



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Risk & Sustainable Management Group

Risk & Uncertainty Program Working Paper: R06_1

Research supported by an Australian Research Council Federation Fellowship
http://www.arc.gov.au/grant_programs/discovery_federation.htm

Dual Approaches to the Analysis of Risk Aversion

Robert G. Chambers

Dept of Agricultural and Resource Economics, University of Maryland,
College Park

and

John Quiggin

Australian Research Council Federation Fellow, University of Queensland

Schools of Economics and Political Science
University of Queensland
Brisbane, 4072
rsmg@uq.edu.au
<http://www.uq.edu.au/economics/rsmg>



**THE UNIVERSITY
OF QUEENSLAND**
AUSTRALIA

Dual Approaches to the Analysis of Risk Aversion

Robert G. Chambers and John Quiggin

June 15, 2006

DUAL APPROACHES TO THE ANALYSIS OF RISK AVERSION

Dual approaches have proved their value in many areas of economic analysis. Until recently, however, they have been virtually ignored in the analysis of choice under uncertainty. Instead reliance has been placed almost exclusively on primal methods, and, in particular, on the expected-utility model. Perhaps the best explanation of the expected-utility model's continuing endurance, in spite of its well-known weaknesses, is its ability to yield predictions about economic behavior. Much of its 'predictive bite', however, comes from what many regard as its Achilles heel, the independence axiom and its associated structural property — additive separability. A prominent challenge for choice theory is to develop a model that retains this predictive bite while dispensing with the unpleasant characteristics associated with the independence axiom.

Additively separable preferences were discarded as a reasonable representation of preferences in standard consumer theory decades ago. Instead, when extra precision is required, reliance is usually placed on direct assumptions about the nature of the decisionmaker's preference map. In particular, notions of translation homotheticity, homotheticity and quasi-homotheticity have proven very useful in both empirical and theoretical analyses. These restrictions have percolated into expected-utility theory, albeit in a disguised form, as the notions of constant absolute risk aversion (CARA), constant relative risk aversion (CRRA), and linear risk tolerance (LRT).

In a series of classic papers (Yaari, 1965; Yaari, 1969; Peleg and Yaari, 1975), Yaari and Peleg observe that choice under uncertainty, just like choice under certainty, can be reasonably modelled in terms of convex preference sets and their supporting hyperplanes. Convex preferences and supporting hyperplanes are the natural stuff of modern duality theory.

In this paper, we exploit these observations by presenting a dual formulation of choice under uncertainty based on a few simple assumptions about preferences, namely, continuity, monotonicity and convexity of preference sets. Particular emphasis is given to showing that the additive separability restriction, key to expected-utility theory, on preferences can be dropped with little loss of analytic power for a broad class of choice problems.

The analysis commences by representing convex preference sets over uncertain outcomes in terms of the translation function, which was originally developed in the theories of inequality measurement and consumer preferences under certainty (Blackorby and Donaldson, 1980; Luenberger, 1992), and its concave conjugate, which we refer to as the expected-value function. Subjective probabilities are interpreted as normalized supporting hyperplanes in the neighborhood of the sure thing, and more generally, for risk-averse individuals, marginal rates of substitution between state-contingent

incomes are interpreted as relative ‘risk-neutral’ probabilities.

We next consider risk aversion, beginning with Yaari’s (1969) concept of risk aversion. A dual definition of risk aversion with respect to a probability vector is then offered. This dual definition of risk aversion yields dual versions of the Pratt–Arrow absolute and relative risk premiums as functions of the probabilities. For an individual risk-averse with respect to a given probability vector, these dual risk premiums take their maximum values (zero and one) at that vector, just as the corresponding primal measures are minimized at certainty.

We then examine the concepts of CARA, CRRA, and LRT in a dual framework. These concepts are interpreted as homotheticity properties, and each is shown to be characterized by an invariance condition on the risk-neutral probabilities and on the dual risk premiums. We illustrate the power of the dual approach in analyzing preferences over uncertain outcomes by showing that linear risk tolerance is simply characterized as quasi-homotheticity. And, even though, for general quasi-concave LRT preferences, there typically do not exist closed-form preference functionals, the dual formulation offers a simple characterization.

Next, a dual analysis of the class of constant risk averse preferences studied by Safra and Segal (1998) is provided. The associated expected-value function is derived and shown to imply that the only quasi-concave preference structures belonging to this class are the maxmin expected value (MMEV) preferences identified by Safra and Segal (1998). Our demonstration in terms of risk-neutral probabilities and the expected-value function leads to the further observation that the plunging behavior observed by Yaari (1987) for his dual preference structure is characteristic of the entire class of constant-risk-averse, quasi-concave preferences.

Our methods are then applied to a generic convex choice problem. Equilibrium conditions for such problems are characterized by a preference analogue to the Peleg–Yaari (1975) efficient set result, and comparative static results for LRT and constant risk averse preferences are presented.

1 Notation and Basic Concepts

For a proper concave function $f : \mathbb{R}^S \rightarrow \mathbb{R}$, its *superdifferential* at \mathbf{x} is the closed, convex set:

$$\partial f(\mathbf{x}) = \{\mathbf{v} \in \mathbb{R}^S : f(\mathbf{x}) + \mathbf{v}(\mathbf{z} - \mathbf{x}) \geq f(\mathbf{z}) \text{ for all } \mathbf{z}\}. \quad (1)$$

The elements of $\partial f(\mathbf{x})$ are referred to as *supergradients*. If f is differentiable at \mathbf{x} , $\partial f(\mathbf{x})$ is a singleton. If $\partial f(\mathbf{x})$ is a singleton, f is differentiable at \mathbf{x} (Rockafellar, 1970). $\frac{\partial^+}{\partial x}$ denotes the right-hand partial derivative with respect to x .

We consider preferences over random variables represented as mappings from a state space Ω to an outcome space $Y \subseteq \mathbb{R}$.¹ We refer to the outcomes as income. Our focus is on the case where Ω is a finite set $\{1, \dots, S\}$, and the space of random variables is $Y^S \subseteq \mathbb{R}^S$. The unit vector is denoted $\mathbf{1} = (1, 1, \dots, 1)$, and $\mathcal{P} \subset \mathbb{R}_{++}^S$ denotes the probability simplex. Define \mathbf{e}_i as the i -th row of the $S \times S$ identity matrix, $\mathbf{e}_i = (0, \dots, 1, 0, \dots, 0)$.

Preferences over state-contingent incomes are given by a continuous, nondecreasing, and quasi-concave certainty equivalent $e : Y^S \rightarrow \mathbb{R}$. Quasi-concavity ensures that the the preference mapping's least-as-good sets

$$V(e) = \{\mathbf{y} : e(\mathbf{y}) \geq e\}$$

are convex, and that the individual is averse to risk in the sense of Yaari (1969).²

2 The Translation Function and the Expected-Value Function

The *translation function*, $B : \mathbb{R} \times Y^S \rightarrow \mathbb{R}$, is defined:

$$B(e, \mathbf{y}) = \max\{\beta \in \mathbb{R} : \mathbf{y} - \beta \mathbf{1} \in V(e)\}$$

if $\mathbf{y} - \beta \mathbf{1} \in V(e)$ for some β , and $-\infty$ otherwise (Blackorby and Donaldson, 1980; Luenberger, 1992).³ The properties of $B(e, \mathbf{y})$ are well known (Blackorby and Donaldson, 1980; Luenberger, 1992; Chambers, Chung, and Färe, 1996), and are summarized for later use in the following lemma:

Lemma 1 *$B(e, \mathbf{y})$ satisfies:*

- a) *$B(e, \mathbf{y})$ is nonincreasing in e and nondecreasing and concave in \mathbf{y} ;*
- b) *$B(e, \mathbf{y} + \alpha \mathbf{1}) = B(e, \mathbf{y}) + \alpha$, $\alpha \in \mathbb{R}$ (the translation property);*
- c) *$B(e, \mathbf{y}) \geq 0 \Leftrightarrow \mathbf{y} \in V(e)$;*
- d) *$B(e, \mathbf{y})$ is jointly continuous in \mathbf{y} and e in the interior of the region $\mathbb{R} \times Y^S$ where $B(e, \mathbf{y})$ is finite.*

¹We allow the outcome space to include negative reals so that our preference representation is general enough to cover instances where the random variables studied may involve actual losses (for example, lost bets) to the decisionmaker. In many decision problems, however, the random variable of interest may only involve nonnegative outcomes. In this case, with largely technical modifications to the arguments, most of the results that have been presented below follow with little change.

²The original version of this paper defined preferences in terms of an ordinal, nondecreasing, continuous and quasi-concave preference function and induced the certainty equivalent from its least-as-good sets. We thank an anonymous reviewer for suggesting the approach taken in the current version.

³The translation function is a special case of the benefit function defined by Luenberger (1992).

