E_{\mathbf{p}}[(\mathbf{p}/\mathbf{q})^{\beta-1}] - 1}{\beta(\beta-1)}$$

- Pseudospherical divergence

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

And what's really really new…

- We introduce a novel parameterization of the power family of utility functions, with the following properties
  - It is indexed by a single parameter, β, which is the agent's *risk tolerance coefficient*, i.e., his risk tolerance ("r") at wealth x is $r = \beta x$.
  - It is normalized to a utility of zero and a marginal utility of 1 at $x = 0$, where x is interpreted as a deviation from the status quo.
  - This turns out to correspond exactly to the conventional parameterizations of the power and pseudospherical divergences.

This means his risk premium for a normally distributed asset with variance σ^2 is $\frac{1}{2}\sigma^2 / r$, i.e., $1/2r$ is his rate of substitution between mean and variance.

It looks like this:

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

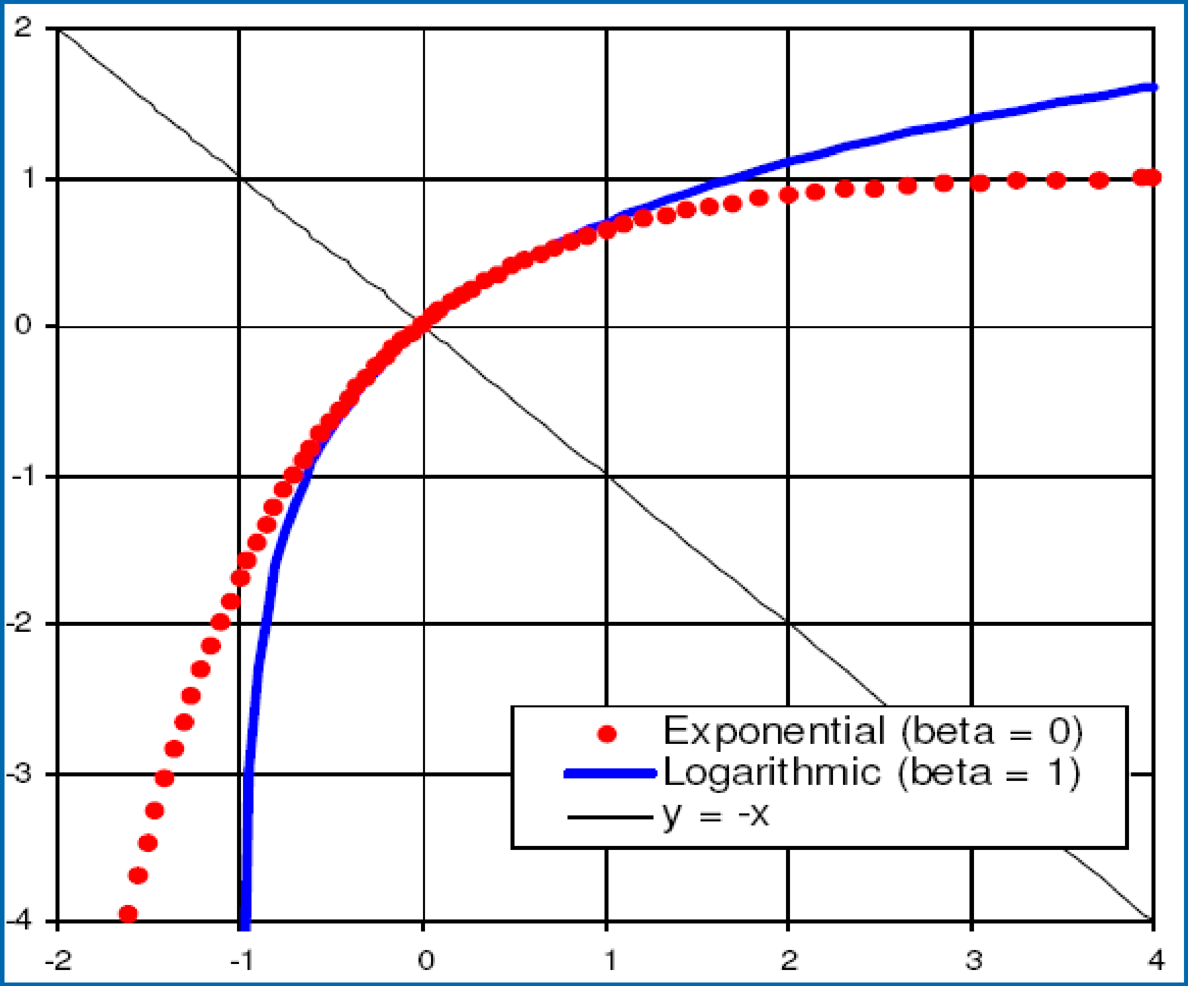

- Its graph passes through the origin with a slope of 1.
- The graph of u_β is the reflection of $u_{1-\beta}$ around the line $y = -x$.
- In particular, the exponential utility function (β=0) is the reflection of the log utility function (β=1) around $y = -x$
- The reciprocal utility function (β= ½) is its own reflection.

