Recursive Preferences[*]

David K. Backus,[†] Bryan R. Routledge,[‡] and Stanley E. Zin[§]

Revised: December 5, 2005

Abstract

We summarize the class of recursive preferences. These preferences fit naturally with recursive solution methods and hold the promise of generating new insights into familiar problems. Portfolio choice is used as an example.

---

1 Introduction

Recursive methods have become a standard tool for studying economic behavior in dynamic stochastic environments. In this chapter, we characterize the class of preferences that is the natural complement to this framework, namely *recursive preferences.*

Why model preferences rather than behavior? Preferences plays two critical roles in economic models. First, preferences provide, in principle, an unchanging feature of a model in which agents can be confronted with a wide range of different environments, institutions, or policies. For each environment, we derive behavior (decision rules) from the same preferences. If we modeled behavior directly, we would also have to model how it adjusted to changing circumstances. The second role played by preferences is to allow us to evaluate the welfare effects of changing policies or circumstances. Without the ranking of opportunities that a model of preferences provides, it's not clear how we should distinguish good policies from bad.

Why recursive preferences? Recursive preferences focus on the tradeoff between current-period utility and the utility to be derived from all future periods. Since an agent's actions today can affect the evolution of opportunities in the future, summarizing the future consequences of these actions with a single index, *i.e.*, future utility, allows multi-period decision problems to be reduced to a series of two-period problems, and in the case of a stationary infinite-horizon problem, a single, time-invariant two-period decision problem. As we will see, this logic applies equally well to environments in which current actions affect the values of random events for all future periods. In this case, the two-period tradeoff is between current utility and a *certainty equivalent* of random future utility. This recursive approach not only allows complicated dynamic optimization problems to be characterized as much simpler and more intuitive two-period problems, it also lends itself to straightforward computational methods. Since many computational algorithms for solving stochastic dynamic models themselves rely on recursive methods, numerical versions of recursive utility models can be solved and simulated using standard algorithms.

2 The Stationary Recursive Utility Function

Assume time is discrete, with dates $t = 0, 1, 2, \ldots$. At each $t > 0$, an event z_t is drawn from a finite set $\mathcal{Z}$, following an initial event z_0. The t-period history of events is denoted by $z^t = (z_0, z_1, \ldots, z_t)$ and the set of possible t-histories by $\mathcal{Z}^t$. Environments like this, involving time and uncertainty, are the starting point for much of modern economics. A typical agent in such a setting has preferences over payoffs $c(z^t)$ for each possible history. A general set of preferences might be represented by a utility function $U(\{c(z_t)\})$. In what follows, we will think of consumption as a scalar. This is purely for exposition since the extension to a vector of consumption at each point in time is straightforward.

Consider the structure of preferences in this dynamic stochastic environment. We define the class of stationary recursive preferences by

$$U_t = V[c_t, \mu_t(U_{t+1})], \tag{1}$$

where U_t is short-hand for utility starting at some date-t history z^t, U_{t+1} refers to utilities for histories $z^{t+1} = (z^t, z_{t+1})$ stemming from z^t, V is a time aggregator, and μ_t is a certainty-equivalent function based on the conditional probabilities $p(z_{t+1}|z^t)$. As with other utility functions, increasing functions of U, with suitable adjustment of μ, imply the same preferences. This structure of preferences leads naturally to recursive solutions of economic problems, with (1) providing the core of a Bellman equation.

In general, the properties of U_t depend on both the properties of the time aggregator and the certainty equivalent. Since the certainty equivalent will be scaled such that $\mu(x) = x$ when x is a perfect certainty, the time aggregator V is all that matters in deterministic settings. Similarly, for a purely static problem with uncertainty, the certainty-equivalent function μ is all that matters. We consider the specification of each of these components in turn.

It is important to note that the utility functions presented in this chapter are not *ad hoc*, but rather have clear axiomatic foundations, and can be derived from more primitive assumptions on preference orderings. Since utility functions are the typical starting point for applied research, we skip this step and refer the interested reader to the axiomatic characterizations of recursive preferences in the papers cited at the end of this chapter.