We refer to the concave conjugate of the translation function, $B(e, \mathbf{y})$, as the *expected-value function* $E : \mathcal{P} \times \mathfrak{R} \rightarrow \mathfrak{R}$. It is defined by

$$E(\boldsymbol{\pi}, e) = \inf_{\mathbf{y}} \{ \boldsymbol{\pi} \mathbf{y} - B(e, \mathbf{y}) \} \quad \boldsymbol{\pi} \in \mathcal{P}.$$

Chambers (2001) shows that, as a consequence of Lemma 1.b, if $\mathbf{y}(\boldsymbol{\pi}, e) \in \arg \inf \{ \boldsymbol{\pi} \mathbf{y} - B(e, \mathbf{y}) \}$, then $\mathbf{y}(\boldsymbol{\pi}, e) + \delta \mathbf{1} \in \arg \inf \{ \boldsymbol{\pi} \mathbf{y} - B(e, \mathbf{y}) \}$ for $\delta \in \mathfrak{R}$. This indeterminacy in the optimizing set can be resolved by a convenient and familiar normalization. Because, for $\boldsymbol{\pi} \in \mathcal{P}$, $\boldsymbol{\pi} \mathbf{y} - B(e, \mathbf{y}) = \boldsymbol{\pi} [\mathbf{y} - B(e, \mathbf{y}) \mathbf{1}]$, and because Lemma 1.c implies $B(e, \mathbf{y} - B(e, \mathbf{y}) \mathbf{1}) \geq 0$, the expected-value function is equivalently expressed as:

$$\begin{aligned} E(\boldsymbol{\pi}, e) &= \inf_{\mathbf{y}} \{ \boldsymbol{\pi} \mathbf{y} : B(e, \mathbf{y}) \geq 0 \} \quad \boldsymbol{\pi} \in \mathcal{P} \\ &= \inf_{\mathbf{y}} \{ \boldsymbol{\pi} \mathbf{y} : \mathbf{y} \in V(e) \} \end{aligned}$$

if there exists some $\mathbf{y} \in V(e)$, and ∞ otherwise. This representation of $E(\boldsymbol{\pi}, e)$ justifies its name because it shows that $E(\boldsymbol{\pi}, e)$ gives the lowest expected value of any random variable that is consistent with a certainty equivalent of e . It also shows that the expected-value function has an alternative interpretation as the expenditure function for $V(e)$ in the state-claim prices $\boldsymbol{\pi}$.

If $V(e)$ is nonempty, $B(e, \mathbf{y})$ is a continuous and nondecreasing proper concave function, and thus $E(\boldsymbol{\pi}, e)$ is a closed, proper concave⁴ function nondecreasing on \mathcal{P} (Theorem 12.2, Rockafellar, 1970). It is also continuous and nondecreasing in e in the region where it is finite. And because $e(e \mathbf{1}) = e$, $E(\boldsymbol{\pi}, e) \leq e$.

Figure 1 illustrates the relationship between the expected value function and the certainty equivalent. Because the expected-value function is an expenditure function in terms of the Arrow–Debreu state-claim prices $\boldsymbol{\pi}$, the difference between $e(\mathbf{y})$ and $E(\boldsymbol{\pi}, e)$ measures the cost savings in achieving e that can be realized by operating in a complete Arrow–Debreu contingent claims economy at state-claims prices $\boldsymbol{\pi}$.

By basic results on conjugate duality (Theorem 12.2, Rockafellar, 1970), the translation function can be reconstructed from the expected-value function by applying the following conjugacy relationship

$$B(e, \mathbf{y}) = \inf_{\boldsymbol{\pi} \in \mathcal{P}} \{ \boldsymbol{\pi} \mathbf{y} - E(\boldsymbol{\pi}, e) \}.$$

⁴A concave function, $g(x)$, is proper if there is at least one x such that $g(x) > -\infty$, and $g(x) < \infty$ for all x . A concave function is closed if and only if it is upper semi-continuous (Rockafellar, 1970, p. 52).

A well-known implication of the conjugacy of the translation function and the expected-value function (Corollary 23.5.1, Rockafellar, 1970) is

$$\boldsymbol{\pi} \in \partial B(e, \mathbf{y}) \iff \mathbf{y} \in \partial E(\boldsymbol{\pi}, e) \quad (2)$$

in the relative interior of their domains.⁵ Expression (2) is a general statement of Shephard's Lemma, familiar from standard consumer and producer theory, for superdifferentiable structures. More formally, by (2), in the relative interior of their domains

$$\mathbf{y}(\boldsymbol{\pi}, e) \in \arg \inf \{ \boldsymbol{\pi} \mathbf{y} - B(e, \mathbf{y}) \} \Rightarrow \mathbf{y}(\boldsymbol{\pi}, e) \in \partial E(\boldsymbol{\pi}, e),$$

because

$$\begin{aligned} \boldsymbol{\pi} \mathbf{y}(\boldsymbol{\pi}, e) - B(e, \mathbf{y}(\boldsymbol{\pi}, e)) &\leq \boldsymbol{\pi} \mathbf{y} - B(e, \mathbf{y}), \quad \text{for all } \mathbf{y} \\ &\Downarrow \\ \boldsymbol{\pi} &\in \partial B(e, \mathbf{y}(\boldsymbol{\pi}, e)) \\ &\Downarrow \\ \mathbf{y}(\boldsymbol{\pi}, e) &\in \partial E(\boldsymbol{\pi}, e), \end{aligned}$$

where the first \Rightarrow follows by the definition of the superdifferential, and the second by (2). By a parallel argument, $\mathbf{p}(e, \mathbf{y}) \in \arg \inf \{ \boldsymbol{\pi} \mathbf{y} - E(\boldsymbol{\pi}, e) \} \Rightarrow \mathbf{p}(e, \mathbf{y}) \in \partial B(e, \mathbf{y})$. Thus, the supergradients of the translation function are interpretable as compensated state-claim-price-dependent demand functions for the state-claim vector \mathbf{y} .

2.1 Risk-neutral probabilities

Because the translation function and the expected-value function form a conjugate pair, they offer a natural method for defining and generating subjective notions of probability in terms of their superdifferentials. Because there is no requirement for smoothness, this allows for the analysis of both first-order and second-order risk aversion (Epstein and Zinn 1990; Segal and Spivak 1990; Machina, 2001).

Yaari (1969) identifies subjective probabilities with a supporting hyperplane to $V(e)$ along the sure-thing vector. Any such supporting hyperplane must belong to $\partial B(e, e\mathbf{1})$ and is interpretable as a vector of state-claim prices which support the constant portfolio, $e\mathbf{1}$. Nau (2001) has confirmed the

⁵Here, as elsewhere in the paper, these are understood to be the superdifferential of B in terms of \mathbf{y} and the superdifferential of E in terms of $\boldsymbol{\pi}$.

importance of considering supporting hyperplanes for the indifference set away from the sure-thing vector by noting that these correspond to the ‘risk-neutral probabilities’ central to the construction of consumption-based pricing kernels (state-price densities) in finance theory.

We start our analysis of these issues by stating a lemma, which ensures that the superdifferential of B has the convenient property that its elements belong to the unit simplex and that it is invariant to translations in the direction of the constant portfolio, $e\mathbf{1}$.

Lemma 2 *Let $\mathbf{p}(e, \mathbf{y}) \in \partial B(e, \mathbf{y})$. Then $\sum_{s \in \Omega} p_s(e, \mathbf{y}) = 1$, and $\mathbf{p}(e, \mathbf{y} + \delta \mathbf{1}) = \mathbf{p}(e, \mathbf{y})$, for all $\delta \in \Re$.*

Proof By Lemma 1.b, $B(e, \mathbf{y} + \delta \mathbf{1}) = B(e, \mathbf{y}) + \delta$. Let $\mathbf{v} \in \partial B(e, \mathbf{y})$ and $\mathbf{z} = \mathbf{y} + \delta \mathbf{1}$, $\mathbf{z}^* = \mathbf{y} - \delta \mathbf{1}$, then

$$\begin{aligned} B(e, \mathbf{y}) + \mathbf{v}(\mathbf{z} - \mathbf{y}) &\geq B(e, \mathbf{z}) \\ B(e, \mathbf{y}) + \mathbf{v}(\mathbf{z}^* - \mathbf{y}) &\geq B(e, \mathbf{z}^*) \end{aligned}$$

which implies $\delta \mathbf{v} \mathbf{1} \geq \delta \geq \delta \mathbf{v} \mathbf{1}$. For the second part,

$$\begin{aligned} \partial B(e, \mathbf{y} + \delta \mathbf{1}, \mathbf{1}) &= \left\{ \mathbf{v} : B(e, \mathbf{y} + \delta \mathbf{1}) + \mathbf{v}(\mathbf{z} - \mathbf{y} - \delta \mathbf{1}) \right. \\ &\quad \left. \geq B(e, \mathbf{z}) \text{ for all } \mathbf{z} \right\} \\ &= \{ \mathbf{v} : B(e, \mathbf{y}) + \mathbf{v}(\mathbf{z} - \mathbf{y}) \geq B(e, \mathbf{z}) \text{ for all } \mathbf{z} \} \\ &= \partial B(e, \mathbf{y}), \end{aligned}$$

where the second equality follows by Lemma 1 and the first part of this lemma. ■

It now seems natural to refer to the elements of any vector $\mathbf{p}(e, \mathbf{y}) \in \partial B(e, \mathbf{y}) \subset \Re_+^S$ for mnemonic purposes as *risk-neutral probabilities* for the certainty equivalent e . If the translation function is differentiable, these probabilities are unique and given by the gradient, $\nabla B(e, \mathbf{y})$. We define the set of *risk-neutral probabilities*, $\boldsymbol{\pi}(\mathbf{y}) \subset \Re_+^S$, which correspond to the supporting hyperplanes for the indifference set, by

$$\boldsymbol{\pi}(\mathbf{y}) = \partial B(e(\mathbf{y}), \mathbf{y}) = \mathbf{p}(e(\mathbf{y}), \mathbf{y}).$$

Thus, $\boldsymbol{\pi}(\mathbf{y})$ is intuitively analogous to an inverse demand correspondence for the Arrow commodities. When preferences are smooth (differentiable), $\boldsymbol{\pi}(\mathbf{y})$ is a singleton.

These risk-neutral probabilities are the preference counterpart to the shadow probabilities developed by Peleg and Yaari (1975), who consider, for a given choice set C , the probabilities that would lead a risk-neutral decision-maker to choose \mathbf{y} as the optimal element of C . We return to this observation in our analysis of choice over convex choice sets.

