

NBER WORKING PAPER SERIES

EXOTIC PREFERENCES FOR MACROECONOMISTS

David Backus
Bryan Routledge
Stanley Zin

Working Paper 10597
<http://www.nber.org/papers/w10597>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2004

Prepared for the 2004 NBER Macroeconomics Annual. We are unusually grateful to the many people who made suggestions and answered questions while we worked through this enormous body of work. Deserving special mention are Mark Gertler, Lars Hansen, Tom Sargent, Rob Shimer, Tony Smith, Ivan Werning, Noah Williams, and especially Martin Schneider. We also thank Kfir Eliaz, Larry Epstein, Anjela Kniazeva, Per Krusell, David Laibson, John Leahy, Tom Tallarini, Stijn Van Nieuwerburgh, and seminar participants at CMU and NYU. The views expressed herein are those of the author(s) and not necessarily those of the National Bureau of Economic Research.

©2004 by David Backus, Bryan Routledge, and Stanley Zin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Exotic Preferences for Macroeconomists
David Backus, Bryan Routledge, and Stanley Zin
NBER Working Paper No. 10597
June 2004
JEL No. D81, D91, E1, G12

ABSTRACT

We provide a user's guide to “exotic” preferences: nonlinear time aggregators, departures from expected utility, preferences over time with known and unknown probabilities, risk-sensitive and robust control, “hyperbolic” discounting, and preferences over sets (“temptations”). We apply each to a number of classic problems in macroeconomics and finance, including consumption and saving, portfolio choice, asset pricing, and Pareto optimal allocations.

David Backus
Stern School of Business
NYU
44 West 4th Street
New York, NY 10012-1126
and NBER
dbackus@stern.nyu.edu

Bryan Routledge
Carnegie Mellon University
Pittsburgh, PA 15213-3890
routledge@cmu.edu

Stanley Zin
Graduate School of
Industrial Administration
Carnegie Mellon University
Pittsburgh, PA 15213-3890
and NBER
zin@cmu.edu

1 Introduction

Applied economists (including ourselves) are generally content to study theoretical agents whose preferences are additive over time and across states of nature. One version goes like this: Time is discrete, with dates $t = 0, 1, 2, \dots$. 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 $z^t = (z_0, z_1, \dots, z_t)$ and the set of possible t -histories by \mathcal{Z}^t . The evolution of events and histories is conveniently illustrated by an event tree, as in Figure 1, with each branch representing an event and each node a history or state. Environments like this, involving time and uncertainty, are the starting point for most of modern macroeconomics and finance. 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 $U(\{c(z_t)\})$. More common, however, is to impose the additive expected utility structure

$$U(\{c(z_t)\}) = \sum_{t=0}^{\infty} \beta^t \sum_{z^t \in \mathcal{Z}^t} p(z^t) u[c(z^t)] = E_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \quad (1)$$

where $0 < \beta < 1$, $p(z^t)$ is the probability of history z^t , and u is a period/state utility function. These preferences are remarkably parsimonious: behavior over time and across states depends solely on the discount factor β , the probabilities p , and the function u .

Although (1) remains the norm throughout economics, there has been extraordinary theoretical progress over the last fifty years (and particularly the last twenty five) in developing alternatives. Some of these alternatives were developed to account for the anomalous predictions of expected utility in experimental work. Others arose from advances in the pure theory of intertemporal choice. Whatever their origin, they offer greater flexibility along several dimensions, often with only a modest increase in analytical difficulty.

What follows is a user's guide, intended to serve as an introduction and instruction manual for economists studying problems in which the structure of preferences may play an important role. Our goal is to describe exotic preferences to mainstream economists: preferences over time, preferences across states or histories, and (especially) combinations of the two. We take an overtly practical approach, downplaying or ignoring altogether the many technical issues that arise in specifying preferences in dynamic stochastic settings, including their axiomatic foundations. (References are provided in Appendix A for those who are interested.) We generally assume without comment that preferences can be represented by increasing, (weakly) concave functions, with enough smoothness and boundary conditions to generate interior solutions to optimizations. We focus instead on applications, using tractable functional forms to revisit some classic problems: consumption and saving, portfolio choice, asset pricing, and Pareto optimal allocations. In most cases, we use utility functions that are homogeneous of degree one (hence invariant to scale) with constant elasticities (think power utility). These functions are the workhorses of macroeconomics and finance, so little

is lost by restricting ourselves in this way.

You might well ask: Why bother? Indeed, we will not be surprised if most economists continue to use (1) most of the time. Exotic preferences, however, have a number of potential advantages that we believe will lead to much wider application than we've seen to date. One is more flexible functional forms for approximating features of data — the equity premium, for example. Another is the ability to ask questions that have no counterpart in the additive model. How should we make decisions if we don't know the probability model that generates the data? Can preferences be dynamically inconsistent? If they are, how do we make decisions? What is the appropriate welfare criterion? Can we think of some choices as tempting us away from better ones? Each of these advantages raises further questions: Are exotic preferences observationally equivalent to additive preferences? If not, how do we identify their parameters? Are they an excuse for free parameters? Do we even care whether behavior is derived from preferences?

These questions run through a series of non-additive preference models. In Section 2, we discuss time preference in a deterministic setting, comparing Koopmans' time aggregator to the traditional time-additive structure. In Section 3, we describe alternatives to expected utility in a static setting, using a certainty-equivalent function to summarize preference toward risk. We argue that the Chew-Dekel class extends expected utility in useful directions without sacrificing analytical and empirical convenience. In Section 4, we put time and risk preference together in a Kreps-Porteus aggregator, which leads to a useful separation between time and risk preference. Dynamic extensions of Chew-Dekel preferences follow the well-worn path of Epstein and Zin. In Section 5, we consider risk-sensitive and robust control, whose application to economics is associated with the work of Hansen and Sargent. Section 6 is devoted to ambiguity, in which agents face uncertainty over probabilities as well as states. We describe Gilboa and Schmeidler's "max-min" utility for static settings and Epstein and Schneider's recursive extension to dynamic settings. In Section 7, we turn to "hyperbolic discounting" and provide an interpretation based on Gul and Pesendorfer's "temptation" preferences. The final section is devoted to a broader discussion of the role and value of exotic preferences in economics.

A word on notation and terminology: We typically denote parameters by Greek letters and functions and variables by Latin letters. We denote derivatives with subscripts; thus V_2 refers to the derivative of V with respect to its second argument. In a stationary dynamic programming problem, J is a value function and a prime ($'$) distinguishes a future value from a current value. The abbreviation "iid" means independent and identically distributed and $NID(x, y)$ means normally and independently distributed with mean x and variance y .

2 Time

Time preference is a natural starting point for macroeconomists, since so much of our subject is concerned with dynamics. Suppose there is no risk and (for this paragraph only) 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 the “45-degree line.”) If $\beta(c)$ is less than one, the agent is said to be impatient: she requires more than one unit of consumption at $t + 1$ to induce her to give up 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. The rest of this section is concerned with preferences in which the discount factor can vary with the level of consumption.

Koopmans’ time aggregator

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 . Our approach and terminology follow Johnsen and Donaldson (1985). Preferences at all dates come from the same “date-zero” utility function U . As a result, they are dynamically consistent by construction: preferences over consumption streams starting at any future date t are consistent with U . Following Koopmans, let ${}_t c \equiv (c_t, c_{t+1}, \dots)$ be an infinite consumption sequence starting at t . Then we might write utility from date $t = 0$ on as

$$U({}_0 c) = U(c_0, {}_1 c).$$

Koopmans’ first condition is *history-independence*: preferences over sequences ${}_t c$ do not depend on consumption at dates prior to t . Without this condition, an agent making sequential decisions would need to keep track of the history of consumption choices to be able to make

future choices consistent with U . The marginal rate of substitution between consumption at two arbitrary dates could depend, in general, on consumption at all dates past, present, and future. History-independence rules out dependence on the past. With it, the utility function can be expressed in the form

$$U({}_0c) = V[c_0, U_1({}_1c)]$$

for some *time aggregator* V . As a result, choices over ${}_1c$ do not depend on c_0 . (Note, for example, that marginal rates of substitution between elements of ${}_1c$ do not depend on c_0 .) Koopmans' second condition is *future independence*: preferences over c_t do not depend on ${}_{t+1}c$. (In Koopmans' terminology, the first and second conditions together imply that preferences over the present (c_t) and future (${}_{t+1}c$) are *independent*.) This is trivially true if c_t is a scalar, but a restriction on preferences otherwise. The two conditions together imply that utility can be written

$$U({}_0c) = V[u(c_0), U_1({}_1c)]$$

for some functions V and u , which defines u as a composite commodity for consumption at a specific date. Koopmans' third condition is that preferences are stationary (the same at all dates). The three conditions together imply that utility can be written in the stationary recursive form,

$$U({}_t c) = V[u(c_t), U({}_{t+1}c)] \tag{3}$$

for all dates t . This is a generalization of the traditional utility function (2), where (evidently) $V(u, U) = u + \beta U$ or the equivalent. As in traditional utility theory, preferences are unchanged when we apply a monotonic transformation to U : if $\hat{U} = f(U)$ for f increasing, then we replace the aggregator V with $\hat{V}(u, \hat{U}) = f(V[u, f^{-1}(\hat{U})])$.

