

Behavior under risk: recent developments in theory and applications

Larry G. Epstein*

1 INTRODUCTION

In the last decade a number of new theories have been proposed to explain individual behavior under risk where, following Knight (1921), risk is defined as randomness with a known probability distribution. Some of these theories are formally atemporal and generalize the classical expected utility model of choice. Their development was inspired primarily by the growing body of laboratory evidence regarding static or one-shot choices, that has cast doubt upon the descriptive validity of the expected utility model. Other theories are explicitly intertemporal and generalize the time-additive expected utility model which is standard in capital theory. The noted laboratory evidence also provides some motivation for this work since it is clearly desirable that a theory of intertemporal utility, when restricted to static gambles, be consistent with the evidence.

The capacity to explain behavior in the laboratory is one criterion that might be applied to a theory of choice under risk. At least as important, however, is that the theory be useful as an engine of inquiry into standard market-based economic questions. The expected utility model has proven extremely successful in this respect. By application of now standard techniques, modelers have been able to derive a rich set of predictions in a variety of contexts. Moreover, a substantial body of non-experimental evidence has been shown to conform well to the expected utility hypothesis. The ultimate influence of the new theories on the profession at large will probably depend on whether they can match the elegance and power of expected utility as a tool of analysis and whether they can significantly improve the explanation of non-experimental evidence. In order to cast light upon these questions, I will survey some of the new theories of choice and their applications to a number of standard problems in macroeconomics, finance, and game theory.

The limitations and objectives of this chapter should be made clear at the outset. First, I will emphasize economic applications and, where it exists, market-based evidence. The theories themselves will be described only to the extent necessary to understand the applications. This is particularly so for the formally atemporal theories, since there already exist a number of excellent surveys dealing with them and their relation to laboratory evidence (Machina, 1983a, 1987; Sugden, 1986; Weber and Camerer, 1987; Fishburn, 1988; Karni and Schmeidler, forthcoming; Camerer, 1989b). Second, a number of papers have demonstrated that many results, in diverse areas of economic analysis, are robust to generalizations of expected utility. In contrast, I will focus below primarily on those instances in the literature where it has been shown that the added flexibility provided by the generalization of expected utility “makes a difference” because it delivers either

- 1 added analytical power and consequently new theoretical insights, or
- 2 interesting new testable implications regarding market behavior.

It will be useful to distinguish between *static* or 1-shot choice problems and *dynamic* choice problems. Typically, in dynamic problems a number of decisions are made subject to different information as a result of the resolution of some risks, and thus the issue of consistent choice arises. Within dynamic problems distinguish further between *intertemporal* choice, where consumption sequences are the objects of choice and source of utility – as in capital theory – and *sequential* choice, where terminal wealth is the source of utility. In a sequential problem there are several stages at which decisions are made but the entire process involves so little time that consumption/savings plans may be reasonably viewed as fixed. Alternatively, the multistage nature of a decision problem could be due exclusively to the way in which the problem is perceived by the individual. In that case, no real time passes as the stages are traversed, but the problem is still sequential. Many games fit into the sequential choice framework.

In order to better achieve the objective stated earlier, I will emphasize the literature on dynamic choice with non-expected utility preferences. Static choice will be discussed primarily in order to facilitate understanding of the theory and applications of dynamic choice. Specifically, the chapter proceeds as follows. The next section describes some atemporal theories and applications. Intertemporal utility theory is reviewed in section 3 and then section 4 describes applications to consumption theory and asset pricing. Finally, section 5 considers sequential choice and then how received game theory is affected by generalized specifications of utility. The first part of section 5 parallels closely and relies upon section 3.

2 STATIC CHOICE

A decision problem under risk is static if all decisions must be made at a single point in time and dynamic otherwise. The bulk of the laboratory evidence cited earlier is based on static choice behavior. Accordingly, the generalizations of expected utility that have been developed to explain that evidence have been formulated and analyzed primarily from the perspective of static choice problems. In fact, these new theories are strictly speaking inapplicable to dynamic choice problems, since they invariably specify a **single** utility function rather than, as would be needed, a **sequence** of utility functions, one for each decision time. This section describes some of these static theories and some results that they have produced. A primary objective, however, is to lay the groundwork for the following sections where it is shown how these theories may be extended and applied to dynamic choice settings.

The objects of choice are lotteries, where a lottery is represented by a cumulative distribution function (cdf) over R^n , or by the underlying random variable, or by a probability measure on some more general space of outcomes. Utility functions are defined over lotteries. I will concentrate on the case of real-valued outcomes; extensions to more general prizes are available in the cited references.

2.1 Generalized expected utility analysis

The standard description of choice from a set of feasible lotteries is based on the maximization of an expected utility function, or equivalently, of a function that is “linear in probabilities.” Machina (1982b) points out that, just as in ordinary calculus, analysis of linear functions can be readily extended to the analysis of “locally linear” or “smooth” functions in such a way that many of the techniques of expected utility analysis apply to suitably smooth utility functions.

Formally, Machina considers the set $D[a,b]$ of cdf’s defined on a bounded interval $[a,b]$ and restricts the utility function V to be Fréchet differentiable with respect to the L^1 metric, $\|F - F_0\| \equiv \int_a^b |F(x) - F_0(x)| dx$. Such differentiability is equivalent to V being “locally expected utility or linear in probabilities.” To be precise, there exists $u: [a,b] \times D[a,b] \rightarrow R^1$ with $u(\cdot; F)$ absolutely continuous for all F , such that throughout the domain:

$$V(F) - V(F_0) = \int_a^b u(x; F_0) d(F(x) - F_0(x)) + o(\|F - F_0\|), \quad (2.1.1)$$

where $o(|x|)/|x| \rightarrow 0$ as $|x| \rightarrow 0$. Consider the probability simplex for three-outcome lotteries (figure 1.1), where the outcomes are fixed and




Figure 1.1 Indifference curves for gambles with outcomes $x_1 < x_2 < x_3$ and probabilities $p_1, p_2, p_3, p_2 = 1 - p_1 - p_3$.

different points in the triangle represent different probability vectors and hence different lotteries. For expected utility functions indifference curves in the simplex are parallel straight lines, while for generalized expected utility functions indifference curves are generally non-linear but they have unique tangent lines everywhere.

The function $u(\cdot; F_0)$ is V 's local utility function at F_0 . If it does not depend on F_0 then expected utility is obtained globally. Even in general, however, there is a strong similarity to expected utility analysis in that the monotonicity and concavity of $u(\cdot; F_0)$ for all F_0 are equivalent to the global increasingness and risk aversion of V respectively. These and other results derived by Machina are based on a fundamental insight of generalized expected utility analysis: as in ordinary calculus, one can use integration to "piece together" qualitative results regarding differential (local) situations

to derive qualitative non-differential (global) results.¹ Formally, let $\{F_\alpha: \alpha \in [0,1]\}$ be a path such that for each $\alpha^* \in [0,1]$, $\|F_\alpha - F_{\alpha^*}\|/|\alpha - \alpha^*|$ is bounded in α for α near α^* , e.g., $\|F_\alpha - F_{\alpha^*}\|$ is differentiable in α at $\alpha = \alpha^*$. Then:

$$\left. \frac{d}{d\alpha} V(F_\alpha) \right|_{\alpha^*} = \frac{d}{d\alpha} \left[\int u(x; F_{\alpha^*}) dF_\alpha(x) \right] \Big|_{\alpha^*} \quad (2.1.2)$$

and hence, under the conditions of the Fundamental Theorem of Calculus:

$$V(F_1) - V(F_0) = \int_0^1 \frac{d}{d\alpha} \left[\int u(x; F_{\alpha^*}) dF_\alpha(x) \right] \Big|_{\alpha^*} d\alpha^*. \quad (2.1.3)$$

If the sign of the integrand on the right side is uniform across α^* , then the same sign will be shared by the left side.

There exist theoretical arguments for postulating a Fréchet differentiable but non-linear utility function. Such a specification can be justified if F represents a delayed risk (Machina, 1984) or if V represents group preference and F represents a gamble which is to be shared optimally among members of the group, each of whom may conform with the expected utility model (Machina, 1989b).

On the other hand, the assumption of Fréchet differentiability is not innocuous. For example, the disappointment averse utility functions defined in section 2.3 are not Fréchet differentiable. For another example, Chew, Karni, and Safra (1987) show that rank-dependent expected utility functions (see Quiggin, 1982; Yaari, 1987; Segal, 1989 for definitions and axiomatizations) are generally not Fréchet differentiable. Fortunately, the former authors show that the weaker property of Gâteaux differentiability is sufficient to provide a notion of a local utility function and subsequently Machina's (1982b) main results. (Other notions of smoothness are exploited in Chew and Nishimura, forthcoming; Chew, Epstein, and Zilcha, 1988; Wang, 1991.)

Another important contribution of Machina (1982b) is the formulation of a hypothesis regarding utility, called Hypothesis II, that is shown to be closely connected to consistency with Allais-type behavior and with the laboratory-based empirical patterns that have come to be known as the common consequence and common ratio effects. Hypothesis II states that indifference curves *fan out* as shown in figure 1.1, or that they become steeper as one moves upward along any vertical line. The opposite pattern is referred to as *fanning in*. A natural and important question, to be addressed in section 4.5, is whether Hypothesis II is useful in explaining market data.

2.2 Betweenness-conforming utility

For the sake of greater specificity, it is desirable to restrict utility functions by more than a smoothness requirement. There are a number of alternative axiomatically based generalizations of expected utility theory that have been developed, but the one which seems to me to strike the optimal balance between generality and tractability, at least for the applications that I will consider, is the *betweenness* theory due to Chew (1983, 1989), Fishburn (1983), and Dekel (1986).² The cornerstone axiom imposes the following requirement on a preference ordering of cumulative distribution functions, where \sim denotes indifference:

Betweenness: If $F \sim G$, then $\alpha F + (1 - \alpha)G \sim F$ for all $\alpha \in (0, 1)$.

The relation between this axiom and the independence axiom ($F \sim G \Rightarrow \alpha F + (1 - \alpha)H \sim \alpha G + (1 - \alpha)H$) is clarified by their respective implications for indifference curves in the probability simplex corresponding to three outcome lotteries. Under betweenness, indifference curves are straight lines but not necessarily parallel (see the two examples in the bottom half of figure 1.1).³

Betweenness delivers (given some other specified axioms) an elegant and tractable functional form for the representing utility function V . In sections 3 and 4 it will be convenient to work with the *certainty equivalent* representation of utility μ , where, for any F , $\mu(F)$ is defined implicitly by:

$$V(\delta_{\mu(F)}) = V(F). \quad (2.2.1)$$

Here δ_x denotes the cdf corresponding to the lottery with the certain outcome x . Consequently, $\mu(F)$ equals that amount of money which, if received with certainty, would be indifferent to the lottery F . If the preference ordering satisfies betweenness and other standard assumptions, then there exists a function $H: \mathbb{R}^2 \rightarrow \mathbb{R}^1$, with $H(x, x) \equiv 0$, such that $\mu(F)$ is defined implicitly as the unique solution to:

$$\int H(x, \mu(F)) dF(x) = 0, \quad (2.2.2)$$

for each cumulative distribution function F in the domain of μ . If H is increasing and concave in its first argument and decreasing in its second argument, then μ is increasing in the sense of first-order stochastic dominance and risk averse in the sense of being averse to mean-preserving spreads (see footnote 4). Expected utility is obtained if H is specialized to $H(x, z) \equiv v(x) - v(z)$.

The above functional structure provides a perspective on the relation with expected utility functions. Suppose that F^* is an optimum according to

μ in some feasible set D . If H is decreasing in its second argument, then $\mu(F) \leq \mu(F^*) \Rightarrow$

$$\int H(x, \mu(F^*)) dF(x) \leq 0 = \int H(x, \mu(F^*)) dF^*(x) \quad \forall F \in D,$$

or

$$F^* \in \operatorname{argmax} \left\{ \int H(x, \mu(F^*)) dF(x) : F \in D \right\}. \quad (2.2.3)$$

That is, F^* is also optimal according to the expected utility function with von Neumann–Morgenstern index $H(\cdot, \mu(F^*))$. Of course (2.2.3) does not imply that betweenness conforming functions are empirically indistinguishable from expected utility functions, since $H(\cdot, \mu(F^*))$ generally varies with F^* and hence with the feasible set D . However (2.2.3) and the fact that $H(\cdot, \mu(F^*))$ depends on F^* only through its utility level help to explain why betweenness functions retain much of the tractability of expected utility as demonstrated in both sections 4 and 5.

When H is sufficiently differentiable, a similar point can be made by reference to (2.1.2)–(2.1.3) since then μ is Fréchet differentiable with local utility function:

$$u(x, F) = -H(x, \mu(F)) / \int H_2(y, \mu(F)) dF(y), \quad (2.2.4)$$

for which the local measure of risk aversion $-u_{11}(x, F)/u_1(x, F)$ depends on F via $\mu(F)$. But, as mentioned earlier, there are interesting examples of betweenness-conforming utility functions (see (2.3.6)) that are not Fréchet differentiable and thus the more general perspective provided by (2.2.3) is useful. In conjunction with the underlying functional form (2.2.2), it also suggests the alternative and possibly more illuminating name “implicit expected utility functions,” first used by Dekel (1986), for betweenness-conforming utility functions.

2.3 Parametric functional forms

Two parametric specializations of (2.2.2) will be described here, both in order to clarify further the nature and variety of betweenness-satisfying utility functions and also because they will be applied in section 4.5. For empirical tractability and consistent with the existing empirical asset pricing literature, the hypothesis of *constant relative risk aversion* is adopted for μ , i.e.:

$$\mu(F_{\lambda \tilde{x}}) = \lambda \mu(F_{\tilde{x}}) \text{ for all } \lambda > 0, \quad (2.3.1)$$

where $F_{\lambda \tilde{x}}$ and $F_{\tilde{x}}$ denote the cdf's for the random variables $\lambda \tilde{x}$ and \tilde{x} .

Restrict attention to cdf's on R_{++}^1 . Given the above assumption, the functional structure in (2.2.2) simplifies to:

$$\int \phi(x/\mu(F))dF(x) = 0, \quad (2.3.2)$$

where $\phi(x) \equiv H(x, 1)$ and so ϕ is increasing and concave on an appropriate subset of R_{++}^1 with $\phi(1) = 0$.⁴

For the first parametric functional form let:⁵

$$\phi(x) = x^\delta(x^\alpha - 1)/\alpha, \quad \alpha \neq 0, \quad \alpha + 2\delta < 1. \quad (2.3.3)$$

This leads to the constant relative risk averse specialization of Chew's (1983) *weighted utility* having the explicit representation:

$$\mu_{wu}(F) = \left\{ \int x^{\alpha+\delta} dF(x) / \int x^\delta dF(x) \right\}^{1/\alpha}. \quad (2.3.4)$$

Expected utility is obtained if $\delta = 0$. (To obtain the limiting form corresponding to $\alpha = 0$, use the fact that $(x^\alpha - 1)/\alpha \rightarrow \log x$ as $\alpha \rightarrow 0$; similarly for other functional forms below.)

Under the restriction $\alpha + 2\delta < 1$, ϕ in (2.3.3) is both increasing and concave in a neighborhood of 1, from which it follows (see footnote 4) that there exists an interval $[a, b]$ containing 1 such that μ_{wu} is monotone and risk averse on $D[a, b]$. Under some auxiliary assumptions, μ_{wu} is well-behaved in this sense for all cdf's having finite mean and support in the positive real line. Those assumptions are:

- (i) $\delta = 0$ and $\alpha < 1$, or
 - (ii) $\delta < 0$ and $0 < \alpha + \delta < 1$, or
 - (iii) $0 < \delta < 1$ and $\alpha + \delta < 0$.
- (2.3.5)

For weighted utility, indifference curves in the three-outcome probability simplex all emanate from a single point which is at infinity in the expected utility special case. The case of fanning out is shown in figure 1.1 and corresponds to $\delta < 0$, but fanning in occurs if $\delta > 0$ when the point of projection is on the north-east side of the triangle.

An alternative generalization of expected utility is obtained by choosing ϕ to satisfy:

$$\phi(x) = \begin{cases} (x^\alpha - 1)/\alpha, & x \leq 1 \\ A(x^\alpha - 1)/\alpha, & x \geq 1 \end{cases} \quad (2.3.6)$$

where $\alpha < 1$ and $0 < A \leq 1$. This leads to the constant relative risk averse specialization of Gul's (1991) *disappointment averse* utility functions. The certainty equivalent μ_{da} is defined implicitly by:

$$\mu_{da}^\alpha(F)/\alpha = \int x^\alpha dF(x)/\alpha + (A^{-1} - 1) \int_{x < \mu_{da}(F)} [(x^\alpha - \mu_{da}^\alpha(F))/\alpha] dF(x). \quad (2.3.7)$$

If $A = 1$, one obtains the common homogeneous expected utility function. Smaller values for A reflect an aversion to disappointment in the following sense: refer to an outcome as disappointing if it is worse than expected in the sense of being smaller than the certainty equivalent of the lottery. When $A < 1$ disappointing outcomes generate negative values for the second integral on the right side of (2.3.7), producing a smaller certainty equivalent value than if $A = 1$.

The indifference map for μ_{da} , shown in figure 1.1, features two distinct sources of projection Q and Q' both of which recede to infinity as $A \rightarrow 1$. Indifference curves fan out in the lower part of the triangle and otherwise fan in. It is interesting to note in this regard that the behavioral evidence supporting fanning out is weaker in the upper triangle than in the lower region. (See Conlisk, 1989, for example.) Finally, with regard to the domain of μ_{da} note that ϕ is increasing and concave on $(0, \infty)$ and refer to footnote 4.

The weighted utility and disappointment averse functional forms have several properties in common. First, they have both been used to explain laboratory-based evidence against expected utility theory. Secondly, they both constitute single parameter extensions of the common constant relative risk averse expected utility specification. Moreover, they retain the tractability of the latter for empirical work (section 4.5) and thus allow some econometric evidence to be brought to bear upon the significance of the independence axiom or the fanning out property for market data. Finally, for each functional form there is a simple qualitative relation between the parameters and the degree of risk aversion. To elaborate, say that the certainty equivalent μ^* is *more risk averse than* μ if $\mu^*(F) \leq \mu(F)$ everywhere.⁶ Then for μ_{wu} risk aversion decreases with α and δ and for μ_{da} it decreases with α and A .

However, the functions μ_{wu} and μ_{da} differ from one another in the way in which they evaluate small gambles. Since this difference will be of some importance in the subsequent discussion of asset pricing models, I move now to an examination of risk premia for small gambles.

2.4 First-order risk aversion

Let π be the risk premium associated with an actuarially fair gamble $t\tilde{\epsilon}$ and initial non-stochastic wealth x . Thus the decision-maker is indifferent between the gamble $x + t\tilde{\epsilon}$ and the sure prospect $x - \pi$. Following Segal and Spivak (1990), say that the utility function exhibits *first-order* (resp.

second-order) risk aversion if, for all x and for all \tilde{e} with non-zero and finite variance, π is of the order of t (resp. t^2) as $t \rightarrow 0$. In that sense, the risk premium for a small gamble is proportional to the standard deviation of the gamble under first-order risk aversion, rather than to its variance as in the more familiar case of second-order risk aversion. A useful graphical representation is possible in outcome space in the case of binary gambles. As portrayed in figure 1.2, given second-order risk aversion the indifference curve is tangent to the actuarially fair market line at certainty, reflecting risk neutrality to the first order. In the case of first-order risk aversion, there is a kink at the certainty line.