Following Yaari (1969), the risk-neutral probabilities associated with outcomes along the sure-thing vector are of particular interest. Because $e(e\mathbf{1}) = e$, $E(\pi, e) \leq e$. And, because preferences are quasi-concave, $\pi \in \partial B(e, e\mathbf{1}) \iff E(\pi, e) = e$. We, thus, define the set of *subjective probabilities* $\pi(\mathbf{1}) \subset \mathbb{R}_+^S$ as

$$\pi(\mathbf{1}) = \cap_e \{\partial B(e, e\mathbf{1})\}.$$

Typically, we shall assume that $\pi(\mathbf{1})$ is non-empty, although in general it need not be. In the case of smooth preferences, non-emptiness implies that indifference surfaces are parallel along the sure-thing vector (a form of ray homotheticity). The set will be empty, however, if there is any systematic tendency for the indifference surfaces to ‘tilt’ as one moves out the sure-thing vector. It is easy to see that this can happen for state-dependent preferences implying, for example, that these subjective probabilities are not the ‘probabilities’ associated with the state-dependent expected-utility model. The set of subjective probabilities satisfies $\pi(\mathbf{1}) = \cap_e \arg \sup_{\pi \in \mathcal{P}} \{E(\pi, e) - e\}$.

Example For an expected utility maximizer with subjective probabilities π , $\{\pi\} = \partial B(e, e\mathbf{1}) \forall e$.

3 Risk aversion

Yaari’s (1969) approach to the definition of risk aversion was first to define concepts of risk neutrality and comparative risk aversion, and then to derive a definition of risk aversion by saying that any decisionmaker who was more risk averse than a ‘risk neutral’ decisionmaker was risk averse. This neatly allows the treatment of concepts of ‘more risk averse’ and ‘decreasing risk aversion’ in a common framework. The standard risk-neutral normalization is the class of preferences which evaluate stochastic outcomes only in terms of their expected outcomes. Karni (1985) and others have criticized this normalization, but in what follows we shall adopt it as the norm in defining risk aversion.

More formally, we have upon recognizing that $V(e)$ corresponds to Yaari’s (1965) acceptance set for the wealth level e :

Definition 1 *A is more risk-averse in the Yaari sense than B if for all e , $V^A(e) \subseteq V^B(e)$.*

It then follows naturally from this definition that a decisionmaker can be said to be *risk-averse* for the probability vector π^0 if for all e

$$V(e) \subseteq \{\mathbf{y} : \pi^0 \mathbf{y} \geq e\}.$$

The definition of risk aversion requires that an individual can be risk averse with respect to π^0 only if $\pi^0 \in \pi(\mathbf{1})$. The definition of risk aversion implies that an individual is risk-averse with respect to π^0 if, from an initial position of certainty represented by some $e\mathbf{1}$, he rejects all bets \mathbf{z} that are fair in the sense that $\pi^0\mathbf{z} = 0$ and, *a fortiori*, all bets that are unfavorable in the sense that $\pi^0\mathbf{z} < 0$. In the case where $\pi(\mathbf{1})$ is empty, there exists no probability vector with this property for all e .

Dually, we can define a notion of relative riskiness of probability vectors and then deduce a notion of risk aversion with respect to a particular probability vector:

Definition 2 *For a given expected value function E , π is less risky than π' at e , denoted $\pi \preceq_e \pi'$, if $E(\pi', e) \geq E(\pi, e)$.*

Intuitively, $\pi \preceq_e \pi'$ implies that π' is ‘closer’ to the set of maximally risk-neutral probabilities (those for which $E(\pi, e) = e$) than π . Thus, the more risky are the probabilities, the ‘closer’ will be $\mathbf{y} \in \partial E(\pi, e)$ to the constant portfolio, $e\mathbf{1}$. The riskiest probabilities are the supporting state-claim prices for the constant portfolio, $e\mathbf{1}$.

Lemma 3 *An individual is risk-averse with respect to probabilities π^0 if and only if $E(\pi^0, e) = e \quad \forall e$, and $\pi \preceq_e \pi^0$ for all (π, e) .*

There are several immediate consequences of these definitions. We summarize them in the following theorem:

Theorem 1 *The following are equivalent:*

- (a) *A is more risk averse than B;*
- (b) *$B^A(e, \mathbf{y}) \leq B^B(e, \mathbf{y})$ for all \mathbf{y} and e ;*
- (c) *$E^A(\pi, e) \geq E^B(\pi, e)$ for all π and e ; and*
- (d) *for all \mathbf{y} , $e^A(\mathbf{y}) \leq e^B(\mathbf{y})$.*

Moreover, if A is more risk-averse than B, and B is risk-averse with respect to probabilities π^0 , so is A.

Proof *(a) \Rightarrow (c) is immediate. (c) \Rightarrow (b) follows by applying $E^A(\pi, e) \geq E^B(\pi, e)$ for all π and e in the conjugacy mapping. (b) \Rightarrow (d) follows because $e(\mathbf{y})$ is determined by $\max\{e : B(e, \mathbf{y}) \geq 0\}$. (d) \Rightarrow (a) is immediate from the definition of V . The second part of the theorem is trivial.■*

An easy corollary to part (d) is the well-known result that for individuals A and B with expected-utility preferences, A is more risk averse than B if and only if A 's *ex post* utility function is a concave transformation of B 's.

3.1 Dual Measures of risk aversion

We introduce an absolute and a relative dual measure of risk aversion. The *dual absolute risk premium* is

$$a(\boldsymbol{\pi}, e) = E(\boldsymbol{\pi}, e) - e,$$

and the *dual relative risk premium* (defined only for $e > 0$) is

$$r(\boldsymbol{\pi}, e) = \frac{E(\boldsymbol{\pi}, e)}{e}.$$

These risk premiums provide exact indexes of the cost saving that a decisionmaker can realize in achieving e by operating in a complete contingent claims market at $\boldsymbol{\pi}$. Notice that $a(\boldsymbol{\pi}, e) \leq 0$ and $r(\boldsymbol{\pi}, e) \leq 1$. Moreover, because E is concave in $\boldsymbol{\pi}$, so are a and r . Thus, they achieve their maximal values at the maximally risky $\boldsymbol{\pi}$. These two measures are directly related in the case $e > 0$ by $a(\boldsymbol{\pi}, e) = e(r(\boldsymbol{\pi}, e) - 1)$.

Lemma 4 *The following conditions are equivalent:*

- (1) A is more risk-averse than B ;
- (2) $a^A(\boldsymbol{\pi}, e) \geq a^B(\boldsymbol{\pi}, e) \quad \forall \boldsymbol{\pi}, e$; and
- (3) for all $e > 0$ $r^A(\boldsymbol{\pi}, e) \geq r^B(\boldsymbol{\pi}, e) \quad \forall \boldsymbol{\pi}$.

An individual is risk-averse with respect to probabilities $\boldsymbol{\pi}^0$ if and only if $a(\boldsymbol{\pi}^0, e) = 0$ and $r(\boldsymbol{\pi}^0, e) = 1$.

Example If preferences are risk-neutral with respect to $\boldsymbol{\pi}^0$,⁶

$$a(\boldsymbol{\pi}, e) = \begin{cases} -\infty & \boldsymbol{\pi} \neq \boldsymbol{\pi}^0 \\ 0 & \boldsymbol{\pi} = \boldsymbol{\pi}^0 \end{cases}.$$

⁶Here, for the sake of a simplified notation, we set

$$\inf \{ \boldsymbol{\pi} \mathbf{y} : \mathbf{y} \in Y^S \} = -\infty.,$$

which is correct if $Y = \Re$. For more restrictive domains in the risk-neutral case

$$a(\boldsymbol{\pi}, e) = \begin{cases} \inf \{ (\boldsymbol{\pi} - \boldsymbol{\pi}^0) \mathbf{y} : \mathbf{y} \in Y^S \} & \boldsymbol{\pi} \neq \boldsymbol{\pi}^0 \\ 0 & \boldsymbol{\pi} = \boldsymbol{\pi}^0 \end{cases}.$$

The decisionmaker makes an unboundedly large saving by operating in a contingent claims market if the state-claim prices depart from his subjective probabilities. This reflects his willingness to take arbitrarily large short or long positions in the pursuit of expected return. For completely risk averse preferences,

$$e(\mathbf{y}) = \min \{y_1, y_2, \dots, y_S\},$$

$a(\boldsymbol{\pi}, e) = 0$, for all $\boldsymbol{\pi}$. Because the individual is completely risk averse, he realizes no cost savings by operating in a complete contingent claims market over holding e units of the riskless asset. The ability to take a short or long position is valueless to such a decisionmaker.

4 Constant Absolute and Relative Risk Aversion and Linear Risk Tolerance

Preferences exhibit CARA if, for all $\boldsymbol{\pi}$,

$$a(\boldsymbol{\pi}, e) = a(\boldsymbol{\pi}, e') \quad \text{all } e, e'.$$

In dual terms, this implies that the decisionmaker's absolute cost saving from operating in a complete contingent claims market only depends on the state-claims prices. Preferences exhibit CRRA if, for all $\boldsymbol{\pi}$,

$$r(\boldsymbol{\pi}, e) = r(\boldsymbol{\pi}, e') \quad \text{all } e, e' > 0,$$

and thus the relative cost saving from operating in a complete contingent claims market is independent of the level of e .

Our next result shows that these dual notions of CARA and CRRA are equivalent to the more familiar notions. It also characterizes the risk-neutral probabilities for both classes of preferences.