In the Koopmans class of preferences represented by (3), time preference is a property of the time aggregator V . Consider our measure of time preference for the composite commodity u . If U_t and u_t represent $U({}_t c)$ and $u(c_t)$, respectively, then

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

The marginal rate of substitution between u_t and u_{t+1} is therefore

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

A constant consumption path with period utility u is defined by $U = V(u, U)$, implying $U = g(u) = V[u, g(u)]$ for some function g . (Koopmans calls g the *correspondence function*.) The discount factor is therefore $\beta(u) = V_2[u, g(u)]$. You might verify for yourself that V_2 is invariant to increasing transformations of U .

In modern applications, we generally work in reverse order: we specify a period utility function u and a time aggregator V and use them to characterize the overall utility function

U . Any U constructed this way defines preferences that are dynamically consistent, history independent, future independent, and stationary. In contrast to time-additive preferences (2), discounting depends on the level of utility u . To get a sense of how this works, consider the behavior of V_2 . If preferences are increasing in consumption, u must be increasing in c and V must be increasing in both arguments. If we consider sequences with constant consumption, U must be increasing in u , so that

$$g_1(u) = V_1[u, g(u)] + V_2[u, g(u)]g_1(u) = \frac{V_1[u, g(u)]}{1 - V_2[u, g(u)]} > 0.$$

Since $V_1 > 0$, $0 < V_2[u, g(u)] < 1$: the discount factor is between zero and one and depends (in general) on u . Many economists impose an additional condition of *increasing marginal impatience*: $V_2[u, g(u)]$ is decreasing in u , or

$$V_{21}[u, g(u)] + V_{22}[u, g(u)]g_1(u) = V_{21}[u, g(u)] + V_{22}[u, g(u)]\frac{V_1[u, g(u)]}{1 - V_2[u, g(u)]} < 0.$$

In applications, this condition is typically used to generate stability of steady states.

Two variants of Koopmans' structure have been widely used by macroeconomists. One was proposed by Uzawa (1968), who suggested a continuous-time version of

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

(In his model, $\beta(u) = \exp[-\delta(u)]$.) Since $V_{21} = 0$, increasing marginal impatience is simply $\beta_1(u) < 0$ (equivalently, $\delta_1(u) > 0$). Another is used by Epstein and Hynes (1983), Lucas and Stokey (1984), and Shi (1993), who generalize Koopmans by omitting the future independence condition. The resulting aggregator is $V(c, U)$, rather than $V(u, U)$, which allows choice over c to depend on U . If c is a scalar, this is equivalent to (3) (set $u(c) = c$), but otherwise need not be. An example is

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

where there is no particular relationship between the functions u and β .

Examples

Example 1 (growth and fiscal policy). In the traditional growth model, Koopmans preferences can change both the steady state and the short-run dynamics. Suppose the period utility function is $u(c)$ and the time aggregator is $V(u, U') = u + \beta(u)U'$, with u increasing and concave and $\beta_1(u) < 0$. Gross output y is produced with capital k using an increasing concave technology f . The resource constraint is $y = f(k) = c + k' + g$, where c is consumption, k' is

tomorrow's capital stock, and g is government purchases (constant). The Bellman equation is

$$J(k) = \max_{k'} u[f(k) - k' - g] + \beta(u[f(k) - k' - g])J(k').$$

The first-order and envelope conditions are

$$\begin{aligned} u_1(c)\{1 + \beta_1[u(c)]J(k')\} &= \beta[u(c)]J_1(k') \\ J_1(k) &= u_1(c)f_1(k)\{1 + \beta_1[u(c)]J(k')\}, \end{aligned}$$

which together imply $J_1(k) = \beta[u(c)]J_1(k')f_1(k)$. In a steady state, $1 = \beta(u[f(k) - k - g])f_1(k)$.

One clear difference from the traditional model is the role of preferences in determining the steady state. With constant β , the steady state capital stock solves $\beta f_1(k) = 1$; u is irrelevant. With recursive preferences, the steady state solves $\beta(u[f(k) - k - g])f_1(k) = 1$, which depends on u through its impact on β . Consider the impact of an increase in g . With traditional preferences, the steady state capital stock doesn't change, so any increase in g is balanced by an equal decrease in c . With recursive preferences and increasing marginal impatience, an increase in g reduces current utility and therefore raises the discount factor. The initial drop in c is therefore larger than in the traditional case. In the resulting steady state, the increase in g leads to an increase in k and a decline in c that is smaller than the increase in g . The magnitude of the decline depends on β_1 , the sensitivity of the discount factor to current utility. [Adapted from Dolmas and Wynne (1998).]

Example 2 (optimal allocations). Time preference affects the optimal allocation of consumption among agents over time. Consider an economy with a constant aggregate endowment y of a single good, to be divided between two agents with Koopmans preferences, represented here by the aggregators V (the first agent) and W (the second). A Pareto optimal allocation is summarized by the Bellman equation

$$J(w) = \max_{c, w'} V[y - c, J(w')]$$

subject to

$$W(c, w') \geq w.$$

Note that both consumption c and promised utility w pertain to the second agent. If λ is the Lagrange multiplier on the constraint, the first-order and envelope conditions are

$$\begin{aligned} V_1[y - c, J(w')] &= \lambda W_1(c, w') \\ V_2[y - c, J(w')]J_1(w') + \lambda W_2(c, w') &= 0 \\ J_1(w) &= -\lambda. \end{aligned}$$

If agents' preferences are additive with the same discount factor β , then the second and third equations imply $J_1(w')/J_1(w) = W_2(c, w')/V_2[y - c, J(w')] = \beta/\beta = 1$: an optimal

allocation places the same weight $\lambda = -J_1(w)$ on the second agent's utility at all dates and promised utility w is constant. If preferences are additive and $\beta_2 > \beta_1$ (the second agent is more patient), then $J_1(w')/J_1(w) = \beta_2/\beta_1 > 1$: an optimal allocation increases the weight over time on the second, more patient agent and raises her promised utility ($w' > w$). In the more general Koopmans setting, the dynamics depend on the time aggregators V and W . The allocation converges to a steady state if both aggregators exhibit increasing marginal impatience and future utility is a normal good. [Adapted from Lucas and Stokey (1984).]

Example 3 (long-run properties of a small open economy). Small open economies with perfect capital mobility raise difficulties with the existence of a steady state that can be resolved by endogenizing the discount factor. We represent preferences over sequences of consumption c and leisure $1 - n$ with a period utility function $u(c, 1 - n)$ and a time aggregator $V(c, 1 - n, U) = u(c, 1 - n) + \beta(c, 1 - n)U$. Let output be produced with labor using the linear technology $y = \theta n$, where θ is a productivity parameter. The economy's resource constraint is $y = c + x$, where x is net exports. The agent can borrow and lend in international capital markets at gross interest rate r , giving rise to the budget constraint $a' = r(a + x) = r(a + \theta n - c)$. The Bellman equation is

$$J(a) = \max_{c,n} u(c, 1 - n) + \beta(c, 1 - n)J[r(a + \theta n - c)].$$

The first-order and envelope conditions are:

$$\begin{aligned} u_1 + \beta_1 J(a') &= \beta J_1(a') \\ u_2 + \beta_2 J(a') &= \beta J_1(a')\theta \\ J_1(a) &= \beta J_1(a')r. \end{aligned}$$

The last equation tells us that in a steady state, $\beta(c, 1 - n)r = 1$. With constant discounting, there is no steady state, but with more general discounting schemes the form of discounting determines the steady state and its response to changes in the environment. Here the long-run impact of a change in (say) θ (the "wage") depends on the form of β . Suppose β is a function of n only. Then the steady state condition $\beta(1 - n)r = 1$ determines n independently of θ ! More generally, the long-run impact on n of a change in θ depends on the form of the discount function $\beta(c, 1 - n)$. [Adapted from Epstein and Hynes (1983), Mendoza (1991), Obstfeld (1981), Schmitt-Grohe and Uribe (2002), and Shi (1994).]