3   The Time Aggregator

Time preference is a natural starting point. Suppose there is no risk and c_t is one-dimensional. Preferences might then be characterized by a general utility function $U(\{c_t\})$. A common measure of time preference in this setting is the marginal rate of substitution between consumption at two consecutive dates (c_t and c_{t+1}, say) along a constant consumption path ($c_t = c$ for all t). If the marginal rate of substitution is

$$\text{MRS}_{t,t+1} = \frac{\partial U/\partial c_{t+1}}{\partial U/\partial c_t},$$

then time preference is captured by the discount factor

$$\beta(c) \equiv \text{MRS}_{t,t+1}(c).$$

(Picture the slope, $-1/\beta$, of an indifference curve along along the "45-degree line"). If $\beta(c)$ is less than one, the agent is said to be impatient: along a constant consumption path (*i.e.*, in the absence of diminishing marginal utility considerations), the agent requires more than one unit of consumption at $t+1$ to induce a sacrifice of one unit at t.

For the traditional time-additive utility function,

$$U(\{c_t\}) = \sum_{t=0}^{\infty} \beta^t u(c_t), \tag{2}$$

$\beta(c) = \beta < 1$ regardless of the value of c, so impatience is built in and constant. A popular and useful special case of this utility function implies a constant elasticity of intertemporal

substitution by assuming $u(c) = c^\rho/\rho$ for $\rho < 1$. Note that we can define the utility function in (2) recursively:

$$U_t = u(c_t) + \beta U_{t+1}, \tag{3}$$

for $t = 1, 2, \ldots$. The constant elasticity version can be expressed

$$U_t = [(1-\beta)c_t^\rho + \beta U_{t+1}^\rho]^{1/\rho}, \tag{4}$$

where $\rho < 1$ and $\sigma = 1/(1-\rho)$ is the intertemporal elasticity of substitution. (To put this in additive form, use the transformation $\widehat{U} = U^\rho/\rho$.) Note that U_t is homothetic and that the scaling we have chosen measures utility on the same scale as consumption:

$$U(c, c, c, \ldots) = c.$$

More generally, impatience summarized by the discount factor, $\beta(c)$, could vary with the level of consumption. Koopmans (1960) derives a class of stationary recursive preferences by imposing conditions on a general utility function U for a multi-dimensional consumption vector c. In the Koopmans class of preferences, time preference is a property of the time aggregator V. Consider our measure of time preference:

$$U_t = V(c_t, U_{t+1}) = V[c_t, V(c_{t+1}, U_{t+2})].$$

The marginal rate of substitution between c_t and c_{t+1} is therefore

$$\text{MRS}_{t,t+1} = \frac{V_2(c_t, U_{t+1})V_1(c_{t+1}, U_{t+2})}{V_1(c_t, U_{t+1})}.$$

A constant consumption path at c is defined by $U = V(c, U)$, implying $U = g(c) = V[c, g(c)]$ for some function g.

In modern applications, we typically work in reverse order: we specify a time aggregator V and use it to characterize the overall utility function U. Any U constructed this way defines preferences that are stationary and dynamically consistent. In contrast to time-additive preferences, discounting depends on the level of consumption c.

There most common example of Koopmans' structure in applications is a generalization of equation (3):

$$V(c, U) = u(c) + \beta(c)U,$$

where there is no particular relationship between the functions u and β. For this example, the intertemporal tradeoff is given by

$$\text{MRS}_{t,t+1} = \beta(c_t)\left[\frac{u'(c_{t+1}) + \beta'(c_{t+1})U_{t+2}}{u'(c_t) + \beta'(c_t)U_{t+1}}\right].$$

When $\beta'(c) \neq 0$, optimal consumption plans will depend on the level of future utility. And along a constant consumption path, discounting is decreasing (increasing) in consumption when $\beta'(c) < 0$ ($\beta'(c) > 0$). Also note that U_t in this example is not homothetic.