Theorem 2 *Preferences exhibit CARA if and only if $E(\boldsymbol{\pi}, e) = \hat{a}(\boldsymbol{\pi}) + e$, where $\hat{a}(\boldsymbol{\pi}) \leq 0$ is a closed, nondecreasing proper concave function, $B(e, \mathbf{y}) = B(0, \mathbf{y}) - e$, and $\boldsymbol{\pi}(\mathbf{y} + \beta \mathbf{1}) = \boldsymbol{\pi}(\mathbf{y})$, $\beta \in \mathbb{R}$. Preferences exhibit CRRA if and only if $E(\boldsymbol{\pi}, e) = \hat{r}(\boldsymbol{\pi}) e$ where $\hat{r}(\boldsymbol{\pi}) \leq 1$ is a closed proper concave function, $B(e, \mathbf{y}) = eB(1, \frac{\mathbf{y}}{e})$, and $\boldsymbol{\pi}(\mu \mathbf{y}) = \boldsymbol{\pi}(\mathbf{y})$, $\mu > 0$.*

Proof *The proof is for CARA. The proof for CRRA is parallel. By CARA $a(\boldsymbol{\pi}, e) = \hat{a}(\boldsymbol{\pi})$, with $\hat{a}(\boldsymbol{\pi}) \leq 0$ a nondecreasing, closed proper concave function by the properties of the expected-*

value function. Hence, $E(\pi, e) = \hat{a}(\pi) + e$. By conjugacy,

$$\begin{aligned} B(e, \mathbf{y}) &= \min_{\pi} \{\pi \mathbf{y} - \hat{a}(\pi)\} - e \\ &= B(0, \mathbf{y}) - e, \end{aligned}$$

where $B(0, \mathbf{y})$ is the concave conjugate of $\hat{a}(\pi)$. Because $B(e, \mathbf{y}) = B(0, \mathbf{y}) - e$, it follows that $\mathbf{p}(e, \mathbf{y}) = \mathbf{p}(0, \mathbf{y})$ for all \mathbf{y} . By the second part of Lemma 2, $\mathbf{p}(0, \mathbf{y} + \beta \mathbf{1}) = \partial B(0, \mathbf{y} + \beta \mathbf{1}) = \partial B(0, \mathbf{y}) = \mathbf{p}(0, \mathbf{y})$. Conversely, if $B(e, \mathbf{y}) = B(0, \mathbf{y}) - e$,

$$\begin{aligned} E(\pi, e) &= \min \{\pi \mathbf{y} - B(0, \mathbf{y}) + e\} \\ &= \min \{\pi \mathbf{y} - B(0, \mathbf{y})\} + e. \blacksquare \end{aligned}$$

Corollary 1 If preferences exhibit CARA $\pi \in \cap_e \{\partial B(e, e\mathbf{1})\} \iff \hat{a}(\pi) = 0$. If preferences exhibit CRRA $\pi \in \cap_e \{\partial B(e, e\mathbf{1}; \mathbf{1})\} \iff \hat{r}(\pi) = 1$. In both cases, $\pi(\mathbf{1})$ is nonempty.

A direct consequence of Theorem 2 is that for CARA preferences, $e(\mathbf{y}) = B(0, \mathbf{y})$. Thus, by Lemma 1.b, $e(\mathbf{y} + \beta \mathbf{1}) = e(\mathbf{y}) + \beta$. This is the standard primal definition of CARA for general preferences (Chambers and Quiggin, 2000). Hence, any ordinal transformation of the certainty equivalent must be translation homothetic (Blackorby and Donaldson, 1980; Chambers and Färe, 1998). Similarly, for CRRA preferences, $B(1, \frac{\mathbf{y}}{e}) = 0$, whence $e(\mu \mathbf{y}) = \mu e(\mathbf{y})$ $\mu > 0$ implying that any ordinal representation of the certainty equivalent is homothetic. By observing that under CARA $e(\mathbf{y}) = B(0, \mathbf{y})$ and using Lemma 1.a, one obtains:

Corollary If preferences exhibit CARA, $e(\mathbf{y})$ is concave in \mathbf{y} .

Example Expected utility preferences, risk-averse for the probabilities π^0 , exhibit CARA if and only if

$$e(\mathbf{y}) = -\frac{1}{r} \ln \left[\sum_s \pi_s^0 \exp(-r y_s) \right] = B(0; \mathbf{y}),$$

with $e(\mathbf{y} + \delta \mathbf{1}) = -\frac{1}{r} \ln [\sum_s \pi_s^0 \exp(-r(y_s + \delta))] = e(\mathbf{y}) + \delta$, and

$$E(\pi, e) = e - \frac{1}{r} \sum_s \pi_s \ln \left(\frac{\pi_s^0}{\pi_s} \right).$$

We use as our notion of decreasing absolute risk aversion that E be sub-additive in e and for decreasing relative risk aversion that E be sub-homogeneous in e .

Definition 3 *Preferences display decreasing absolute risk aversion (DARA) if for all π , $E(\pi, e + e^*) \leq E(\pi, e) + e^*$, $e^* > 0$.*

Definition 4 *Preferences display decreasing relative risk aversion (DRRA) if for $e > 0$ and all π , $E(\pi, \mu e) \leq \mu E(\pi, e)$, $\mu > 1$.*

Under DARA, $\frac{\partial^+}{\partial e} E(\pi, e) \leq 1$, while under DRRA, $\frac{\partial^+}{\partial \ln e} \ln E(\pi, e) \leq 1$. Thus, DARA requires that the marginal cost of increasing the certainty equivalent (the marginal utility of income) is always less (greater) than one. Hence, the π -weighted average of income effects across state-claims is never greater than one under DARA. (For CARA, all state-claim income effects are one.) DRRA implies that the marginal cost of increasing the certainty equivalent is always less than the average cost of the certainty equivalent. More familiarly, in terminology borrowed from basic firm theory, DRRA requires that the average cost of the certainty equivalent be increasing in e . (CRRA requires that the average cost of the certainty equivalent is constant in e and equals marginal cost.)

An immediate consequence of these definitions and Theorem 2 is that:

Corollary If preferences exhibit CRRA, they also exhibit DARA. If preferences exhibit CARA, they exhibit increasing relative risk aversion (IRRA).

Using as the primal definition of DARA that

$$e(\mathbf{y} + \delta \mathbf{1}) \geq e(\mathbf{y}) + \delta, \quad \delta > 0,$$

and as the primal definition of DRRA that

$$e(\mu \mathbf{y}) \geq \mu e(\mathbf{y}), \quad \mu > 1,$$

Chambers and Quiggin (2000) have derived a version of this Corollary for strictly quasi-concave primal preferences. The corollary, thus, weakens the requirement to quasi-concavity because, as we now establish, our dual definition of DARA and DRRA are equivalent to the primal definitions used by Chambers and Quiggin (2000).

Theorem 3 *Preferences display DARA if and only if for $\delta > 0$, $B(e + \delta, \mathbf{y}) \geq B(e, \mathbf{y}) - \delta$, and $e(\mathbf{y} + \delta \mathbf{1}) \geq e(\mathbf{y}) + \delta$. Preferences display DRRA if and only if for $e > 0$ $\mu > 1$, $B(\mu e, \mathbf{y}) \geq \mu B\left(e, \frac{\mathbf{y}}{\mu}\right)$ and $e(\mu \mathbf{y}) \geq \mu e(\mathbf{y})$.*

Proof The proof is for DARA. By DARA

$$\begin{aligned} B(e + e^*, \mathbf{y}) &= \inf \{ \boldsymbol{\pi} \mathbf{y} - E(\boldsymbol{\pi}, e + e^*) \} \\ &\geq \inf \{ \boldsymbol{\pi} \mathbf{y} - E(\boldsymbol{\pi}, e) - e^* \} \\ &= B(e, \mathbf{y}) - e^*. \end{aligned}$$

Now apply Lemma 1.c. The converse follows by duality. The proof for DRRA is parallel. ■

Because CRRA corresponds to homotheticity and CARA corresponds to translation homotheticity, it is natural to speculate that the class of quasi-homothetic preferences, which contains both CRRA and CARA preferences as subsets, will prove useful for choice over uncertain prospects. Quasi-homothetic preferences possess linear income-expansion paths (Gorman, 1953). In the expected-utility literature, this characteristic is associated with preferences that exhibit LRT (Brennan and Kraus, 1976; Milne, 1979), and for which two-fund spanning applies (Cass and Stiglitz, 1970). Therefore, we say that preferences exhibit LRT if E assumes the Gorman polar form:

$$E(\boldsymbol{\pi}, e) = E^0(\boldsymbol{\pi}) + E^1(\boldsymbol{\pi})e$$

where $E^0(\boldsymbol{\pi})$ and $E^1(\boldsymbol{\pi})$ are expected-value functions for least-as-good sets that are independent of the certainty equivalent and $E^1(\boldsymbol{\pi}) \geq 0$. CARA is the special case of LRT where $E^1(\boldsymbol{\pi}) = \boldsymbol{\pi} \mathbf{1} = 1$ for all $\boldsymbol{\pi}$, while CRRA is the special case of LRT where $E^0(\boldsymbol{\pi}) = \boldsymbol{\pi} \mathbf{0} = 0$ for all $\boldsymbol{\pi}$.

CRRA and CARA preferences are tractable in either their dual or their primal formulations. This partially explains their popularity in models based on primal representations of preferences, such as expected utility. Preferences exhibiting LRT are simply expressed in terms of $E(\boldsymbol{\pi}, e)$ or $V(e)$. Both $e(\mathbf{y})$ and B , however, assume very inconvenient forms for general LRT preferences.

It is well-known that dual to an expected-value function exhibiting LRT there must exist a $V(e)$ of the form $V(e) = V^0 + eV^1$, where V^0 is a least-as-good set dual to E^0 , and V^1 is a least-as-good set dual to E^1 . However, it is also well-known that quasi-homothetic preferences generally do not have a closed form certainty equivalent. The manifestation of this in terms of B is a special case of a result originally due to Chambers, Chung, and Färe (1996) in the producer context.

Theorem 4 (*Chambers, Chung, and Färe*) *Preferences exhibit LRT if and only if*

$$B(e, \mathbf{y}) = \sup \left\{ \min \left\{ B^0(\mathbf{y}^0), eB^1\left(\frac{\mathbf{y}^1}{e}\right) \right\} : \mathbf{y}^0 + \mathbf{y}^1 = \mathbf{y} \right\},$$

where B^0 is the translation function conjugate to E^0 , and B^1 is the translation function conjugate to E^1 .