As opposed to the usual mess…

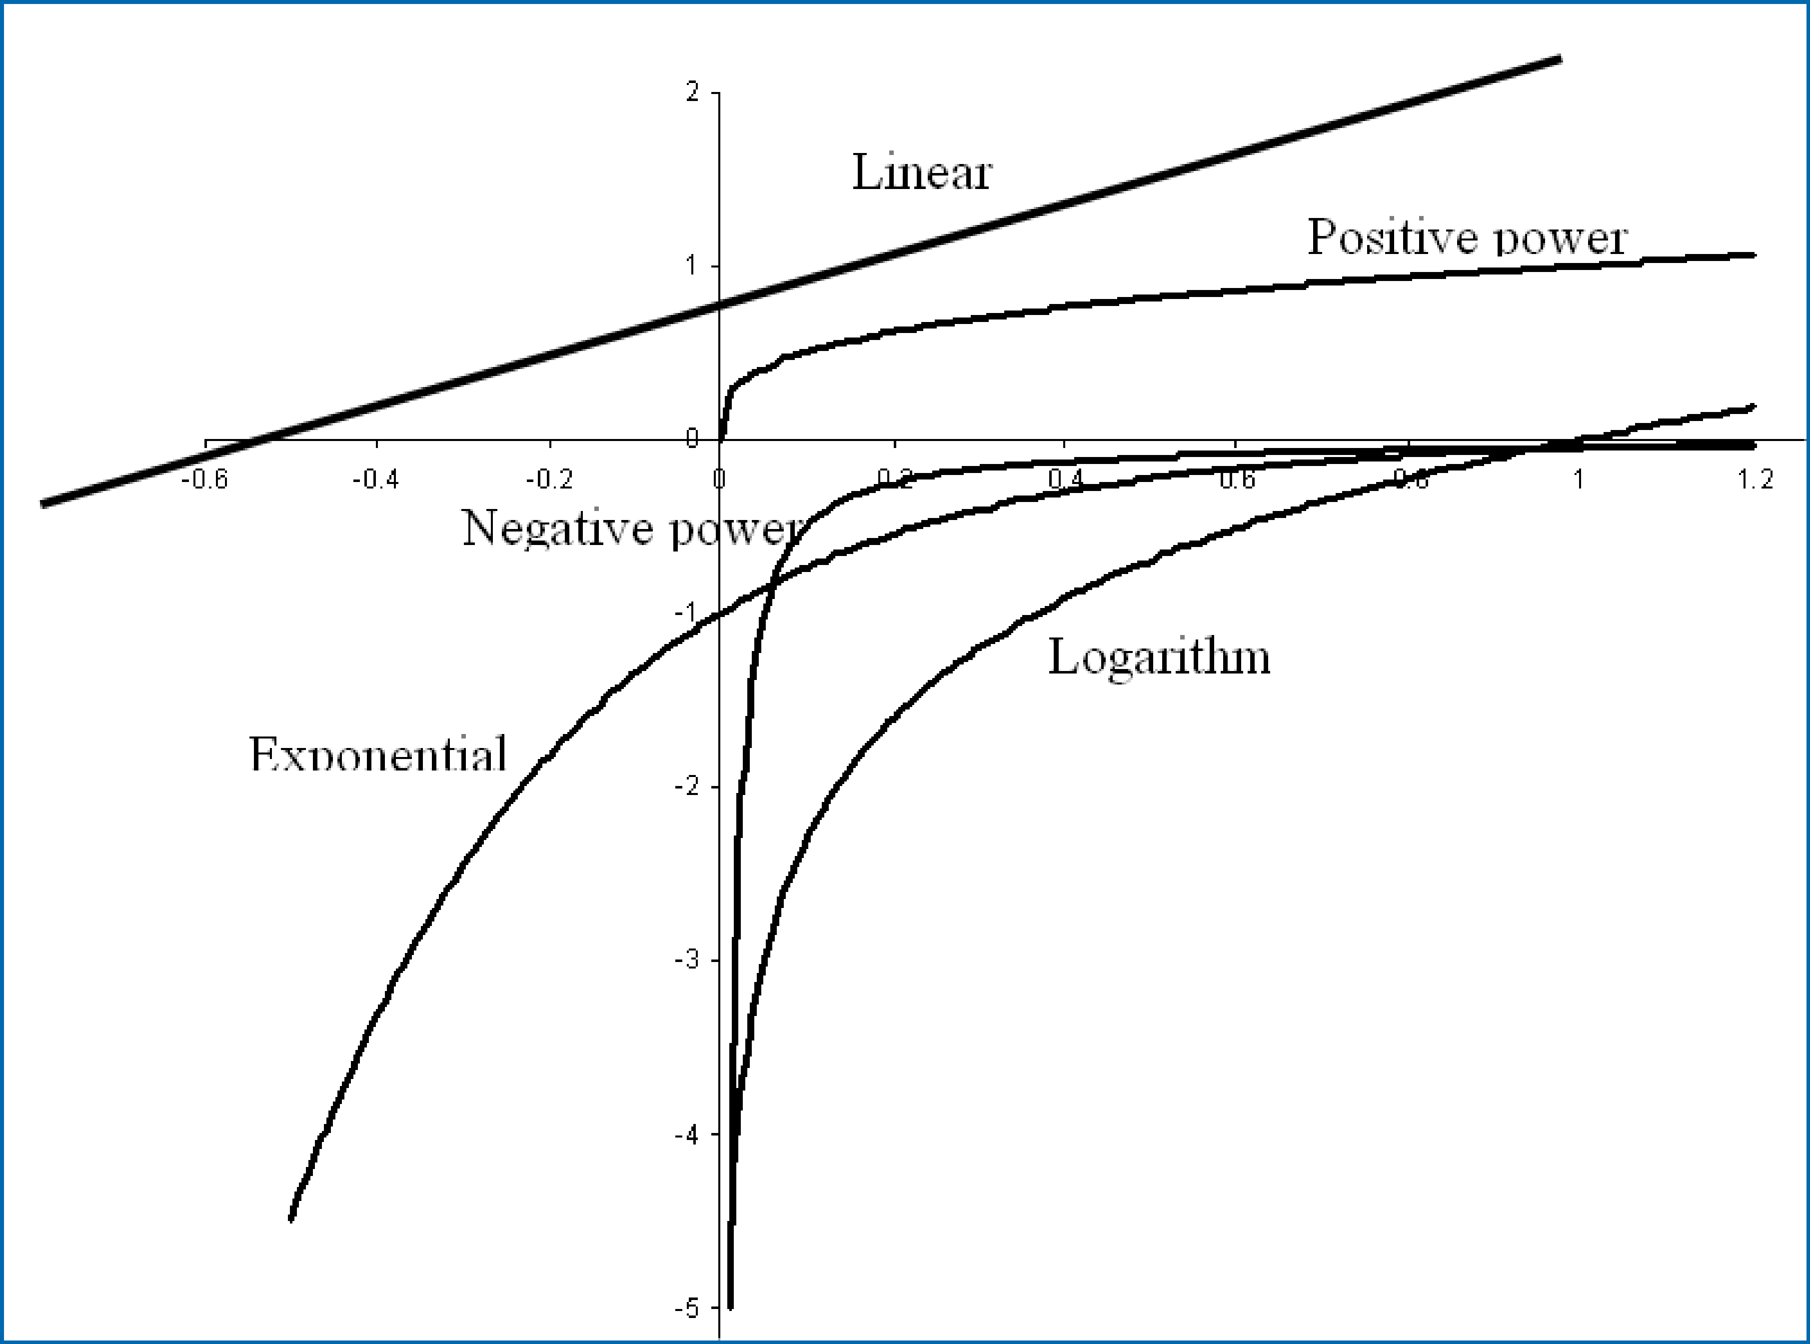

Special cases

$\beta = -1$	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	$u_2(x) = \sqrt{1+2x} - 1$