Dual Approaches to the Analysis of Risk Aversion

by

Robert G. Chambers and John Quiggin

WP 02-04

Department of Agricultural and Resource Economics

The University of Maryland, College Park

DUAL APPROACHES TO THE ANALYSIS OF RISK AVERSION

Robert G. Chambers
Professor of Agricultural and Resource Economics, University of Maryland
Adjunct Professor of Agricultural and Resource Economics, University of Western Australia
rchambers@arec.umd.edu

John Quiggin
Australian Research Council Senior Fellow, Australian National University

May 2002

Dual Approaches to the Analysis of Risk Aversion

Robert G. Chambers and John Quiggin

May 15, 2002

DUAL APPROACHES TO THE ANALYSIS OF RISK AVERSION

A series of classic papers (Yaari, 1965; Yaari, 1969; Peleg and Yaari, 1975) observe that choice under uncertainty is formally equivalent to choice over commodity bundles in conventional consumer theory. Although this is an old observation, its implications are yet to be fully exploited. Using dual concepts familiar from standard producer-consumer theory, this paper presents a systematic analysis of preferences over uncertain outcomes in terms of convex sets and their supporting hyperplanes. The aim is to extend the insights of Peleg and Yaari in the light of the dual approaches developed since their work was published.

The analysis commences by representing convex preference sets over uncertain outcomes in terms of the translation function, 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 the notion of 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 gives rise to 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.

The paper then examines the concepts of constant absolute risk aversion (CARA), constant relative risk aversion (CRRA), and linear risk tolerance (LRT) in a dual framework. These concepts are interpreted as homotheticity properties, and each is shown to be characterized by an invariance property of the risk-neutral probabilities. Homotheticity conditions of various kinds play a central role in consumer and producer theory. However, relatively little attention has been given to these properties for preferences over uncertain outcomes. 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 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. Finally, our methods are 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. The final section concludes.

1 Notation and Basic Concepts

For a proper concave function $f : \Re^S \to \Re$, its *superdifferential* at x is the closed, convex set:

(1) $$\partial f(\mathsf{x}) = \left\{\mathsf{v} \in \Re^S : f(\mathsf{x}) + \mathsf{v}(\mathsf{z} - \mathsf{x}) \geq f(\mathsf{z}) \text{ for all } \mathsf{z}\right\}.$$

The elements of $\partial f(\mathsf{x})$ are referred to as *supergradients*. If f is differentiable at x, $\partial f(\mathsf{x})$ is a singleton. If $\partial f(\mathsf{x})$ is a singleton, f is differentiable at 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 $Y \subseteq \Re$. We refer to the outcomes as income. Our focus is on the case where Ω is a finite set $\{1, ...S\}$, and the space of random variables is $Y^S \subseteq \Re^S$. The unit vector is denoted $\mathsf{1} = (1, 1, ...1)$, and $\mathcal{P} \subset \Re^S_{++}$ denotes the probability simplex. Define e_i as the i-th row of the $S \times S$ identity matrix , $\mathsf{e}_i = (0, ..., 1, 0, ..., 0)$.

Preferences over state-contingent incomes are given by an ordinal mapping $W : \Re^S \to \Re$. W is continuous, nondecreasing, and quasi-concave in y. Quasi-concavity ensures that the least-as-good sets of the preference mapping

$$V(w) = \{\mathsf{y} : W(\mathsf{y}) \geq w\}$$

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

2 The Translation Function and the Expected-Value Function

The *translation function*, $B : \Re \times Y^S \to \Re$, is defined:

$$B(w, \mathsf{y}; \mathsf{1}) = \max\{\beta \in \Re : \mathsf{y} - \beta\mathsf{1} \in V(w)\}$$

if $\mathsf{y} - \beta\mathsf{1} \in V(w)$ for some β, and $-\infty$ otherwise (Blackorby and Donaldson, 1980; Luenberger, 1992).[1] The properties of $B(w, \mathsf{y}; \mathsf{1})$ 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 $B(w, \mathsf{y}; \mathsf{1})$ *satisfies:*

a) $B(w, \mathsf{y}; \mathsf{1})$ *is nonincreasing in* w *and nondecreasing and concave in* y*;*

b) $B(w, \mathsf{y} + \alpha\mathsf{1}; \mathsf{1}) = B(w, \mathsf{y}; \mathsf{1}) + \alpha$, $\alpha \in \Re$ *(the translation property);*

c) $B(w, \mathsf{y}; \mathsf{1}) \geq 0 \Leftrightarrow \mathsf{y} \in V(w)$*;*

d) $B(w, \mathsf{y}; \mathsf{1})$ *is jointly continuous in* y *and* w *in the interior of the region* $\Re \times Y^S$ *where* $B(w, \mathsf{y}; \mathsf{1})$ *is finite.*

An important special case is the certainty-equivalent representation of preferences,

$$e(\mathsf{y}) = -B(W(\mathsf{y}), \mathsf{0}; \mathsf{1}).$$

Because the preference functional is ordinal, no loss of generality is involved in replacing $W(\mathsf{y})$ by its ordinal transform $e(\mathsf{y})$, and in what follows we always do so, writing $B(e, \mathsf{y}; \mathsf{1})$ in place of $B(w, \mathsf{y}; \mathsf{1})$.

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

We refer to the concave conjugate of the translation function, $B(e, \mathsf{y}; \mathsf{1})$, as the *expected-value function* $E : \mathcal{P} \times \Re \rightarrow \Re$. It is defined by

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

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

$$\begin{aligned} E(\boldsymbol{\pi}, e) &= \inf_{\mathsf{y}} \{\boldsymbol{\pi}\mathsf{y} : B(e, \mathsf{y}; \mathsf{1}) \geq 0\} \quad \boldsymbol{\pi} \in \mathcal{P} \\ &= \inf_{\mathsf{y}} \{\boldsymbol{\pi}\mathsf{y} : \mathsf{y} \in V(e)\} \end{aligned}$$

if there exists some $\mathsf{y} \in V(e)$, and ∞ otherwise. Hence, 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, $B(e, \mathsf{y}; \mathsf{1})$ is a continuous and nondecreasing proper concave function, and thus $E(\boldsymbol{\pi}, e)$ is a closed, proper concave[2] 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 $e(e\mathsf{1}) = e$, $E(\boldsymbol{\pi}, e) \leq e$.