Proof By LRT

$$\begin{aligned}
B(e; \mathbf{y}; \mathbf{1}) &= \sup \{ \beta : \mathbf{y} - \beta \mathbf{1} \in V^0 + eV^1 \} \\
&= \sup \{ \beta : \mathbf{y}^0 - \beta \mathbf{1} \in V^0, \mathbf{y}^1 - \beta \mathbf{1} \in eV^1 : \mathbf{y}^0 + \mathbf{y}^1 = \mathbf{y} \} \\
&= \sup \left\{ \min \left\{ B^0(\mathbf{y}^0), eB^1\left(\frac{\mathbf{y}^1}{e}\right) \right\} : \mathbf{y}^0 + \mathbf{y}^1 = \mathbf{y} \right\},
\end{aligned}$$

where the last equality follows by monotonicity of preferences. ■

It seems unlikely, therefore, that much information can be gleaned directly from examining $\partial B(e; \mathbf{y})$ for general LRT preferences. However, some things are apparent from the $E(\boldsymbol{\pi}, e)$ formulation. For example, if LRT preferences are risk-averse with respect to the probability vector, $\boldsymbol{\pi}^0$, then $E^0(\boldsymbol{\pi}^0) = 0$, $E^1(\boldsymbol{\pi}^0) = 1$ and

$$\begin{aligned}
0 &\geq E^0(\boldsymbol{\pi}), \\
1 &\geq E^1(\boldsymbol{\pi}),
\end{aligned}$$

$\boldsymbol{\pi} \in \mathcal{P}$. This allows us to conclude:

Theorem 5 *If LRT preferences are risk-averse with respect to a probability vector $\boldsymbol{\pi}^0$, they exhibit both IRRA and DARA for all $\boldsymbol{\pi} \in \mathcal{P}$.*

Proof $E(\boldsymbol{\pi}, e + e^*) = E(\boldsymbol{\pi}, e) + E^1(\boldsymbol{\pi}) e^*$, and $E(\boldsymbol{\pi}, \mu e) = \lambda E(\boldsymbol{\pi}, e) + (1 - \lambda) E^0(\boldsymbol{\pi})$. ■

Increased tractability can be obtained by imposing functional structure beyond LRT. For example, an important special case of LRT preferences are the *affinely homothetic* preferences (Milne, 1979, Färe and Lovell, 1984) given by, $E^0(\boldsymbol{\pi}) = \boldsymbol{\pi} \mathbf{v}$ $\mathbf{v} \in \mathbb{R}^S$. These preferences have linear expansion paths emanating from \mathbf{v} . In a standard consumer context, \mathbf{v} is usually interpreted as a vector of subsistence demands. Perhaps the best known member of the LRT class is the Stone–Geary utility structure, which underlies the linear-expenditure system. Expected-utility LRT preferences are also affinely homothetic (Milne, 1979). Thus, results for general LRT preferences also apply to the expected-utility subclass of LRT preferences.

Another special case, which has received relatively less attention in literature on portfolio choice, is the class of preferences that are translation homothetic in an arbitrary direction \mathbf{u} (Chambers and Färe, 1998). This class, which has played a role in the empirical modelling of labor demand and consumer preferences (Blackorby, Boyce, Russell, 1978; Dickinson, 1980), is defined by $E^1(\boldsymbol{\pi}) = \boldsymbol{\pi} \mathbf{u}$, where $\mathbf{u} \in \mathbb{R}^S$. CARA is the special case where $\mathbf{u} = \mathbf{1}$.

Preferences satisfying CARA, CRRA, and LRT can all be characterized in terms of the notion of demand rank for asset demands for individuals facing complete contingent claims markets. Demand rank corresponds to the dimension of the function space spanned by the individual's Engel curves in budget-share form (Lewbel, 1991). By Theorem 1 of Lewbel (1991), CRRA corresponds to a rank-one demand system, while LRT corresponds to a rank-two demand system. Using the general results of Lewbel and Perraudin (1995), this establishes that each of these preference structures satisfy the conditions for portfolio separation associated with the theory of mutual funds. Lewbel and Perraudin (1995) show that a necessary and sufficient condition for portfolio separation, with smooth preferences, is that $E(\boldsymbol{\pi}, e) = E'(\rho^1(\boldsymbol{\pi}), \dots, \rho^K(\boldsymbol{\pi}), e)$ where $K < S$.

Constant relative risk aversion, thus, implies that preferences have a dual representation in terms of a composite of the state-claims. The corresponding holdings of the respective state-claims per unit of real income are given by the gradient of $\hat{r}(\boldsymbol{\pi})$. Constant absolute risk aversion has a dual representation in terms of two such composites. One is degenerate and corresponds to the traditionally safe asset, $\mathbf{1}$. The holding of the degenerate composite is proportional to real wealth, while the holding of the other composite is independent of real wealth and only depends on the state-claim prices. It is this characteristic of CARA which yields the well-known result that changes in real wealth do not affect the individual's holding of the risky asset. LRT generalizes the rank-two case to allow the composite dependent on real wealth to be risky.

4.1 Constant Risk Aversion

Safra and Segal (1998) investigated the class of preferences exhibiting both CARA and CRRA. They refer to this class of preferences as *constant risk averse*. Among other results, they have demonstrated that the only class of quasi-concave preferences which can exhibit constant risk aversion are the MMEV class.

Quiggin and Chambers (1998), who do not impose quasi-concavity, show that preferences defined over a finite state space exhibit constant risk aversion if and only if

$$B(e, \mathbf{y}) = g(\mathbf{y} - \text{Min}\{y_1, \dots, y_S\}\mathbf{1}) + \text{Min}\{y_1, \dots, y_S\} - e,$$

where g is positively linearly homogeneous. Maxmin, linear mean-standard deviation, and risk-neutral preferences are all special cases of this preference structure. The expected value function for this class of preferences can be derived as

$$E(\boldsymbol{\pi}, e) = \inf_{\mathbf{y}} \{ \boldsymbol{\pi}\mathbf{y} - \text{Min}\{y_1, \dots, y_S\} - g(\mathbf{y} - \text{Min}\{y_1, \dots, y_S\}\mathbf{1}) \} + e$$

$$\begin{aligned}
&= \inf_{\mathbf{y}} \{ \pi (\mathbf{y} - \text{Min}\{y_1, \dots, y_S\} \mathbf{1}) - g(\mathbf{y} - \text{Min}\{y_1, \dots, y_S\} \mathbf{1}) \} + e \\
&= \inf_{\hat{\mathbf{y}}} \{ \pi \hat{\mathbf{y}} - g(\hat{\mathbf{y}}) \} + e.
\end{aligned}$$

Because g is positively linearly homogeneous, $\inf_{\hat{\mathbf{y}}} \{ \pi \hat{\mathbf{y}} - g(\hat{\mathbf{y}}) \}$ equals either 0 or $-\infty$. This observation and conjugacy leads to the following compact demonstration of the Safra and Segal (1998) result, and its extension to the associated dual structures.⁷

Theorem 6 (*Safra and Segal*): *Preferences exhibit constant risk aversion if and only if*

$$E(\pi, e) = \begin{cases} e & \pi \in \hat{\mathcal{P}} \\ -\infty & \pi \notin \hat{\mathcal{P}} \end{cases},$$

and $B(e, \mathbf{y}) = \inf \{ \pi \mathbf{y} : \pi \in \hat{\mathcal{P}} \} - e$, for $\hat{\mathcal{P}} \subseteq \mathcal{P}$ closed and convex.

Proof By Theorem 2, preferences exhibit CARA if and only if $E(\pi, e) = \hat{a}(\pi) + e$, where $\hat{a}(\pi) \leq 0$ is a closed, proper concave function. To satisfy CRRA, it further follows from Theorem 2 that $\mu \hat{a}(\pi) = \hat{a}(\pi) \quad \mu > 0$. There are three possibilities: either $\hat{a}(\pi) = 0$; $\hat{a}(\pi) = \infty$; or $\hat{a}(\pi) = -\infty$. If $\hat{a}(\pi) = \infty$, there is no \mathbf{y} such that $B(e, \mathbf{y}; \mathbf{1}) \geq 0$, and hence $V(e)$ is empty. If $\hat{a}(\pi) = -\infty$ for all π , preferences are not well defined, and that case is ruled out. The only closed, proper concave function remaining is

$$\hat{a}(\pi) = \begin{cases} 0 & \pi \in \hat{\mathcal{P}} \\ -\infty & \pi \notin \hat{\mathcal{P}} \end{cases},$$

for some $\hat{\mathcal{P}} \subseteq \mathcal{P}$ closed. This establishes necessity of the first part. By the conjugacy of the translation and expected-value functions:

$$B(e, \mathbf{y}; \mathbf{1}) = \inf_{\pi \in \mathcal{P}} \{ \pi \mathbf{y} - E(\pi, e) \}.$$

For all $\pi \notin \hat{\mathcal{P}}$, $\pi \mathbf{y} - E(\pi, e) = \infty$, and thus

$$B(e, \mathbf{y}) = \inf_{\pi} \{ \pi \mathbf{y} - E(\pi, e) : \pi \in \hat{\mathcal{P}} \} < \infty,$$

if it is to be finite. Conversely, by Corollary 4, the certainty equivalent for quasi-concave CARA preferences is concave in \mathbf{y} . CRRA requires positive linear homogeneity of the certainty

⁷Here again we set $\inf \{ \pi \mathbf{y} : \mathbf{y} \in Y^S \} = -\infty$ for the sake of a streamlined proof. Using dual methods from the standard consumer problem, the statement of the result and the proof can be adapted to more restrictive domains at the expense of more notational complexity.

equivalent in \mathbf{y} . Thus, the certainty equivalent must be superlinear as a function of \mathbf{y} , and it must be the lower support function for some closed convex set (Rockafellar, 1970), whence

$$e(\mathbf{y}) = \inf \{ \boldsymbol{\pi} \mathbf{y} : \boldsymbol{\pi} \in \mathcal{P} \}.$$

The remainder is trivial. ■

Besides exhaustively characterizing the class of constant risk averse preferences, Theorem 6 has an interesting consequence for portfolio theory.