Expected utility functions generally exhibit second-order risk aversion. In the twice differentiable case, this is reflected by the famous Arrow–Pratt formula for the premium associated with a small gamble. More generally, kinks along the certainty line are rare in that an increasing and concave function can fail to be differentiable at only a countable number of points. In particular, second-order risk aversion applies if the von Neumann–Morgenstern index is homogeneous of any degree. More generally, the constant relative risk averse weighted utility function μ_{wu} also exhibits second-order risk aversion. On the other hand, the disappointment averse form is first-order risk averse if $A < 1$; the absolute slopes of the indifference curve in figure 1.2 on either side of the certainty point c are $A^{-1}p_1/p_2$ and Ap_1/p_2 where p_1 and p_2 are the probabilities of the two outcomes.

In terms of usefulness of first-order risk aversion, Segal and Spivak (1990) point out a number of empirical implications which seem consistent with observation. For example, there is a lesser tendency to diversify towards a risky asset while holding a safe asset.⁷ More precisely, such diversification is optimal under second-order risk aversion if the mean excess return to the risky asset is positive, while a sufficiently large mean excess return is necessary for diversification given first-order risk aversion. More generally, only for second-order risk aversion is it true that any favorable bet is desirable at a sufficiently small scale, i.e., any favorable market line through the certainty point c in figure 1.2 lies above the indifference curve shown somewhere near c .

There is a sense in which kinks along the certainty line are typical for community utility functions in a multiperson model with incomplete markets. Consider two individuals with monotonic preferences and an aggregate gamble which is shared according to a given inefficient allocation rule. Let the representative aggregate gamble have equally likely outcomes e_1 and e_2 . Individual preference orderings over individual gambles induce an ordering of aggregate gambles for each agent. Thus for each wealth level e , we can define $I(e)$ as the boundary of $\{(e_1, e_2): \text{each individual prefers}$




Figure 1.2 First and second-order risk aversion

(e_1, e_2) to (e, e) . Then $I(e)$ is a form of community indifference curve through (e, e) , such that (under specified assumptions) points that lie above it are Pareto superior to (e, e) . Because of the inefficiency of the given allocation rule, it will typically be the case that $I(e)$ will be kinked at certainty as in figure 1.2, even if both agents are expected utility maximizers. (This will be true, for example, if one agent receives the fraction $\alpha_i \in (0, 1)$ of the aggregate endowment in state i , $i = 1, 2$, $\alpha_1 \neq \alpha_2$, if the remaining endowments are allocated to the other agent and if their common von Neumann–Morgenstern index is not logarithmic.)

Finally, I would like to describe a separate argument, based on functional form flexibility, for being interested in first-order risk aversion. As has already been noted, for reasons of tractability the assumption of constant relative risk aversion is common in the asset pricing literature, particularly in empirical studies such as those cited in section 4. Thus let a decision-maker have constant relative risk aversion and suppose further that he is an expected utility maximizer. Following Kandel and Stambaugh (1990), consider how such an individual would evaluate binary symmetric gambles with outcomes $x \pm \varepsilon$. For concreteness, let $x = 75,000$. If relative risk aversion is 2, the individual would pay “only” 0.83 to avoid the “small” gamble corresponding to $\varepsilon = 250$, while the willingness to pay rises to 12.48 if relative risk aversion equals 30. But with the larger degree of risk aversion he would pay 1,091.17 to avoid the moderately sized gamble with $\varepsilon = 2,500$ and 23,790.52 to avoid the large gamble having $\varepsilon = 25,000$. The “reason” for these implications of the common homogeneous expected utility function is that it is risk neutral to the first order. Thus it can generate a plausible risk premium for small gambles only if the coefficient of relative risk aversion is so large that the implied risk premium for large gambles is unrealistically large. Epstein and Zin (1990b) show how a first-order risk

averse utility function is better able to model “plausible” risk attitudes over a broad range of gamble sizes. For a numerical illustration, let $v(x) = \log x$ and $A = 5/6$. Then to avoid the three gambles above, the decision-maker would be willing to pay 22.50, 265.50, 6,373.57 respectively. Collectively, these certainty equivalent values seem more plausible on introspective grounds than those obtainable from any homogeneous expected utility function.

First-order risk aversion is of interest below in attempting to fit aggregate US time series data. That is because the smoothness of aggregate consumption data implies that the evaluation of small gambles is critical, while frequently the “plausibility” of a utility function is judged informally on the basis of its ranking of “real-life” moderate or large gambles. To develop this intuition, it is first necessary to consider the problem of intertemporal choice and the integration of the above atemporal theories into a temporal framework. Then their potential role in helping to explain time series data can be considered. Before proceeding with that principal thrust of this survey, however, I will mention briefly some results which have been derived using generalized utility functions in static choice settings.

2.5 Results

Many of the results which have been derived with the above utility functions in static choice settings are direct extensions of well-known propositions from expected utility theory. For the reasons indicated earlier through (2.2.3) for betweenness functions and (2.1.2)–(2.1.3) for “smooth” functions, many of the useful theoretical properties and behavioral implications of expected utility are preserved. These include the Arrow–Pratt characterization of comparative risk aversion across individuals (Machina, 1982b), the Ross characterization of comparative risk aversion (Machina and Nielsen, 1987), the Kihlstrom–Mirman characterization of comparative multivariate risk aversion (Karni, 1989), and the Rothschild–Stiglitz and Diamond–Stiglitz comparative statistics predictions regarding the effects of changing risk (Machina, 1989c and Chew and Nishimura, forthcoming).

Such robustness led Machina (1982b, p. 279) to state that even without the independence axiom “the implications and predictions of theoretical studies which use expected utility analysis typically will be valid, provided preferences are smooth.” Subsequently, the displayed robustness of expected utility predictions has been interpreted as a defense of expected utility analysis; for example, see the finance textbook by Huang and Litzenberger (1988, p. 16). However, there are important examples where

there is a payoff to considering a broader class of preferences. Sections 4 and 5 show that in dynamic choice settings, including an intertemporal asset pricing framework, generalizations of the expected utility model can be advantageous for both theoretical and empirical work.

In static choice models, advantageous applications of generalized utility functions have been made to address the Friedman–Savage observations on insurance and lotteries (Machina, 1982b), formulate and apply strong notions of declining risk aversion (Machina, 1982a and Epstein, 1985), provide an explanation of individual investor behavior in the stock market (Shefrin and Statman, 1984, 1985), clarify the foundations of axiomatic Nash bargaining theory (Rubinstein, Safra, and Thomson, 1990), and contribute to the theory of social choice and inequality measurement (Fishburn, 1988; Yaari, 1988; and Chew and Epstein, 1989b).

3 INTERTEMPORAL UTILITY

This section is concerned with utility functions defined on infinite horizon stochastic consumption programs, with a primary focus being on recursive utilities. When restricted to *timeless gambles*, i.e., those for which all uncertainty is resolved before further consumptions/savings decisions are made, recursive utility coincides with one of the atemporal certainty equivalent functions discussed in section 2. Even if the ranking of timeless gambles conforms with expected utility theory, however, the temporal resolution of consumption risk may matter. Hence recursive utility functions generally do not agree with intertemporal expected utility theory on the domain of consumption programs.

A primary motivation for this work is the desire to disentangle intertemporal substitution from risk aversion as explained in section 3.1. One way in which this has been accomplished, via recursive intertemporal utility functions, is described in section 3.3, after the important issue of intertemporal consistency has been clarified in section 3.2. Some normative issues are considered next.

3.1 Risk aversion and intertemporal substitution

In much of received capital theory it is assumed that the ranking of intertemporal stochastic consumption programs $\tilde{c} = (\tilde{c}_0, \tilde{c}_1, \dots)$ may be represented by a function of the form:

$$V(\tilde{c}) = E_0 \sum_{t=0}^{\infty} \beta^t u(\tilde{c}_t). \quad (3.1.1)$$

Here \tilde{c}_t denotes random scalar consumption at time t , $0 < \beta < 1$ is the rate of

discount and E_0 denotes the expected value operator conditional upon period 0 information (see section 3.2 for some measure-theoretic details). A property of the felicity function u which is of particular interest is its curvature as measured by $-cu''(c)/u'(c)$. This elasticity is often referred to as the measure of relative risk aversion (with respect to consumption gambles in any single period). It is also inversely related to the willingness to substitute consumption across time. For example, in the case of the homogeneous specification:

$$u(c) = \begin{cases} c^\rho/\rho, & \rho \neq 0 \\ \log c, & \rho = 0, \end{cases} \quad (3.1.2)$$

the constant elasticity of intertemporal substitution σ equals $(1 - \rho)^{-1}$, while the measure of risk aversion equals $(1 - \rho)$. Thus a precise inverse relation between intertemporal substitutability and risk aversion is imposed a priori.

Such a restriction is unfortunate firstly because risk aversion and substitution correspond to two conceptually distinct aspects of preference – one concerns the attitude toward the variation in consumption across states of the world (at a given time) while the other is concerned with variations across time (in the absence of risk). Some economists might conjecture that the pairings (high risk aversion, low substitution) and (low risk aversion, high substitution) would be more common empirically than the other possibilities, i.e., either one strongly dislikes “change” or one does not. But the validity of the conjecture cannot be determined unless the a priori constraint can be relaxed. Besides, the standard functional form (3.1.1) goes even further since it imposes a quantitative relation between the measures of substitution and risk aversion.

A consequence of this inflexibility of the additive expected utility specification is that the effects of increased risk aversion in the model under study cannot be determined. A comparative statics analysis based upon a change in the curvature of the felicity function u does not admit an unambiguous interpretation since both intertemporal substitutability and risk aversion are changed in this way. Lucas (1978, p. 1441) points this out in attempting to understand the determinants of equilibrium asset prices. There are many other theoretical contexts in which the effects of greater risk aversion are of interest.

In the empirical literature also, this inflexibility has been of concern. As described further in section 4.1, expected utility, representative agent optimizing models have not performed well in explaining asset returns and aggregate consumption data and the noted inflexibility of the utility specification has been suggested as one possible reason. For example,

Grossman, Melino, and Shiller (1987) claim that the data suggest that two parameters are needed in place of the single parameter ρ . The situation is reminiscent of that prevailing in demand theory when the Cobb-Douglas or even CES functional forms were the dominant specifications for utility. The severe constraints imposed by these specifications on the pattern of substitution across commodities led to the adoption of more “flexible” forms such as the Generalized Leontief or Translog. Similarly, it would be desirable to have more flexible functional forms that are empirically tractable for the study of consumption and asset return data.

Move from considering risks that are confined to consumption in a single period to the more relevant case of multiperiod risks. For the reasons given by Kihlstrom and Mirman (1974, pp. 365–6), comparisons of risk aversion should be restricted to preferences with a given ordering of deterministic programs. Thus the best that can be hoped for is a class of utility functions for which it is possible to change the degree of risk aversion without affecting the preference ranking of deterministic consumption paths.⁸ One possibility is to apply the Kihlstrom and Mirman (1974) approach to multicommodity risk aversion that retains the expected utility framework and considers ordinally equivalent von Neumann–Morgenstern indices. For example, let h be increasing and concave and define:

$$V^*(\tilde{c}) \equiv E_0 h \left[\sum_0^{\infty} \beta^t u(\tilde{c}_t) \right].$$

The functions V^* and V agree ordinally on deterministic programs but V^* is *more risk averse than* V in the sense that any gamble $(\tilde{c}_0, \tilde{c}_1, \dots)$ that would be rejected by V in favor of some deterministic program, would also be rejected by V^* . (This notion of comparative risk aversion for intertemporal utility functions is adopted throughout the paper; it is consistent with the definition of “more risk averse than” for certainty equivalents described in section 2.3, but see footnote 6.) Moreover, if h is strictly concave then there exist gambles that would be rejected by V^* but accepted by V .

However, this approach encounters serious difficulties in a temporal framework with discounting. Thus I feel that a separation of substitution from risk aversion that is likely to be useful in the applied contexts considered below is achievable only outside the expected utility framework.

To see the nature of the difficulties, consider an individual with the function V^* who arrives at period T and contemplates the remaining future. If previous consumption levels were $\tilde{c}_0, \dots, \tilde{c}_{T-1}$ and if tastes are not changing (see below for a definition), then the utility function for the remaining future is:

$$V^{*T} = E_T h \left[\sum_0^{T-1} \beta^t u(\bar{c}_t) + \sum_T^\infty \beta^t u(\tilde{c}_t) \right],$$

where E_T is the expected value operator conditional upon period T information. Suppose that h does not have constant absolute risk aversion (that case will be covered below), then risk attitudes at T depend implausibly upon the past in that the risk premium for a small gamble in c_T is affected more by a small change in \bar{c}_0 than by a corresponding change in \bar{c}_{T-1} , locally near $\bar{c}_0 = \dots = \bar{c}_{T-1}$. This follows readily upon examination of the Arrow–Pratt absolute risk aversion measure for period T consumption risks (see Epstein and Zin, 1989, p. 951). Such implausible dependence upon the past is not limited to the case where the von Neumann–Morgenstern index is additively separable across time. For example, in a similar fashion one can show that it prevails also if the felicity function u in each period t depends upon finitely many lagged values of consumption, or if it is a function of c_t and z_t , where $z_t = \delta \sum_0^\infty (1 - \delta)^i c_{t-i}$ and $1 - \delta < \beta$.

Further limitations of the Kihlstrom and Mirman approach emerge if we suppose that time begins at $-\infty$. Let $s_t = (\dots, c_{t-2}, c_{t-1})$ denote the history at t . Suppose that the von Neumann–Morgenstern index $v(\cdot; s)$ represents the intertemporal ordering \succeq_s at any time t if $s_t = s$. In such an environment it is both natural and common to assume that preferences are *stationary* – they can vary through time but only because the consumption history does.⁹ It can be shown that the von Neumann–Morgenstern index $v(c_t, c_{t+1}, \dots; s_t)$ implies stationarity of preferences if and only if there exist suitable functions A and B , $B > 0$, such that for all t :

$$v(c_t, c_{t+1}, \dots; s_t) = A(c_t; s_t) + B(c_t; s_t) v(c_{t+1}, \dots; s_{t+1}). \quad (3.1.3)$$

Suppose also that there is positive discounting along globally constant paths, in the sense that for all c :

$$B(c, s(c)) < 1,$$

where $s(c) \equiv (c, c, \dots)$. Such discounting could be inferred from common continuity assumptions.

Now consider whether comparative risk-aversion analysis is possible within the class of stationary expected utility functions; that is, if v satisfies (3.1.3) and if $v^*(\cdot) = h(v(\cdot))$ for some increasing and strictly concave h , can v^* satisfy a similar relation? Assuming twice differentiability, an implication of (3.1.3) is that for $\tau > t$:

$$\frac{-v_{c_{\tau}c_{\tau}}}{v_{c_{\tau}}}(c_{\tau}, \dots, c_{\tau}, \dots, s_{\tau}) = \frac{-v_{c_{\tau}c_{\tau}}}{v_{c_{\tau}}}(c_{\tau}, c_{\tau+1}, \dots, s_{\tau}) \quad (3.1.4)$$

and similarly for v^* . If the Arrow–Pratt measure for period τ risks is computed for $v^*(\cdot) = h(v(\cdot))$ via the chain rule, if the appropriate forms of (3.1.4) are applied, and if a constant path at level c is considered, then $B(c, s(c)) = 1$ is implied, contradicting positive discounting.

Though these arguments do not prove that the expected utility framework is inadequate in all temporal settings, I believe they do provide a *prima facie* case for exploring more general utility functions. Selden (1978) was the first to propose a generalization that achieves a separation between risk and certainty preferences. Some discussion of his approach appears in the next section.

3.2 Intertemporal consistency

A central and often misunderstood property of intertemporal utility functions is intertemporal consistency, which I now clarify following Johnsen and Donaldson (1985). This discussion will be useful also for understanding the work on sequential choice with non-expected utility preferences (section 5).

First define consumption programs somewhat more formally. Let (Ω, F, P) be a probability space and $\{F_t : t \geq 0\}$ an increasing filtration of sub σ -algebras of F that represents the information structure. A consumption program $(\tilde{c}_0, \tilde{c}_1, \dots, \tilde{c}_T, \dots)$ is a sequence of random variables such that each \tilde{c}_t is F_t -measurable. Frequently, consumption will be restricted to be a positive scalar but occasionally a vector of consumption goods within each period will be allowed. When the tilde is deleted, the corresponding variable is deterministic. Give $T \geq 0$, an event $I_T \in F_T$ and a consumption program as above, $(\tilde{c}_T, \tilde{c}_{T+1}, \dots | I_T)$ denotes the consumption program whose i^{th} component, $i = 0, 1, \dots$, is $\tilde{c}_{T+i} | I_T$, the restriction of \tilde{c}_{T+i} to $I_T \subseteq \Omega$.

Consider a time $T > 0$, an event $I_T \in F_T$ of positive probability and two programs $\tilde{c} = (\tilde{c}_0, \tilde{c}_1, \dots, \tilde{c}_T, \tilde{c}_{T+1}, \dots)$ and $\tilde{c}' = (\tilde{c}_0, \tilde{c}_1, \dots, \tilde{c}'_T, \tilde{c}'_{T+1}, \dots)$ that agree on $\Omega \setminus I_T$ (see figure 1.3 for a simple example). If \tilde{c} is ranked better than \tilde{c}' at time 0, then *intertemporal consistency* of preferences requires that the continuation $(\tilde{c}_T, \tilde{c}_{T+1}, \dots | I_T)$ be preferred to $(\tilde{c}'_T, \tilde{c}'_{T+1}, \dots | I_T)$ at (T, I_T) . Otherwise, the choice made at 0 would with positive probability be reversed. Evidently, whether or not consistency prevails depends upon which utility function dictates choice at (T, I_T) .

Consider the corresponding situation when there is no risk and let V be the utility function at time 0. The natural specification for the utility function at T , in the absence of changing tastes, is the restriction




Figure 1.3 Two consumption programs

$V^T(\cdot) \equiv V(c_0, \dots, c_{T-1}, \cdot)$, where c_0, \dots, c_{T-1} are the consumption levels actually experienced in the interim periods. Intertemporal consistency obtains trivially since:

$$V(c_0, \dots, c_{t-1}, c_T, \dots) > V(c_0, \dots, c_{T-1}, c_T, \dots)$$

and

$$V^T(c_T, \dots) > V^T(c_T, \dots)$$

are equivalent statements.

This argument has a counterpart in the case of risk. For example, consider the two programs represented by the probability trees in figure 1.3. The specification of utility at the lower node at $t=1$ that corresponds to constant tastes is the restriction $V^{1,b}(\cdot) = V(c_0, y^a, \cdot)$, where equality is modulo ordinal equivalence and where the notation introduced above has been modified in the obvious way. (In general, *constant tastes* will refer to the case where intermediate utility functions are defined as the obvious restrictions of the initial V . Otherwise, I'll speak of *changing tastes*.) Precisely as in the certainty case, the choice between the two consumption programs shown will be identical whether it is made at $t=0$ or at the lower node at $t=1$; that is, intertemporal consistency is ensured if tastes are constant.