Example 4 (dynamically inconsistent preferences). Suppose preferences "from date t on" are given by:

$$U_t(t, c) = u(c_t) + \delta\beta u(c_{t+1}) + \delta\beta^2 u(c_{t+2}) + \delta\beta^3 u(c_{t+3}) + \dots,$$

with $0 < \delta \leq 1$. When $\delta = 1$ this reduces to the time-additive utility function (2). Otherwise, we discount utility in periods $t + 1, t + 2, t + 3, \dots$ by $\delta\beta, \delta\beta^2, \delta\beta^3, \dots$. A little effort should convince you that these preferences cannot be put into stationary recursive form. In fact,

they are dynamically inconsistent in the sense that preferences over (say) (c_{t+1}, c_{t+2}) at date t are different from preferences at $t + 1$. (Note, for example, the marginal rates of substitution between c_{t+1} and c_{t+2} at t and $t + 1$.) This structure is ruled out by Koopmans, who begins with the presumption of a consistent set of preferences. We'll return to this example in Section 7. [Adapted from Harris and Laibson (2003) and Phelps and Pollack (1968).]

3 Risk

Our next topic is risk, which we consider initially in a static setting. Our theoretical agent makes choices that have risky consequences or payoffs and has preferences 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 $\mathcal{Z} = \{1, 2, \dots, Z\}$. Consequences (c , say) depend on the state. Having read Debreu's *Theory of Value* or the like, we might guess that with the appropriate technical conditions the agent's preferences can be represented by a utility function of state-contingent consequences ("consumption"):

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

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 , as they do in (1). In this section, we regard the probabilities as known, which you might think of as an assumption of "risk" or "rational expectations." We consider unknown probabilities ("ambiguity") in Sections 5 and 6.

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:

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

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

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

a special case of (1). Preferences of this form are used in virtually all macroeconomic theory, but decades of experimental research have documented numerous difficulties with it. Among them: people seem more averse to bad outcomes than expected utility implies. See, for example, the summaries in Kreps (1988, ch 14) and Starmer (2000). We suggest the broader Chew-Dekel class of preferences, which allows us to account for some of the empirical anomalies of expected utility without giving up its analytical tractability.

The Chew-Dekel risk aggregator

Chew (1983, 1989) and Dekel (1986) derive a class of risk preferences that generalizes expected utility, yet leads to first-order conditions that are linear in probabilities, hence easily solved and amenable to econometric analysis. In the Chew-Dekel class, the certainty equivalent function μ for a set of payoffs and probabilities $\{c(z), p(z)\}$ is defined implicitly by a *risk aggregator* M satisfying

$$\mu = \sum_z p(z) M[c(z), \mu]. \quad (4)$$

(This is Epstein and Zin's (1989) equation (3.10) with $M \equiv F + \mu$.) Chew (1983, 1989) and Dekel (1986, Section 2) show that such preferences satisfy a weaker condition than the notorious independence axiom that underlies expected utility. 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 (4) in probabilities.

In the examples that follow, we focus our attention on the following tractable members of the Chew-Dekel class:

- *Expected utility.* A version with constant relative risk aversion 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 (4), 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.* Chew (1983) suggests a relatively easy way to generalize expected utility given (4): 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.* Gul (1991) proposes another model that increases sensitivity to bad events (“disappointments”). Preferences are 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. 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).

Each of these models is described in Epstein and Zin (2001). Other tractable preferences include semi-weighted utility (Epstein and Zin, 2001), generalized disappointment aversion (Routledge and Zin, 2003), and rank-dependent preferences (Epstein and Zin, 1990). All but the last one are members of the Chew-Dekel class.

One source of intuition about these preferences is their state-space indifference curves, examples of which are pictured in Figure 2. For the purpose of illustration, suppose there are two equally likely states ($Z = 2$, $p(1) = p(2) = 1/2$). The 45-degree line represents certainty ($c(1) = c(2)$). Since preferences are linear homogeneous, the unit indifference curve ($\mu = 1$) completely characterizes preferences. For expected utility, the unit indifference curve is

$$\mu(\text{EU}) = [0.5c(1)^\alpha + 0.5c(2)^\alpha]^{1/\alpha} = 1.$$

This is the usual convex arc with a slope of minus one (the odds ratio) at the 45-degree line. As we decrease α , the arc becomes more convex. For weighted utility, the unit indifference curve is

$$\mu(\text{WU}) = \left[\frac{c(1)^{\gamma+\alpha} + c(2)^{\gamma+\alpha}}{c(1)^\gamma + c(2)^\gamma} \right]^{1/\alpha} = 1.$$

Drawn for the same value of α and a modest negative value of γ , it is more convex than expected utility, suggesting greater risk aversion. With disappointment aversion, the equation governing the indifference curve depends on whether $c(1)$ is larger or smaller than $c(2)$. If it's smaller (so that $z = 1$ is the bad state), the indifference curve is

$$\mu(\text{DA}) = \left[\left(\frac{1+\delta}{2+\delta} \right) c(1)^\alpha + \left(\frac{1}{2+\delta} \right) c(2)^\alpha \right]^{1/\alpha} = 1.$$

If it's larger, we switch the two states around. To express this more compactly, define sets of transformed probabilities, $\hat{p}_1 = [(1+\delta)/(2+\delta), 1/(2+\delta)]$ (when $z = 1$ is the bad state) and $\hat{p}_2 = [1/(2+\delta), (1+\delta)/(2+\delta)]$ (when $z = 2$ is the bad state). Then the indifference curve can be expressed

$$\left[\min_i \sum_z \hat{p}_i(z) c(z)^\alpha \right]^{1/\alpha} = 1.$$

We'll see something similar in Section 6. For now, note that the indifference curve is the upper envelope of two curves based on different sets of probabilities. The envelope is denoted by a solid line, and the extensions of the two curves by dashed lines. The result is an indifference curve with a kink at the 45-degree line, where the bad state switches. (As we cross from below, the bad state switches from 2 to 1.)

Another source of intuition is the sensitivity of certainty equivalents to small risks. For the two-state case discussed above, consider the certainty equivalent of the outcome $c(1) = 1 - \sigma$ and $c(2) = 1 + \sigma$ for small $\sigma > 0$, thereby defining the certainty equivalent as a function

of σ . How much does a small increase in σ reduce μ ? For expected utility, a second-order Taylor series expansion of $\mu(\sigma)$ around $\sigma = 0$ is

$$\mu(\text{EU}) \approx 1 - (1 - \alpha)\sigma^2/2.$$

This familiar bit of mathematics suggests $1 - \alpha$ as a measure of risk aversion. For weighted utility, a similar approximation yields

$$\mu(\text{WU}) \approx 1 - (1 - \alpha - 2\gamma)\sigma^2/2,$$

which suggests $1 - \alpha - 2\gamma$ as a measure of risk aversion. Note that neither expected utility nor weighted utility has a linear term: agents with these preferences are effectively indifferent to very small risks. For disappointment aversion, however, the Taylor series expansion is

$$\mu(\text{DA}) \approx 1 - \left(\frac{\delta}{2 + \delta}\right)\sigma - (1 - \alpha) \left(\frac{4 + 4\delta}{4 + 4\delta + \delta^2}\right)\sigma^2/2.$$

The linear term tells us that disappointment aversion exhibits *first-order risk aversion*, a consequence of the kink in the indifference curve.

Examples

Example 5 (certainty equivalents for log-normal risks). We illustrate the behavior of Chew-Dekel preferences in an environment in which the impact of risk on utility is particularly transparent. Define the risk premium on a risky consumption distribution by $rp \equiv \log[E(c)/\mu(c)]$, the logarithmic difference between consumption's expectation and its certainty equivalent. Suppose consumption is log-normal: $\log c(z) = \kappa_1 + \kappa_2^{1/2}z$, with z distributed $N(0,1)$. Recall that if $\log x \sim N(a, b)$, then $\log E(x) = a + b/2$ ("Ito's lemma," equation (42) of Appendix B). Since $\log c \sim N(\kappa_1, \kappa_2)$, expected consumption is $\exp(\kappa_1 + \kappa_2/2)$. Similarly, the certainty equivalent for expected utility is $\mu = \exp(\kappa_1 + \alpha\kappa_2/2)$ and the risk premium is $rp = (1 - \alpha)\kappa_2/2$. The proportionality factor $(1 - \alpha)$ is the traditional coefficient of relative risk aversion. Weighted utility is not quite kosher in this context (M is concave only in a neighborhood of μ), but the example nevertheless gives us a sense of its properties. Using similar methods, we find that the certainty equivalent is $\mu = \exp(\kappa_1 + (\alpha + 2\gamma)\kappa_2/2)$ and the risk premium is $rp = (1 - \alpha - 2\gamma)\kappa_2/2$. Note that the risk premium is the same as expected utility with parameter $\alpha' = \alpha + 2\gamma$. This equivalence of expected utility and weighted utility doesn't extend to other distributions, but it suggests that we might find some difficulty distinguishing between the two in practice. For disappointment aversion, we find the certainty equivalent using mathematics much like that underlying the Black-Scholes formula:

$$\mu^\alpha = e^{\alpha\kappa_1 + \alpha^2\kappa_2/2} + \delta \left[e^{\alpha\kappa_1 + \alpha^2\kappa_2/2} \Phi\left(\frac{\log \mu - \kappa_1 - \alpha\kappa_2}{\kappa_2^{1/2}}\right) - \Phi\left(\frac{\log \mu - \kappa_1}{\kappa_2^{1/2}}\right) \right],$$

where Φ is the standard normal distribution function; see equation (41) in Appendix B. Apparently the risk premium is no longer proportional to κ_2 . We show this in Figure 3, where we graph rp against κ_2 for all three preferences using the same parameter values as Figure 2 ($\alpha = \delta = 0.5$, $\gamma = -0.25$). As you might expect, disappointment aversion implies proportionately greater aversion to small risks than large ones; in this respect it is qualitatively different from expected utility and weighted utility. Routledge and Zin's (2003) generalized disappointment aversion does the reverse: it generates greater aversion to large risks. Different sensitivity to large and small risks provides a possible method to distinguish such preferences from expected utility.