Corollary 2 *Preferences exhibit constant risk aversion if and only if either $\partial E(\boldsymbol{\pi}, e) = e\mathbf{1}$ or $\partial E(\boldsymbol{\pi}, e)$ is undefined.*

This corollary generalizes Yaari's (1987) observation that preferences in his dual model display 'plunging' behavior. That is, either the individual will reject a given risk entirely and adopt a non-stochastic portfolio, or he will accept an amount of the risk that is either unbounded or fixed by the constraints of the choice problem. Corollary 2 establishes the more general result that plunging behavior characterizes the entire class of quasi-concave, constant risk averse preferences.

5 An Application to Convex Choice Sets

Following Peleg and Yaari (1975), we consider an individual faced with a closed, bounded, convex choice set $Y \subseteq \mathbb{R}^S$. Such choice problems may arise, for example, from the standard portfolio choice problem, the production decisions of a firm under uncertainty, or as an investment allocation problem with a nonlinear but appropriately convex tax structure. We endow Y with the properties that $\mathbf{0} \in Y$ and $Y \cap \mathbb{R}_{++}^S \neq \emptyset$.

The decisionmaker's choice problem is $\max_{\mathbf{y}} \{e(\mathbf{y}) : \mathbf{y} \in Y\}$. Upon defining

$$R(\boldsymbol{\pi}, Y) = \max \{ \boldsymbol{\pi} \mathbf{y} : \mathbf{y} \in Y \},$$

and by restricting attention to the region where E is finite (see below for more on this assumption), given state-claim prices, an equivalent dual formulation of the individual's choice problem is

$$\max_e \{e : E(\boldsymbol{\pi}, e) \leq R(\boldsymbol{\pi}, Y)\}.$$

$R(\boldsymbol{\pi}, Y)$ is the revenue function dual to Y that is associated with the state-claim prices $\boldsymbol{\pi}$. Observe that

$$R(\boldsymbol{\pi}, Y + \delta \mathbf{u}) = R(\boldsymbol{\pi}, Y) + \delta \boldsymbol{\pi} \mathbf{u},$$

$$R(\pi, \mu Y) = \mu R(\pi, Y), \quad \mu > 0.$$

The optimization problem applies both when predetermined state-claim prices exist, as they would, for example, in the presence of complete markets, or in the absence of any predetermined state-claim prices. In the former case, optimization is over e , and equilibrium e is determined by $E(\pi, e) = R(\pi, Y)$. In a complete market, the decisionmaker maximizes income given the state-claim prices, and then uses this income to purchase the bundle of state-claims which maximize his preferences. This is analogous to equilibrium determination for a small-open economy with a representative consumer.⁸

In the latter case, which is analogous to autarkic price determination in general equilibrium with a representative consumer, state-claim prices are chosen so that the individual's internal market clears. Here it is convenient to recall E 's interpretation as a cost function. Picking state-claim prices is thus equivalent to picking state-claim demands for $E(\pi, e)$ and state-claim supplies for R . Hence, the market clearing conditions require that there exist a $y \in \partial E(\pi, e)$ such that

$$y \in \partial R(\pi, Y),$$

where the notation ∂R denotes the subdifferential of R in π . By Walras' Law and the basic properties of cost and revenue functions, one of the S market clearing conditions is redundant (alternatively $E(\pi, e) = R(\pi, Y)$ is redundant in the presence of the S market clearing conditions).

Define the maximal nonstochastic income consistent with Y as

$$y^Y = \max \{c : c\mathbf{1} \in Y\}.$$

Dual to $y^Y \mathbf{1}$ is a set of 'risk-neutral probabilities', \mathcal{P}^Y , which correspond to the supporting hyperplanes of Y at $y^Y \mathbf{1}$,

$$\mathcal{P}^Y = \{\pi : y^Y \mathbf{1} \in \partial R(\pi, Y)\}.$$

$R(\pi^Y, Y) = y^Y, \pi^Y \in \mathcal{P}^Y$. Besides offering the highest sure income that the decisionmaker can realize from Y , because y^Y is always feasible, $R(\pi^Y, Y) = y^Y$ also places a lower bound on equilibrium e and $E(\pi, e)$.

Theorem 7 *For any $\hat{\pi}$ consistent with the decisionmaker's choice equilibrium, $\pi^Y \preceq_e \hat{\pi} \preceq_e \pi^0, \pi^Y \in \mathcal{P}^Y$, where π^0 are the maximally risk-neutral probabilities.*

⁸In this context, $E(\pi, e) - R(\pi, Y)$ is exactly analogous to a trade expenditure function.

Proof Because $R(\pi^Y, Y)$ is feasible, $E(\hat{\pi}, e) \geq y^Y = R(\pi^Y, Y)$, $\pi^Y \in \mathcal{P}^Y$. By the definition of equilibrium, $\hat{y} \in \partial E(\hat{\pi}, e) \in Y$, and hence $R(\pi^Y, Y) \geq \pi^Y \hat{y}$ for any such y . But it is also true that $\hat{y} \in V(e)$, whence $\pi^Y \hat{y} \geq E(\pi^Y, e)$. ■

In interpreting Theorem 7, one might think of π^Y as the set of ‘minimally risky’ risk-neutral probabilities determined by the structure of Y . They are the probabilities that would lead a risk-neutral decisionmaker facing Y to choose the non-stochastic outcome. Theorem 7 is the preference analogue of the famous Peleg and Yaari (1975) result characterizing the set of risk averse efficient points over a general convex choice set. It implies that decisionmakers choose state-contingent income allocations so that their equilibrium risk-neutral probabilities are ranked between the minimally risky probabilities \mathcal{P}^Y and the maximally risky π^0 . This means that their optimal state-contingent income vector must fall ‘between’ the nonstochastic portfolio, $e\mathbf{1}$, and the portfolio that would be picked if the decisionmaker were forced to make trades at $\pi^Y \in \mathcal{P}^Y$. Figure 2 illustrates. Hence, just as the Peleg–Yaari notion of risk-averse efficiency constrains optimal choices of risk averters to lie in a particular subset of a convex choice set, Theorem 7 restricts choice associated with Y to lie within the subset of $V(e)$ determined by these supporting hyperplanes.

Theorem 7 is true for general monotonic preferences and does not require convexity of $V(e)$. It is a basic consequence of choice over convex sets. Among other things, the result implies that individuals create perfect insurance in the face of such a convex choice problem if and only if $y^Y \mathbf{1} \in \partial E(\pi^Y, y^Y)$ for some $\pi^Y \in \mathcal{P}^Y$, or, in other words, if and only if the choice set permits them to create fair insurance at their maximally risk-neutral probabilities.

Now consider the class of preferences for which there exists a unique probability measure, that is, for which $\pi(\mathbf{1})$ is a singleton. Suppose that, for some initial choice set Y , equilibrium is characterized by $\pi \neq \pi(\mathbf{1})$, so that the optimal y is not equal to $e\mathbf{1}$. Then because translating Y in the direction of $\mathbf{1}$ or radially expanding or shrinking Y has no effect on \mathcal{P}^Y , we conclude:

Theorem 8 *If for some initial choice set, Y , equilibrium $\pi \neq \pi(\mathbf{1})$, then translating Y in the direction of $\mathbf{1}$ or radially expanding or shrinking Y can only lead the decisionmaker to adopt the nonstochastic portfolio if \mathcal{P}^Y is not a singleton.*

Thus, such shifts in Y can lead to the decisionmaker fully insuring only if Y exhibits a kink at $y^Y \mathbf{1}$. A special case of this theorem is the well-known result that decisionmakers with unique subjective probabilities will never fully insure if $\mathcal{P}^Y = \{\pi\}$ is a singleton such that π provides what the decisionmaker views as unfair odds. It is the choice set analogue of the result, derived

by Segal and Spivak (1990), that decisionmakers with first-order risk aversion may fully insure at unfair odds. Those results can be derived by an analogous argument, which is left to the reader.

5.1 Comparative Statics for Linear Risk Tolerance and Constant Risk Aversion

In the cases where preferences exhibit CARA and CRRA, the decisionmaker's choice problem is particularly transparent. In the former, Theorem 2 implies that, for given state-claim prices, the decisionmaker equilibrium is characterized by

$$\begin{aligned} e^A(Y) &= \max \{e : e \leq R(\pi, Y) - \hat{a}(\pi)\} \\ &= R(\pi, Y) - \hat{a}(\pi), \end{aligned}$$

and in the latter by

$$e^R(Y) = \frac{R(\pi, Y)}{\hat{r}(\pi)}.$$

Because

$$\begin{aligned} R(\pi, Y + \delta \mathbf{1}) &= R(\pi, Y) + \delta, \\ R(\pi, \mu Y) &= \mu R(\pi, Y), \quad \mu > 0 \end{aligned}$$

one obtains the well-known results that a sure increase of wealth of δ dollars increases a CARA individual's equilibrium e by δ , while a radial increase or decrease in wealth leads to a proportionate change in the individual's equilibrium certainty equivalent. Similarly, for the class of preferences translation homothetic in the direction of \mathbf{u} ,

$$e^T(\mathbf{y}) = \frac{R(\pi, Y) - E^0(\pi)}{\pi \mathbf{u}},$$

whence:

Theorem 9 *If preferences are translation homothetic in the direction of \mathbf{u} , replacing Y by $Y + \delta \mathbf{u}$ with $\delta > 0$ raises the equilibrium certainty equivalent by δ with no effect on equilibrium π .*

For the case of LRT, and given state-claim prices, the equilibrium e is defined by

$$e^L(Y) = \frac{R(\pi, Y) - E^0(\pi)}{E^1(\pi)}.$$

Denote equilibrium π here by $\hat{\pi}$ and note that, since $E^1(\hat{\pi}) \leq 1$,

$$\begin{aligned} e^L(Y + \delta \mathbf{1}) &= \frac{R(\pi, Y) + \delta - E^0(\pi)}{E^1(\pi)} \\ &\geq \frac{R(\hat{\pi}, Y) + \delta - E^0(\hat{\pi})}{E^1(\hat{\pi})} \\ &\geq e^L(Y) + \delta, \end{aligned}$$

which is to be expected in light of Theorem 5.