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 $e(\mathsf{y})$ and $E(\boldsymbol{\pi}, e)$ thus 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

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

[2] A concave function, $\mathsf{g}(\mathsf{x})$, is proper if there is at least one x such that $\mathsf{g}(\mathsf{x}) > -\infty$, and $\mathsf{g}(\mathsf{x}) < \infty$ 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

$$\boldsymbol{\pi} \in \partial B\left(e, \mathsf{y}; \mathsf{1}\right) \Longleftrightarrow \mathsf{y} \in \partial E\left(\boldsymbol{\pi}, e\right) \tag{2}$$

in the relative interior of their domains.[3] 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 interior of their domains

$$\mathsf{y}\left(\boldsymbol{\pi}, e\right) \in \arg\inf\left\{\boldsymbol{\pi}\mathsf{y} - B\left(e, \mathsf{y}; \mathsf{1}\right)\right\} \Rightarrow \mathsf{y}\left(\boldsymbol{\pi}, e\right) \in \partial E\left(\boldsymbol{\pi}, e\right),$$

because

$$\begin{array}{rcl} \boldsymbol{\pi}\mathsf{y}\left(\boldsymbol{\pi}, e\right) - B\left(e, \mathsf{y}\left(\boldsymbol{\pi}, e\right); \mathsf{1}\right) & \leq & \boldsymbol{\pi}\mathsf{y} - B\left(e, \mathsf{y}; \mathsf{1}\right), \quad \text{for all } \mathsf{y} \\ & \Downarrow & \\ \boldsymbol{\pi} & \in & \partial B\left(e, \mathsf{y}\left(\boldsymbol{\pi}, e\right); \mathsf{1}\right) \\ & \Downarrow & \\ \mathsf{y}\left(\boldsymbol{\pi}, e\right) & \in & \partial E\left(\boldsymbol{\pi}, e\right), \end{array}$$

where the first $\Rightarrow$ follows by the definition of the superdifferential, and the second by (2). By a parallel argument, $\mathsf{p}\left(e, \mathsf{y}\right) \in \arg\inf\left\{\boldsymbol{\pi}\mathsf{y} - E\left(\boldsymbol{\pi}, e\right)\right\} \Rightarrow \mathsf{p}\left(e, \mathsf{y}\right) \in \partial B\left(e, \mathsf{y}; \mathsf{1}\right)$. Thus, the supergradients of the translation function are interpretable as compensated state-claim-price-dependent demand functions for the state-claim vector 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, which has proven central in the generalized expected utility literature (Epstein and Zinn 1990; Segal and Spivak 1990; Machina, 2001).

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

Yaari (1969) identifies subjective probabilities with the supporting hyperplane to $V(e)$ along the sure-thing vector, which are given by $\partial B(e, e\mathsf{1}, \mathsf{1})$. These are the normalized state-claim prices which support the constant portfolio, $e\mathsf{1}$. Nau (2001), who restricts attention to differentiable preferences, has confirmed the 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 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, $e\mathsf{1}$.

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

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

$$B(e, \mathsf{y}, \mathsf{1}) + \mathsf{v}(\mathsf{z}-\mathsf{y}) \geq B(e, \mathsf{z}, \mathsf{1})$$
$$B(e, \mathsf{y}, \mathsf{1}) + \mathsf{v}(\mathsf{z}^*-\mathsf{y}) \geq B(e, \mathsf{z}^*, \mathsf{1})$$

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

$$\begin{aligned} \partial B(e, \mathsf{y}+\boldsymbol{\delta}\mathsf{1}, \mathsf{1}) &= \left\{ \begin{array}{c} \mathsf{v} : B(e, \mathsf{y}+\boldsymbol{\delta}\mathsf{1}, \mathsf{1}) + \mathsf{v}(\mathsf{z}+\boldsymbol{\delta}\mathsf{1} - [\mathsf{y}+\boldsymbol{\delta}\mathsf{1}]) \\ \geq B(e, \mathsf{z}+\boldsymbol{\delta}\mathsf{1}, \mathsf{1}) \text{ for all } \mathsf{z}+\boldsymbol{\delta}\mathsf{1} \end{array} \right\} \\ &= \{\mathsf{v} : B(e, \mathsf{y}, \mathsf{1}) + \mathsf{v}(\mathsf{z}-\mathsf{y}) \geq B(e, \mathsf{z}, \mathsf{1}) \text{ for all } \mathsf{z}\} = \partial B(e, \mathsf{y}; \mathsf{1}). \end{aligned}$$

where the second equality follows by Lemma 1.¥

It now seems natural to refer to the elements of any vector $\mathsf{p}(e, \mathsf{y}) \in \partial B(e, \mathsf{y}, \mathsf{1}) \subset \Re^S_+$ as *risk-neutral probabilities* for the certainty equivalent e. If the translation function is differentiable, these probabilities are unique and given by the gradient, $\nabla B(e, \mathsf{y}, \mathsf{1})$. We define the set of *risk-neutral probabilities* $\boldsymbol{\pi}(\mathsf{y}) \subset \Re^S_+$, which correspond to the supporting hyperplanes for the indifference set, by

$$\boldsymbol{\pi}(\mathsf{y}) = \partial B(e(\mathsf{y}), \mathsf{y}; \mathsf{1}) = \mathsf{p}(e(\mathsf{y}), \mathsf{y}).$$

Thus, $\boldsymbol{\pi}(\mathsf{y})$ is analogous to an inverse demand correspondence for the Arrow commodities. When preferences are smooth, $\boldsymbol{\pi}(\mathsf{y})$ 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 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 $e(e\mathbf{1}) = e$, $E(\boldsymbol{\pi}, e) \leq e$. And, because preferences are quasi-concave, $\boldsymbol{\pi} \in \partial B(e, e\mathbf{1}; \mathbf{1}) \Longleftrightarrow E(\boldsymbol{\pi}, e) = e$. We, thus, define the set of *subjective probabilities* $\boldsymbol{\pi}(\mathbf{1}) \subset \Re_{+}^{S}$ as

$$\boldsymbol{\pi}(\mathbf{1}) = \cap_{e}\{\partial B(e, e\mathbf{1}; \mathbf{1})\}.$$

Typically, we shall assume that $\boldsymbol{\pi}(\mathbf{1})$ is non-empty, although in general it need not be. In the case of smooth preferences, nonemptiness 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. The set of subjective probabilities satisfies $\boldsymbol{\pi}(\mathbf{1}) = \cap_{e} \arg\sup_{\pi \in \mathcal{P}} \{E(\boldsymbol{\pi}, e) - e\}$.