Example 6 (portfolio choice with Chew-Dekel preferences). One strength of the Chew-Dekel class is that it leads to first-order conditions that are easily solved and used in econometric work. Consider an agent with initial net assets a_0 who invests fractions w in a risky asset with (gross) return $r(z)$ in state z and $1 - w$ in a risk-free asset with return r_0 . For an arbitrary choice of w , consumption in state z is $c(z) = a_0[r_0 + w(r(z) - r_0)]$. The portfolio choice problem might then be written

$$\max_w \mu[a_0\{r_0 + w(r(z) - r_0)\}] = a_0 \max_w \mu[r_0 + w(r(z) - r_0)],$$

the second equality stemming from the linear homogeneity of μ . The direct approach to this problem is to choose w to maximize μ , and in some cases we'll do that. For the general Chew-Dekel class, however, we may not have an explicit expression for the certainty equivalent function. In those cases, we use equation (4):

$$\max_w \mu[\{r_0 + w(r(z) - r_0)\}] = \max_w \sum_z p(z) M[r_0 + w(r(z) - r_0), \mu^*],$$

where μ^* is the maximized value of the certainty equivalent function. The problem on the right-hand side has first-order condition

$$\sum_z p(z) M_1[r_0 + w(r(z) - r_0), \mu^*][r(z) - r_0] = E[M_1(r_0 + w(r - r_0), \mu^*)(r - r_0)] = 0. \quad (5)$$

(There are M_2 terms, too, but you might verify for yourself that they can be eliminated.) We find the optimal portfolio by solving the first-order condition and (4) simultaneously for w and μ^* . The same conditions can also be used in econometric work to estimate preference parameters.

To see how you might use (5) to determine w , consider a numerical example with two equally-likely states and returns $r_0 = 1.01$, $r(1) = 0.90$, and $r(2) = 1.24$ (the "equity premium" is 6%). With expected utility, the first-order condition is

$$(\mu^*)^{\alpha-1} (1 - \beta) \sum_z p(z) (r_0 + w[r(z) - r_0])^{\alpha-1} [r(z) - r_0] = 0.$$

μ^* drops out and we can solve for w independently. For $\alpha = 0.5$, the solution is $w = 4.791$, which implies $\mu^* = 1.154$. The result is the dual of the equity premium puzzle: with modest

risk aversion, the observed equity premium induces a huge long position in the risky asset, financed by borrowing. With disappointment aversion, the first-order condition is

$$(1 + \delta)p(1) (r_0 + w[r(1) - r_0])^{\alpha-1} [r(1) - r_0] \\ + p(2) (r_0 + w[r(2) - r_0])^{\alpha-1} [r(2) - r_0] = 0,$$

since $z = 1$ is the bad state. For $\delta = 0.5$, $w = 2.147$ and $\mu^* = 1.037$. [Adapted from Epstein and Zin (1989, 2001).]

Example 7 (portfolio choice with rank-dependent preferences). Rank-dependent preferences are an interesting alternative to the Chew-Dekel class. We rank states so that the payoffs $c(z)$ are increasing in z and define the certainty equivalent function by

$$\mu = u^{-1} \left(\sum_z (g[P(z)] - g[P(z-1)]) u[c(z)] \right) = u^{-1} \left(\sum_z \hat{p}(z) u[c(z)] \right),$$

where g is an increasing function satisfying $g(0) = 0$ and $g(1) = 1$, $P(z) = \sum_{u=1}^z p(u)$ is the cumulative distribution function, and $\hat{p}(z) = g[P(z)] - g[P(z-1)]$ is a transformed probability. If $g(p) = p$, this is simply expected utility. If g is concave, these preferences exhibit risk aversion even if u is linear. However, since μ is nonlinear in probabilities it cannot be expressed in Chew-Dekel form. At the end of this section, we discuss the difficulties this raises for econometric estimation. In the portfolio choice problem, the first-order condition is

$$\sum_z \hat{p}(z) u_1[c(z)] [r(z) - r_0] = 0, \quad (6)$$

which is readily solved if we know the probabilities. [Adapted from Epstein and Zin (1990) and Yaari (1987).]

Example 8 (risk sharing). Consider a Pareto problem with two agents who divide a given risky aggregate endowment $y(z)$. If their certainty equivalent functions are identical and homogeneous of degree one, each agent consumes the same fraction of the aggregate endowment in all states. The problem is more interesting if the agents have different preferences. Let us say that two agents, indexed by i , have certainty equivalent functions $\mu^i[c^i(z)]$. A Pareto optimal allocation solves: choose $\{c^1(z), c^2(z)\}$ to maximize μ^1 subject to $c^1(z) + c^2(z) \leq y(z)$ and $\mu^2 \geq \bar{\mu}$ for some number $\bar{\mu}$. If λ is the Lagrange multiplier on the second constraint, the first-order conditions have the form

$$\frac{\partial \mu^1}{\partial c^1(z)} = \lambda \frac{\partial \mu^2}{\partial c^2(z)}.$$

With Chew-Dekel risk preferences, the derivatives have the form:

$$\frac{\partial \mu^i}{\partial c^i(z)} = p(z) M_1^i[c^i(z), \mu^i] + \sum_x p(x) M_2^i[c^i(x), \mu^i] \frac{\partial \mu^i}{\partial c^i(z)} \\ = p(z) M_1^i[c^i(z), \mu^i] / (1 - \sum_x p(x) M_2^i[c^i(x), \mu^i]).$$

This expression is not particularly user-friendly, but in principle we can solve it numerically for specific functional forms. With expected (power) utility, an optimal allocation solves

$$[\mu^1]^{1-\alpha_1}[y(z) - c^2(z)]^{\alpha_1-1} = \lambda[\mu^2]^{1-\alpha_2}c^2(z)^{\alpha_2-1},$$

which implies allocation rules that we can express in the form $c^i = s^i(y)y$. If we substitute into the optimality condition and differentiate, we find $ds^1/dy > 0$ if $\alpha_1 > \alpha_2$: the less risk averse agent absorbs a disproportionate share of the risk.

Discussion: moment conditions for preference parameters

One of the most useful features of Chew-Dekel preferences is how easily they can be used in econometric work. Since the risk aggregator (4) is linear in probabilities, we can apply method of moments estimators directly to first-order conditions.

In a typical method of moments estimator, a vector-valued function f of data x and a vector of parameters θ of equal dimension satisfies the *moment condition*

$$Ef(x, \theta_0) = 0, \tag{7}$$

where $\theta = \theta_0$ is the parameter vector that generated the data. A method of moments estimator θ_T for a sample of size T replaces the population mean with the sample mean:

$$T^{-1} \sum_{t=1}^T f(x_t, \theta_T) = 0.$$

Under reasonably general conditions, a law of large numbers implies that the sample mean converges to the population mean and θ_T converges to θ_0 . When the environment permits a central limit theorem, we can also derive an asymptotic normal distribution for θ_T . If the number of moment conditions (the dimension of f) is greater than the number of parameters (the dimension of θ), we can apply a generalized method of moments estimator with similar properties; see Hansen (1982).

The portfolio choice problem with Chew-Dekel preferences has exactly this form if the number of preference parameters is no greater than the number of risky assets. For each risky asset i there is a moment condition,

$$f_i(x, \theta) = M_1(c, \mu^*)(r_i - r_0),$$

analogous to equation (5). In the static case, we also need to estimate μ^* , which we do using (4) as an additional moment condition. (In a dynamic setting, a homothetic time aggregator allows us to replace μ^* with a function of consumption growth; see equation (13).)

Outside the Chew-Dekel class, estimation is a more complex activity. First-order conditions are no longer linear in probabilities and do not lead to moment conditions in the

form of equation (7). To estimate, say, equation (6) for rank-dependent preferences, we need a different estimation strategy. One possibility is a simulated method of moments estimator, which involves something like the following: (i) conjecture a probability distribution and parameter values; (ii) given these values, solve the portfolio problem for decision rules; (iii) calculate (perhaps through simulation) moments of the decision rule and compare them to moments observed in the data; (iv) if the two sets of moments are sufficiently close, stop; otherwise, modify parameter values and return to step (i). All of this can be done, but it highlights the econometric convenience of Chew-Dekel risk preferences.