4   The Risk Aggregator

Turn now to the specification of risk preferences, which we consider initially in a static setting. Choices have risky consequences or payoffs, and agents have preferences defined over those consequences and their probabilities. To be specific, let us say that the state z is drawn with probability $p(z)$ from the finite set $\mathcal{Z} = \{1, 2, \ldots, Z\}$. Consequences ($c$, say) depend on the state and the agent's preferences are be represented by a utility function of state-contingent consequences ("consumption"):

$$U(\{c(z)\}) \;=\; U[c(1), c(2), \ldots, c(Z)].$$

At this level of generality there is no mention of probabilities, although we can well imagine that the probabilities of the various states will show up somehow in U. We regard the probabilities as known, which you might think of as an assumption "rational expectations."

We prefer to work with a different (but equivalent) representation of preferences. Suppose, for the time being, that c is a scalar; very little of the theory depends on this, but it streamlines the presentation. We define the *certainty equivalent* of a set of consequences as a certain consequence μ that gives the same level of utility:

$$U(\mu, \mu, \ldots, \mu) \;=\; U[c(1), c(2), \ldots, c(Z)].$$

If U is increasing in all its arguments, we can solve this for the certainty-equivalent function $\mu(\{c(z)\})$. Clearly μ represents the same preferences as U, but we find its form particularly useful. For one thing, it expresses utility in payoff ("consumption") units. For another, it summarizes behavior toward risk directly: since the certainty equivalent of a sure thing is itself, the impact of risk is simply the difference between the certainty equivalent and expected consumption.

The traditional approach to preferences in this setting is expected utility, which takes the form

$$U(\{c(z)\}) \;=\; \sum_z p(s)u[c(z)] \;=\; Eu(c),$$

or

$$\mu(\{c(z)\}) \;=\; u^{-1}\left(\sum_z p(z)u[c(z)]\right) \;=\; u^{-1}\left[Eu(c)\right].$$

Preferences of this form have been used in virtually all economic theory. The utility function of Kreps and Porteus employs a general time aggregator and an expected utility certainty equivalent. Following Epstein and Zin, many recent applications, particularly in dynamic asset pricing models, use the homothetic version of this utility function which combines the constant elasticity time aggregator in (4) with a linear homogeneous (constant relative risk aversion) expected utility certainty equivalent.

Empirical research both in the laboratory and in the field has documented a variety of difficulties with the predictions of expected utility models. In particular, people seem more averse to bad outcomes than implied by expected utility. In response to this evidence, there is a growing body of work that looks at decision making under uncertainty outside

of the traditional expected utility framework. For example, a popular recent innovation is Gilboa and Schmeidler's "max-min" preferences, which have been extended to recursive dynamic problems in a number of different ways. Without surveying all of these extensions, we demonstrate the basic mechanics of recursive utility with non-expected utility certainty equivalents by studying one particular analytically convenient class of preferences, the Chew-Dekel class, in detail.

The Chew-Dekel certainty equivalent function μ for a set of payoffs and probabilities $\{c(z), p(z)\}$ is defined implicitly by a *risk aggregator* M satisfying

$$\mu = \sum_z p(z) M[c(z), \mu]. \tag{5}$$

Such preferences satisfy a weaker condition than the notorious independence axiom that underlies expected utility, yet like expected utility, they lead to first-order conditions in decision problems that are linear in probabilities, hence easily solved and amenable to econometric analysis. We assume M has the following properties: (i) $M(m, m) = m$ (sure things are their own certainty equivalents), (ii) M is increasing in its first argument (first-order stochastic dominance), (iii) M is concave in its first argument (risk aversion), and (iv) $M(kc, km) = kM(c, m)$ for $k > 0$ (linear homogeneity). Most of the analytical convenience of the Chew-Dekel class follows from the linearity of equation (5) in probabilities. (Note that this implies that indifference curves on the probability simplex are linear, but not necessarily parallel.)