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

3 Risk aversion

Yaari's (1969) approach to the definition of risk aversion was to define first a notion of comparative risk aversion and then induce absolute 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

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

The definition of risk aversion requires that an individual can be risk averse with respect to $\boldsymbol{\pi}^0$ only if $\boldsymbol{\pi}^0 \in \boldsymbol{\pi}(\mathbf{1})$. The definition of risk aversion implies that an individual is risk-averse with respect to $\boldsymbol{\pi}^0$ 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 $\boldsymbol{\pi}^0\mathbf{z} = 0$ and, *a fortiori*, all bets that are unfavorable in the sense that $\boldsymbol{\pi}^0\mathbf{z} < 0$. In the case where $\boldsymbol{\pi}(\mathbf{1})$ 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 $\boldsymbol{\pi}$ *is less risky than* $\boldsymbol{\pi}'$ *at* e, *denoted* $\boldsymbol{\pi} \preceq_e \boldsymbol{\pi}'$, *if* $E(\boldsymbol{\pi}', e) \geq E(\boldsymbol{\pi}, e)$.

Intuitively, $\boldsymbol{\pi} \preceq_e \boldsymbol{\pi}'$ implies that $\boldsymbol{\pi}'$ is 'closer' to the set of maximally risk-neutral probabilities (those for which $E(\boldsymbol{\pi}, e) = e$) than $\boldsymbol{\pi}$. Thus, the more risky are the probabilities, the 'closer' will be $\mathbf{y} \in \partial E(\boldsymbol{\pi}, e)$ 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* $E(\boldsymbol{\pi}^0, e) = e \quad \forall e$, *and* $\boldsymbol{\pi} \preceq_e \boldsymbol{\pi}^0$ *for all* $(\boldsymbol{\pi}, e)$.

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

Theorem 4 *The following are equivalent:*

(a) A is more risk averse than B;

(b) $B^A(e, \mathsf{y}; \mathsf{1}) \leq B^B(e, \mathsf{y}; \mathsf{1})$ for all y and e;

(c) $E^A(\boldsymbol{\pi}, e) \geq E^B(\boldsymbol{\pi}, e)$ for all π and e;and

(d) for all y, $e^A(\mathsf{y}) \leq e^B(\mathsf{y})$.

Moreover, if A is more risk-averse than B, and B is risk-averse with respect to probabilities $\boldsymbol{\pi}^0$, so is A.

Proof *(a)⇒(c) is immediate. (c)⇒(b) follows by applying $E^A(\boldsymbol{\pi}, e) \geq E^B(\boldsymbol{\pi}, e)$ for all π and e in the conjugacy mapping. (b)⇒(d) follows because $e(\mathsf{y})$ is determined by* $\max\{e : B(e, \mathsf{y}; \mathsf{1}) \geq 0\}$ *.(d)⇒(a) is immediate from the definition of V. The second part of the theorem is trivial.*¥

An easy corollary to part (d) is 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 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 $e > 0$) as

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

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, and they thus 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 5 *The following conditions are equivalent:*

(1) A is more risk-averse than B;

(2) $a^A(\boldsymbol{\pi},e) \geq a^B(\boldsymbol{\pi},e) \quad \forall\pi, e$*; and*

(3) for all $e > 0$ $r^A(\boldsymbol{\pi},e) \geq r^B(\boldsymbol{\pi},e) \quad \forall\pi.$

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 π^0,

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

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,

$$e(\mathsf{y}) = \min\{y_1, y_2, ..., y_S\},$$

$a(\boldsymbol{\pi},e) = 0$, 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 6 *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,* $B(e,\mathsf{y};\mathsf{1}) = B(0,\mathsf{y};\mathsf{1})-e$*, and* $\boldsymbol{\pi}(\mathsf{y}+\beta\mathsf{1}) = \boldsymbol{\pi}(\mathsf{y}), \quad \beta \in \Re$*. 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,* $B(e,\mathsf{y};\mathsf{1}) = eB\left(1,\frac{\mathsf{y}}{e};1\right)$*, and* $\boldsymbol{\pi}(\mu\mathsf{y}) = \boldsymbol{\pi}(\mathsf{y}), \quad \mu > 0$.

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,* $E(\boldsymbol{\pi},e) = \hat{a}(\boldsymbol{\pi}) + e$*. By conjugacy,*

$$\begin{aligned} B(e,\mathsf{y};\mathsf{1}) &= \min_{\boldsymbol{\pi}}\{\boldsymbol{\pi}\mathsf{y}-\hat{a}(\boldsymbol{\pi})\} - e \\ &= B(0,\mathsf{y};\mathsf{1}) - e, \end{aligned}$$

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

Corollary 7 *If preferences exhibit CARA* $\boldsymbol{\pi} \in \cap_e\{\partial B(e,e\mathsf{1};\mathsf{1})\} \Longleftrightarrow \hat{a}(\boldsymbol{\pi}) = 0$*. If preferences exhibit CRRA* $\boldsymbol{\pi} \in \cap_e\{\partial B(e,e\mathsf{1};\mathsf{1})\} \Longleftrightarrow \hat{r}(\boldsymbol{\pi}) = 1$*. In both cases,* $\boldsymbol{\pi}(\mathsf{1})$ *is nonempty.*

A direct consequence of Theorem 6 is that for CARA preferences, $e(\mathsf{y}) = B(0,\mathsf{y};\mathsf{1})$. Thus, by Lemma 1.b, $e(\mathsf{y}+\beta\mathsf{1}) = e(\mathsf{y}) + \beta$. This is the standard primal definition of CARA for general preferences (Chambers and Quiggin, 2000). Hence, $W(\mathsf{y})$ is translation homothetic (Blackorby and Donaldson, 1980; Chambers and Färe, 1998). Similarly, for CRRA preferences, $B\left(1,\frac{\mathsf{y}}{e};1\right) = 0$, whence $e(\mu\mathsf{y}) = \mu e(\mathsf{y}) \quad \mu > 0$ implying that $W(\mathsf{y})$ is homothetic.