4 Time and risk

We are now in a position to describe non-additive preferences in a dynamic stochastic environment like that illustrated by Figure 1. You might guess that the process of specifying preferences over time and states of nature is simply a combination of the two. In fact, the combination raises additional issues that are not be readily apparent. We touch on some of them here; others come up in the next two sections.

Recursive preferences

Consider the structure of preferences in a dynamic stochastic environment. In the tradition of Kreps and Porteus (1978), Johnsen and Donaldson (1985), and Epstein and Zin (1989), we represent a class of recursive preferences by

$$U_t = V[u_t, \mu_t(U_{t+1})], \tag{8}$$

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 , u_t is date- t utility, V is a time aggregator, and μ_t is a certainty-equivalent function based on the conditional probabilities $p(z_{t+1}|z^t)$. This structure is suggested by Kreps and Porteus (1978) for expected utility certainty equivalent functions. Epstein and Zin (1989) extend their work to stationary infinite-horizon settings and propose the more general Chew-Dekel class of risk preferences. As in Section 2, such preferences are dynamically consistent, history independent, future independent, and stationary. They are also *conditionally independent* in the sense of Johnsen and Donaldson (1985): preferences over choices at any history at date t (\bar{z}^t , for example) do not depend on other histories that may have (but did not) occur ($z^t \neq \bar{z}^t$). You can see this in Figure 1: If we are now at the node marked (A), then preferences do not depend on consumption at nodes stemming from (B) denoting histories that can no longer occur.

If equation (8) seems obvious, think again. If you hadn't read the previous paragraph or

its sources, you might just as easily propose

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

another seemingly natural combination of time and risk preference. This combination, however, has a serious flaw: it implies dynamically inconsistent preferences unless it reduces to (1). See Kreps and Porteus (1978) and Epstein and Zin (1989, Section 4). File away for later the idea that the combination of time and risk preference can raise subtle dynamic consistency issues.

We refer to the combination of the recursive structure (8) and an expected utility certainty equivalent as Kreps-Porteus preferences. A popular parametric example consists of the constant elasticity aggregator,

$$V[u, \mu(U)] = [(1 - \beta)u^\rho + \beta\mu(U)^\rho]^{1/\rho}, \quad (9)$$

and the “power certainty equivalent,”

$$\mu(U) = [E(U^\alpha)]^{1/\alpha}, \quad (10)$$

with $\rho, \alpha < 1$. Equations (9) and (10) are homogeneous of degree one with constant discount factor β . This is more restrictive than the aggregators we considered in Section 2, but linear homogeneity rules out more general discounting schemes: it implies that indifference curves have the same slope along any ray from the origin, so their slope along the 45-degree line must be the same, too. If U is constant, the weights $(1 - \beta)$ and β define $U = u$ as the (steady state) level of utility. It is common to refer to $1 - \alpha$ as the coefficient of relative risk aversion and $1/(1 - \rho)$ as the intertemporal elasticity of substitution. If $\rho = \alpha$, the model is equivalent to one satisfying (1) and intertemporal substitution is the inverse of risk aversion. More generally, the Kreps-Porteus structure allows us to specify risk aversion and intertemporal substitution independently. Further, a Kreps-Porteus agent prefers early resolution of risk if $\alpha < \rho$; see Epstein and Zin (1989, Section 4). This separation of risk aversion and intertemporal substitution has proved to be not only a useful empirical generalization, but an important source of intuition about the properties of dynamic models.

We can generate further flexibility by combining (8) with a Chew-Dekel risk aggregator (4), thereby introducing Chew-Dekel risk preferences to dynamic environments. We refer to this combination as Epstein-Zin preferences.

Examples

Example 9 (Weil’s model of precautionary saving). We say consumption-saving models generate *precautionary saving* if risk decreases consumption as a function of current assets. In the

canonical consumption problem with additive preferences, income risk has this effect if the period utility function u has constant $k \equiv u_{111}u_1/(u_{11})^2 > 0$. See, for example, Ljungqvist and Sargent (2000, pp 390-393). Both power utility and exponential utility satisfy this condition. With power utility ($u(c) = c^\alpha/\alpha$), $k = (\alpha - 2)(\alpha - 1)$, which is positive for $\alpha < 1$ and therefore implies precautionary saving. (In the next section we look at quadratic utility, which effectively sets $\alpha = 2$, implying $k = 0$ and no precautionary saving.) Similarly, with exponential utility ($u(c) = -\exp(-\alpha c)$), $k = 1 > 0$. With Kreps-Porteus preferences we can address a somewhat different question: Does precautionary saving depend on intertemporal substitution, risk aversion, or both? To answer this question, consider the problem characterized by the Bellman equation

$$J(a) = \max_c \{(1 - \beta)c^\rho + \beta\mu[J(a')]^\rho\}^{1/\rho}$$

subject to the budget constraint $a' = r(a - c) + y'$, where $\mu(x) = -\alpha^{-1} \log E \exp(-\alpha x)$ and $\{y_t\} \sim \text{NID}(\kappa_1, \kappa_2)$. The exponential certainty equivalent μ is not homogeneous of degree one, but it is analytically convenient for problems with additive risk. The parameters satisfy $\rho \leq 1$, $\alpha \geq 0$, $r > 1$, and $\beta^{1/(1-\rho)}r^{\rho/(1-\rho)} < 1$. Of particular interest are ρ , which governs intertemporal substitution, and α , which governs risk aversion.

The value function in this example is linear with parameters that can be determined by the time-honored guess-and-verify method. We guess (we've seen this problem before) $J(a) = A + Ba$ for parameters (A, B) to be determined. The certainty equivalent of future utility is

$$\mu[J(a')] = \mu[A + Br(a - c) + By'] = A + Br(a - c) + B\kappa_1 - \alpha B^2\kappa_2/2, \quad (11)$$

which follows from equation (42) of Appendix B. The first-order and envelope conditions are

$$\begin{aligned} 0 &= J(a)^{1-\rho} [(1 - \beta)c^{\rho-1} - \beta\mu^{\rho-1}Br] \\ J_1(a) &= B = J(a)^{1-\rho}\beta\mu^{\rho-1}Br, \end{aligned}$$

which imply

$$\begin{aligned} \mu &= (\beta r)^{1/(1-\rho)}J(a) = (\beta r)^{1/(1-\rho)}(A + Ba) \\ c &= [(1 - \beta)/B]^{1/(1-\rho)}J(a) = [(1 - \beta)/B]^{1/(1-\rho)}(A + Ba). \end{aligned}$$

The latter tells us that the decision rule is linear, too. If we substitute both equations into (11), we find that the parameters of the value function must be

$$A = (r - 1)^{-1}(\kappa_1 - B\alpha\kappa_2/2)B, \quad B = \left[\frac{(1 - \beta)^{1/(1-\rho)}}{1 - \beta^{1/(1-\rho)}r^{\rho/(1-\rho)}} \right]^{(1-\rho)/\rho}.$$

They imply the decision rule

$$c = \left(1 - \beta^{1/(1-\rho)} r^{\rho/(1-\rho)}\right) \left(a + (r-1)^{-1}[\kappa_1 - B\alpha\kappa_2/2]\right).$$

The last term is the impact of risk. Apparently a necessary condition for precautionary saving is $\alpha > 0$, so the parameter controlling precautionary saving is risk aversion. [Adapted from Weil (1993).]

Example 10 (Merton-Samuelson portfolio model). Our next example illustrates the relation between consumption and portfolio decisions in iid environments. The model is similar to the previous example, and we use it to address a similar issue: the impact of asset return risk on consumption. At each date t a theoretical agent faces the budget constraint

$$a_{t+1} = (a_t - c_t) \sum_i w_{it} r_{it+1},$$

where w_{it} is the share of post-consumption wealth invested in asset i and r_{it+1} is its return. Returns $\{r_{it+1}\}$ are iid over time. Preferences are characterized by the constant elasticity time aggregator (9) and an arbitrary linearly homogeneous certainty equivalent function. The Bellman equation is

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

subject to

$$a' = (a - c) \sum_i w_i r'_i = (a - c)r'_p$$

and $\sum_i w_i = 1$, where r_p is the portfolio return. Since the time and risk aggregators are linear homogeneous, so is the value function, and the problem decomposes into separate portfolio and consumption problems. The portfolio problem is:

$$\max_w \mu[J(a')] = (a - c) \max_w \mu[J(r'_p)].$$

Since returns are iid, the portfolio problem is the same at all dates and can be solved using methods outlined in the previous section. Given a solution μ^* to the portfolio problem, the consumption problem is:

$$J(a) = \max_c \{(1-\beta)c^\rho + \beta[(a-c)\mu^*]^\rho\}^{1/\rho}.$$

The first-order condition implies the decision rule $c = [A/(1+A)]a$, where

$$A = [(1-\beta)/\beta]^{1/(1-\rho)} (\mu^*)^{-\rho/(1-\rho)}.$$

The impact of risk is mediated by μ^* and involves the familiar balance of income and substitution effects. If $\rho < 0$, the intertemporal elasticity of substitution is less than one and smaller μ^* (larger risk premium) is associated with lower consumption (the income effect).