Examples of tractable members of the Chew-Dekel class are:

- *Expected utility.* A version with constant relative risk aversion (*i.e.*, linear homogeneity) is implied by
  $$M(c, m) = c^\alpha m^{1-\alpha}/\alpha + m(1 - 1/\alpha).$$

  If $\alpha \leq 1$, M satisfies the conditions outlined above. Applying (5), we find

  $$\mu = \left(\sum_z p(z) c(z)^\alpha \right)^{1/\alpha},$$

  the usual expected utility with a power utility function.

- *Weighted utility.* A relatively easy way to generalize expected utility given (5): weight the probabilities by a function of outcomes. A constant-elasticity version follows from

  $$M(c, m) = (c/m)^\gamma c^\alpha m^{1-\alpha}/\alpha + m[1 - (c/m)^\gamma/\alpha].$$

  For M to be increasing and concave in c in a neighborhood of m, the parameters must satisfy either (a) $0 < \gamma < 1$ and $\alpha + \gamma < 0$ or (b) $\gamma < 0$ and $0 < \alpha + \gamma < 1$. Note that (a) implies $\alpha < 0$, (b) implies $\alpha > 0$, and both imply $\alpha + 2\gamma < 1$. The associated certainty equivalent function is

  $$\mu^\alpha = \frac{\sum_z p(z) c(z)^{\gamma+\alpha}}{\sum_x p(x) c(x)^\gamma} = \sum_z \hat{p}(z) c(z)^\alpha,$$

where
$$\hat{p}(z) = \frac{p(z)c(z)^\gamma}{\sum_x p(x)c(x)^\gamma}.$$

This version highlights the impact of bad outcomes: they get greater weight than with expected utility if $\gamma < 0$, less weight otherwise.

- *Disappointment aversion.* Another model that increases sensitivity to bad events ("disappointments") is defined by the risk aggregator

$$M(c,m) \;=\; \begin{cases} c^\alpha m^{1-\alpha}/\alpha + m(1-1/\alpha) & c \geq m \\ c^\alpha m^{1-\alpha}/\alpha + m(1-1/\alpha) + \delta(c^\alpha m^{1-\alpha}-m)/\alpha & c < m \end{cases}$$

with $\delta \geq 0$. When $\delta = 0$ this reduces to expected utility. Otherwise, disappointment aversion places additional weight on outcomes worse than the certainty equivalent. The certainty equivalent function satisfies

$$\mu^\alpha \;=\; \sum_z p(z)c(z)^\alpha + \delta \sum_z p(z)I[c(z) < \mu][c(z)^\alpha - \mu^\alpha] \;=\; \sum_z \hat{p}(z)c(z)^\alpha,$$

where $I(x)$ is an indicator function that equals one if x is true and zero otherwise and

$$\hat{p}(z) \;=\; \left(\frac{1 + \delta I[c(z) < \mu]}{1 + \delta \sum_x p(x)I[c(x) < \mu]} \right) p(z).$$

It differs from weighted utility in scaling up the probabilities of all bad events by the same factor, and scaling down the probabilities of good events by a complementary factor, with good and bad defined as better and worse than the certainty equivalent. (This implies a "kink" in state-space indifference curves at certainty, which is referred to as "first-order" risk aversion.) All three expressions highlight the *recursive* nature of the risk aggregator M: we need to know the certainty equivalent to know which states are bad so that we can compute the certainty equivalent (and so on).