It is noteworthy, however, that $V^{1,b}$ may depend through y^a on consumption alternatives that were *ex ante* possible at $t=0$ but which were never realized. Such dependence is not irrational, e.g., it could arise through feelings of disappointment or relief at the failure of some potential outcomes to be realized. Nevertheless, ruling out dependence on unrealized alternatives provides an intuitively appealing way to narrow down the class of admissible intertemporal utility functions and thereby add both predictive power and tractability. Therefore, say that V is *weakly recursive* if

$V^{1,b}$ above, and more generally the restrictions of V at each intermediate time and state, are independent (up to ordinal equivalence) of unrealized alternatives. This property is weaker than the separability across states exhibited by intertemporal expected utility functions.¹⁰ Under weak recursivity, the constant tastes specification and the concomitant intertemporal consistency are uncontroversial.

With the above terminology in place, it may be useful at this point to refer to figure 1.4 which outlines three approaches to the specification of intertemporal utility functions along with their main features. Two approaches feature constant tastes and thus dynamic consistency. The route corresponding to the middle branch, that assumes weak recursivity, has been by far the most productive to date and will be the focus of the remaining discussion of intertemporal utility and applications.

Changing tastes, corresponding to the branch on the left, usually arises from the assumption that at each time T the individual acts as though time begins anew – she disregards past and unrealized parts of the consumption program and uses the original utility function V to evaluate the future, i.e., in figure 1.3, $V^{1,b}(\cdot)$ is ordinally equivalent to $V(\cdot)$. As first elucidated by Strotz (1956), changing tastes, or the intertemporal inconsistency of preferences, poses problems for the modeler in describing behavior. The notion of sophisticated planning (Pollak, 1968) may be adapted to the present context to describe a consistent course of action (Chew and Epstein, 1990), but this approach has not yet delivered any interesting new empirical implications for consumption and asset returns.

As a further illustration, consider the following functional form due to Selden and Stux (1978) which is a multiperiod extension of Selden (1978):

$$V(\tilde{c}_0, \tilde{c}_1, \dots) = \sum_0^{\infty} \beta^t u(\hat{c}_t), \quad \hat{c}_t \equiv v^{-1} E_0 v(\tilde{c}_t). \quad (3.2.1)$$

This functional structure is appealing on several grounds. It reflects a seemingly natural algorithm for computing utility – first, each random consumption level is replaced by its certainty equivalent and second, the intertemporal utility of the sequence of certainty equivalents is computed by the common additive function. Moreover, a separation between certainty and risk preferences is possible in that aversion to multiperiod risks can be increased by suitably changing v while keeping β and u fixed. In addition, note that the certainty equivalent could be redefined using a betweenness-conforming function, for example, in order to accommodate Allais-type behavior.

However, the Selden and Stux function violates weak recursivity. Thus the constant tastes assumption would necessarily impose dependence of

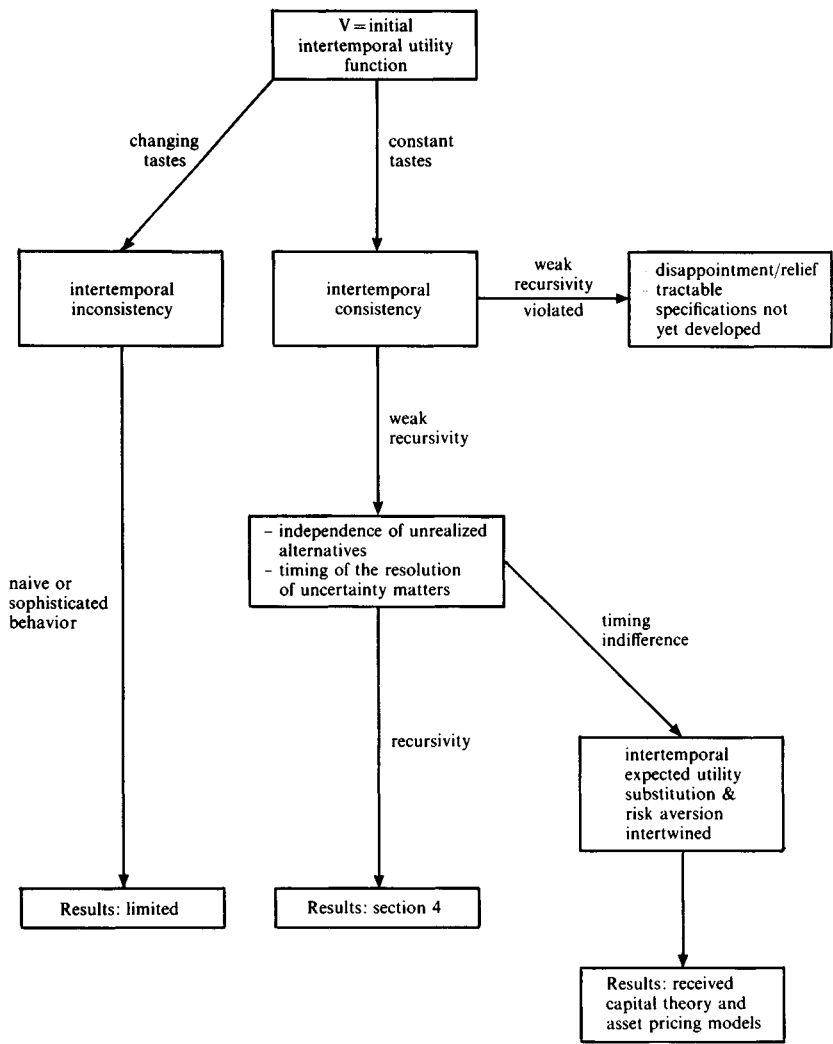


Figure 1.4 Intertemporal Utility

preferences upon unrealized alternatives. In particular, under constant tastes, period T preferences would continue to be based upon certainty equivalents \hat{c}_t computed according to period 0 information. In fact, where the function has been employed (Attansio and Weber, 1989; Hall, 1985; Zin, 1987) it is assumed that the utility function at any time T is computed as above except that the expected values are computed conditional upon period T information, i.e., tastes change as risk resolves. Moreover, the

hypothesis of naive behavior is adopted in these studies – inconsistencies are assumed to be ignored by the decision-maker who continually revises plans.

3.3 Recursive utility

Under weak recursivity and constant tastes, preferences are dynamically consistent and independent of unrealized alternatives. In the interest of still greater specificity and tractability, restrict utility further. Require that preferences be independent also of realized past consumption levels and further that V itself dictate choice at any intermediate time, e.g., in figure 1.3, require that $V(c_0, y^a, \cdot)$ be ordinally equivalent to $V(\cdot)$. Call V recursive if it satisfies these requirements.¹¹ Under recursivity, stationary dynamic programming techniques are applicable to specified optimization problems in such a way that state variables reflecting past consumption are unnecessary (see section 4).

Recursive utility functions may be constructed by means of the following recursive functional relation:¹²

$$V_t = W(c_t, \mu_t), \quad t \geq 0, \quad (3.3.1)$$

where $V_t = V(c_t, \tilde{c}_{t+1}, \dots | I_t)$ is intertemporal utility beginning at t , $\mu_t = \mu(\tilde{V}_{t+1} | I_t)$ is the certainty equivalent of the distribution of future utility \tilde{V}_{t+1} conditional upon period t information, and W is called an *aggregator* function since it aggregates current consumption c_t with an index of the future to determine current utility. The function W is such that $W(c, \cdot)$ is increasing.

The certainty equivalent function μ assigns to each real valued random variable \tilde{x} a certainty equivalent value. Require that μ be increasing in the sense of first-order stochastic dominance and that $\mu(\tilde{x}) = x$ if \tilde{x} equals x with certainty. Then (3.3.1) restricted to deterministic consumption programs coincides with the Koopmans (1960) structure:

$$V_t = W(c_t, V_{t+1}), \quad t \geq 0,$$

which generalizes the common intertemporally additive utility function (corresponding to $W(c, z) = u(c) + \beta z$) by endogenizing the discount factor $W_2(c_t, V_{t+1})$.

A broader special case applicable to stochastic programs arises if μ has the expected utility form:

$$\mu_{eu}(\tilde{x}) = h^{-1} E h(\tilde{x}), \quad (3.3.2)$$

for some increasing h . Then (3.3.1) corresponds to a special case of the structure studied in a finite horizon framework by Kreps and Porteus

(1978). A further parametric specialization (see (4.2.2) and (4.2.3) below) has been proposed and applied by several researchers (Epstein and Zin, 1989; Weil, 1990a; Farmer, 1990; Kocherlakota (1987)). Note that generally, even given (3.3.2), preferences over consumption programs do not conform with intertemporal expected utility theory since they are not indifferent to the way in which risk resolves over time (see section 3.4).

It is apparent from (3.3.1) how a degree of separation is achieved between substitution and risk aversion. Certainty preferences are determined by W alone. Thus only risk attitudes are affected by a change in μ . In particular, let $\mu^*(\tilde{x}) \leq \mu(\tilde{x})$ for all \tilde{x} and suppose that V^* is the intertemporal utility function corresponding to W and μ^* . Then V^* is more risk averse than V (Chew and Epstein, 1990b).¹³

Though μ is applied to utility gambles, it is intimately related to the induced preference ordering over timeless wealth gambles. For example, if intertemporal utility is linearly homogeneous as is the case if W is linearly homogeneous and μ exhibits constant relative risk aversion, then $\mu(\tilde{x})$ represents the preference ordering over timeless wealth gambles \tilde{x} (Epstein and Zin, 1989, p. 956). For another example, if μ satisfies betweenness then so does the ordering of timeless wealth gambles.¹⁴ Moreover, choices amongst timeless gambles constitute the laboratory evidence regarding atemporal expected utility and its generalizations. Thus it is appropriate to use μ as the route through which these atemporal theories are integrated into the temporal framework and the possible links between the laboratory evidence and intertemporal market behavior are investigated. Any of the functions discussed in section 2 are admissible.

In the next section, I will offer a normative argument for restricting μ to the betweenness class. From a descriptive perspective, Duffie and Epstein (1991a) provide some insight into the specifications for μ that might prove useful. They formulate the continuous time counterpart to (3.3.1), under a smoothness assumption for μ . (Fréchet differentiability, with a sufficiently smooth local utility function, is sufficient to imply such smoothness.) In (3.3.1), μ is applied to the conditional distribution of \tilde{V}_{t+1} given period t information. In a continuous environment, this conditional distribution has a small variance if the period corresponds to a short interval of time. Thus the only aspect of μ that is relevant is how it evaluates small gambles about certainty. The continuous time result is that smooth certainty equivalents are *observationally equivalent* to one another if they agree in their evaluations of infinitesimal gambles about certainty, and if only choices between Brownian consumption processes are observable. In particular, a Fréchet differentiable specification for μ , which satisfies the additional regularity conditions typically assumed by Machina (1982b), is empirically indistinguishable from an expected utility form (3.3.2) within

the standard diffusion models of asset prices. The above suggests also that there may be little empirical gain in generalizing from μ_{eu} to another smooth specification for μ if one is dealing with time series data, such as for aggregate U.S. consumption, having a small conditional variance. In terms of intertemporal utility functions, it is being suggested that there may be little advantage in generalizing from the structure studied by Kreps and Porteus (1978) to the class of recursive utility functions having a suitability differentiable certainty equivalent. The economic essence of the smoothness property that underlies the noted observational equivalence appears to be that of second-order risk aversion, but that remains a conjecture.

3.4 Some normative considerations

The next section will demonstrate the tractability and power of recursive utility for addressing some central questions in macroeconomics and finance. Here I examine further the rationality of recursive utility, which examination was started above with the description of its underlying axioms. I do this to assuage concerns that recursive utility may have logical implications for behavior that are obviously contradicted by observation and also to describe a normative argument for restricting the class of recursive utility functions.

Recursive utility functions are generally not indifferent to the way in which risk resolves over time in the sense of Kreps and Porteus (1978).¹⁵ Given the expected utility certainty equivalent (3.3.2), define $U_t(\cdot) \equiv h(V_t(\cdot))$ and the new aggregator $\hat{W}(c, \cdot) \equiv h(W(c, h^{-1}(\cdot)))$. Then (see footnote 13), $U_t = \hat{W}(c_t, E_t \tilde{U}_{t+1})$, from which it follows that late (early) resolution is preferred if $\hat{W}(c, \cdot)$ is concave (convex). (This is readily verified for the consumption programs in figure 1.5; for the general case see Kreps and Porteus, 1978.) Thus recursive utilities generally distinguish between the consumption programs in figure 1.5, unlike the case for (3.1.1) or for any other intertemporal expected utility function. Moreover, such indifference to timing is exactly the property of preference which is being dropped in generalizing from expected to recursive (or weakly recursive) utility functions (Kreps and Porteus, 1978; Chew and Epstein, 1989a, 1991), as indicated in figure 1.4.¹⁶

Introspection suggests that one might care about the temporal resolution of risk even in the absence of any implications for planning. For example, early resolution might be preferred by a “nervous” or “edgy” person who does not like living with risk, while an affinity for surprises could lead to the opposite preference. An individual might also prefer to defer resolution in order to continue to “consume” the hope for or illusion of a favorable outcome for a risky prospect. Kreps and Porteus (1978) offer other




Figure 1.5 Temporal resolution

supporting arguments. Supporting anecdotal evidence exists, as does some laboratory-based evidence (Cook, 1989). Unfortunately, casual observation of market behavior is not informative. Market choices generally reflect both the planning advantages of early resolution, e.g., of income risk, and the psychic costs or benefits of the corresponding early resolution of consumption risk which are the focus here.

Suppose that the risk to be resolved is which of two indifferent alternatives will be realized, e.g., $V(y^a) = V(y^b)$ in figure 1.5. Then the psychic costs or benefits of early resolution are less apparent and timing indifference is plausible. Say that utility satisfies *quasi-timing indifference* if in all such circumstances there is indifference to the timing of resolution. This reasonable property restricts the class of recursive utility functions to those for which the certainty equivalent μ in (3.3.1) satisfies betweenness (Chew and Epstein, 1989a, 1991). (See section 4.5 for a separate argument in support of betweenness based upon empirical tractability.)

To this point the discussion has dealt with individual preferences and behavior and thus, ideally, individual level data should be used to investigate and apply recursive (or other) utility functions empirically. However, since aggregate data are much more readily available and also call for “explanation,” the applications in the coming section are confined to them. Consequently, the utility function in question is hypothesized for a representative agent. Since the latter is fictional, the relevance of rationality and the normative properties of various functional form specifications becomes clouded. The existence of the representative agent and her utility should ideally be deduced via an appropriate aggregation theorem. In principle, the derived utility function need not possess all of the rationality properties considered important at the level of the individual. However, aggregation can be justified theoretically only under stringent conditions

and, moreover, tractability is a serious concern. For example, recall the complete markets justification for aggregation (Constantinides, 1982) and suppose that individual utility functions are recursive. The representative agent's function V is an optimized weighted average of the individual functions V^1, \dots, V^N ; more precisely, for some utility weights $\lambda_1, \dots, \lambda_N$:

$$V(\tilde{c}_0, \tilde{c}_1, \dots) = \max \left\{ \sum_{i=1}^N \lambda_i V^i(\tilde{c}_0^i, \tilde{c}_1^i, \dots) : \sum_{i=1}^N \tilde{c}_t^i \leq \tilde{c}_t \forall t \right\}.$$

Then V is generally not even weakly recursive. Moreover, the representative agent "works" in the complete markets case only if constant tastes are assumed for her. But then the dependence of preferences on unrealized alternatives might render the model untractable.

We are left with the familiar "excuse" for representative agent modeling, namely the current lack of a superior alternative. Moreover, it is useful to treat the representative agent as though she were real, since that allows us to organize observations in terms of familiar microeconomic principles and notions.

4 CONSUMPTION AND ASSET RETURNS

This section describes applications of recursive utility to both theoretical and empirical issues in macroeconomics and finance dealing with aggregate consumption and asset pricing. It begins by considering the consumption/savings and portfolio behavior of an agent having a recursive utility function and operating in a standard competitive and stationary environment. The Euler equations for the optimal intertemporal plan serve as the basis for all of the results below. When the agent is taken to be a representative agent for the economy, the Euler equations define relations between aggregate consumption and rates of return that must hold in a rational expectations equilibrium. In particular, they determine the equilibrium prices of the assets. Before describing the applications of recursive utility to the issues that have been of concern in this literature, the first subsection provides as background a brief review of the relevant results and limitations of the standard expected utility model.

4.1 Background

In the standard model with intertemporal expected utility function (3.1.1)–(3.1.2), the Euler equations take the following familiar form:

$$\beta E_t \left[\left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\rho-1} \tilde{r}_{i,t+1} \right] = 1, \quad i = 1, \dots, N, \quad (4.1.1)$$

where $\tilde{r}_{i,t+1}$ is the gross real return to the i th asset (see Lucas, 1978, for example). These equations generalize to the stochastic context the equality of the marginal rate of substitution and marginal rate of transformation that characterizes optimal consumption/savings behavior under certainty. Subtracting the i^{th} equation from the j^{th} leads to the following relations that characterize an optimal portfolio allocation:

$$E_t \left[\left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\rho-1} (\tilde{r}_{i,t+1} - \tilde{r}_{j,t+1}) \right] = 0, \quad i, j = 1, \dots, N.$$

If, for expository convenience, the first asset is taken to be riskless, then one obtains a form of Breeden's (1979) consumption-based restrictions on asset returns, whereby:

$$E_t[\tilde{r}_{i,t+1} - r_{1,t+1}] = K_t \text{cov}_t \left[- \left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\rho-1}, \tilde{r}_{i,t+1} \right], \quad (4.1.2)$$

where $K_t^{-1} \equiv E_t[\tilde{c}_{t+1}/c_t]^{\rho-1}$ and $\text{cov}_t(\cdot, \cdot)$ is the covariance conditional upon period t information. Thus, conditional on available information, the mean excess return to each asset is proportional to its systematic risk, which in turn is measured by the covariance of its return with an increasing function of the consumption growth rate.

Unfortunately, (4.1.1) and (4.1.2) do not accurately represent the relation between asset returns and aggregate per capita consumption observed in the U.S. For example, when confronted with time series data, (4.1.1) has been rejected statistically by Hansen and Singleton (1983) and Grossman, Melino and Shiller (1987) and on less formal grounds by Grossman and Shiller (1981) and Mehra and Prescott (1985). Furthermore, Mankiw and Shapiro (1986) have shown that, contrary to (4.1.2), cross-sectional mean returns are related much more closely to covariances with the return on the aggregate stock market than to covariances with aggregate consumption.