If $\rho > 0$, the opposite happens. In contrast to the previous example, the governing parameter is ρ ; the impact of risk parameters is imbedded in μ^* . Note, too, that the impact on consumption of a change in μ^* can generally be offset by a change in β that leaves A unchanged. This leads to an identification issue that we discuss at greater length in the next example. Farmer and Gertler use a similar result to motivate setting $\alpha = 1$ (risk neutrality) in the Kreps-Porteus preference models, which leads to linear decision rules even with risk to income, asset returns, and length of life. [Adapted from Epstein and Zin (1989), Farmer (1990), Gertler (1999), and Weil (1990).]

Example 11 (asset pricing). The central example of this section is an exploration of time and risk preference in the traditional exchange economy of asset pricing. Preferences are governed by the constant elasticity time aggregator (9) and the Chew-Dekel risk aggregator (4). We characterize asset returns for general recursive preferences and discuss the identification of time and risk preference parameters. We break the argument into a series of steps.

Step (i) (consumption and portfolio choice). Consider a stationary Markov environment with states z and conditional probabilities $p(z'|z)$. A dynamic consumption/portfolio problem for this environment is characterized by the Bellman equation

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

subject to the budget 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 r_p is the portfolio return. 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 .

As in the previous example, the 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 mechanics are similar to Example 6. The portfolio first-order conditions are

$$\sum_{z'} p(z'|z) M_1[L(z')r_p(z, z'), \mu] L(z') [r_i(z, z') - r_j(z, z')] = 0 \quad (12)$$

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

$$(1 - \beta)b^{\rho-1} = \beta(1 - b)^{\rho-1}\mu^\rho. \quad (13)$$

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

$$\begin{aligned} \mu &= [(1 - \beta)/\beta]^{1/\rho} [b/(1 - b)]^{(\rho-1)/\rho} \\ L &= (1 - \beta)^{1/\rho} b^{(\rho-1)/\rho}. \end{aligned} \quad (14)$$

The first-order condition (13) and value function (14) allow us to express the relation between consumption and returns in almost 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

$$\mu(x'r'_p) = \mu \left(\left[\beta(c'/c)^{\rho-1} r'_p \right]^{1/\rho} \right) = 1, \quad (15)$$

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 (12) 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

$$E \left[M_1(x'r'_p, 1)x'r'_i \right] = E \left[M_1(x'r'_p, 1)x'r'_p \right]. \quad (16)$$

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 is helpful depends on M . [Adapted from Epstein and Zin (1989).]

Step (ii) (equilibrium). Now shift focus to an exchange economy in which output growth follows a stationary Markov process: $g' = y'/y = g(z')$. In equilibrium, consumption equals output and the optimal portfolio is a claim to the stream of future output. We denote the price of this claim by q and the price-output ratio by $Q = q/y$. Its return is therefore

$$r'_p = (q' + y')/q = (Q'y' + y')/(Qy) = g'(Q' + 1)/Q. \quad (17)$$

With linear homogeneous preferences, the equilibrium price-output ratio is a stationary function of the current state, $Q(z)$. Asset pricing then consists of these steps: (a) Substitute (17) into (15) and solve for Q :

$$\mu \left(\left[\beta(g')^\rho (Q' + 1) \right]^{1/\rho} \right) = Q^{1/\rho}.$$

(b) Compute the portfolio return r'_p from (17). (c) Use (16) to derive returns on other assets.

Step (iii) (the iid case). If the economy is iid, we cannot generally identify separate time and risk parameters. Time and risk parameters are intertwined in (16), but suppose we were somehow able to estimate the risk parameters. How might we estimate the time preference

parameters β and ρ from observations of r_p (returns) and b (the consumption-wealth ratio)? Formally, equations (13) and (14) imply the intertemporal optimality condition

$$(1 - b)^{1-\rho} = \beta\mu(r'_p)^\rho.$$

If r_p is iid, μ and b are constant. With no variation in μ or b , the optimality condition cannot tell us both ρ and β : for any value of ρ , we can satisfy the condition by adjusting the discount factor β . The only limit to this is the restriction $\beta < 1$. Evidently a necessary condition for identifying separate time and risk parameters is that risk varies over time. The issue doesn't arise with additive preferences, which tie time preference to risk preference. [Adapted from Kocherlakota (1990) and Wang (1993).]

Step (iv) (extensions). With Kreps-Porteus preferences and non-iid returns, the model does somewhat better in accounting for asset returns. It nevertheless fails to provide an entirely persuasive account of observed relations between asset returns and aggregate consumption. Roughly speaking, the same holds for more general risk preference specifications, although the combination of exotic preferences and time-varying risk shows promise. [See Bansal and Yaron (2003), Epstein and Zin (1991), Lettau, Ludvigson, and Wachter (2003), Routledge and Zin (2003), Tallarini (2000), and Weil (1989).]

Example 12 (risk sharing). With additive preferences and equal discount factors, Pareto problems generate constant weights on agents' utilities over time and across states of nature, even if period/state utility functions differ. With Kreps-Porteus preferences, differences in risk aversion lead to systematic drift in the weights. To be concrete, suppose states z follow a Markov chain with conditional probabilities $p(z'|z)$. Aggregate output is $y(z)$. Agents have the same aggregator, $V(c, \mu) = (c^\rho + \beta\mu^\rho)/\rho$, but different certainty equivalent functions,

$$\mu^i[x(z')] = \left(\sum_{z'} p(z'|z)x(z')^{\alpha_i} \right)^{1/\alpha_i}$$

for state-dependent "utility" x . The Bellman equation for the Pareto problem is

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

subject to

$$(c^\rho + \beta\mu^2[w_{z'}]^\rho) / \rho \geq w.$$

Here c and $w_{z'}$ refer to consumption and promised future utility of the second agent. The first-order and envelope conditions imply

$$\begin{aligned} (y(z) - c)^{\rho-1} &= \lambda c^{\rho-1} \\ (\mu^1)^{\rho-\alpha_1} J(w_{z'}, z')^{\alpha_1-1} J_1(w_{z'}, z') &= J_1(w, z) (\mu_{z'}^2)^{\rho-\alpha_2} w_{z'}^{\alpha_2-1} \\ J_1(w, z) &= -\lambda. \end{aligned}$$

The first equation leads to the familiar allocation rule $c = [1 + \lambda^{1/(\rho-1)}]^{-1}y(z)$. If $\alpha_1 \neq \alpha_2$, the weight λ will generally vary over time. [Adapted from Anderson (2004) and Kan (1995).]

Example 13 (habits, disappointment aversion, and conditional independence). Habits and disappointment aversion both assess utility by comparing consumption to a benchmark. With disappointment aversion, the benchmark is the certainty equivalent. With habits, the benchmark is a function of past consumption. Despite this apparent similarity, there are a number of differences between them. One is timing: the habit is known and fixed when current decisions are made, while the certainty equivalent generally depends on those decisions. Another is that disappointment aversion places restrictions on the benchmark that have no obvious analog in the habit model. A third is that habits take us outside the narrowly-defined class of recursive preferences summarized by equation (8): they violate the assumption of conditional independence. Why? Because preferences at any node in the event tree depend on past consumption through the habit, which in turn depends on nodes that can no longer be reached. In Figure 1, for example, decisions at node (A) depend on the habit, which was chosen at (say) the initial node z_0 and therefore depends on anything that could have happened from there on, including (B) and its successors. The solution, of course, is to define preferences conditional on a habit state variable and proceed in the natural way.

Discussion: distinguishing time and risk preference

The defining feature of this class of preferences is the separation of time preference (summarized by the aggregator V) and risk preference (summarized by the certainty equivalent function μ). In the functional forms used in this section, time preference is characterized by a discount factor and an intertemporal substitution parameter. Risk preference is characterized by risk aversion and possibly other parameters indicated by the Chew-Dekel risk aggregator. We have therefore added one or more parameters to the conventional additive utility function (1). Examples suggest that the additional parameters may be helpful in explaining precautionary saving, asset returns, and the intertemporal allocation of risk.

A critical question in applications is whether these additional parameters can be identified and estimated from a single time series realization of all the relevant variables. If so, we can use the methods outlined in the previous section: apply a method of moments estimator to the first-order conditions of the problem of interest. Identification hinges on the nature of risk. If risk is iid, we cannot identify separate time and risk parameters. This is clear in examples, but the logic is both straightforward and general: we need variation over time to identify time preference. A more formal statement is given by Wang (1993).