5 Optimization and the Bellman Equation

For an illustrative application of recursive utility, we turn to the classic Merton-Samuelson consumption/portfolio-choice problem. Consider a stationary Markov environment with states z and conditional probabilities $p(z'|z)$. Preferences are represented by a constant-discounting/constant-elasticity aggregator and a general linear homogeneous certainty equivalent. A dynamic consumption/portfolio problem for this environment is characterized by the Bellman equation which implicitly defines the value function:

$$J(a,z) \;=\; \max_{c,w} \{(1-\beta)c^\rho + \beta\mu[J(a',z')]^\rho\}^{1/\rho},$$

subject to the wealth constraint, $a' = (a-c)\sum_i w_i r_i(z,z') = (a-c)\sum_i w_i r_i' = (a-c)r_p'$, where a denotes wealth, r_p is the return on the portfolio $(w_1, w_2, \ldots, w_{N-1}, 1 - \sum_{i=1}^{N-1} w_i)$, of assets with risky returns $(r_1, r_2, \ldots, r_N)$. The budget constraint and linear homogeneity

of the time and risk aggregators imply linear homogeneity of the value function: $J(a, z) = aL(z)$ for some scaled value function L. The scaled Bellman equation is

$$L(z) = \max_{b,w} \{(1-\beta)b^\rho + \beta(1-b)^\rho \mu[L(z')r_p(z, z')]^\rho\}^{1/\rho},$$

where $b \equiv c/a$. Note that $L(z)$ is the marginal utility of wealth in state z.

This problem divides into separate portfolio and consumption decisions. The portfolio decision solves: choose $\{w_i\}$ to maximize $\mu[L(z')r_p(z, z')]$. The portfolio first-order conditions are

$$\sum_{z'} p(z'|z)M_1[L(z')r_p(z, z'), \mu]L(z')[r_i(z, z') - r_j(z, z')] = 0 \qquad (6)$$

for any two assets i and j.

Given a maximized μ, the consumption decision solves: choose b to maximize L. The intertemporal first-order condition is

$$(1-\beta)b^{\rho-1} = \beta(1-b)^{\rho-1}\mu^\rho. \qquad (7)$$

If we solve for μ and substitute into the (scaled) Bellman equation, we find

$$\begin{aligned} \mu &= [(1-\beta)/\beta]^{1/\rho}[b/(1-b)]^{(\rho-1)/\rho} \\ L &= (1-\beta)^{1/\rho}b^{(\rho-1)/\rho}. \end{aligned} \qquad (8)$$

The first-order condition (7) and value function (8) allow us to express the relation between consumption and returns in a familiar form. Since μ is linear homogeneous, the first-order condition implies $\mu(x'r'_p) = 1$ for

$$x' = L'/\mu = \left[\beta(c'/c)^{\rho-1}(r'_p)^{1-\rho}\right]^{1/\rho}.$$

The last equality follows from $(c'/c) = (b'/b)(1-b)r'_p$, a consequence of the budget constraint and the definition of b. The intertemporal first-order condition can therefore be expressed

$$\mu(x'r'_p) = \mu\left(\left[\beta(c'/c)^{\rho-1}r'_p\right]^{1/\rho}\right) = 1, \qquad (9)$$

a generalization of the tangency condition for an optimum (set the marginal rate of substitution equal to the price ratio). Similar logic leads us to express the portfolio first-order conditions (6) as

$$E\left[M_1(x'r'_p, 1)x'(r'_i - r'_j)\right] = 0.$$

If we multiply by the portfolio weight w_j and sum over j we find

$$E\left[M_1(x'r'_p, 1)x'r'_i\right] = E\left[M_1(x'r'_p, 1)x'r'_p\right]. \qquad (10)$$

Euler's theorem for homogeneous functions allows us to express the right side as

$$E\left[M_1(x'r'_p, 1)x'r'_p\right] = 1 - EM_2(x'r'_p, 1).$$

Whether this expression is helpful depends on the precise form of M. For example, with disappointment aversion, (10) is

$$E\left[z^{\alpha-1}(1+\delta I[z<1])\frac{r_i'}{r_p'}\right] = 1 + \delta E I[z<1],$$

where $z = \left[\beta(c'/c)^{\rho-1}r_p'\right]^{1/\rho}$. This reduces to the Kreps-Porteus model when $\delta = 0$, and to the time-additive expected utility model when, in addition, $\rho = \alpha$.