To a large degree, the research on intertemporal non-expected utility functions has been motivated by the desire to see whether the above empirical shortcomings can be ameliorated by a more general specification of utility for the representative agent.¹⁷ Thus in this section the standard model is modified by specifying that the utility function of the agent is recursive. A consequence is that the associated Euler equations generalize (4.1.1) and take the form:

$$E_t[\text{IMRS}_{t,t+1} \tilde{r}_{i,t+1}] = 1, \quad i = 1, \dots, N, \quad (4.1.3)$$

where $\text{IMRS}_{t,t+1}$ is a suitably defined intertemporal marginal rate of substitution between consumption at t and $t+1$ (see Hansen and Jagan-

nathan, 1991 for a general formulation of models of asset returns in terms of such marginal rates of substitution). For several models described below, the associated IMRS depends on more than just aggregate consumption. However, it can be computed from data given a parametrization of intertemporal utility and thus (4.1.3) is empirically tractable. Moreover, (4.1.2) may be adjusted to reflect the corresponding new model of the cross-sectional variation of excess mean returns given by:

$$E_t[\tilde{r}_{i,t+1} - r_{1,t+1}] = K_t^* \text{cov}_t[-\text{IMRS}_{t,t+1}, \tilde{r}_{i,t+1}], \quad (4.1.4)$$

where $K_t^{*-1} \equiv E_t[\text{IMRS}_{t,t+1}]$.

As discussed in section 3, there are also theoretical considerations that motivate generalizations of intertemporal expected utility theory. Below (see section 4.7) the theoretical gains from the disentanglement of substitution and risk aversion are demonstrated in the context of a Lucas (1978) style general equilibrium endowment economy. Under the expected utility specification (3.1.1)–(3.1.2), the ex-dividend price P_t of an asset paying dividend \tilde{d}_t in each period t , can be derived by substituting $(\tilde{P}_{t+1} + \tilde{d}_{t+1})/P_t$ for $\tilde{r}_{i,t+1}$ in (4.1.1), yielding the recursive relation:

$$P_t = \beta E_t \left[\left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\rho-1} (\tilde{P}_{t+1} + \tilde{d}_{t+1}) \right]. \quad (4.1.5)$$

Recursive substitution and imposition of a transversality condition deliver the formula:

$$P_t = E_t \sum_{s=1}^{\infty} \beta^s (\tilde{c}_{t+s}/c_t)^{\rho-1} \tilde{d}_{t+s},$$

whereby the price is the expected value of the discounted sum of future dividends. Comparative statics analysis of these formulae typically have ambiguous interpretations because of the dual role played by the parameter ρ , thus preventing a clear understanding of the determinants of equilibrium asset prices. In contrast, for recursive utility, (4.1.3) produces the following generalization of (4.1.5):

$$P_t = E_t[\text{IMRS}_{t,t+1}(\tilde{P}_{t+1} + \tilde{d}_{t+1})]. \quad (4.1.6)$$

For some parametric forms of recursive utility, the marginal rate of substitution involves separate substitution and risk aversion parameters, thus permitting more illuminating comparative statics analyses to be conducted.

The standard utility function (3.1.1) is separable across time and across states of the world. The results in the coming subsection are due primarily to the relaxation of state separability, though time non-separability

underlies the restriction (4.6.3) on asset returns. It should be noted that time non-separability due to habits or the durability of goods has been examined extensively in the empirical macro/finance literature (see Singleton, 1990; Constantinides, 1990; Heaton, 1991, for example). The utility functions used in these studies extend (3.1.1) in that the felicity function in period t depends also on lagged values of consumption; consequently they are weakly recursive but not recursive. A separation between substitution and risk aversion is *not* delivered in this way (see footnote 9 and the surrounding discussion in section 3.1), nor are any of the new asset pricing models described below. Of course, it is an empirical question, that is still unresolved, as to whether time or state non-separability is more useful in explaining aggregate data. Neither is it clear yet which form of time non-separability is most useful.

4.2 Euler equations

Suppose there are N assets with the vector of gross real returns $\tilde{r}_{t+1} = (\tilde{r}_{1,t+1}, \dots, \tilde{r}_{N,t+1})$ over the interval $[t, t+1]$. Denote the fraction of total real wealth held in the j^{th} asset in period t by $\omega_{j,t}$ and the N -vector of portfolio weights by ω_t . Let wealth evolve according to the process:

$$\tilde{x}_{t+1} = (x_t - c_t)\omega_t' \tilde{r}_{t+1}, \quad x_0 > 0 \text{ given.} \quad (4.2.1)$$

Denote by $J(I, x)$ the value of the agent's intertemporal optimization problem beginning with wealth x and information variable I used to predict future rates of return. Suppose the recursive utility function is represented by the aggregator W and certainty equivalent function μ . The latter is taken in the beginning of this section to be of the expected utility form (3.3.2). Then J should solve the Bellman equation:

$$J(I_t, x_t) = \text{Max } W(c_t, h^{-1} E_t[hJ(\tilde{I}_{t+1}, (x_t - c_t)\omega_t' \tilde{r}_{t+1})]), \quad (4.2.2)$$

because of the recursive relation (3.3.1).¹⁸ If J , W , and h are suitably smooth, then the envelope theorem and first-order conditions for the Bellman equation can be used to derive Euler equations for the intertemporal optimization problem. Such a straightforward attack at the above level of generality proves unsatisfactory, however, since the Euler equations (more particularly, the appropriate IMRS from (4.1.3)), invariably involve the unobservable value function. To generate empirically useful results, the utility function must be restricted further.

One powerful assumption is that intertemporal utility is homothetic, i.e., the common rescaling of two consumption programs does not affect their relative ranking. (An alternative assumption is described in the context of the multicommodity asset pricing model below.) Duffie and Epstein (1991b) derive the implications of homotheticity in a continuous-time

setting. Here, to simplify and facilitate interpretation, specialize further to the following convenient functional forms:

$$W(c, z) = \begin{cases} ((1 - \beta)c^\rho + \beta z^\rho)^{1/\rho}, & 0 \neq \rho < 1 \\ \exp((1 - \beta)\log c + \beta \log z), & \rho = 0, \end{cases} \quad (4.2.3)$$

$$\mu(\tilde{x}) = \begin{cases} (E\tilde{x}^\alpha)^{1/\alpha}, & 0 \neq \alpha < 1 \\ \exp(E\log \tilde{x}), & \alpha = 0. \end{cases} \quad (4.2.4)$$

The utility of deterministic programs is evaluated by a CES function with elasticity of substitution $\sigma = (1 - \rho)^{-1}$ and rate of time preference $\beta^{-1} - 1 > 0$. Aversion toward intertemporal consumption gambles increases as α falls. In addition, $1 - \alpha$ equals the degree of relative risk aversion with respect to timeless wealth gambles (see footnote 14 and the discussion leading to it). The common specification (3.1.1)–(3.1.2) corresponds to the special case $\alpha = \rho$.

If $\alpha > (<) \rho$, then a preference for late (early) resolution of consumption risk is implied (section 3.4). However, even if there is a psychic cost associated with the early resolution of consumption risk, early resolution of rate of return risks provides planning advantages that could outweigh the associated psychic costs. As an example where closed-form solutions are possible, consider the case of a single asset ($N = 1$) with rates of return \tilde{r}_t that are identically and independently distributed like \tilde{r} over time. The value function for the associated planning problem coincides with that implied by the deterministic problem where the rate of return is constant at the certainty equivalent level $(E\tilde{r}^\alpha)^{1/\alpha}$. Thus maximum intertemporal utility is $x_0(1 - \beta)^{1/\rho} [1 - \beta^{1/(1 - \rho)} (E\tilde{r}^\alpha)^{\rho/\alpha(1 - \rho)}]^{(\rho - 1)/\rho}$. Next suppose that the entire sequence of rates of return (r_1, \dots, r_t, \dots) is revealed at $t = 0$ before any consumption decisions are made and evaluate maximum utility from the perspective of an instant preceding 0 when the sequence that will be revealed is not yet known. Maximum utility in this case of early resolution

is $x_0(1 - \beta)^{1/\rho} \left[E \left(\sum_{t=0}^{\infty} \beta^{t/(1 - \rho)} (\tilde{r}_1 \tilde{r}_2 \dots \tilde{r}_{t+1})^{\rho/(1 - \rho)} \right)^{\alpha(1 - \rho)/\rho} \right]^{1/\alpha}$.

It follows that the early resolution of rate of return risk is preferred if and only if $\rho(1 + \alpha) \geq \alpha$, or equivalently, $\sigma + (1 - \alpha) \geq 2$. This is compatible with the condition $\sigma(1 - \alpha) < 1$, which corresponds to a preference for the late resolution of consumption risk.

With the above functional forms, the Euler equations for $\rho \neq 0$ are:

$$\begin{aligned} E_t \left[\left(\beta^{1/\rho} \left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{(\rho - 1)/\rho} \tilde{M}_{t+1}^{1/\rho} \right)^\alpha - 1 \right] / \alpha = 0 \text{ and} \\ E_t \left[\left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\alpha(\rho - 1)/\rho} \tilde{M}_{t+1}^{(\alpha - \rho)/\rho} (\tilde{r}_{i,t+1} - \tilde{r}_{j,t+1}) \right] = 0 \quad i, j = \dots, N, \end{aligned} \quad (4.2.5)$$

where $\tilde{M}_{t+1} \equiv \omega_t^* \tilde{r}_{t+1}$ is the return to the optimal portfolio and where $(x^\alpha - 1)/\alpha$ stands for $\log x$ when $\alpha = 0$.¹⁹ The familiar Euler equations (4.1.1), analysed by Hansen and Singleton (1983) for example, are obtained if $\alpha = \rho$. A corresponding set of Euler equations involving only “observables,” is not available without additional assumptions if $\rho = 0$ and $\alpha \neq 0$.²⁰ Henceforth, $\rho \neq 0$ is assumed unless explicitly stated otherwise. Note that, if $\alpha \neq 0$, the Euler equations can be rewritten in the equivalent and sometimes more convenient form:

$$E_t \left[\beta^{\alpha/\rho} \left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\alpha(\rho-1)/\rho} \tilde{M}_{t+1}^{(\alpha-\rho)/\alpha} \tilde{r}_{t+1} \right] = 1, \quad i = 1, \dots, N, \quad (4.2.5')$$

which is the special case of (4.1.3) in which:

$$\text{IMRS}_{t,t+1} = \beta^{\alpha/\rho} (\tilde{c}_{t+1}/c_t)^{\alpha(\rho-1)/\rho} \tilde{M}_{t+1}^{(\alpha-\rho)/\alpha}. \quad (4.2.6)$$

Since the utility function is homothetic, if it is assumed that it is also common to each agent in the economy and if common information is also assumed, then demand aggregation in the sense of Gorman (1953) holds. In this way theoretical justification can be provided for the existence of a representative agent and for the application of (4.2.5) to aggregate data.

Some applications of the above Euler equations will now be described.

4.3 Consumption

The study of consumption in the rational expectations school of macroeconomics is commonly based on the Euler equation characterization of optimal consumption and empirical tests of it. Hall's (1978) random walk hypothesis and its generalizations (Hall, 1988; Hansen and Singleton, 1983) have assumed expected utility function specifications. The resulting inability to separately identify substitution and risk aversion poses a problem since, as Hall (1988, p. 339) notes, the magnitude of the intertemporal substitution effect of a change in the expected real interest “is one of the central questions of macroeconomics.” Identification is not an issue if it is believed that the elasticity of substitution σ and the degree of relative risk aversion are in fact reciprocals of one another as in the conventional homogeneous specification. But then estimates of σ near 0 such as found by Hall imply an incredibly large degree of risk aversion.

Epstein and Zin (1991a) show that the Euler equations (4.2.5) identify both α and ρ when the Generalized Method of Moments (GMM) estimation procedure is applied to monthly post-war US data on NYSE stock returns, treasury bills, and various measures of per capital consumption expenditures. The performance of the model varies with the choice of

instrumental variables and also with the measure adopted for consumption. For some choices the expected utility restriction $\alpha = \rho$ is rejected and the general model cannot be rejected. For a study based on Israeli data see Bufman and Leiderman (1990), who report results favorable to the recursive utility model.

Hall and other researchers have based their empirical analyses on a linear relation between consumption growth and individual asset returns and have interpreted the parameter multiplying the asset return as σ . Qualified justification for such an analysis can be provided as follows (see also Attanasio and Weber, 1989): define $\tilde{z}_{t+1} \equiv \beta^{1/\rho} (\tilde{c}_{t+1}/c_t)^{(\rho-1)/\rho} \tilde{M}_{t+1}^{1/\rho}$. After taking logarithms, the first Euler equation above can be transformed into the relation:

$$\log(\tilde{c}_{t+1}/c_t) = k + \sigma \log \tilde{M}_{t+1} + (1 - \sigma) E_t \log \tilde{z}_{t+1} + \tilde{\varepsilon}_{t+1},$$

where k is a constant, $E_t[(\tilde{z}_{t+1}^\alpha - 1)/\alpha] = 0$ and $E_t \tilde{\varepsilon}_{t+1} = 0$. If $E_t \log \tilde{z}_{t+1}$ is constant through time, then the desired linear regression equation is obtained and σ can be estimated consistently by instrumental variables techniques. A similar equation, where the return to any single asset replaces the market return on the right side, can be obtained by analogous arguments.

It is intuitive that the relationship between consumption growth and the real interest rate is governed by the intertemporal substitution aspect of preferences. Intuition may not be clear or correct in slightly modified models, however. Consider the framework of Hall (1978) where the return to saving r is deterministic and constant and suppose that there is an exogenous stochastic stream of payments to inelastically supplied labour. The conventional Euler equation is:

$$\beta r E_t[u'(\tilde{c}_{t+1})/u'(c_t)] = 1. \quad (4.3.1)$$

Does the curvature of u represent intertemporal substitution or risk aversion? Under a functional form specification for recursive utility, based on exponential rather than power functions as in (4.2.3)–(4.2.4), Epstein (1991) points out that an appropriate form of (4.3.1) applies and that it is risk aversion, rather than substitution, that is embodied in u .

Weil (1990c) uses yet another parametric specification of recursive utility to study behavior given a constant interest rate and undiversifiable labor income risk. He shows that a CES form for the aggregator W and a constant absolute risk aversion form for μ are particularly well suited to the analysis of the determinants of precautionary savings.

Finally, it is worth noting a difficulty with the implementation of the extended Euler equations or a linear regression equation derived from them – the measurement of consumption. In the conventional model, the use of

any component of consumption, such as non-durables consumption for example, can be justified theoretically by the assumption that the intertemporal von Neumann–Morgenstern index is additively separable between the selected component and the rest of consumption. A comparable argument does not exist in the recursive utility model, however, and the use of a comprehensive measure of consumption is required by the theory.

4.4 A two-factor asset pricing model

The two principal models in finance for explaining excess mean returns are the consumption-based CAPM of Lucas (1978) and Breeden (1979) and the market-portfolio-based CAPM. In one systematic risk of an asset is measured by covariance of its return with consumption growth and in the other by covariance with the return to the market. It is evident from the portfolio allocation equation in (4.2.5) (see also (4.1.4) and (4.2.6)) that for recursive utility functions, both consumption and the market return enter into the covariance that defines systematic risk. Roughly speaking, \tilde{c}_{t+1} and \tilde{M}_{t+1} both enter into the marginal rate of substitution in (4.2.6) because the marginal rate of substitution between c_t and consumption in a specified state I_{t+1} at $t+1$ depends upon both consumption levels and, given the non-separabilities present when $\alpha \neq \rho$, upon future intertemporal utility $J(I_{t+1}, x_{t+1})$. The market return acts as a proxy for J .

If there is a riskless asset with return r_{t+1}^f over the interval $[t, t+1]$, then for each risky asset:

$$\begin{aligned} E_t[\tilde{r}_{i,t+1} - r_{t+1}^f] &= \frac{\alpha(1-\rho)}{\rho} \text{cov}_t\left(\frac{\tilde{c}_{t+1}}{c_t}, \tilde{r}_{i,t+1}\right) \\ &\quad + \frac{\rho-\alpha}{\rho} \text{cov}_t(\tilde{M}_{t+1}, \tilde{r}_{i,t+1}). \end{aligned} \quad (4.4.1)$$

Duffie and Epstein (1991b) derive the counterpart of this equation in a continuous time framework. Alternatively, the equation can be derived from (4.1.2) as an approximation via a joint lognormality assumption for consumption growth and asset returns (Epstein and Zin, 1991a); in that case, one half of the conditional variance of $\tilde{r}_{i,t+1}$ must be subtracted from the right side. The consumption-CAPM follows if $\alpha = \rho$, but in general a linear combination of the two common asset pricing models obtains. For empirical evidence regarding this 2-factor model based on a cross-sectional analysis of various securities, see Mankiw and Shapiro (1986) and Giovannini and Weil (1989); for evidence that considers time series data as well, see Bollerslev, Engle, and Wooldridge (1988), Epstein and Zin

(1991a), and Bufman and Leiderman (1990). Overall the evidence suggests that the market covariance is the more important, but that both factors are statistically significant.

Motivated by the problems in measuring consumption, Campbell (1990) adopts auxiliary assumptions in order to substitute out consumption from the asset pricing model. In his model, covariance with consumption growth is replaced by the covariance with news about the discounted value of all future market returns.

4.5 Non-expected utility for timeless wealth gambles

To this point the results described have been based mostly on the expected utility certainty equivalent (4.2.4). Since, as discussed in section 3.3, the certainty equivalent μ represents the ranking of timeless wealth gambles and since choices amongst such gambles constitute the laboratory evidence regarding atemporal expected utility theory, it is of interest to determine the theoretical and/or empirical gains from admitting more general specifications for μ . In particular, can such generalizations improve the explanation of the consumption and asset return data that have been studied and do the market data and laboratory evidence indicate similar deviations from an expected utility form for μ ?

These questions can be addressed by taking μ to be a betweenness-conforming function (sections 2.2 and 2.3). Suppose further that μ has constant relative risk aversion and is represented by the function ϕ as in (2.3.2). If the CES form for the aggregator is also retained, then Epstein and Zin (1989) show that the Euler equations (4.2.5) generalize to:

$$\begin{aligned} E_t \left[\phi \left(\beta^{1/\rho} \left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{(\rho-1)/\rho} \tilde{M}_{t+1}^{1/\rho} \right) \right] &= 0 \quad \text{and} \\ E_t \left[\phi' \left(\beta^{1/\rho} \left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{(\rho-1)/\rho} \tilde{M}_{t+1}^{1/\rho} \right) \cdot \left[\frac{\tilde{c}_{t+1}}{\tilde{M}_{t+1}} \right]^{(\rho-1)/\rho} \cdot (\tilde{r}_{i,t+1} - \tilde{r}_{j,t+1}) \right] &= 0. \end{aligned} \quad (4.5.1)$$

It is noteworthy that these equations are as tractable as the earlier ones; for example, generalized method of moments estimation is applicable once a functional form is specified for ϕ . This feature seems to be due to the considerable affinity already noted between expected utility and betweenness orderings. While other suitably differentiable specifications for μ can generate first-order conditions for the portfolio allocation, those equations are much less attractive for estimation purposes since they generally require functional form assumptions about the distribution describing consumption and asset returns.²¹

Two convenient parametric forms for ϕ , corresponding to weighted utility and disappointment averse utility, are given by (2.3.3) and (2.3.6).²² Each defines a single parameter extension of the expected utility certainty equivalent (4.2.4), but they differ from one another in that only disappointment averse utility satisfies first-order risk aversion (section 2.4). In light of the observational equivalence result of Duffie and Epstein (1991a) described in section 3.3 and the smoothness of US aggregate consumption data, one might therefore expect that the generalization to weighted utility would be of little help in explaining US monthly data, but that the disappointment averse functional form could move the theory significantly closer to observed behavior. Preliminary evidence reported by Epstein and Zin (1991b) supports these conjectures.