5 Risk-sensitive and robust control

Risk-sensitive and robust control emerged in the engineering literature in the 1970s and were brought to economics and developed further by Hansen and Sargent, their many coauthors, and a few other brave souls. The most popular version of risk-sensitive control is based on Kreps-Porteus preferences with an exponential certainty equivalent function. Robust control considers a new issue: decision making when the agent does not know the probability model generating the data. The agent considers instead a range of models, and makes decisions that maximize utility given the worst possible model. The same issue is addressed from a different perspective in the next section. Much of this work deals with linear-quadratic-gaussian (LQG) problems, but the ideas are applicable more generally. We start by describing risk-sensitive and robust control in a static scalar LQG setting, where the insights are less cluttered by algebra. We go on to consider dynamic LQG problems, robust control problems outside the LQG universe, and challenges of estimating, and distinguishing between, models based on risk-sensitive and robust control.

Static control

Many of the ideas behind risk-sensitive and robust control can be illustrated with a static, scalar example. We consider traditional optimal control, risk-sensitive control, and robust control as variants of the same underlying problem. The striking result is the equivalence of optimal decisions made under risk-sensitive and robust control.

In our example, an agent maximizes some variant of a quadratic “return” function,

$$u(v, x) = -[Qv^2 + Rx^2],$$

subject to the linear constraint,

$$x = Ax_0 + Bv + C(w + \varepsilon), \tag{18}$$

where v is a control variable chosen by the agent, x is a state variable that is controlled indirectly through v , x_0 is a fixed “initial” value, $(Q, R) > 0$ are preference parameters, (A, B, C) are nonzero parameters describing the determination of x , $\varepsilon \sim N(0, 1)$ is noise, and w is a distortion of the model that we’ll describe in greater detail when we get to robust control. The problem sets up a tradeoff between the cost (Qv^2) and potential benefit (Rx^2) of nonzero values of v . If you’ve seen LQG control problems before, most of this should look familiar.

Optimal control. In this problem and the next one we set $w = 0$, thereby ruling out distortions. The control problem is: choose v to maximize Eu given the constraint (18). Since

$$Eu = -[Qv^2 + R(Ax_0 + Bv)^2] - RC^2, \tag{19}$$

the objective functions with and without noise differ only by a constant. Noise therefore has no impact on the optimal choice of v . For both problems, the optimal v is

$$v = -(Q + B^2R)^{-1}(ABR)x_0.$$

This solution serves as a basis of comparison for the next two.

Risk-sensitive control. Continuing with $w = 0$, we consider an alternative approach that brings risk into the problem in a meaningful way: we maximize an exponential certainty equivalent of u :

$$\mu(u) = -\alpha^{-1} \log E \exp(-\alpha u),$$

where $\alpha \geq 0$ is a risk aversion parameter. (This is more natural in a dynamic setting, where we would compute the certainty equivalent of future utility à la Kreps and Porteus.) We find $\mu(u)$ by applying formula (43) of Appendix B:

$$\mu(u) = -(1/2) \log(1 - 2\alpha RC^2) - [Qv^2 + [R/(1 - 2\alpha RC^2)](Ax_0 + Bv)^2] \quad (20)$$

as long as $1 - 2\alpha RC^2 > 0$. This condition places an upper bound on the risk aversion parameter α . Without it, the agent can be so sensitive to risk that her objective function is negative infinity regardless of the control. The first term on the right side of (20) does not depend on v or x , so it has no effect on the choice of v . The important difference from (19) is the last term: the coefficient of $(Ax_0 + Bv)^2$ is larger than R , making the agent more willing to tolerate nonzero values of v to bring x close to zero. The optimal v is

$$v = -(Q + B^2R - \alpha QRC^2)^{-1}(ABR)x_0.$$

If $\alpha = 0$ (risk neutrality) or $C = 0$ (no noise), this is the same as the optimal control solution. If $\alpha > 0$ and $C \neq 0$, the optimal choice of v is larger in absolute value because risk aversion increases the benefit of driving x to zero.

Robust control. Our third approach is conceptually different. We bring back the distortion w and tell the following story: We are playing a game against a malevolent nature, who chooses w to minimize our objective function. If our objective were to maximize Eu , then w would be infinite and our objective function would be minus infinity regardless of what we do. Let us therefore add a penalty (to nature) of θw^2 , making our objective function

$$\min_w Eu + \theta w^2.$$

The parameter $\theta > 0$ has the effect of limiting how much nature distorts the model, with small values of θ implying weaker limits on nature. The minimization implies

$$w = (\theta - RC^2)^{-1}R(Ax_0 + Bv),$$

making the robust control objective function

$$\min_w Eu + \theta w^2 = -[Qv^2 + [R/(1 - \theta^{-1}RC^2)](Ax_0 + Bv)^2] - RC^2. \quad (21)$$

The remarkable result: if we set $\theta^{-1} = 2\alpha$, the robust control objective differs from the risk-sensitive control objective (20) only by a constant, so it leads to the same choice of v . As in risk-sensitive control, the choice of v is larger in absolute value, in this case to offset the impact of w . There is, once again, a limit on the parameter: where α was bounded above, θ is bounded below. An infinite value of θ reproduces the optimal control objective function and solution.

A further result applies to the example: risk-sensitive and robust control are observationally equivalent to the traditional control problem with suitably adjusted R . That is, if we replace R in equation (19) with

$$\hat{R} = R/(1 - 2\alpha RC^2) = R + 2\alpha R^2 C^2 / (1 - 2\alpha RC^2) > R, \quad (22)$$

then the optimal control problem is equivalent to risk-sensitive control, which we've seen is equivalent to robust control. If Q and R are functions of more basic parameters it may not be possible to adjust R in this way, but the exercise points to the qualitative impact on the control: be more aggressive. This result need not survive beyond the scalar case, but it's suggestive.

Although risk-sensitive and robust control lead to the same decision, they are based on different preferences and give the decision different interpretations. With risk-sensitive control, we are concerned with risk for traditional reasons and the parameter α measures risk aversion. With robust control, we are concerned with model uncertainty (possible nonzero values of w). To deal with it, we make decisions that maximize given the worst possible specification error. The parameter θ controls how bad the error can be.

Entropy constraints. One of the most interesting developments in robust control is a procedure for setting θ : namely, choose θ to limit the magnitude of model specification error, with specification error measured by *entropy*. We define the entropy of transformed probabilities \hat{p} relative to reference probabilities p by

$$I(\hat{p}; p) \equiv \sum_z \hat{p}(z) \log[\hat{p}(z)/p(z)] = \hat{E} \log(\hat{p}/p), \quad (23)$$

where the expectation is understood to be based on \hat{p} . Note that $I(\hat{p}; p)$ is non-negative and equals zero when $\hat{p} = p$. Since the likelihood is the probability density function expressed as a function of parameters, entropy can be viewed as the expected difference in log-likelihoods between the reference and transformed models, with the expectation based on the latter.

In a robust control problem, we can limit the amount of specification error faced by an agent by imposing an upper bound on I : consider (say) only transformations \hat{p} such

that $I(\hat{p}; p) \leq I_0$ for some positive number I_0 . This *entropy constraint* takes a particularly convenient form in the normal case. Let \hat{p} be the density of x implied by equation (18) and p the density with $w = 0$:

$$\begin{aligned}\hat{p}(x) &= (2\pi C^2)^{-1/2} \exp[-(x - Ax_0 - Bv - Cw)^2/2C^2] = (2\pi C^2)^{-1/2} \exp[-\varepsilon^2/2] \\ p(x) &= (2\pi C^2)^{-1/2} \exp[-(x - Ax_0 - Bv)^2/2C^2] = (2\pi C^2)^{-1/2} \exp[-(w + \varepsilon)^2/2].\end{aligned}$$

Relative entropy is

$$I(\hat{p}; p) = \widehat{E}(w^2/2 + w\varepsilon) = w^2/2.$$

If we add the constraint $w^2/2 \leq I_0$ to the optimal control objective (19), the new objective is

$$\min_w -[Qv^2 + R(Ax_0 + Bv + Cw)^2] - RC^2 + \theta(w^2 - 2I_0),$$

where θ is the Lagrange multiplier on the constraint. The only difference from the robust control problem we discussed earlier is that θ is determined by I_0 . Low values of I_0 (tighter constraints) are associated with high values of θ , so the lower bound on θ is associated with an upper bound on I_0 .

Example 14 (Kyddland and Prescott's inflation game). A popular macroeconomic policy game goes like this: the government chooses inflation q to maximize the quadratic return function,

$$u(q, y) = -[q^2 + Ry^2],$$

subject to the Phillips curve,

$$y = y_0 + B(q - q^e) + C(w + \varepsilon),$$

where y is the deviation of output from its social optimum, q^e is expected inflation, (R, B, C) are positive parameters, y_0 is the noninflationary level of output, and $\varepsilon \sim N(0, 1)$. We assume $y_0 < 0$, which imparts an inflationary bias to the economy.