Example Expected utility preferences, risk-averse for the probabilities $\boldsymbol{\pi}^0$, exhibit CARA if and only if

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

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

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

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, $E(\boldsymbol{\pi}, e+e^*) \leq E(\boldsymbol{\pi}, e) + e^*$, $e^* > 0$.

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

Under DARA, $\frac{\partial^+}{\partial e}E(\boldsymbol{\pi}, e) \leq 1$, while under DRRA, $\frac{\partial^+}{\partial \ln e}\ln E(\boldsymbol{\pi}, e) \leq 1$. 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 $\boldsymbol{\pi}$-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 at marginal cost.)

An immediate consequence of these definitions and Theorem 6 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

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

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

Proof The proof is for DRRA, the proof for DARA is parallel. By definition, $E(\boldsymbol{\pi}, \mu e) \leq \mu E(\boldsymbol{\pi}, e)$. Hence

$$\boldsymbol{\pi}\mathbf{y} - E(\boldsymbol{\pi}, \mu e) \geq \boldsymbol{\pi}\mathbf{y} - \mu E(\boldsymbol{\pi}, e) = \mu\left[\boldsymbol{\pi}\frac{\mathbf{y}}{\mu} - E(\boldsymbol{\pi}, e)\right].$$

Taking the infimum of both sides yields $B(\mu e, \mathbf{y}; \mathbf{1}) \geq \mu B\left(e, \frac{\mathbf{y}}{\mu}; \mathbf{1}\right)$ by conjugacy. By Lemma 1.c and this form $B\left(\mu e\left(\frac{\mathbf{y}}{\mu}\right), \mathbf{y}; \mathbf{1}\right) \geq \mu B\left(e\left(\frac{\mathbf{y}}{\mu}\right), \frac{\mathbf{y}}{\mu}; \mathbf{1}\right) \geq 0$. The converse follows by the conjugacy relationship.¥

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:

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

with $E^0(\boldsymbol{\pi})$ and $E^1(\boldsymbol{\pi}) \geq 0$ expected-value functions for least-as-good sets that are independent of the certainty equivalent. 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 $e(\mathsf{y})$ 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 9 *(Chambers, Chung, and Färe) Preferences exhibit LRT if and only if*

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

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

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

where the last equality follows by monotonicity of preferences.¥

It seems unlikely, therefore, that much information can be gleaned directly from examining $\partial B(e; \mathsf{y}; \mathbf{1})$ 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 10 *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 $E(\boldsymbol{\pi}, \mu e) = \lambda E(\boldsymbol{\pi}, e) + (1 - \lambda) E^0(\boldsymbol{\pi})$.¥

Further tractability for LRT preferences, can be had by imposing further functional structure. For example, an important special case of LRT preferences are the *affinely homothetic* preferences (Milne, 1979). Affine homotheticity is the special case of LRT given by, $E^0(\boldsymbol{\pi}) = \boldsymbol{\pi}\mathsf{v} \quad \mathsf{v} \in \Re^S$. These preferences have linear expansions paths emanating from a common point that is independent of state-claim prices. In a standard consumer context, v is usually interpreted as a vector of subsistence demands, and perhaps the best known member of the LRT class is the Stone–Geary utility structure, which underlies the linear-expenditure systems. Expected-utility LRT preferences are also affinely homothetic (Milne, 1979). Thus, it is a trivial corollary that results established for the general class of LRT preferences also apply to the expected-utility subclass of LRT preferences.

Another special case, which appears not to have been considered in the literature on portfolio choice, is the class of preferences that are translation homothetic in an arbitrary direction 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 $E^1(\boldsymbol{\pi}) = \boldsymbol{\pi}\mathsf{u}$, where $\mathsf{u} \in \Re^S$. CARA is the special case where $\mathsf{u} = \mathsf{1}$. We briefly return to this class below in our discussion of comparative statics for LRT preferences.

Corollary Affinely homothetic preferences of the form, $\boldsymbol{\pi}\mathsf{v} + E^1(\boldsymbol{\pi}) e$ are risk averse with respect to $\boldsymbol{\pi}^0$ only if $\mathsf{v} \leq \mathsf{0}$. Preferences translation homothetic in the direction of u are risk averse with respect to $\boldsymbol{\pi}^0$ only if $\mathsf{u} \leq \mathsf{1}$.

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 CARA, and linear risk-tolerance each correspond to rank-two demand systems. Further, 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 $E(\boldsymbol{\pi},e) = E'\left(\rho^1(\boldsymbol{\pi}), ..., \rho^K(\boldsymbol{\pi}), e\right)$where $K < S$.

Constant relative risk aversion, thus, implies that preferences can be represented indirectly in terms of a composite of the state-claims, and 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 is associated with preferences that can be represented indirectly in terms of two composites, one of which is degenerate and corresponds to the traditionally safe asset. 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 in the portfolio allocation problem. 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

$$B(e, \mathsf{y}; \mathsf{1}) = g(\mathsf{y} - Min\{y_1, ..., y_S\}\mathsf{1}) + Min\{y_1, ..., y_S\} - e,$$

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

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

Because g is positively linearly homogeneous, $\inf_{\hat{\mathbf{y}}} \{\boldsymbol{\pi}\hat{\mathbf{y}} - g(\hat{\mathbf{y}})\}$ 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.

Theorem 11 *(Safra and Segal): Preferences exhibit constant risk aversion if and only if*

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

and $B(e, \mathbf{y}; \mathbf{1}) = \inf \left\{ \boldsymbol{\pi}\mathbf{y} : \boldsymbol{\pi} \in \hat{\mathcal{P}} \right\} - e$, *for* $\hat{\mathcal{P}} \subseteq \mathcal{P}$ *closed and convex.*

Proof *By Theorem 6, 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, proper concave function. To satisfy CRRA, it further follows from Theorem 6 that* $\mu\hat{a}(\boldsymbol{\pi}) = \hat{a}(\boldsymbol{\pi}) \quad \mu > 0$. *There are three possibilities: either* $\hat{a}(\boldsymbol{\pi}) = 0; \hat{a}(\boldsymbol{\pi}) = \infty$ *; or* $\hat{a}(\boldsymbol{\pi}) = -\infty$. *If* $\hat{a}(\boldsymbol{\pi}) = \infty$, *there is no* $\mathbf{y}$ *such that* $B(e, \mathbf{y}; \mathbf{1}) \geq 0$, *and hence* $V(e)$ *is empty. If* $\hat{a}(\boldsymbol{\pi}) = -\infty$ *for all* $\boldsymbol{\pi}$, *preferences are not well defined, and that case is ruled out. The only closed, proper concave function remaining is*

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

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

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

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

$$B(e, \mathbf{y}; \mathbf{1}) = \inf_{\boldsymbol{\pi}} \left\{ \boldsymbol{\pi}\mathbf{y} - E(\boldsymbol{\pi}, e) : \boldsymbol{\pi} \in \hat{\mathcal{P}} \right\} < \infty,$$

if it is to be finite.¥

Besides exhaustively characterizing the class of constant risk averse preferences, Theorem 11 , when combined with the interpretation of the expected-value function as an expenditure function in the presence of complete contingent claims, has an interesting consequence.