6   Conclusion

A recursive utility function can be constructed from two components: (1) a time aggregator that completely characterizes preferences in the absence of uncertainty and (2) a risk aggregator that defines the certainty equivalent function that characterizes preferences over static gambles and is used to aggregate the risk associated with future utility. We looked at natural candidates for each of these components and gave an example of how Bellman's equation can be used to characterize optimal plans in a dynamic stochastic environment when agents have recursive preferences.

Literature Guide: If you'd like to see more on this subject, see Backus, Routledge, and Zin (2005) and the references cited there. Much of the material in this chapter builds from Epstein and Zin (1989), who extend the preferences in Kreps and Porteus (1978) to allow for a stationary infinite-horizon model and for non-expected utility certainty equivalents. They also derive the consumption/portfolio-choice results of Section (5). For more on time aggregators, see Koopmans (1960), Uzawa (1968), Epstein and Hynes (1983), Lucas and Stokey (1984), and Shi (1994). Common departures from expected utility are documented in Kreps (1988, ch 14) and Starmer (2000). Epstein and Schneider (2003) and Hansen and Sargent (2004) propose different dynamic and recursive extensions of the max-min risk preference of Gilboa and Schmeidler (1993). The Chew-Dekel risk aggregator was proposed by Chew (1983, 1989) and Dekel (1986). Examples within this class: weighted utility (Chew, 1983), disappointment aversion (Gul, 1991), semi-weighted utility (Epstein and Zin, 2001), and generalized disappointment aversion (Routledge and Zin, 2003).

References

Backus, David K., Bryan R. Routledge, and Stanley E. Zin, 2004, "Exotic preferences for macroeconomists," *NBER Macroeconomics Annual*.

Chew, Soo Hong, 1983, "A generalization of the quasi-linear mean with applications to the measurement of inequality and decision theory resolving the Allais paradox," *Econometrica* 51,1065-1092.

Chew, Soo Hong, 1989, "Axiomatic utility theories with the betweenness property," *Annals of Operations Research* 19, 273-298.

Dekel, Eddie, 1986, "An axiomatic characterization of preferences uncer uncertainty: Weakening the independence axiom," *Journal of Economic Theory* 40, 304-318.

Epstein, Larry G., and J. Allan Hynes, 1983, "The rate of time preference and dynamic economic analysis," *Journal of Political Economy* 91, 611-635.

Epstein, Larry G., and Martin Schneider, 2003, "Recursive multiple-priors," *Journal of Economic Theory* 113, 1-31.

Epstein, Larry G., and Stanley E. Zin, 1989, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework," *Econometrica* 57, 937-969.

Epstein, Larry G., and Stanley E. Zin, 2001, "The independence axiom and asset returns," *Journal of Empirical Finance* 8, 537-572.

Gilboa, Itzhak, and David Schmeidler, 1993, "Updating ambiguous beliefs," *Journal of Economic Theory* 59, 33-49.

Gul, Faruk, 1991, "A theory of disappointment aversion," *Econometrica* 59, 667-686.

Hansen, Lars Peter, and Thomas J. Sargent, 2004, *Misspecification in Recursive Macroeconomic Theory* manuscript, January.

Koopmans, Tjalling C., 1960, "Stationary ordinal utility and impatience," *Econometrica* 28, 287-309.

Kreps, David M., 1988, *Notes on the Theory of Choice*, Boulder, CO: Westview Press.

Kreps, David M., and Evan L. Porteus, 1978, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica* 46, 185-200.

Lucas, Robert E., and Nancy L. Stokey, 1984, "Optimal growth with many consumers," *Journal of Economic Theory* 32, 139-171.

Routledge, Bryan R., and Stanley E. Zin, 2003, "Generalized disappointment aversion and asset prices," NBER Working Paper 10107, November.

Shi, Shouyong, 1994, "Weakly nonseparable preferences and distortionary taxes in a small open economy," *International Economic Review* 35, 411-428.

Starmer, Chris, 2000, "Developments in non-expected utility theory," *Journal of Economic Literature* 38, 332-382.

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 IL: Aldine Publishing Company.