Now consider the archetypal comparative static changes: the replacement of the choice set Y by tY for some $t > 1$ and the replacement of Y by $Y + \delta \mathbf{1}$ for some $\delta > 0$. The first arises, for example, in the case of the firm under uncertainty facing a proportional increase in all input and output prices. The second arises in the wealth allocation problem from an exogenous, non-taxable increase in income.

Theorem 10 *Suppose preferences exhibit LRT and are risk-averse with respect to some π^0 . Replacement of Y by $Y + \delta \mathbf{1}$ for $\delta > 0$ cannot lead to the choice of a less risky π , and replacement of Y by tY for $t > 1$ cannot lead to the choice of a more risky π .*

Proof The proof is for the replacement of Y by $Y + \delta \mathbf{1}$. Let $\hat{\pi}$ denote the originally optimal choice of π and π^δ the optimal choice for $Y + \delta \mathbf{1}$.

$$\begin{aligned} e^L(Y + \delta \mathbf{1}) &= \frac{R(\pi^\delta, Y) + \delta - E^0(\pi^\delta)}{E^1(\pi^\delta)} \\ &\geq \frac{R(\hat{\pi}, Y) + \delta - E^0(\hat{\pi})}{E^1(\hat{\pi})}. \end{aligned}$$

Now since

$$\frac{R(\hat{\pi}, Y) - E^0(\hat{\pi})}{E^1(\hat{\pi})} \geq \frac{R(\pi^\delta, Y) - E^0(\pi^\delta)}{E^1(\pi^\delta)},$$

we must have

$$\frac{\delta}{E^1(\pi^\delta)} \geq \frac{\delta}{E^1(\hat{\pi})},$$

that is, $E^1(\pi^\delta) \leq E^1(\hat{\pi})$. Hence, it cannot be true that π^δ is more risky than $\hat{\pi}$. A similar argument yields the result for tY . ■

Theorem 10, in conjunction with Theorem 7, implies that a sure increase in income leads a LRT decisionmaker to adopt a state-claim portfolio that is ‘closer’ to the optimal portfolio for \mathcal{P}^Y than his original portfolio. On the other hand, radial changes in the choice set lead a LRT decisionmaker to adopt a state-claim portfolio that is ‘closer’ to the riskless portfolio than his original portfolio.

The results of Theorem 10 may be combined to derive comparative statics for upward shifts in mean returns, multiplicative increases in the riskiness of assets, and so on. The results are consistent with those derived using the primal approach to characterize comparative statics in the presence of decreasing absolute risk aversion, as in Sandmo (1971), Feder (1977) and Milgrom (1994). However, the results of Theorem 10 are more general because these earlier papers were confined to the case of a scalar choice variable and relied on the restrictive assumption of expected-utility maximization.

Now consider constant risk averse preferences. Recall that in the dual equilibrium formulation, it was required that E be restricted to the region where it is finite. This requirement reflects a need for sufficient continuity to permit ‘market’ equilibration in the dual structure. The class of constant risk averse preferences, for which E is only finite on \mathcal{P}^* , neatly illustrates the requirement for such an assumption. For that class of preferences,

$$E(\pi, e) = \begin{cases} e & \pi \in \mathcal{P}^* \\ -\infty & \text{otherwise} \end{cases},$$

and, by Corollary 2,

$$\partial E(\pi, e) = \begin{cases} e\mathbf{1} & \pi \in \mathcal{P}^* \\ \emptyset & \text{otherwise} \end{cases}.$$

Suppose that $\mathcal{P}^Y \cap \mathcal{P}^* \neq \emptyset$, then equilibrium is determined by $e\mathbf{1} = y^Y\mathbf{1}$. The individual creates complete full insurance. On the other hand if $\mathcal{P}^Y \cap \mathcal{P}^* = \emptyset$, this method is not applicable. Well defined demand correspondences for state claims, which match the supplies generated from $R(\pi, Y)$, do not exist. Instead, equilibrium, is determined by the decisionmaker ‘plunging’ to the bounds of the choice set as

$$\inf \{R(\pi, Y) : \pi \in \mathcal{P}^*\}.$$

The choice problem reduces to finding the least favorable $R(\pi, Y)$ consistent with $\pi \in \mathcal{P}^*$. Figure 3 illustrates plunging behavior for the standard portfolio problem, with one safe asset, one risky asset, and no short selling. This latter characterization of equilibrium behavior always holds under constant risk aversion. We say that *plunging exists* when the dual equilibration process approach outlined above cannot be used in place of this latter characterization. By observing that \mathcal{P}^Y is invariant to either radial changes in Y or translations of Y in the direction of the sure thing, we can characterize comparative statics compactly under constant risk aversion by:

Theorem 11 *If the decisionmaker has constant risk averse preferences, her equilibrium certainty equivalent is given by $\inf \{R(\pi, Y) : \pi \in \mathcal{P}^*\}$. Replacing Y by $Y + \delta\mathbf{1}$ shifts the equilibrium certainty equivalent by δ , and replacing Y by tY for $t > 0$ rescales the equilibrium certainty equivalent by t . If the decisionmaker plunges before Y is replaced by $Y + \delta\mathbf{1}$ or by tY , she will plunge after the replacement. If she does not plunge before these replacements, she will not plunge after these replacements.*

6 Uncertainty and Risk⁹

The approach to preferences over stochastic outcomes developed in this paper is decidedly not decision theoretic. That is intentional. Our goal is to demonstrate that standard tools from microeconomic theory can be usefully and informatively applied to preferences over stochastic outcomes, and that many concepts familiar from standard consumer and producer theory have exact counterparts in decision making under uncertainty. Thus, the only requirements that we have placed on preferences is that they be monotonic and associated with convex indifference maps. While we consider the resulting generality the primary strength of the paper, it does have its costs. In particular, we forgo the benefits of additional structure that may be obtained if we assume that preferences are defined with respect to probability distributions over outcomes.

To illustrate, suppose we denote by Σ the set of all subsets of Ω , that is, 2^Ω . Then (Ω, Σ) can be thought of as a measurable space, and if we endow it with a particular probability measure, call it \mathbf{p} , then $(\Omega, \Sigma, \mathbf{p})$ is a measure or probability space. Elements of Σ are referred to as ‘events’.

Following in the tradition of de Finetti and Savage, it is traditional to identify this probability measure, which is taken as subjective, with an individual’s willingness to pay for lottery tickets associated with each of the events in the neighborhood of the degenerate random variable, $\mathbf{1}$. A lottery ticket for event $A \in \Sigma$ is the special case of our random variable mapping from $\Omega \rightarrow Y$ that pays \$1 if event A occurs. So, for example, the lottery tickets associated with the primitive events $s \in \Omega$ are the random variables denoted by \mathbf{e}_s in our notation.

Because the primitive lottery tickets span the space of random variables, any random variable can be built up constructively from these primitive lotteries. Visually, the probabilities so derived can be identified with the normals to the supporting hyperplanes for $V(e)$ in the neighborhood of the degenerate random variable. It is the invocation of the Savage axioms (or some closely related set of axioms) that permits this identification of subjective probabilities in formal decision theoretic models.

If preferences satisfy the Savage axioms, they may be described as probabilistically sophisticated (Machina and Schmeidler 1992). This means that any two lotteries yielding the same probability distributions over outcomes are judged as indifferent. In particular, if $\mathbf{p}(A) = \mathbf{p}(B)$ then lottery tickets for A and B are equally valued. If Ω contains S events, each with equal probability, then the certainty equivalent function e is symmetric.

⁹We thank an anonymous reviewer for suggesting this section to us.

But there is a problem. Starting with Ellsberg's famous paradoxes, considerable evidence has emerged to suggest that individual attitudes towards uncertain outcomes cannot be identified with a single probability space of the form $(\Omega, \Sigma, \mathbf{p})$ for a unique probability measure in the presence of ambiguity. In particular, Ellsberg's paradox suggests that individual behavior cannot be accurately described in terms of a single probability measure. Visually, this suggests that there may exist multiple supporting hyperplanes for $V(e)$ in the neighborhood of the degenerate random variables. Indifference maps would then appear to be 'kinked' in the neighborhood of the sure thing.

As we have illustrated with the example of CRA preferences, which also have a kink in the neighborhood of the degenerate random variable, such a preference structure is perfectly consistent with our general set up. And because we work in terms of superdifferentials, it presents no inherent analytic difficulties. However, as has been illustrated by a number of authors, it is easy in such circumstances to confuse attitudes to risk and to ambiguity.

An example illustrates the problem. One of the most popular models capable of coping with the theoretical problems that arise from Ellsberg's paradox is the multiple-prior maximin expected utility model of Gilboa and Schmeidler (1989). Here an individual's evaluation of a stochastic outcome is represented by a preference function of the form (note the similarity to CRA)

$$u(e(\mathbf{y})) = \inf \left\{ \sum_{s \in \Omega} \pi_s u(y_s) : \pi \in \hat{\mathcal{P}} \right\}$$

where $\hat{\mathcal{P}} \subseteq \mathcal{P}$ is closed and convex, and u is concave and increasing. As usually interpreted, the curvature of u measures the individual's attitude towards risk, while the practice of evaluating outcomes in terms of the least favorable probability measure reflects aversion to ambiguity.