It is conceivable that the weighted utility form would prove more valuable in explaining data from more volatile economies. In any case, it is evident that a number of new functional forms have been added to the tool kit of the empiricist interested in fitting consumption and asset return data.

Another demonstration of the usefulness of a non-expected utility certainty equivalent is provided by the discussion of the equity premium puzzle in section 4.7.

4.6 Multicommodity asset pricing

The Euler equations (4.2.5) were generalized above by adopting more general specifications for the certainty equivalent. Now consider an alternative generalization which has the following motivation: according to the wealth accumulation equation (4.2.1), all income is investment income. Thus the agent's portfolio presumably includes human capital and other non-traded assets, which renders the return to the aggregate portfolio \tilde{M}_{t+1} unobservable, i.e., the Roll (1977) critique of CAPM is relevant here. (In practice, the market return is sometimes measured as an index of the returns to stocks traded on the New York Stock Exchange.)

Epstein and Zin (1990b) in discrete time and Duffie and Epstein (1991b) in continuous time describe a model in which endogenous labor supply may be incorporated and in which returns to non-traded assets are not needed for empirical implementation. Reconsider a recursive utility function with general aggregator W and expected utility certainty equivalent μ_{eu} from (3.3.2), but let consumption at time t be a vector $(c_t, q_t) \in R_+^1 \times R_+^L$, $L \geq 1$. The first good is the numeraire; the remaining relative prices are denoted p_t . One of the components of q_t could represent leisure and the corresponding component of p_t would be the real wage rate. The real wealth accumulation equation is:

$$\tilde{x}_{t+1} = (x_t - c_t - p_t' q_t + y_t) \omega_t' \tilde{r}_{t+1},$$

where y_t represents exogenous income which includes the value of skill and time endowments, and where ω_t and \tilde{r}_{t+1} refer to traded assets only.

As mentioned earlier, recursive utility models generate first-order conditions that may be difficult to implement empirically because they contain the unobserved value function. But in a multicommodity context that unobservability may be overcome as follows: an intratemporal optimality condition for the agent is:

$$\frac{W_{q_l}}{W_c}(c_t, q_t, \mu_t) = P_{l,t}, l = 1, \dots, L, \quad (4.6.1)$$

where μ_t is the certainty equivalent of the $(t + 1)$ period value function given period t information. The ratio on the left is the marginal rate of substitution between c_t and $q_{l,t}$. In the standard specification (replacing c by (c, q) in (3.1.1)), this marginal rate of substitution is independent of μ_t since (c_t, q_t) is weakly separable from consumption at all times $t' \neq t$. Consider the opposite assumption in the form of the following: there exists at least one $l \in \{1, \dots, L\}$ such that for each (c, q) the function:

$$\Psi(\cdot) \equiv \frac{W_{q_l}}{W_c}(c, q, \cdot)$$

is strictly monotonic and hence invertible. In that case refer to q_l as being *invertibly non-separable from c* . Given such non-separability, μ_t can be uniquely recovered from the l^{th} intratemporal equation in (4.6.1).

Epstein and Zin (1990b) show how this permits the derivation of Euler equations of the form:

$$E_t[f(z_t, \tilde{z}_{t+1})\tilde{r}_{i,t+1}] = 1, \quad i = 1, \dots, N, \quad (4.6.2)$$

for some function f derived from W and μ , where $z_t \equiv (c_t, q_t, p_{l,t})$. In continuous time, Duffie and Epstein derive the counterpart of the following model of excess mean returns:

$$\begin{aligned} E_t[\tilde{r}_{i,t+1} - r_{t+1}^f] = & K_c(t) \text{cov}_t\left(\frac{\tilde{c}_{t+1}}{c_t}, \tilde{r}_{i,t+1}\right) + \sum_{k=1}^L K_k(t) \text{cov}_t\left(\frac{\tilde{q}_{k,t+1}}{q_{k,t}}, \tilde{r}_{i,t+1}\right) \\ & + K_l^*(t) \text{cov}_t\left(\frac{\tilde{p}_{l,t+1}}{p_{l,t}}, \tilde{r}_{i,t+1}\right), \quad i = 1, \dots, N, \end{aligned} \quad (4.6.3)$$

where l is any index such that q_l is invertibly non-separable from c , which non-separability permits the price p_l to serve as a proxy for future utility. Since the coefficients K_c , K_k , and K_l^* do not depend on the asset i , it might be possible to apply (4.6.3) to a cross-section of securities. The significance of the price covariance term would provide support for the model based on invertible non-separability and against the standard utility specification

where (4.6.3) applies with $K_l^* \equiv 0$ for any l . There is no empirical evidence available yet regarding (4.6.2) or (4.6.3).

Since the assumption of invertible non-separability is an assumption regarding W , it restricts certainty preferences rather than μ or attitudes towards timeless wealth gambles. In fact, the cited papers show that it can be accommodated by an intertemporal expected utility ordering in which the von Neumann–Morgenstern index has the Uzawa (1968) functional form. Such a specification has:

$$W(c, q, \phi) = u(c, q) + B(c, q)\phi, \quad \mu(\tilde{x}) = E\tilde{x}.$$

Thus a non-expected utility function would not be implied (nor would it be contradicted) by the significance of a price beta in the regression suggested by (4.6.3). However, I will outline a modified multicommodity model in which a non-expected utility ordering for intertemporal gambles is an essential ingredient.

A feature of (4.6.3) is that only a single price beta appears on the right side, though l can correspond to any good that satisfies the requisite invertible non-separability. The modified model requires $L \geq 2$ and provides justification for having two price betas appearing simultaneously. Thus the joint significance of two price betas in the obvious cross-sectional regression would indicate a non-expected utility ordering. It will be evident to the reader how to extend the model further to admit more than two price betas.

Consider an extension of (3.3.1) whereby for some function W with $W(c, q, \cdot, \cdot)$ increasing:

$$V_t = W(c_t, q_t, \mu_t, v_t), \quad t \geq 0, \quad (4.6.4)$$

where $\mu_t = \mu(\tilde{V}_{t+1})$ and $v_t = v(\tilde{V}_{t+1})$ represent conditional certainty equivalents of \tilde{V}_{t+1} computed according to the two distinct certainty equivalent functions μ and v . If the last two arguments of W are weakly separable from (c_t, q_t) , then μ_t and v_t can be aggregated into a single certainty equivalent and the recursive structure (3.3.1) is obtained. In the absence of such weak separability, recursivity is violated by V .²³ Weak recursivity is always satisfied, however, and under the constant tastes hypothesis intertemporal consistency prevails and dynamic programming is applicable with the obvious Bellman equation.

The intratemporal optimality conditions are the counterparts of (4.6.1). They involve the two unobservables μ_t and v_t . Thus it is natural to introduce the following terminology and ensuing assumption: say that $(q_t, q_{t'})$ is *invertibly non-separable from c* if for each (c, q) the vector-valued function Ψ :

$$\Psi(\cdot) \equiv \left(\frac{W_{q_t}}{W_c}(c, q, \cdot), \frac{W_{q_{t'}}}{W_c}(c, q, \cdot) \right)$$

is invertible. Suppose that $\exists l, l' \in \{1, \dots, L\}$ such that this non-separability obtains. Then μ_t and v_t can be recovered from the appropriate intratemporal optimality conditions. Consequently, by extending the arguments in Epstein and Zin (1990b) and in Duffie and Epstein (1991b), counterparts to (4.6.2) and (4.6.3) can be derived.²⁴ In the latter, precisely two price betas appear, corresponding to l and l' .

As an example which is particularly relevant to continuous time, restrict attention to non-negative random variables and let:

$$\mu(\tilde{x}) = (E\tilde{x}^2)^{1/2} \quad \text{and} \quad v(\tilde{x}) = E\tilde{x}.$$

Note that the mean $m(\cdot)$ and variance $\text{var}(\cdot)$ are expressible as:

$$m(\tilde{x}) = v(\tilde{x}), \quad \text{var}(\tilde{x}) = \mu^2(\tilde{x}) - v^2(\tilde{x}).$$

Since the latter equations define a one-to-one relation between $(\mu(\tilde{x}), v(\tilde{x}))$ and $(m(\tilde{x}), \text{var}(\tilde{x}))$, (4.6.4) can be expressed equivalently as:

$$V_t = \tilde{W}(c, q_t, m(\tilde{V}_{t+1}), \text{var}(\tilde{V}_{t+1})), \quad (4.6.4')$$

for suitable \tilde{W} . It is immediate that $(q_t, q_{t'})$ is invertibly non-separable from c if:

$$\begin{aligned} \frac{\tilde{W}_{q_t}}{\tilde{W}_c}(c, q, \cdot, \cdot) \nearrow m \text{ and } \nearrow \text{var}, \\ \frac{\tilde{W}_{q_{t'}}}{\tilde{W}_c}(c, q, \cdot, \cdot) \nearrow m \text{ and } \searrow \text{var}. \end{aligned}$$

In words, a brighter expected future or increased uncertainty about the future shift preferences toward q_t and away from the numeraire, given fixed levels of the other commodities. On the other hand, the mean and variance of future prospects work in opposite directions in affecting the marginal rate of substitution between $q_{t'}$ and c .

4.7 Asset prices in general equilibrium

Now consider some general equilibrium issues. First, it seems that recursive utility defines a natural class of preferences for dynamic general equilibrium theory. Since intertemporal consistency obtains, so does the non-market reopening property typically assumed in the standard Arrow–Debreu model of contingent commodity markets. Donaldson and Selden (1981) and Johnsen and Donaldson (1985) have discussed the link between preferences and the reopening of markets.

Next consider Lucas' (1978) endowment model modified so that the utility function V of the representative agent is recursive and corresponds to

the functional form specifications (4.2.3)–(4.2.4). If an asset pays dividend \tilde{d}_t in period t , then substitution of $(\tilde{P}_{t+1} + \tilde{d}_{t+1})/P_t$ for $\tilde{r}_{t,t+1}$ in the Euler equation (4.2.5') implies that the equilibrium exdividend price P_t must satisfy:

$$P_t = \beta^{\alpha/\rho} E_t \left[\left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\alpha(\rho-1)/\rho} \tilde{M}_{t+1}^{(\alpha-\rho)/\rho} (\tilde{P}_{t+1} + \tilde{d}_{t+1}) \right], \text{ for all } t. \quad (4.7.1)$$

In particular, the price p_t^* of aggregate equity or claims to the endowment process satisfies:

$$(p_t^*)^{\alpha/\rho} = \beta^{\alpha/\rho} E \left[\left[\frac{\tilde{c}_{t+1}}{c_t} \right]^{\alpha(\rho-1)/\rho} (\tilde{p}_{t+1}^* + \tilde{c}_{t+1})^{\alpha/\rho} \right], \quad (4.7.2)$$

where $\tilde{M}_{t+1} = (\tilde{p}_{t+1}^* + \tilde{c}_{t+1})/p_t^*$ has been substituted.²⁵

Consider the effects on the price of equity of a change in the degree of risk aversion as represented by α . Following Epstein (1988), if the endowment and thus also consumption levels are i.i.d., then $p_t^* = K c_t^{1-\rho}$, where K is the unique solution to $K^{\alpha/\rho} = \beta^{\alpha/\rho} E[(\tilde{c}^\rho + K)^{\alpha/\rho}]$. Thus the period t price increases (falls) as risk aversion increases if the elasticity of intertemporal substitution σ is less (greater) than 1. The intuition is clear. At fixed prices, an increase in risk aversion acts to reduce the certainty equivalent of the return \tilde{M}_{t+1} to savings. The effect on behavior is similar to the consequence of a lower rate of return in a deterministic model. If $\sigma < (> 1)$, the dominant income (substitution) effect implies reduced (enhanced) present consumption and an increased (reduced) demand for securities. Thus the price of equity is forced to rise (fall).²⁶ Results of comparable sharpness and intuitive clarity are possible for other comparative statics exercises, e.g., the consequences of changes in the consumption endowment process. However, the clarity is lost if α and ρ are constrained to be equal as in the conventional model.

Next consider the term structure of interest rates. In order to obtain closed-form solutions, assume that $\rho = 0$ (i.e., $\sigma = 1$). The pricing formula (4.7.1) does not apply in this case but an alternative formula can be derived if it is assumed that consumption growth rates $\tilde{y}_{t+1} \equiv \tilde{c}_{t+1}/c_t$ follow the autoregressive process:

$$\log \tilde{y}_{t+1} = \lambda \log y_t + \delta + v \tilde{\varepsilon}_{t+1},$$

where $\beta|\lambda| < 1$, and the ε_t 's are white-noise normal variates. Consider the appropriate form of the Bellman equation (4.2.2) for the representative agent's optimization problem. Since the intertemporal utility function V is linearly homogeneous, the corresponding value function is linearly homo-

geneous in wealth. Then the first-order conditions for the Bellman equation imply that optimal consumption equals the fraction $(1 - \beta)$ of current wealth. Since wealth evolves according to $\tilde{x}_{t+1} = (x_t - c_t)\tilde{M}_{t+1}$, it follows that $\tilde{y}_{t+1} = \beta\tilde{M}_{t+1}$. Thus the equilibrium market return follows the process:

$$\log \tilde{M}_{t+1} = \lambda \log M_t + \delta + (\lambda - 1) \log \beta + v\tilde{e}_{t+1}.$$

The implied lognormality can be exploited to derive an explicit solution for the value function, with $I_t = M_t$, of the form:

$$J(M_t, x_t) = kM_t^\theta x_t, \quad \theta \equiv \lambda\beta/(1 - \lambda\beta).$$

Finally, we obtain the following counterpart to the first-order conditions in (4.2.5) for optimal portfolio allocation:

$$E_t[\tilde{M}_{t+1}^{\alpha\theta + \alpha - 1}(\tilde{r}_{i,t+1} - \tilde{r}_{j,t+1})] = 0.$$

The latter equation can be used to price any asset. In particular, let $P_{t,t+1}$ be the price in period t of a one period pure discount bond. Then:

$$P_{t,t+1} = E_t[\tilde{M}_{t+1}^{\alpha\theta + \alpha - 1}] / E_t[\tilde{M}_{t+1}^{\alpha\theta + \alpha}],$$

which implies, after integrating, that the interest rate $r_t \equiv P_{t,t+1}^{-1}$ satisfies the autoregressive model:

$$\log \tilde{r}_{t+1} = \lambda \log r_t + K + \alpha v^2(1 - \lambda)(1 - \lambda\beta)^{-1} + \lambda v\tilde{e}_{t+1},$$

for some constant K that does not depend on α . Under expected utility, $\alpha = 0$ necessarily. With α free, however, it can be deduced that: (i) the conditional mean of the interest rate decreases (increases) with risk aversion if $\lambda < (>) 1$ and is independent of risk aversion in the unit root case $\lambda = 1$; (ii) the conditional variance of the interest rate does not depend upon the degree of risk aversion.

Similarly, if $P_{t,T}$ denotes the price at t of a discount bond promising one unit of consumption at T , then:

$$P_{t,T} = E_t[\tilde{M}_{t+1}^{\alpha\theta + \alpha - 1} \tilde{P}_{t+1,T}] / E_t[\tilde{M}_{t+1}^{\alpha\theta + \alpha}],$$

which implies that:

$$P_{t,T} = B(t, T)M_t^{A(t, T)},$$

where A and B are explicitly determinable functions of t and T . Define by $(T - t) \log R(r_t, T) = -\log P_{t,T}$ the yield-to-maturity of a bond maturing at T given the interest rate r_t at t . It is a straightforward matter to show that an increase in risk aversion reduces the yield for bonds of all maturities if $\lambda \geq -1$. In addition, the slope of the yield curve decreases with risk aversion if $\lambda \geq 0$.

A separate and important question concerning the Lucas economy is

whether preferences, including both intertemporal substitutability and risk aversion, are recoverable from a *single* dynamic equilibrium; in particular, could an observer of equilibrium consumption and asset prices in a *given* economy distinguish between the intertemporal expected utility case ($\alpha = \rho$) and the more general recursive case $\alpha \neq \rho$? If not, then the more general model does not provide any additional power for explaining time series from a single economy. Kocherlakota (1990) has shown that, if consumption growth rates are i.i.d., then indeed none of the parameters β, ρ , and α is uniquely determined and the recursive utility model based on (4.2.3)–(4.2.4) is observationally equivalent to the standard model.²⁷ In fact, if consumption growth rates are i.i.d., then the price of aggregate equity satisfies $p_t^* = Kc_t$, where K is a constant that combines β, α , and ρ . The intuition underlying such observational equivalence is clear. Asset prices at t reflect marginal rates of substitution at the conditional consumption program faced by the agent at time t . If that program does not vary sufficiently with t , as in the i.i.d. case, then marginal rates of substitution will be delivered only on a limited domain. The i.i.d. case is analogous to the situation in demand theory where a single price/quantity data point cannot be used to pin down the underlying utility function. This intuition suggests that observational equivalence should be a problem only for a “small” set of consumption processes. Wang (1990) proves a number of results that confirm this intuition. In particular, he shows that the following is true generically in the space of finite state, first-order Markov processes for consumption growth rates: the function describing the equilibrium price of equity in an economy with the Kreps–Porteus utility (4.2.3)–(4.2.4) and $\alpha \neq \rho$ is distinct from that implied by any intertemporally additive expected utility function satisfying specified technical conditions.

Finally, in an empirical vein, consider the equity premium puzzle posed by Mehra and Prescott (1985). They argue that the general equilibrium implications of the representative agent, expected additive utility model, sensibly restricted with regard to the degree of risk aversion, are inconsistent with the historically observed low average real rate of return on debt and large risk premium for equity in the US. Weil (1989) and Kocherlakota (1990) have shown that the generalization of utility to the recursive form (4.2.3)–(4.2.4) does not improve matters. It is intuitive, once again in light of the smoothness of aggregate consumption, that the “order” of risk aversion should be important here. Indeed, Epstein and Zin (1990a) have shown that a partial resolution of the puzzle is achieved with a first-order risk-averse specification for the certainty equivalent μ . As noted earlier, such a specification can produce a sizeable risk premium for a small gamble without implying implausibly large premia for larger gambles. Note that first-order risk aversion precludes μ from being expected utility based. Thus

it **requires** that choices amongst timeless wealth gambles violate expected utility. First-order risk aversion can be accommodated readily, however, within the betweenness theory via the disappointment averse functional form, for example (see sections 2.3 and 2.4).

5 SEQUENTIAL CHOICE AND GAME THEORY

Many dynamic decision problems are reasonably viewed as taking place over short intervals of time during which consumption/savings plans can be taken to be fixed. The source of utility is terminal wealth rather than a consumption sequence. Such decision problems are called sequential. It may be convenient to think of such problems as being faced at $t \geq 0$ and extending over a small subinterval of $[t, t + 1]$. In contrast, recall the static problems considered earlier of the one-shot choice at t between timeless wealth lotteries. Alternatively, the multistage nature of a decision problem could be due exclusively to the way in which the problem is perceived by the individual. In that case no real time passes as the stages are traversed but the problem is still sequential.