Corollary 12 *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 12 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 $Y \subseteq \Re^S$. 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 nonlinear but appropriately convex tax structures. We endow Y with the following properties $\mathbf{0} \in Y$, $Y \cap \Re^S_{++} \neq \emptyset$.

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), an equivalent dual formulation of the individual's choice problem is

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

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

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

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 $E(\boldsymbol{\pi},e) = R(\boldsymbol{\pi},Y)$. 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.[4]

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 $E(\boldsymbol{\pi},e)$ and state-claim supplies for R. Hence, the market clearing conditions require that there exist a $\mathsf{y} \in \partial E(\boldsymbol{\pi},e)$ such that

$$\mathsf{y} \in \partial R(\boldsymbol{\pi},Y),$$

where the notation ∂R denotes the subdifferential of R in $\boldsymbol{\pi}$. By Walras' Law and the basic properties of cost and revenue functions, one of the S market clearing conditions is redundant (alternatively $E(\boldsymbol{\pi},e) = R(\boldsymbol{\pi},Y)$ 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 $y^Y\mathbf{1}$,

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

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

Theorem 13 *For any $\hat{\boldsymbol{\pi}}$ consistent with the decisionmaker's choice equilibrium, $\boldsymbol{\pi}^Y \preceq_e \hat{\boldsymbol{\pi}} \preceq \boldsymbol{\pi}^0, \boldsymbol{\pi}^Y \in \mathcal{P}^Y$, where $\boldsymbol{\pi}^0$ are the maximally risk-neutral probabilities.*

[4] In this context, $\mathsf{E}(\boldsymbol{\pi},\mathsf{e}) - \mathsf{R}(\boldsymbol{\pi},\mathsf{Y})$ is exactly analogous to a trade expenditure function

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

In interpreting Theorem 13, one might think of $\boldsymbol{\pi}^{Y}$ 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 13 is the preference analogue of the famous Peleg and Yaari (1975) result characterizing the set of risk aversely 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 $\boldsymbol{\pi}^{0}$. 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 $\boldsymbol{\pi}^{Y} \in \mathcal{P}^{Y}$. 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 13 restricts choice associated with Y to lie within the subset of $V(e)$ determined by these supporting hyperplanes.

Theorem 13 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 $y^{Y}\mathbf{1} \in \partial E\left(\boldsymbol{\pi}^{Y}, y^{Y}\right)$ for some $\boldsymbol{\pi}^{Y} \in \mathcal{P}^{Y}$, 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 $\boldsymbol{\pi}(\mathbf{1})$ is a singleton. Suppose that, for some initial choice set Y, equilibrium is characterized by $\boldsymbol{\pi} \neq \boldsymbol{\pi}(\mathbf{1})$, so that the optimal $\mathbf{y}$ 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 14 *If for some initial choice set, Y, equilibrium $\boldsymbol{\pi} \neq \boldsymbol{\pi}(\mathbf{1})$, 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 $\mathcal{P}^Y$ is a singleton but provides what the decisionmaker views as unfair odds. It is the choice set analogue of the result, derived by Segal and Spivak (1991), 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 6 implies that the decisionmaker equilibrium is characterized by

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

and in the latter by

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

Because

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

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 $\mathbf{u}$,

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

whence:

Theorem 15 *If preferences are translation homothetic in the direction of* u*, replacing* Y *by* $Y+\delta\mathsf{u}$ *with* $\delta>0$ *raises the equilibrium certainty equivalent by* δ *with no effect on equilibrium* $\boldsymbol{\pi}$*.*

For the case of LRT, the equilibrium e is defined by

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

Denote the optimal choice of $\boldsymbol{\pi}$ by $\hat{\pi}$ and note that, since $E^{1}(\hat{\pi})\leq 1$,

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

which is to be expected in light of Theorem 10.

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 $Y+\delta\mathsf{1}$ 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 16 *Suppose preferences exhibit LRT and are risk-averse with respect to some* $\boldsymbol{\pi}^{0}$*. Replacement of* Y *by* $Y+\delta\mathsf{1}$ *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 $Y+\delta\mathsf{1}$. Let $\hat{\pi}$ denote the originally optimal choice of π and $\boldsymbol{\pi}^{\delta}$ the optimal choice for $Y+\delta\mathsf{1}$.

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

Now since

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

we must have

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

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

Theorem 16, in conjunction with Theorem 13, 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 16 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 16 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,

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

and, by Corollary 12,

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

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 $R(\boldsymbol{\pi},Y)$, do not exist. Instead, equilibrium, is determined by the decisionmaker 'plunging' to the bounds of the choice set as

$$\max_{\mathsf{y}}\{\inf\{\boldsymbol{\pi}\mathsf{y}:\boldsymbol{\pi}\in\mathcal{P}^*\}:\mathsf{y}\in Y\}=\inf\{R(\boldsymbol{\pi},Y):\boldsymbol{\pi}\in\mathcal{P}^*\},$$

The choice problem reduces to finding the least favorable $R(\boldsymbol{\pi},Y)$ consistent with $\boldsymbol{\pi}\in\mathcal{P}^*$. 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 17 *If the decisionmaker has constant risk averse preferences, her equilibrium certainty equivalent is given by* $\inf\{R(\boldsymbol{\pi},Y):\boldsymbol{\pi}\in\mathcal{P}^*\}$. *Replacing Y by $Y+\delta\mathsf{1}$ shifits the equilibrium certainty equivalent by δ, and replacing Y by tY for $t>0$ rescales the equilbrium certainty equivalent by t. If the decisionmaker plunges before Y is replaced by $Y+\delta\mathsf{1}$ 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 Concluding comments

Dual approaches have proved their value in many areas of economic analysis. Until recently, however, the analysis of choice under uncertainty has made little use of duality concepts. Instead reliance has been placed almost exclusively on primal methods, and, in particular, on the expected-utility model. Perhaps the best explanation of the endurance of the expected-utility model, in spite of its well-known weaknesses, is its ability to yield predictions about economic behavior. Much of the predictive 'bite' of the expected utility model comes from what many regard as its Achilles heel, the independence axiom and its consequent additive separability.