Let us take the special case where

$$\hat{\mathcal{P}} = \mathcal{P}.$$

Intuitively, this is equivalent to saying that the decision maker is willing to entertain all possible probabilities as underlying the stochastic outcomes that he or she faces. Also assume, however, that if the individual were presented with risk, a known probability space, he or she would be risk neutral so that $u(y) = y$. It follows that

$$\begin{aligned} e(\mathbf{y}) &= \inf \{ \pi \mathbf{y} : \pi \in \mathcal{P} \} \\ &= \min \{ y_1, y_2, \dots, y_S \}. \end{aligned}$$

This individual would behave as though he or she were what we described as 'perfectly risk averse'.

However, this apparent risk averse behavior does not emerge from aversion to risk in the usual sense. Rather it emerges from concerns that are typically attributed to ambiguity aversion.

6.1 Distinguishing risk and ambiguity

There is no generally agreed way to distinguish between risk and ambiguity aversion in the primal setting. However, a distinction may be drawn for the special case of maxmin expected utility preferences where the utility function satisfies CARA.¹⁰ Moreover, this opens the way to a general treatment for the dual premium.

Given an expected-value function $E(\pi, e) : \mathcal{P} \times \mathfrak{R} \rightarrow \mathfrak{R}$, we can define the generalization $E(P, e) : C \times \mathfrak{R} \rightarrow \mathfrak{R}$ where C consists of convex subsets of \mathcal{P} , given by

$$E(P, e) = \sup_{\pi \in P} E(\pi, e).$$

Given any particular π^0 in P , it is always true that

$$E(P, e) = \left(\sup_{\pi \in P} E(\pi, e) - E(\pi^0, e) \right) + E(\pi^0, e), \quad (3)$$

so that relative to π^0 , one can partition expenditure into an ambiguity term and a risk term.¹¹ The ambiguity term,

$$\left(\sup_{\pi \in P} E(\pi, e) - E(\pi^0, e) \right),$$

has the desirable property that, in the absence of ambiguity, P is a singleton, that is $P = \{\pi^0\}$, and thus

$$\left(\sup_{\pi \in P} E(\pi, e) - E(\pi^0, e) \right) = 0.$$

Although the partition (3) is applicable for any choice of π^0 in P , two choices seem of particular interest. The first is

$$\pi^0 = \arg \min_{\pi \in P} E(\pi, e)$$

which yields

$$E(\pi^0, e) = \inf_{\pi \in P} E(\pi, e).$$

The effect of this choice is to maximize the component of expected-value imputed to the ambiguity represented by P .

¹⁰We thank a referee for suggesting this approach.

¹¹This decomposition was inspired by the comments of an anonymous reviewer. We note, in passing, that a number of alternative risk-uncertainty decompositions have been offered elsewhere in the literature (for example, Epstein, 1999). To date, no single one appears to have gained universal acceptance.

Alternatively, π^0 may be a salient or ‘anchor’ reference probability vector, as in the work of Gajdos, Tallon and Vergnaud (2004). Gajdos et al. note that, in most of the classic cases of ambiguity, such as those of the Ellsberg problems there is a natural reference probability vector π^0 . In the Ellsberg two-urn problem, for example, where the second urn contains unknown numbers of black and white balls, the natural reference probability for the event ‘a white ball is drawn’ is 0.5.

In this and many other cases, the reference probability may be derived from symmetry considerations. However, as Gajdos et al. observe, the reference probability vector may also be derived from the point estimates of parameters in an econometric or physical model, while the set P corresponds to a confidence interval. If the model is non-linear, P need not be symmetric about π^0 .

In the special case of CARA, we have

$$E(\pi, e) = \hat{a}(\pi) + e$$

so we get

$$\begin{aligned} E(P, e) &= \sup_{\pi \in P} (\hat{a}(\pi) + e) \\ &= \sup_{\pi \in P} \hat{a}(\pi) + e. \end{aligned}$$

Looking directly at the premium, write

$$a(P, e) = E(P, e) - e$$

so with CARA

$$a(P, e) = \sup_{\pi \in P} \hat{a}(\pi) - e$$

and we have the partition

$$a(P, e) = \left(\sup_{\pi \in P} \hat{a}(\pi) - \hat{a}(\pi^0) \right) + \hat{a}(\pi^0).$$

Thus, under CARA, the partition is independent of the choice of e , and may be regarded as being determined entirely by the decision-maker’s beliefs.

7 Concluding comments

In this paper, we have attempted to show that a dual treatment of choice under uncertainty is both tractable and informative. Particular emphasis has been placed on characterizing dual risk premiums and showing that various invariance restrictions placed on these risk premiums

lead naturally to the generalizations of the concepts of CARA, CRRA, and LRT familiar from expected-utility theory. Each of these concepts conforms to a notion of homotheticity familiar from the literature on consumer preferences, and each has a tractable, and intuitive, dual formulation.

References

- Blackorby, C., and D. Donaldson. "A Theoretical Treatment of Indices of Absolute Inequality." *International Economic Review* 21, no. 1 (February 1980): 107–36.
- Blackorby, C., R. Boyce, and R. R. Russell. "Estimation of Demand Systems Generated by the Gorman Polar Form: A Generalization of the S-Branch Utility Tree." *Econometrica* 46 (1978): 345–64.
- Brennan, M. J., and A. Kraus. "The Geometry of Separation and Myopia." *Journal of Financial and Quantitative Analysis* 11 (1976): 171–93.
- Cass, D., and J. E. Stiglitz. "The Structure of Investor Preferences and Asset Returns, and Separability in Portfolio Allocation: A Contribution to the Pure Theory of Mutual Funds." *Journal of Economic Theory* 2 (1970): 122–60.
- Chambers, R. G. "Consumers' Surplus As an Exact and Superlative Cardinal Welfare Measure." *International Economic Review* 42 (2001): 105–20.
- Chambers, R. G., and J. Quiggin. *Uncertainty, Production, Choice, and Agency: The State-Contingent Approach*. New York: Cambridge University Press, 2000.
- Chambers, R. G., and R. Färe. "Translation Homotheticity." *Economic Theory* 11 (1998): 629–41.
- Chambers, Robert G., Y. Chung, and R. Fare. "Benefit and Distance Functions." *Journal of Economic Theory* 70 (1996): 407–19.
- Dickinson, J. G. "Parallel Preference Structures in Labour Supply and Commodity Demand: An Adaptation of the Gorman Polar Form." *Econometrica* 48 (November 1980): 1711–25.
- Epstein, L.G. "A Definition of Uncertainty Aversion". *Review of Economic Studies* 66 (1999): 579-608.
- Epstein, L. G., and S. E. Zin. "First-Order Risk Aversion and the Equity Premium Puzzle." *Journal of Monetary Economics* 26 (1990): 387–407.
- Färe, R. and C.A.K. Lovell. "Affinely Homothetic Production Technology", *Methods of Operations Research* 48 (1984): 197-205.
- Feder, G. "The Impact of Uncertainty on a Class of Objective Functions." *Journal of Economic Theory* 16 (1977): 504–12.
- Gajdos, T., J-M. Tallon, & J-C. Vergnaud. "Decision making with imprecise probabilistic information." *Journal of Mathematical Economics*, 40(6), 2004: 647-681.
- Gilboa, I. and Schmeidler, D. (1989), "Maxmin expected utility with non-unique prior", *Journal*

of *Mathematical Economics*, 18(2), 141–53.

Gorman, W. M. “Community Preference Fields.” *Econometrica* 21 (1953): 63–80.

Karni, E. *Decision Making Under Uncertainty: The Case of State-Dependent Preferences*. Cambridge: Harvard University Press, 1985.

Lewbel, A. “The Rank of Demand Systems: Theory and Nonparametric Estimation.” *Econometrica* 59 (May1991): 711–30.

Lewbel, A., and W. Perraudin. “A Theorem on Portfolio Selection With General Preferences.” *Journal of Economic Theory* 65 (1995): 624–26.

Luenberger, D. G. “Benefit Functions and Duality.” *Journal of Mathematical Economics* 21 (1992): 461–81.

Machina, M. J. “Payoff Kinks in Preferences Over Lotteries.” *Journal of Risk and Uncertainty* 23 (2001): 207–60.

Machina, M.J., and D. Schmeidler. “A More Robust Definition of Subjective Probability”. *Econometrica* 60 (1992): 745-780.

Milgrom, Paul. “Comparing Optima: Do Simplifying Assumptions Affect Conclusions.” *Journal of Political Economy*, no. 3 (1994): 607–15.

Milne, F. “Consumer Preferences, Linear Demand Functions, and Aggregation in Competitive Asset Markets.” *Review of Economic Studies* 46 (1979): 407–17.

Nau, R. “A Generalization of Pratt–Arrow Measure to Non–Expected Utility Preferences and Inseparable Probability and Utility.” Working Paper, Fuqua School of Business, Duke University, 2001.

Peleg, B., and M. Yaari. “A Price Characterisation of Efficient Random Variables.” *Econometrica* 43 (1975): 283–92.

Quiggin, J., and R. G. Chambers. “Risk Premiums and Benefit Measures for Generalized Expected Utility Theories.” *Journal of Risk and Uncertainty* 17 (1998): 121–38.

Rockafellar, R. T. *Convex Analysis*. Princeton: Princeton University Press, 1970.

Safra, Zvi, and Uzi Segal. “Constant Risk Aversion.” *Journal of Economic Theory* 83, no. 1 (1998): 19–42.

Sandmo, A. “On the Theory of the Competitive Firm Under Price Uncertainty.” *American Economic Review* 61 (1971): 65–73.

Segal, U., and A. Spivak. “First–Order Versus Second–Order Risk–Aversion.” *Journal of Economic Theory* 51, no. 1 (1990): 111–25.

- Yaari, M. “Convexity in the Theory of Choice Under Risk.” *Quarterly Journal of Economics* 79 (May 1965): 278–90.
- “The Dual Theory of Choice Under Risk.” *Econometrica* 55 (1987): 95–115.
- “Some Remarks on Measures of Risk Aversion and on Their Uses.” *Journal of Economic Theory* 1 (1969): 315–29.