This section begins by discussing how agents who do not maximize expected utility behave in sequential choice settings. As in the intertemporal setting, dynamic consistency is an important issue and it can be understood here in similar terms. The discussion of sequential choice behavior provides a useful perspective on the research into game theory with non-expected utility preferences, which is reviewed in the second subsection.

5.1 Sequential choice

Frequently risks are resolved gradually over time and hence the objects of choice are compound or multistage lotteries. An example is the two-stage lottery $(F_1, p_1; \dots; F_n, p_n)$, where each F_i occurs with probability p_i in the first stage and F_i is a cumulative distribution function representing the simple wealth lottery to be played in the second stage. If there is no risk in the first stage, $(F, 1)$ may be written simply as F .

An example of a sequential choice problem, that will serve as the vehicle for the discussion, is portrayed in figure 1.6a by means of the standard decision tree diagram. Circles denote chance nodes and squares denote decision nodes. Think of the agent making a choice at the first decision node and formulating a contingent choice for the second node. The choice is essentially between the compound lotteries $(I, 1)$, $(F, p; H, 1 - p)$ and $(G, p; H, 1 - p)$, and is presumably based on the utility function for compound lotteries. If the second decision node is reached, one may wonder




Figure 1.6 Decision trees

whether the contingent plan will be carried out. If so, preferences will be said to be dynamically consistent. If the contingent choice made at the first node is binding, then the decision problem facing the agent is the normal form problem in figure 1.6b. The latter is a one-shot choice problem for which dynamic consistency is not an issue. It is important for what follows to note that the normal form problem is defined as the choice between the two-stage lotteries $(I,1)$, $(F,p;H,1-p)$, and $(G,p;H,1-p)$ rather than between the simple lotteries I , $pF + (1-p)H$ and $pG + (1-p)H$. With the chosen definition, dynamic inconsistency is the only potential reason for different choices in the normal and extensive decision problems.

Whether or not consistency prevails depends upon the utility function for simple lotteries that determines choice at the second node. Denote by v^1 and v^2 the utility functions that apply at nodes 1 and 2. The usual starting point for their specification is a given utility function v for simple lotteries, which will be assumed to be strictly increasing in the sense of first-degree stochastic dominance; v could belong to the betweenness class or be Fréchet differentiable (section 2). The three principal routes that have been followed to move from v to v^1 and v^2 are outlined in figure 1.7, along with their main features.²⁸ It should be evident that the main points apply much more generally than the particular problem in figure 1.6.

Most commonly it is argued that, since the entire decision problem is assumed to occur over a short time interval, the individual is indifferent to the way in which risk is resolved over time. That is, for all compound lotteries:

$$(F_1, p_1; \dots; F_n, p_n) \sim (\sum p_i F_i, 1).$$

Given this reduction of compound lotteries axiom (ROCLA), the utility function v^1 that applies at the first decision node is simply:

$$v^1(F_1, p_1; \dots; F_n, p_n) \equiv v(\sum p_i F_i). \quad (5.1.1)$$

It remains to define utility functions for subsequent nodes or more particularly, for the second decision node in figure 1.6a. The situation is similar to that for intertemporal utility discussed in section 3.2. Drawing the obvious parallels, refer to the constant tastes specification as the one where v^2 is defined as the restriction of v^1 :

$$v^2(\cdot) = v^1(\cdot, p; H, 1 - p), \quad (5.1.2)$$

where equality is modulo ordinal equivalence. Machina (1989a) provides detailed and forceful arguments to support this specification, as well as references to antecedents in the literature. It implies dynamic consistency and also that, unless v is expected utility, preferences at the second node depend upon the unrealized alternative H , due possibly to psychological considerations such as relief or disappointment. Thus backward induction or "rolling back," where decisions at each node are made without reference to "what might have been," is not applicable.²⁹ Finally, the extensive form decision problem is identified by the agent with the associated normal form (figure 1.6b) which amounts under ROCLA to a static problem of choice between $I, pF + (1 - p)H$ and $pG + (1 - p)H$.

Any deviation from (5.1.2) implies changing tastes along the decision tree. For example, suppose that:

$$v^2(\cdot) \equiv v(\cdot), \quad (5.1.3)$$

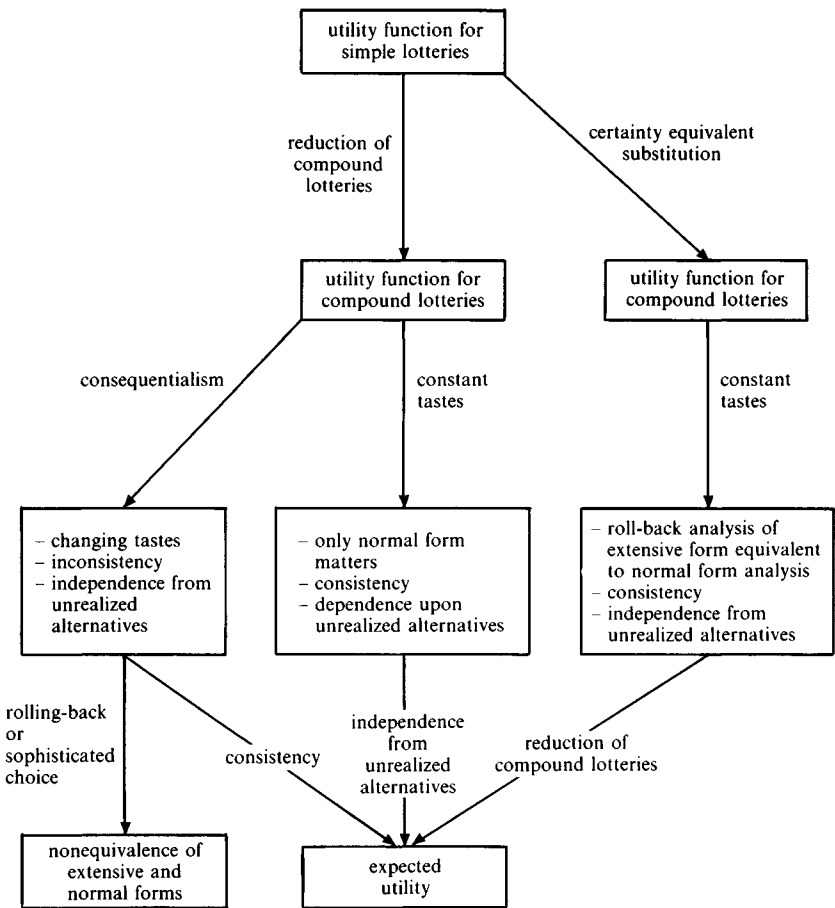


Figure 1.7 Sequential utility

which is the predominant specification for sequential choice problems (e.g., Raiffa, 1968) and is termed *consequentialism* by Machina following Hammond (1988, 1989).³⁰ According to (5.1.3), the choice between F and G at node 2 is made as though the rest of the tree did not exist, whereas from the perspective of decision node 1 and v^1 , the presence of H generally influences the contingent choice between F and G . This change in tastes leads to a dynamic inconsistency for any F, G, H , and p for which the conditions of the independence axiom are violated by v , e.g., if:

$$v(F) > v(G) \text{ and } v(pF + (1 - p)H) < v(pG + (1 - p)H). \tag{5.1.4}$$

This dynamic interpretation of violations of the independence axiom is the basis for the widely held view that non-expected utility preferences must be dynamically inconsistent.

One prescription for behavior in such a situation is to solve the extensive form problem by “rolling back,” letting v^2 determine the contingent choice at the second node, thus determining the opportunity set to which v^1 is applied. (The corresponding behavior in the intertemporal setting was termed sophisticated in section 3.2.) Given (5.1.4), such a procedure implies that the choice at node 1 is between I and $pF + (1 - p)H$. In most cases such rolling back of decision trees does not produce a lottery that is optimal in the normal form (Hazen, 1987). In those cases, the decision-maker would be willing to pay a positive price for the ability to commit to the contingent choices dictated by v^1 , since such commitment would mean that the normal form problem is the one being faced.³¹

There is a third approach to the modeling of the sequential choice behavior of non-expected utility maximizers. In this approach, advocated by Segal (1990), the reduction of compound lotteries axiom is dropped, reflecting the frequent finding in the psychology literature that a compound lottery is perceived differently if it is reduced to a single stage. In a series of papers, referenced in his (1990) paper, Segal argues that dropping ROCLA can lead to a unified explanation of a range of behavioral evidence, thus bolstering the argument against the adoption of ROCLA in descriptive modeling.

As an alternative to (5.1.1), extend v to the domain of compound lotteries by means of *certainty equivalent substitution*, i.e., by declaring the compound lottery $(F_1, p_1; \dots; F_n, p_n)$ to be indifferent to the simple lottery $\Sigma p_i \delta_{x_i}$, where each x_i is the certainty equivalent of F_i with respect to v , i.e.:

$$v(\delta_{x_i}) = v(F_i). \quad (5.1.5)$$

In short, v^1 is defined by (5.1.5) and:

$$v^1(F_1, p_1; \dots; F_n, p_n) \equiv v(\Sigma p_i \delta_{x_i}). \quad (5.1.6)$$

Note the algorithmic appeal of this way of evaluating compound lotteries and the parallel with recursive intertemporal utility. Finally, the function v^2 is specified according to (5.1.2) to reflect constant tastes.

These utility specifications have the following immediate implications. First, $v^2(\cdot) = v(\cdot)$ up to ordinal equivalence, so that unrealized alternatives do not affect preferences. Second, dynamic consistency obtains since:

$$v^2(F) \geq v^2(G) \Leftrightarrow v^1(F, p; H, 1 - p) \geq v^1(G, p; H, 1 - p).$$

(The fact that $v(\cdot)$ is strictly increasing is important here.) Thus roll-back analysis of the extensive form delivers a compound lottery that is optimal in

the normal form. (However, this statement is false if by the normal form one means the choice between the simple lotteries I , $pF + (1-p)H$ and $pG + (1-p)H$.)

A variety of normative arguments have been applied in an attempt to differentiate between the three approaches outlined above. For example, under consequentialism, sophisticated behavior cannot be rationalized by a preference ordering (Hammond, 1976, 1989) unless v is an expected utility function. Machina (1989a) argues in favor of his constant tastes specification in part by drawing the analogy with the implications of non-separability for intertemporal choice under certainty. Finally, when the duration of a choice problem is short, or conceptual rather than real, the reduction of compound lotteries axiom has considerable normative appeal. On the other hand, violation of the axiom might be understood as reflecting a utility or disutility of gambling. It is for this reason that von Neumann and Morgenstern (1953, p. 632) identify the relaxation of the reduction axiom as an important direction for the extension of their theory of games.

One might also try to exploit the obvious analogy with the analysis of intertemporal choice behavior to differentiate between the three approaches. There is some merit to maintaining consistency across the two contexts with respect to constant versus changing tests, for example. On the other hand, if recursive utility is specified in the intertemporal context, it does not follow that its counterpart for sequential choice, Segal's certainty equivalent substitution approach, is the only sensible one to adopt here. It is clearly plausible that the attitude towards the way in which a multistage lottery is resolved as the stages are traversed depends upon whether time is conceptual or real.

It is consistent with the emphasis throughout this chapter to argue that the differentiation being discussed should be made on the basis of "usefulness" in standard sequential choice models. The following section will outline some recent applications of the above three modeling approaches to the study of strategic interactions between agents.³² Unlike the case for intertemporal choice modeling, however, the research on applied sequential choice models has not yet delivered a convincing demonstration of the superiority of any of the three approaches.

One final point regarding the modeling of general sequential choice problems deserves attention. As mentioned, in all three approaches the starting point is a utility function v for simple lotteries. It is interesting to note that none of these approaches directly imply restrictions on the set of acceptable specifications for v . However, some restrictions are forthcoming under additional assumptions regarding the nature of sequential choice. Under consequentialism (and the reduction of compound lotteries) v must be expected utility if dynamic consistency is demanded. In the context of

Machina's approach to sequential choice Gul and Lantto (1990) describe arguments that rule out all but betweenness-satisfying functions v . A complementary argument in support of betweenness, that applies to all three approaches under discussion, follows from the requirement of consistency between the modeling of intertemporal and sequential choice. In sections 3.4 and 4.5, it was suggested that a betweenness-conforming specification for the certainty equivalent is advantageous in the context of recursive intertemporal utility functions. Such a specification corresponds to the hypothesis that the derived utility function for timeless wealth gambles satisfies betweenness. However, that derived utility function is naturally identified with the function v .

5.2 Game theory

The theory of games has been one of the most important areas of application for expected utility theory. Therefore, to demonstrate the comparable usefulness of non-expected utility theories, it is important to show that the theory of strategic interactions between agents can be extended to incorporate broader definitions of individual rationality. The research in this area is still at an early stage in its development. Accordingly the summary to follow is relatively brief.³³

Following Crawford (1990), consider a two-person normal form game of complete information. The players, R and C, have finite sets of pure strategies and corresponding probability simplices Σ_R and Σ_C representing sets of mixed strategies. A combination of mixed strategies induces a probability distribution of monetary outcomes for each player and preferences over such lotteries are represented by the utility functions v_R and v_C . A *Nash equilibrium* is a pair of mixed strategies $(p^*, q^*) \in \Sigma_R \times \Sigma_C$ such that each is a best response to the other. Note that the specification of payoffs is in monetary terms, rather than in utility terms as is customary in the expected utility framework. This is advantageous since it permits the separation of the game structure from players' preferences over monetary gambles.

The primitive utility functions v_R and v_C and the mapping from $\Sigma_R \times \Sigma_C$ into monetary gambles induce utility functions \hat{v}_R and \hat{v}_C on $\Sigma_R \times \Sigma_C$ which represent preferences over strategy pairs. It is well-known (see Debreu, 1952, for example) that a Nash equilibrium exists if \hat{v}_R and \hat{v}_C are each quasiconcave in own probability vectors for each given choice of opponents' strategies. Moreover, $\hat{v}_R(\cdot, q)$ is quasiconcave on $\Sigma_R \forall q \in \Sigma_C$ if and only if v_R is quasiconcave on (a suitable subdomain of) probability distributions over monetary outcomes, i.e., "better-than" sets in probability triangles corresponding to monetary outcomes are convex. A similar statement

applies to v_C and \hat{v}_C . Thus existence of an equilibrium is ensured if v_R and v_C are both quasiconcave and *a fortiori* if both satisfy betweenness, which is the conjunction of quasiconcavity and quasiconvexity. In general, however, an equilibrium may fail to exist because the players may be unwilling to randomize. For example, under strict quasiconvexity, for any mixed strategy of C, R will strictly prefer at least one of his pure strategies. Thus an equilibrium does not exist if there is no pure strategy equilibrium. Technically the difficulty is that best reply correspondences need not be convex-valued.

I have argued in earlier sections that the restriction to betweenness-satisfying utility functions represents an attractive balance between generality and tractability, and the discussion of game theory will bolster this view. Nevertheless, it is desirable to have a theory of strategic interaction that is not limited by assumptions such as the independence axiom, the betweenness axiom or quasiconcavity in probabilities, which should be viewed as empirically refutable hypotheses rather than tenets of individual rationality. Moreover, there exists some laboratory evidence contradicting betweenness, with the violations divided roughly equally between quasiconcavity and quasiconvexity (see footnote 3).

To accommodate more general preferences, Crawford adapts from Aumann (1987) the notion of an *equilibrium in beliefs*, which is a pair (P^*, Q^*) such that: P^* and Q^* are probability distributions on Σ_C and Σ_R respectively, that represent beliefs about opponents' strategy choices; and $P^*(Q^*)$ assigns positive probability only to those mixed strategies of C(R) that are optimal responses for him given his beliefs. The beliefs equilibrium concept has three attractive features. First, such an equilibrium exists whenever utility functions are continuous. This is so because, if players' utility functions are transformed to make them quasiconcave by replacing "better-than" sets by their convex hulls, then each Nash equilibrium of the convexified game defines an equilibrium in beliefs for the original game. For the second feature note that, if $(p^*, q^*) \in \Sigma_R \times \Sigma_C$ is a Nash equilibrium, then the pair of beliefs which assign probability 1 to (p^*, q^*) constitutes a beliefs equilibrium. Conversely, assuming the reduction of compound lotteries axiom, any probability distribution on Σ_R (or Σ_C) defines a unique element of Σ_R (or Σ_C). (Modulo this identification, Crawford shows that any beliefs equilibrium is also a Nash equilibrium given the quasiconcavity of utility.) In this sense, the notion of equilibrium in beliefs coincides with Nash equilibrium given quasiconcavity. Finally, in a beliefs equilibrium explicit randomization on the part of agents is necessary only if it is desired, i.e., if utility is strictly quasiconcave in probabilities. If both players have quasiconvex utility functions, then each chooses a pure strategy, with no attempt to randomize, but there is uncertainty about the opponent's choice (see also Harsanyi, 1973).

Dekel, Safra, and Segal (1991) offer an interesting perspective on Crawford's analysis and an alternative approach to modeling normal form games. They argue that, since the opponent's strategy is generally mixed, each player must rank alternative two-stage lotteries in choosing a mixed strategy. For example, let r_{ij} be the monetary payoff to R if he plays his i^{th} pure strategy and C plays his j^{th} pure strategy, $i = 1, \dots, m$ and $j = 1, \dots, n$. Then, if C plays $q = (q_1, \dots, q_n) \in \Sigma_C$ and if R views himself as playing first, R's choice of $p \in \Sigma_R$ implies the two-stage lottery in which the simple lottery $(r_{i1}, q_1; \dots; r_{in}, q_n)$ is obtained with probability $p_i, i = 1, \dots, m$. Crawford implicitly assumes that R employs the reduction of compound lotteries axiom to reduce the compound lottery to a simple lottery, which is then evaluated by the utility function v_R . Dekel, Safra, and Segal assume instead that each agent uses the certainty equivalent substitution procedure described in section 5.1 to evaluate compound lotteries. Under the assumptions that each player views himself as moving first and that v_R and v_C are continuous and strictly increasing in the sense of first-degree stochastic dominance, they prove that a Nash equilibrium exists. Each best reply set is convex, even if the utility functions are quasiconvex, since it equals the convex hull of the collection of pure strategies that are best replies. On the other hand, if each player perceives himself as moving second, in which case the pair of strategies (p, q) leads to the compound lottery for R in which $(r_{1j}, p_1; \dots; r_{mj}, p_m)$ is encountered with probability $q_j, j = 1, \dots, n$, then convex-valued best response correspondences can be guaranteed only if both players are expected utility maximizers. Thus it remains to find an equilibrium concept (perhaps a variation of the beliefs equilibrium) that can be applied under alternative assumptions about individual perceptions.