Additively separable preferences were discarded as a reasonable representation of preferences in standard consumer theory long ago . Instead, reliance is usually placed on direct assumptions about the nature of the decisionmaker's preference map. In particular, the notions of homotheticity and quasi-homotheticity have proven very useful in both empirical and theoretical analyses. These restrictions have percolated into expected-utility theory, albeit in disguised form, as the notions of CARA, CRRA, and LRT. This paper has taken advantage of this equivalence and the dual formulation to show how behavior can be fully characterized without imposing additive separability, and that a dual formulation of choice under uncertainty is both straightforward and analytically productive.

Reference List

Blackorby, C., and D. Donaldson. "A Theoretical Treatment of Indices of Absolute Inequality." *International Economic Review* 21, no. 1 (February1980): 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 (November1980): 1711–25.

Epstein, L. G., and S. E. Zin. "First–Order Risk Aversion and the Equity Premium Puzzle." *Journal of Monetary Economics* 26 (1990): 387–407.

Feder, G. "The Impact of Uncertainty on a Class of Objective Functions." *Journal of Economic Theory* 16 (1977): 504–12.

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.

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 (May1965): 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.

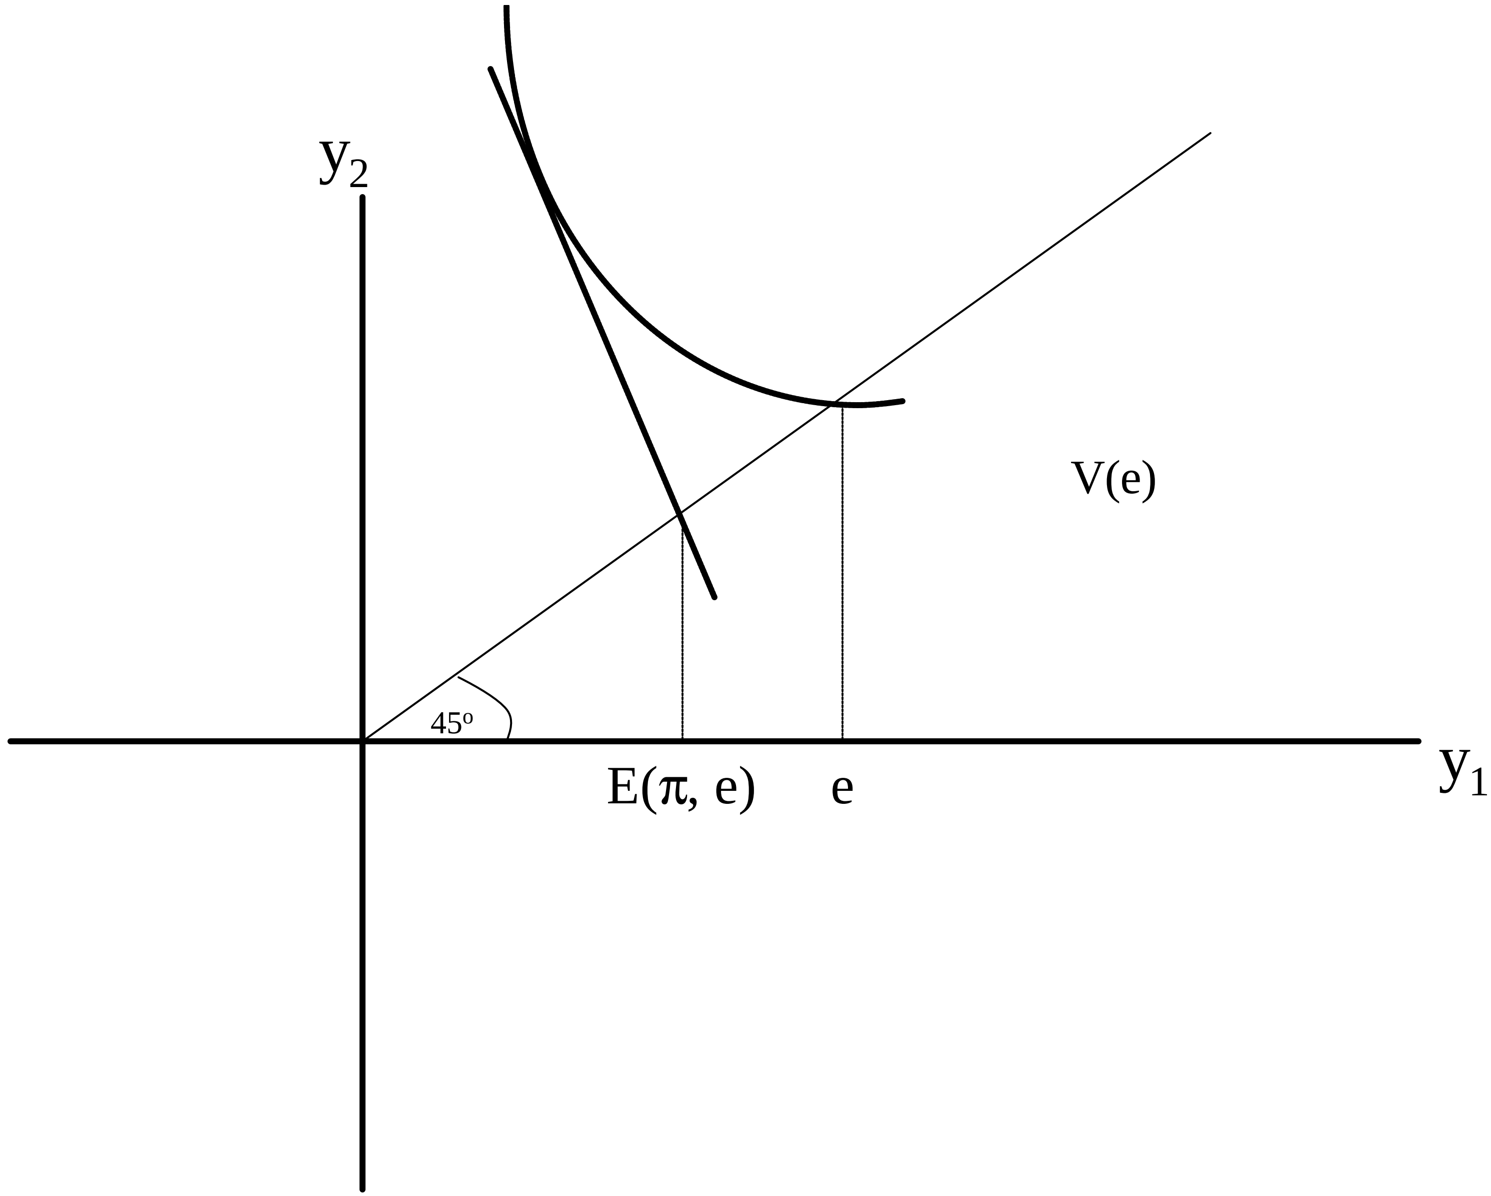

Figure 1: Certainty Equivalent and E(π, e)

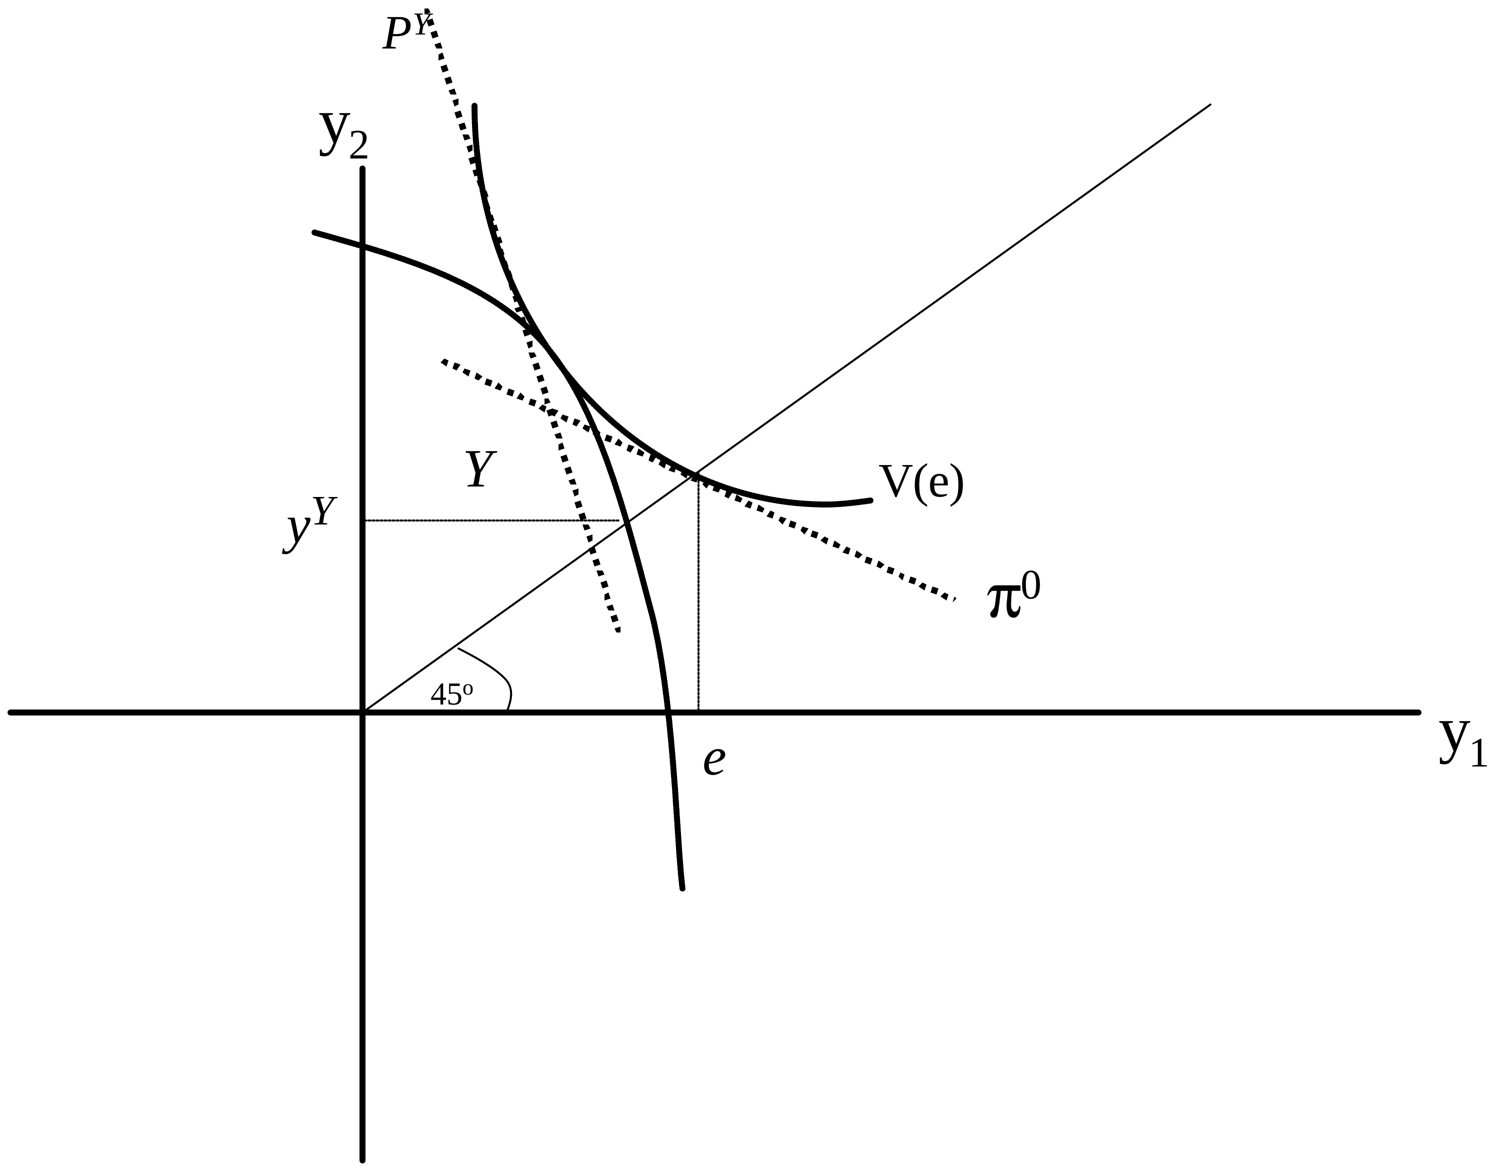

Figure 2: Range of Equilibrium Probabilities