The Dekel, Safra, and Segal analysis draws attention to the distinction between normal and extensive form modeling of games that is one of the principal lessons that has emerged from the extension beyond the expected utility framework. Define the normal form of the above game as in Crawford (1990) whereby, because of the reduction axiom, the strategy pair (p, q) produces a simple monetary lottery for each player. Then the preceding paragraph points to the loss of information in restricting attention to the normal form and consequently to the need for extensive form analysis. On the other hand, if reduction is not imposed in the definition of the normal form, so that (p, q) produces a two-stage monetary lottery, then two different normal forms are being discussed in the preceding paragraph and attention is drawn rather to the proper definition of the normal form. (The parallel issue in the context of individual choice problems was mentioned in section 5.1.) Note that the distinction between the alternative definitions of the normal form is immaterial if reduction of compound lotteries is imposed at the level of preferences.

The explicit formulation of players' choices in terms of compound lotteries calls attention to the sequential nature of individual decision problems and therefore to the issue of dynamic consistency. If explicit randomization is actually undertaken, then each player has two non-trivial decisions to make – first, which mixed strategy to employ and second, whether or not to play the pure strategy that is delivered by the chosen randomization. In light of the discussion in the preceding section, it is evident that consistency is guaranteed in Crawford's analysis if Machina's view of sequential choice behavior is adopted. It is also guaranteed in the analysis of Dekel, Safra, and Segal, since they employ the certainty equivalent substitution approach to sequential choice.

Under consequentialism, however, consistency is potentially a problem if quasiconvexity is violated. This is because a decision-maker with strictly quasiconcave utility may choose to randomize over two non-different lotteries F and G and then renege if confronted with the outcome of the inferior alternative.³⁴ Even if F and G are equally attractive, the decision-maker may prefer to randomize again rather than accepting either F or G . Crawford assumes these problems away by supposing that players commit themselves to abide by the outcome of the mixed strategy. He argues that the ability to make such commitments is implicit in the assumption that the players can use mixed strategies. Viewed in this light, the preceding discussion of consistency is more properly reinterpreted as the comparison of behavior in different extensive form games having the same normal form.

The related issues of dynamic consistency and the equivalence of normal and extensive forms have been considered in the context of auction games. Suppose that bidders behave non-cooperatively and that their valuations of the auctioned object are private and independent. Consider a Dutch or descending-bid auction and a first-price sealed-bid auction. The corresponding games have identical (reduced) normal forms and produce the identical outcomes under expected utility (Milgrom and Weber, 1982; Milgrom, 1989). But more generally, they produce different outcomes assuming consequentialism, because of the dynamic inconsistency present in the Dutch auction. To elaborate, consider a bidder in the Dutch auction with monetary valuation of the auctioned object equal to r . Suppose the price has fallen to $b + 1 < r$ and that the individual is deciding whether to claim the object at that price or at the "next" price b . Either choice implies a lottery which can be represented by a point in the probability simplex corresponding to outcomes 0, $r - b - 1$ and $r - b$. Denote these lotteries by F^{b+1} and F^b (where $F^{b+1} = \delta_{r-b-1}$ and F^b is a mixture of δ_0 and δ_{r-b}). On the other hand, consider the same bidder at the start of the Dutch auction, or in the sealed-bid auction, comparing the same two bids. In this case, either choice entails the additional risk, having probability α say, that someone else will bid more than $b + 1$. Thus the relevant comparison is

between the lotteries $(1-\alpha)F^b + \alpha\delta_0$ and $(1-\alpha)F^{b+1} + \alpha\delta_0$. If the independence axiom is satisfied, F^b is chosen in the former case if and only if $(1-\alpha)F^b + \alpha\delta_0$ is chosen in the latter comparison, but not so more generally. Indeed, if indifference curves fan out as portrayed for weighted utility in figure 1.1, then the individual could offer b in the first-price auction and $b+1$ in the Dutch auction once the price has fallen to that level. (Note that the chord in the probability simplex connecting F^b and F^{b+1} is parallel to that connecting $(1-\alpha)F^b + \alpha\delta_0$ and $(1-\alpha)F^{b+1} + \alpha\delta_0$.) If indifference curves fan in, the lower bid would occur in the Dutch auction.

Karni and Safra (1989a) show that dynamic consistency is a problem also in the English auction under the consequentialist view of sequential choice if the object being sold is a lottery. In (1989b), these authors posit sophisticated behavior and assume that bidders restrict attention to strategies that will actually be carried out. Formally, Karni and Safra adapt Selten's (1975) trick (see his "agent's normal form") of regarding the same bidder at different decision nodes as distinct players. Then they consider the Bayesian-Nash equilibrium of the resulting game of incomplete information. Under the assumption that utility functions are quasiconcave, they establish the existence of an equilibrium, which they show to be value-revealing if and only if betweenness is satisfied.

To this point, the discussion of games has concentrated on the consequentialist and certainty equivalent substitution approaches to sequential choice. If Machina's (1989a) approach is adopted, then the discussion in the last section implies immediately that only the normal form matters and that dynamic consistency generally prevails. However, the approach conflicts with the principle of backwards induction (see also footnote 29).

In summary, the literature on game theory with general preferences has just begun to tackle the problem of formulating a satisfactory equilibrium concept. The modeling of strategic rationality depends upon which of the three approaches to modeling individual sequential choice behavior outlined in figure 1.7 is adopted. Regardless of the approach that is adopted, however, it has been shown that one of the following features of received game theory must be abandoned if utility functions do not conform to expected utility theory: the applicability of backwards induction, or the equivalence of normal and extensive forms in games of perfect recall, at least where the normal form is defined in the usual way in terms of simple rather than multistage payoff lotteries. As for applications, the literature has produced some new predictions regarding the equivalences, efficiency, and demand-revealing properties of common auctions. It remains to be seen whether non-expected utility preferences will deliver interesting new predictions in other contexts.

Notes

- * I am grateful to the Social Sciences and Humanities Research Council of Canada for financial support. I am especially indebted to Darrell Duffie, Angelo Melino, Michael Peters, and Uzi Segal for valuable discussions and comments. This draft has also benefited from suggestions by Eddy Dekel, Ingrid Peters-Fransen, Jerry Green, Carolyn Pitchik, Zvi Safra, A. Siow, Tan Wang, and Philippe Weil.
- 1 For certain quantitative empirical exercises, however, there is a substantial difference in tractability between Fréchet differentiable functions and expected utility functions (see footnote 21).
 - 2 The reader is referred to the surveys cited in the introduction for descriptions of other classes of utility functions and appropriate references. Some notable models include rank-dependent expected or anticipated utility (references given above), the non-transitive regret theory (Bell, 1982; Loomes and Sugden, 1982) and the closely related skew-symmetric-bilinear utility theory (Fishburn, 1982). Some studies that have applied these models are listed below, but these alternative models are not particularly useful for the applications in sections 4 and 5. The same comment applies to prospect theory (Kahneman and Tversky, 1979). The latter also suffers, in comparison with expected utility and the other models mentioned, from more ambiguous predictions because of the lack of a precise theory of the framing and editing processes.
 - 3 See the cited sources on betweenness and the surveys mentioned in the introduction for a discussion of the extent to which the betweenness axiom is compatible with the behavioral evidence against independence. There have also been some attempts to examine directly the descriptive validity of betweenness in the laboratory. See Camerer (1989a,b) and the references and discussions in Machina (1985), Crawford (1990), and Chew, Epstein, and Segal (1991).
 - 4 Conversely, given such a function ϕ defined on $[a,b]$, $0 < a < 1 < b < \infty$, equation (2.3.2) defines a monotone and risk averse μ on $D[a,b]$, the set of cdf's on $[a,b]$. If ϕ is defined and well-behaved on $(0,\infty)$, then μ is well-defined, monotone, and risk averse on the set of all cdf's on R^1_{++} having finite mean. For example, to show risk aversion let G be a mean preserving spread of F . Then:

$$\int \phi(x/\mu(F))dF(x) = 0 = \int \phi(x/\mu(G))dG(x)$$

$$\leq \int \phi(x/\mu(G))dF(x) \Rightarrow \mu(F) \geq \mu(G).$$

See Epstein and Zin (1991b) for further details. These facts can be applied to deduce the domains of the certainty equivalent functions corresponding to (2.3.3) and (2.3.6) below.

- 5 Chew (1989) axiomatizes the family of semi-weighted utility functions, a subset of the betweenness class containing both of the parametric classes to follow.
- 6 This notion of comparative risk aversion for certainty equivalent functions is

employed in the analysis of intertemporal utility (section 3.3). Machina (1982b, p. 299) formulates and analyzes a stronger notion.

- 7 Related inhibitions to trade also appear in Gilboa and Schmeidler (1989), Bewley (1989), and Dow and Werlang (1992). Kinks in indifference curves are critical there also, but the kinks reflect aversion to uncertainty (randomness with unknown probability) rather than aversion to risk.
- 8 If only single period risks are considered, then the functional form (3.2.1), which is based upon research by Larry Selden, provides a complete separation between substitution and risk aversion in that it is possible also to change the ranking of deterministic programs without affecting risk aversion.
- 9 More precisely, stationarity requires that for all t and consumption programs, and for all events $I \in F_{t+1}$ (see the notation introduced in section 3.2):

$$\begin{aligned} (c_t, \tilde{c}_{t+1}, \tilde{c}_{t+2}, \dots | I) &\succeq_{s_t} (c_t, \tilde{c}'_{t+1}, \tilde{c}'_{t+2}, \dots | I) \iff \\ &(\tilde{c}_{t+1}, \tilde{c}_{t+2}, \dots | I) \succeq_{(s_t, c_t)} (\tilde{c}'_{t+1}, \tilde{c}'_{t+2}, \dots | I). \end{aligned}$$

Stationarity is assumed without exception (to my knowledge) in both the theoretical and empirical literature on capital theory that employ the framework of an infinitely lived agent. This is true in particular of the habit-formation literature, for which some references are provided in section 4.1. Epstein (1983) considers stationarity in the case where consumption histories do not influence preferences and derives the appropriate specialization of (3.13).

- 10 Weak recursivity can be defined more formally as follows: let \tilde{c} , \tilde{c}' , \tilde{c}^* and \tilde{c}^{**} be four consumption programs and $T > 0$. For any $I \in F_T$, denote by $(\tilde{c}|I, \tilde{c}^*|\Omega \setminus I)$ the consumption program in which period t consumption is $\tilde{c}_t(\omega)$ if $\omega \in I$ and $\tilde{c}^*_t(\omega)$ if $\omega \in \Omega \setminus I$; and similarly for other combinations of the above programs. Weak recursivity requires that $\forall I_T \in F_T$, if:

$$V(\tilde{c}|I_T, \tilde{c}^*|\Omega \setminus I_T) \geq V(\tilde{c}'|I_T, \tilde{c}^*|\Omega \setminus I_T) \quad \forall I_T \supseteq I \in F_T,$$

then the same should be true if \tilde{c}^* is replaced by \tilde{c}^{**} . An intertemporal expected utility function satisfies the stronger condition that $\forall I_T \in F_T$, if:

$$V(\tilde{c}|I_T, \tilde{c}^*|\Omega \setminus I_T) \geq V(\tilde{c}'|I_T, \tilde{c}^*|\Omega \setminus I_T),$$

then the inequality is true also if \tilde{c}^* is replaced by \tilde{c}^{**} . To see the difference between these conditions in the context of figure 1.3 would require that a third alternative be added at $t=1$. For further elaboration in the context of a two-period model see Johnsen and Donaldson (1985, pp. 1454–5).

- 11 Recursivity requires that $\forall \tilde{c}, \tilde{c}', \tilde{c}^*, T > 0$ and $I_T \in F_T$:

$$V(\tilde{c}|I_T, \tilde{c}^*|\Omega \setminus I_T) \geq V(\tilde{c}'|I_T, \tilde{c}^*|\Omega \setminus I_T) \quad \forall I_T \supseteq I \in F_T,$$

if and only if:

$$V(\tilde{c}_T, \tilde{c}_{T+1}, \dots | I) \geq V(\tilde{c}'_T, \tilde{c}'_{T+1}, \dots | I) \quad \forall I_T \supseteq I \in F_T.$$

- 12 Working in a domain of probability measures rather than random variables, Chew and Epstein (1991) show that (3.3.1) is characterized by a slightly strengthened form of recursivity. Note that (3.3.1) implies stationarity in the sense of section 3.1, footnote 9.
- 13 There are several (W, μ) pairs that represent the same intertemporal preference ordering. If $V^* = g(V)$ is ordinaly equivalent to V then V^* , W^* , and μ^* satisfy the appropriate form of (3.3.1) if $W^*(c, z) \equiv g(W(c, g^{-1}(z)))$ and $\mu^*(\tilde{x}) = g(\mu(g^{-1}(\tilde{x})))$. But if, as here, the aggregator is held fixed, then the certainty equivalent function corresponds uniquely to the intertemporal preference ordering.
- 14 Let J be the value function corresponding to V and the intertemporal optimization problem. Then timeless wealth gambles \tilde{x} are evaluated by $\mu(J(\tilde{x}))$. For most of the axiomatic generalizations of expected utility that have been developed recently, including the betweenness theory, $\mu(\cdot)$ lies in the given axiomatic class if and only if $\mu(J(\cdot))$ does. Finally, if J and its inverse both satisfy Lipschitz conditions, then $\mu(\cdot)$ is Fréchet differentiable if and only if $\mu(J(\cdot))$ is.
- 15 Indeed, Kreps and Porteus studied recursive utility functions **because** they implied non-indifference to temporal resolution. On the other hand, the more recent attention that has been afforded these functions has been motivated more by the separation they deliver between substitution and risk aversion.
- 16 Figure 1.4 also indicates (in the right branch) that dynamic consistency and timing indifference can be achieved simultaneously, but at the cost of allowing dependence upon unrealized alternatives. Any utility function V , which depends only on the joint distribution of the \tilde{z}_t 's, is indifferent to the temporal resolution of risk.
- 17 The utility specification is not the only potential cause of the empirical failures. Some studies have focused on the consequences of generalizing the standard model by incorporating heterogeneous agents and incomplete markets. See Marcet and Singleton (1990) and Weil (1990b), for example.
- 18 A general analysis of dynamic programming with recursive utility is not yet available. Some supporting arguments for the homogeneous functional forms below are provided in Epstein and Zin (1989). Streufert (1990b) and Ma (forthcoming) contain some results for more general cases.
- 19 See Epstein and Zin (1989) for details and also for a proof that there exists a utility function V , defined on a suitable domain, that satisfies recursive relation (3.3.1) for the specified W and μ . Existence theorems comparable in generality to those available in the certainty case (Streufert (1990a)) are not yet available. But see Streufert (1990b) and Ma (forthcoming).
- 20 See the discussion of term structure in section 4.4 for an example of such additional assumptions. It is shown there that optimal consumption is myopic when $\rho = 0$. The non-myopic nature of optimal consumption when $\rho \neq 0$ is essential in the derivation of (4.2.5) (see Epstein and Zin, 1989).
- 21 If μ is Fréchet differentiable with local utility function u , the first-order conditions for portfolio choice take the form:

$$E_t[u_1(\tilde{z}_{t+1}; F_t^*) (\tilde{r}_{i,t+1} - \tilde{r}_{j,t+1})] = 0,$$

where $\tilde{z}_{t+1} \equiv (\tilde{c}_{t+1}/c_t)^{(\rho-1)/\rho} \tilde{M}_{t+1}^{1/\rho}$ and F_t^* is the conditional cdf of \tilde{z}_{t+1} . Generalized method of moments estimation is not applicable since F_t^* is not observed. Only if a parametric form is assumed for F_t^* can the integrand be computed, given specified values for all unknown parameters. On the other hand, if μ satisfies betweenness, then, because of (2.2.4), u_1 can be replaced in the above expectation by $H_1(z_{t+1}, \mu(\tilde{z}_{t+1}|I_t))$ which depends on F_t^* only via $\mu_t^* = \mu(\tilde{z}_{t+1}|I_t)$. Epstein and Zin (1989) show how μ_t^* can be expressed in terms of observables by suitably exploiting the Bellman equation. Thus the specialization to a betweenness-conforming certainty equivalent is advantageous for empirical tractability.

- 22 For the disappointment averse specification, ϕ is not differentiable at 1, but under specified assumptions on the distribution of consumption and asset returns that is of no consequence in (4.5.1). See Epstein and Zin (1991a).
- 23 The following component property of recursivity is violated: for all consumption programs:

$$\begin{aligned} V(y_0, \tilde{y}_1, \tilde{y}_2, \dots) &\geq V(y_0, \tilde{y}_1', \tilde{y}_2', \dots) &<=> \\ V(y_0, \tilde{y}_1, \tilde{y}_2, \dots) &\geq V(y_0, \tilde{y}_1', \tilde{y}_2'', \dots), &\text{where } y_t \equiv (c_t, q_t). \end{aligned}$$

- 24 Theorems guaranteeing that utility is well-defined by the recursive equation (4.6.4) or its continuous time counterpart are not yet available. However, if existence of utility is assumed, the remaining arguments in the cited papers are readily extended.
- 25 The papers cited below provide some results regarding the existence of an equilibrium in the representative agent economy with the CES specialization of recursive utility. Ma (forthcoming) provides an existence result for a heterogeneous agent economy with recursive utilities by extending Duffie, Geanakoplos, Mas-Colell, and McLennan (1988).
- 26 Similar intuition was validated formally by Kihlstrom and Mirman (1974, pp. 378–80) and Selden (1979) with respect to the behavior of an individual facing exogenous prices and operating in a two-period environment.
- 27 Kocherlakota's observational equivalence result for an i.i.d. environment should be distinguished from the Duffie and Epstein (1991a) result for continuous time Brownian (but not necessarily i.i.d.) environments, described in section 3.3. Kocherlakota is concerned with observational equivalence with intertemporal expected utility, while Duffie and Epstein are concerned with whether one could detect if the ordering of timeless wealth gambles (i.e., the certainty equivalent μ) violates expected utility theory.
- 28 There is an evident parallel between the discussion of intertemporal choice (based on figure 1.4) and that of sequential choice (based on figure 1.7). Roughly speaking, the latter can be viewed as the special case of intertemporal choice where there is no discounting, consumption is perfectly substitutable across time, and, possibly, the horizon is finite. But the parallel is imperfect. For example, the usual starting point in the sequential framework is a utility function defined over simple, rather than over multistage, lotteries. Moreover, the associated literatures have developed independently and employ different terminology. Thus a brief separate treatment of sequential choice is provided.
- 29 Mention should be made of the related modeling of psychological factors in

- strategic situations by Geanakoplos, Pearce, and Stacchetti (1989), who also point out the inapplicability of backwards induction in their context.
- 30 There is some disagreement about the proper definition of consequences and hence of consequentialism. See Hammond (1989, pp. 1447–8) and Machina (1989a, section 6.6).
 - 31 In single-person decision problems, the ability to commit, in the sense of guaranteeing that v^1 dictate all choices, is always (weakly) preferable according to v^1 . However, Fershtman, Safra, and Vincent (1990) show that in strategic situations an agent may strictly prefer not to commit in this way. Roughly, entering into a game with utility function v^2 rather than v^1 may produce an equilibrium payoff to the player which is preferred by v^1 .
 - 32 See Karni and Safra (1990) for a single-agent sequential choice problem corresponding to optimal search. They adopt the consequentialist approach and propose a form of sophisticated behavior to resolve the dynamic inconsistency.
 - 33 In addition to the studies referred to below, mention should be made of Fishburn and Rosenthal (1986). They study Nash equilibrium in finite non-cooperative games where players have skew-symmetric-bilinear utility functions and thus are non-transitive. Also, recall the paper on axiomatic bargaining theory by Rubinstein, Safra, and Thomson, cited in section 2.5.
 - 34 See Green (1987) for a related argument that an agent who violates quasiconvexity can be exploited.

References

- Attanasio, O.P. and G. Weber (1989), "Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption," *Economic Journal* (supplement), 99: 59–73.
- Aumann, R.J. (1987), "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55: 1–18.
- Bell, D. (1982), "Regret in Decision Making Under Uncertainty," *Operations Research*, 30: 961–81.
- Bewley, T. (1989), "Market Innovation and Entrepreneurship: A Knightian View," Paper 905, Cowles Foundation.
- Bollerslev, T., R.F. Engle, and J.M. Wooldridge (1988), "A Capital Asset Pricing Model with Time-Varying Covariances," *Journal of Political Economy*, 96: 116–31.
- Breeden, D. (1979), "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment," *Journal of Financial Economics*, 7: 265–96.
- Breeden, D., M.R. Gibbons, and R.H. Litzenberger (1989), "Empirical Tests of the Consumption-Oriented CAPM," *Journal of Finance*, 44: 231–62.
- Bufman, G. and L. Leiderman (1990), "Consumption and Asset Returns Under Non-Expected Utility," *Economic Letters*, 34: 231–5.
- Camerer, C.F. (1989a), "An Experimental Test of Several Generalized Utility Theories," *Journal of Risk and Uncertainty*, 2: 61–104.

- (1989b), "Recent Tests of Generalizations of Expected Utility Theory," Mimeo, Wharton School.
- Campbell, J.Y. (1990), "Intertemporal Asset Pricing Without Consumption," Working Paper 119, Woodrow Wilson School, Princeton University.
- Chew, S.H. (1983), "A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox," *Econometrica*, 51: 1065–92.
- (1989), "Axiomatic Utility Theories with the Betweenness Property," *Annals of Operations Research*, 19: 273–98.
- Chew, S.H. and L.G. Epstein (1989a), "The Structure of Preferences and Attitudes Towards the Timing of the Resolution of Uncertainty," *International Economic Review*, 30: 103–17.
- (1989b), "Axiomatic Rank-Dependent Means," *Annals of Operations Research*, 19: 299–309.
- (1990a), "Non-Expected Utility Preferences in a Temporal Framework with an Application to Consumption-Savings Behavior," *Journal of Economic Theory*, 50: 54–81.
- (1991), "Recursive Utility Under Uncertainty," in A. Khan and N. Yannelis (eds), *Equilibrium Theory with an Infinite Number of Commodities*, Springer Verlag.
- Chew, S.H., L.G. Epstein, and U. Segal (1991), "Mixture Symmetry and Quadratic Utility," *Econometrica*, 59: 139–64.
- Chew, S.H., L.G. Epstein, and I. Zilcha (1988), "A Correspondence Theorem Between Expected Utility and Smooth Utility," *Journal of Economic Theory*, 46: 186–93.
- Chew, S.H., E. Karni, and Z. Safra (1987), "Risk Aversion in the Theory of Expected Utility with Rank Dependent Probabilities," *Journal of Economic Theory*, 42: 370–81.
- Chew, S.H. and N. Nishimura (forthcoming), "Differentiability, Comparative Statics and Non-Expected Utility Preferences," *Journal of Economic Theory*.
- Conlisk, J. (1989), "Three Variants on the Allais Example," *American Economic Review*, 79: 392–407.
- Constantinides, G.M. (1982), "Intertemporal Asset Pricing with Heterogeneous Consumers and Demand Aggregation," *Journal of Business*, 55: 253–67.
- (1990), "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy*, 98: 519–43.
- Cook, V.T. (1989), "The Effects of Temporal Resolution on the Overall Utility and Suspense of Risky Monetary and Survival Gambles," Ph.D. Dissertation, Department of Psychology, McGill University.
- Cox, J., J. Ingersoll, and S. Ross (1985), "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53: 385–408.
- Crawford, V. (1990), "Equilibrium Without Independence," *Journal of Economic Theory*, 50: 127–54.
- Debreu, G. (1952), "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences*, 38: 886–93.

- Dekel, E. (1986), "An Axiomatic Characterization of Preferences under Uncertainty," *Journal of Economic Theory*, 40: 304–18.
- Dekel, E., Z. Safra, and U. Segal (1991), "Existence and Dynamic Consistency of Nash Equilibrium with Non-Expected Utility Preferences," *Journal of Economic Theory*, 55: 229–46.
- Donaldson, J. and L. Selden (1981), "Arrow-Debreu Preferences and the Reopening of Contingent Claims Markets," *Economic Letters*, 8: 209–16.
- Dow, J. and S.R. Werlang (1992), "Uncertainty Aversion and the Optimal Choice of Portfolio," *Econometrica*, 60: 197–204.
- Duffie, D. and L.G. Epstein (1991a), "Stochastic Differential Utility," Research Paper 1078, Graduate School of Business, Stanford University; forthcoming in *Econometrica*.
- (1991b), "Asset Pricing with Stochastic Differential Utility," Working Paper 9122, University of Toronto; forthcoming in *Review of Financial Studies*.
- Duffie, D., J. Geanakoplos, A. Mas-Colell, and A. McLennan (1988), "Stationary Markov Equilibria," mimeo.
- Epstein, L.G. (1983), "Stationary Cardinal Utility and Optimal Growth Under Uncertainty," *Journal of Economic Theory*, 31: 133–52.
- (1985), "Decreasing Risk Aversion and Mean-Variance Analysis," *Econometrica*, 53: 945–61.
- (1988), "Risk Aversion and Asset Prices," *Journal of Monetary Economics*, 22: 179–92.
- (1991), "Discussion of 'Substitution Over Time in Consumption and Work'," in L. McKenzie and S. Zamagni (eds), "Value and Capital" Fifty Years Later, London: Macmillan, pp. 268–78.
- Epstein, L.G. and S.E. Zin (1989), "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57: 937–69.
- (1990a), "'First-Order' Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics*, 26: 387–407.
- (1990b), "Consumption, Labor Supply and Portfolio Choice with Time and State Non-Separable Utility," mimeo, University of Toronto.
- (1991a), "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy*, 99: 263–86.
- (1991b), "The Independence Axiom and Asset Returns," NBER Technical Working Paper 109.
- Farmer, R.E.A. (1990), "RINCE Preferences," *Quarterly Journal of Economics*, 105: 43–60.
- Fershtman, C., Z. Safra, and D. Vincent (1990), "Delayed Agreements and Non-Expected Utility," Discussion Paper 867, Kellogg School, Northwestern University.
- Fishburn, P. (1982), "Nontransitive Measurable Utility," *Journal of Mathematical Psychology*, 26: 31–67.

- (1983), "Transitive Measurable Utility," *Journal of Economic Theory*, 31: 293–317.
- (1988), *Nonlinear Preference and Utility Theory*. Baltimore: Johns Hopkins University Press.
- Fishburn, P. and R. Rosenthal (1986), "Noncooperative Games and Nontransitive Preferences," *Mathematical Social Sciences*, 12: 1–7.
- Geanakoplos, J., D. Pearce, and E. Stacchetti (1989), "Psychological Games and Sequential Rationality," *Games and Economic Behavior*, 1: 60–79.
- Gilboa, I. and D. Schmeidler (1989), "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18: 141–53.
- Giovannini, A. and P. Weil (1989), "Risk Aversion and Intertemporal Substitution in the Capital Asset Pricing Model," NBER Working Paper 2824.
- Gorman, W.M. (1953), "Community Preference Fields," *Econometrica*, 21: 63–80.
- Green, J. (1987), "'Making Book Against Oneself,' The Independence Axiom and Nonlinear Utility Theory," *Quarterly Journal of Economics*, 102: 785–96.
- Grossman, S., A. Melino, and R. Shiller (1987), "Estimating the Continuous-Time Consumption-Based Asset-Pricing Model," *Journal of Business and Economic Statistics*, 5: 315–27.
- Grossman, S. and R. Shiller (1981), "The Determinants of the Variability of Stock Market Prices," *American Economic Review*, 71: 222–7.
- Gul, F. (1991), "A Theory of Disappointment Aversion," *Econometrica*, 59: 667–86.
- Gul, F. and O. Lantto (1990), "Betweenness Satisfying Preferences and Dynamic Choice," *Journal of Economic Theory*, 52: 162–77.
- Hall, Robert E. (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86: 971–87.
- (1985), "Real Interest and Consumption," NBER Working Paper 1694.
- (1988), "Intertemporal Substitution in Consumption," *Journal of Political Economy*, 96: 339–57.
- Hammond, P. (1976), "Changing Tastes and Coherent Dynamic Choice," *Review of Economic Studies*, 43: 159–73.
- (1988), "Consequentialist Foundations for Expected Utility," *Theory and Decision*, 25: 25–78.
- (1989), "Consistent Plans, Consequentialism, and Expected Utility," *Econometrica*, 57: 1445–9.
- Hansen, L.P. and R. Jagannathan (1991), "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99: 225–62.
- Hansen, L.P. and K. Singleton (1982), "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50: 1269–86.
- (1983), "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns," *Journal of Political Economy*, 91: 249–65.
- Harsanyi, J.C. (1973), "Games with Randomly Distributed Payoffs: A New Rationale for Mixed Strategy Equilibrium Points," *International Journal of Game Theory*, 2: 1–23.

- Hazen, G. (1987), "Does Rolling Back Decision Trees Really Require the Independence Axiom?" *Management Science*, 33: 807–9.
- Heaton, J. (1988), "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications," Working Paper 3245-91, Sloan School, MIT.
- Huang, C.F. and R.H. Litzenberger (1988), *Foundations for Financial Economics*, New York: North-Holland.
- Johnsen, T.H. and J.B. Donaldson (1985), "The Structure of Intertemporal Preferences under Uncertainty and Time Consistent Plans," *Econometrica*, 53: 1451–8.
- Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47: 263–91.
- Kandel, S. and R.F. Stambaugh (1990), "Asset Returns and Intertemporal Preferences," NBER Working Paper 3633.
- Karni, E. (1988), "On the Equivalence Between Descending Bid Auctions and First Price Sealed Bid Auctions," *Theory and Decision*, 25: 211–17.
- (1989), "Generalized Expected Utility Analysis of Multivariate Risk Aversion," *International Economic Review*, 30: 297–305.
- Karni, E. and Z. Safra (1989a), "Dynamic Consistency, Revelations in Auctions and the Structure of Preferences," *Review of Economic Studies*, 56: 421–33.
- (1989b), "Ascending Bid Auctions with Behaviorally Consistent Bidders," *Annals of Operations Research*, 19: 435–46.
- (1990), "Behaviorally Consistent Optimal Stopping Rules," *Journal of Economic Theory*, 51: 391–401.
- Karni, E. and D. Schmeidler (forthcoming), "Utility Theory with Uncertainty," in W. Hildenbrand and H. Sonnenschein (eds), *Handbook of Mathematical Economics*, vol. IV.
- Kihlstrom, R.E. and L.J. Mirman (1974), "Risk Aversion with Many Commodities," *Journal of Economic Theory*, 8: 361–88.
- Knight, F. (1921), *Risk, Uncertainty and Profit*, Boston: Houghton Mifflin.
- Kocherlakota, N. (1987), "State Nonseparability: Theory and Empirical Implications," Ph.D. Thesis, University of Chicago.
- (1990), "Disentangling the Coefficient of Relative Risk Aversion from the Elasticity of Intertemporal Substitution: An Irrelevance Result," *Journal of Finance*, 45: 175–90.
- Koopmans, T.C. (1960), "Stationary Ordinal Utility and Impatience," *Econometrica*, 28: 287–309.
- Kreps, D.M. and E.L. Porteus (1978), "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46: 185–200.
- (1979), "Temporal von Neumann-Morgenstern and Induced Preferences," *Journal of Economic Theory*, 20: 81–109.
- Loomes, G. and R. Sugden (1982), "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *Economic Journal*, 92: 805–24.
- Lucas, R.E. (1978), "Asset Prices in an Exchange Economy," *Econometrica*, 46: 1426–45.
- Ma, C. (forthcoming), "Market Equilibrium with Heterogeneous Agents and Recursive Utility," *Economic Theory*.

- Machina, M.J. (1982a), "A Stronger Characterization of Declining Risk Aversion," *Econometrica*, 50: 1069–79.
- (1982b), "'Expected Utility' Analysis Without the Independence Axiom," *Econometrica*, 50: 277–323.
- (1983a), "The Economic Theory of Individual Behavior Towards Risk: Theory, Evidence and New Directions," Stanford University, IMSSS Technical Report #433.
- (1983b), "Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom," in B. Stigum and F. Wenstop (eds), *Foundations of Utility and Risk Theory with Applications*, Dordrecht, Holland: D. Reidel Publishing Co.
- (1984), "Temporal Risk and the Nature of Induced Preferences," *Journal of Economic Theory*, 33: 199–231.
- (1985), "Stochastic Choice Functions Generated from Deterministic Preferences over Lotteries," *Economic Journal*, 95: 575–94.
- (1987), "Choice Under Uncertainty: Problems Solved and Unsolved," *Journal of Economic Perspectives*, 1: 121–54.
- (1989a), "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty," *Journal of Economic Literature*, 27: 1622–68.
- (1989b), "The Behavior of Risk Sharers," unpublished.
- (1989c), "Comparative Statics and Non-Expected Utility Preferences," *Journal of Economic Theory*, 47: 393–405.
- Machina, M.J. and W. Nielsen (1987), "The Ross Measure of Risk Aversion: Strengthening and Extension," *Econometrica*, 55: 1139–49.
- Mankiw, N.G. and M.D. Shapiro (1986), "Risk and Return: Consumption Beta Versus Market Beta," *Review of Economics and Statistics*, 68: 452–9.
- Marcet, A. and K.J. Singleton (1990), "Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Portfolio Constraints," mimeo, Carnegie Mellon University.
- Mehra, R. and E. Prescott (1985), "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15: 145–61.
- Merton, R.C. (1973), "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41: 867–87.
- Milgrom, P.R. (1989), "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, 3: 3–22.
- Milgrom, P.R. and R.J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50: 1089–22.
- von Neumann, J. and O. Morgenstern (1953), *Theory of Games and Economic Behavior*, 3rd edition, Princeton: Princeton University Press.
- Pollak, R.A. (1968), "Consistent Planning," *Review of Economic Studies*, 35: 201–8.
- Quiggin, J. (1982), "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3: 323–43.
- Raiffa, H. (1968), *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*, Reading, MA: Addison-Wesley.
- Roll, R. (1977), "A Critique of the Asset Pricing Theory's Tests: Part I: On Past and Potential Testability of the Theory," *Journal of Financial Economics*, 4: 129–76.

- Rubinstein, A., Z. Safra, and W. Thomson (1990): "On the Interpretation of the Nash Bargaining Solution and Its Extension to Non-Expected Utility Preferences," Working Paper 24-90, Tel Aviv University.
- Segal, U. (1989), "Axiomatic Representation of Expected Utility with Rank-Dependent Probabilities," *Annals of Operations Research*, 19: 359–73.
- (1990), "Two-Stage Lotteries Without the Reduction Axiom," *Econometrica*, 58: 349–77.
- Segal, U. and A. Spivak (1990), "First-Order versus Second-Order Risk Aversion," *Journal of Economic Theory*, 51: 111–25.
- Selden, L. (1978), "A New Representation of Preference over 'Certain \times Uncertain' Consumption Pairs: The 'Ordinal Certainty Equivalent' Hypothesis," *Econometrica*, 46: 1045–60.
- (1979), "An OCE Analysis of the Effect of Uncertainty on Saving Under Risk Preference Independence," *Review of Economic Studies*, 46: 73–82.
- Selden, L. and I. Stux (1978), "Consumption Trees, OCE Utility and the Consumption/Savings Decision," mimeo, Columbia University.
- Selten, R. (1975), "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4: 25–55.
- Shefrin, H.M. and M. Statman (1984), "Explaining Investor Preference for Cash Dividends," *Journal of Financial Economics*, 13: 253–82.
- (1985): "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence," *Journal of Finance*, 40: 777–92.
- Singleton, K.J. (1990), "Specification and Estimation of Intertemporal Asset Pricing Models," in B. Friedman and F. Hahn (eds.), *Handbook in Monetary Economics*, New York: North-Holland, pp. 583–623.
- Streufert, P.A. (1990a), "Stationary Recursive Utility and Dynamic Programming under the Assumption of Biconvergence," *Review of Economic Studies*, 57: 79–98.
- (1990b), "Stochastic Dynamic Programming Under the Assumption of Biconvergence," mimeo, University of Wisconsin.
- Strotz, R.H. (1956), "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23: 165–80.
- Sugden, R. (1986), "New Developments in the Theory of Choice Under Uncertainty," *Bulletin of Economic Reserves*, 38: 1–24. Reprinted in *Surveys in the Economics of Uncertainty*, in J. Hey and P. Lambert (ed.), Oxford: Basil Blackwell.
- Svensson, L.E.O. (1989), "Portfolio Choice with Non-Expected Utility in Continuous Time," *Economic Letters*, 30: 313–17.
- Uzawa, H. (1968), "Time Preference, the Consumption Function and Optimum Asset Holdings," in J.N. Wolfe (ed.), *Value, Capital and Growth: Papers in Honour of Sir John Hicks*. Chicago: Aldine.
- Wang, S.S. (1990), "The Recoverability of Risk Aversion and Intertemporal Substitution," Ph.D. Thesis, University of Toronto.
- Wang, T. (1991), "Lp-Fréchet Differentiable Preference and Local Utility Analysis," mimeo, University of Toronto; forthcoming in *Journal of Economic Theory*.

- Weber, M. and C. Camerer (1987), "Recent Developments in Modelling Preferences Under Risk," *OR Spektrum*, 9: 129–51.
- Weber, R.J. (1982), "The Allais Paradox, Dutch Auctions and Alpha-Utility Theory," Kellogg School of Management, mimeo, Northwestern University.
- Weil, P. (1989), "The Equity Premium Puzzle and the Risk Free Rate Puzzle," *Journal of Monetary Economics*, 24: 401–22.
- (1990a), "Non-Expected Utility in Macroeconomics," *Quarterly Journal of Economics*, 105: 29–42.
- (1990b), "Equilibrium Asset Prices with Undiversifiable Labor Income Risk," mimeo, Harvard University.
- (1990c), "Precautionary Savings and the Permanent Income Hypothesis," mimeo, Harvard University.
- Yaari, M.E. (1987), "The Dual Theory of Choice Under Risk," *Econometrica*, 55: 95–115.
- (1988), "A Controversial Proposal Concerning Inequality Measurement," *Journal of Economic Theory*, 44: 381–96.
- Zin, S. (1987), "Intertemporal Substitution, Risk and the Time Series Behavior of Consumption and Asset Returns," Ph.D. Thesis, University of Toronto.