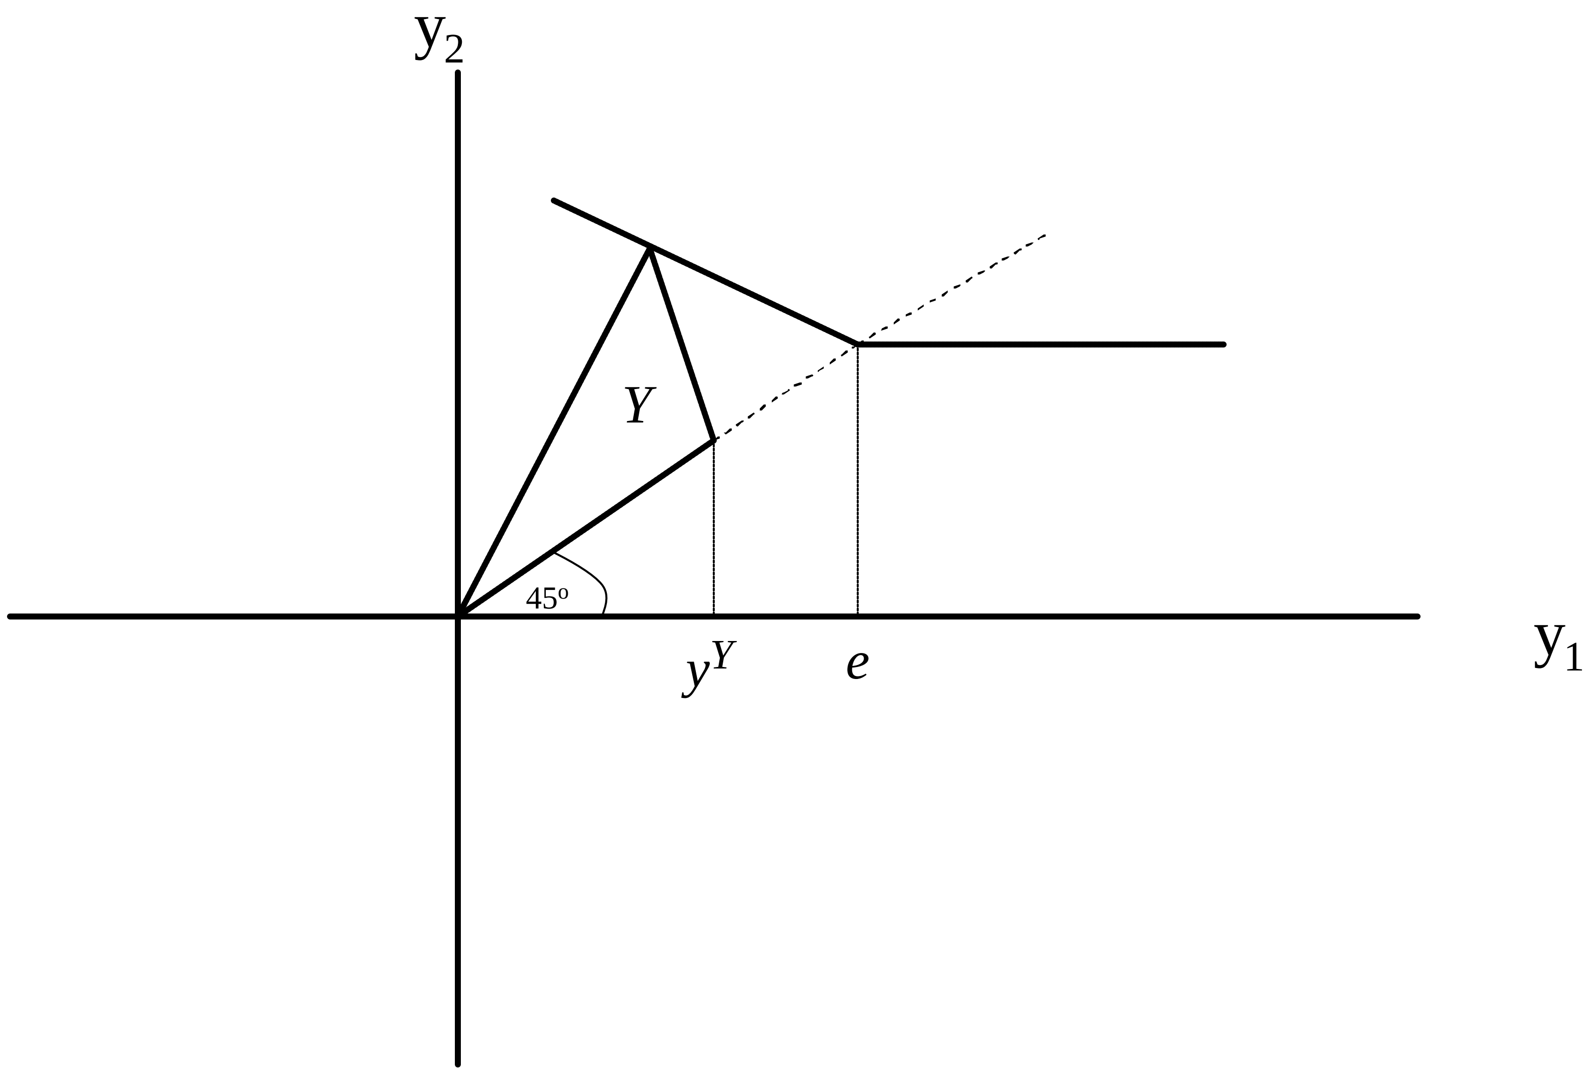

Figure 3: Plunging with Constant Risk Aversion