This problem is similar to our example, with one twist: We assume q^e is chosen by private agents to equal the value of q they expect the government to choose (another definition of rational expectations), but taken as given by the government (and nature). Agents know the model, so they end up setting $q^e = q$. A robust control version of this problem leads to the optimization:

$$\max_q \min_w -E \left(q^2 + R[y_0 + B(q - q^e) + C(w + \varepsilon)]^2 \right) + \theta w^2.$$

Note that we can do the min and max in any order (the min-max theorem). We do both at the same time, which generates the first-order conditions

$$\begin{aligned}q + RB[y_0 + B(q - q^e) + Cw] &= 0 \\ -\theta w + RC[y_0 + B(q - q^e) + Cw] &= 0.\end{aligned}$$

Applying the rational expectations condition $q^e = q$ leads to

$$q = -\left(\frac{RB}{1 - \theta^{-1}RC^2}\right)y_0, \quad w = \left(\frac{\theta^{-1}RC}{1 - \theta^{-1}RC^2}\right)y_0.$$

Take $\theta^{-1} = 0$ as the benchmark. Then $q = -RB y_0 > 0$ (the inflationary bias we mentioned earlier) and $w = 0$ (no distortions). For smaller values of $\theta > RC^2$, inflation is higher. Why? Because negative values of w effectively lower the noninflationary level of output (it becomes $y_0 + Cw$), leading the government to tolerate more inflation. As θ approaches its lower bound of RC^2 , inflation approaches infinity. If we treat this as a “constraint problem” with entropy bound $w^2/2 \leq I_0$, then $w = -(2I_0)^{1/2}$ (recall that $w < 0$) and the Lagrange multiplier θ is related to I_0 by

$$\theta = RC^2 - RCy_0/(2I_0)^{1/2}.$$

The lower bound on θ corresponds to an upper bound on I_0 . All of this is predicated on private agents understanding the government’s decision problem, including the value of θ . [Adapted from Hansen and Sargent (2004, ch 5) and Kydland and Prescott (1977).]

Example 15 (entropy with three states). With three states, the constraint $I(\hat{p}; p) \leq I_0$ is two-dimensional, since the probability of the third state can be computed from the other two. Figure 4 illustrates the constraint for the reference probabilities $p(1) = p(2) = p(3) = 1/3$ (the point marked “+”) and $I_0 = 0.1$. The boundary of the constraint set is the “egg.” By varying I_0 we vary the size of the constraint set. Chew-Dekel preferences can be viewed from the same perspective. Disappointment aversion, for example, is a one-dimensional class of “distortions.” If the first state is the only one worse than the certainty equivalent, the transformed probabilities are $\hat{p}(1) = (1 + \delta)p(1)/[1 + \delta p(1)]$, $\hat{p}(2) = p(2)/[1 + \delta p(1)]$, and $\hat{p}(3) = p(3)/[1 + \delta p(1)]$. Their entropy is

$$I(\delta) = \log[1 + \delta p(1)] - p(1) \log(1 + \delta),$$

a positive increasing function of $\delta \geq 0$. By varying δ subject to the constraint $I(\delta) \leq I_0$, we produce the line shown in the figure. (It hits the boundary at $\delta = 1.5$.) The interpretation of disappointment aversion, however, is different: in the theory of Section 3, the line represents different preferences, not model uncertainty.

Dynamic control

Similar issues and equations arise in dynamic settings. The traditional linear-quadratic control problem starts with the quadratic return function,

$$u(v_t, x_t) = -\left(v_t^\top Q v_t + x_t^\top R x_t + 2x_t^\top S v_t\right),$$

where v is the control and x is the state. Both are vectors, and (Q, R, S) are matrices of suitable dimension. The state evolves according to the law of motion

$$x_{t+1} = Ax_t + Bv_t + C(w_t + \varepsilon_{t+1}), \quad (24)$$

where w is a distortion (zero in some applications) and $\{\varepsilon_t\} \sim \text{NID}(0, I)$ is random noise. We use these inputs to describe optimal, risk-sensitive, and robust control problems. As in the static example, the central result is the equivalence of decisions made under risk-sensitive and robust control. We skip quickly over the more torturous algebraic steps, which are available in the sources listed in Appendix A.

Optimal control. We maximize the objective function,

$$E_0 \sum_{t=0}^{\infty} \beta^t u(v_t, x_t),$$

subject to (24) and $w_t = 0$. From long experience, we know that the value function takes the form

$$J(x) = -x^\top Px - q \quad (25)$$

for a positive semi-definite symmetric matrix P and a scalar q . The Bellman equation is

$$\begin{aligned} -x^\top Px - q = \max_v \left\{ - \left(v^\top Qv + x^\top Rx + 2x^\top Sv \right) \right. \\ \left. - \beta E \left[(Ax + Bv + C\varepsilon')^\top P (Ax + Bv + C\varepsilon') + p \right] \right\}. \end{aligned} \quad (26)$$

Solving the maximization in (26) leads to the Riccati equation

$$P = R + \beta A^\top P A - (\beta A^\top P B + S)(Q + \beta B^\top P B)^{-1}(\beta B^\top P A + S^\top). \quad (27)$$

Given a solution for P , the optimal control is $v = -Fx$, where

$$F = (Q + \beta B^\top P B)^{-1}(\beta B^\top P A + S^\top). \quad (28)$$

As in the static scalar case, risk is irrelevant: the control (28) does not depend on C . You can solve such problems numerically by iterating on the Riccati equation: make an initial guess of P (we use I), plug it into the right side of (27) to generate the next estimate of P , and repeat until successive values are sufficiently close together. See Anderson, Hansen, McGrattan, and Sargent (1996) for algebraic details, conditions guaranteeing convergence, and superior computational methods (the doubling algorithm, for example).

Risk-sensitive control. Risk-sensitive control arose independently, but can be regarded as an application of Kreps-Porteus preferences using an exponential certainty equivalent. The exponential certainty equivalent introduces risk into the decisions without destroying the quadratic structure of the value function. The Bellman equation is

$$J(x) = \max_v \{u(v, x) + \beta \mu[J(x')]\},$$

where the maximization is subject to $x' = Ax + Bv + C\varepsilon'$ and $\mu(J) = -\alpha^{-1} \log E \exp(-\alpha J)$. If the value function has the quadratic form (25), the multivariate analog to (43) gives us:

$$\mu[J(Ax + Bv + C\varepsilon')] = -(1/2) \log |I - 2\alpha C^\top PC| + (Ax + Bv)^\top \hat{P}(Ax + Bv),$$

where

$$\hat{P} = P + 2\alpha PC(I - 2\alpha C^\top PC)^{-1} C^\top P \quad (29)$$

as long as $|I - 2\alpha C^\top PC| > 0$. Each of these pieces has a counterpart in the static case. The inequality again places an upper bound on the risk aversion parameter α ; for larger values, the integral implied by the expectation diverges. Equation (29) corresponds to (22); in both equations, risk sensitivity increases the agent's aversion to non-zero values of the state variable. Substituting \hat{P} into the Bellman equation and maximizing leads to a variant of the Riccati equation,

$$P = R + \beta A^\top \hat{P} A - (\beta A^\top \hat{P} B + S)(Q + \beta B^\top \hat{P} B)^{-1} (\beta B^\top \hat{P} A + S^\top), \quad (30)$$

and associated control matrix,

$$F = (Q + \beta B^\top \hat{P} B)^{-1} (\beta B^\top \hat{P} A + S^\top).$$

A direct (if inefficient) solution technique is to iterate on (29,30) simultaneously. We describe another method shortly.

Robust control. As in our static example, the idea behind robust control is that a malevolent nature chooses distortions w that reduce our utility. A recursive version has the Bellman equation:

$$J(x) = \max_v \min_w \left\{ u(v, x) + \beta \left(\theta w^\top w + EJ(x') \right) \right\}$$

subject to the law of motion $x' = Ax + Bv + C(w + \varepsilon')$. The value function again takes the form (25), so the Bellman equation can be expressed

$$\begin{aligned} -x^\top Px - q &= \max_v \min_w \left\{ - \left(v^\top Qv + x^\top Rx + 2v^\top Sx \right) + \beta \theta w^\top w \right. \\ &\quad \left. - \beta E \left([Ax + Bv + C(w + \varepsilon')]^\top P(Ax + Bv + C\varepsilon') + p \right) \right\}. \end{aligned} \quad (31)$$

The minimization leads to

$$w = (\theta I - C^\top PC)^{-1} C^\top P(Ax + Bv)$$

and

$$\theta w^\top w - (Ax + Bv + Cw)^\top P(Ax + Bv + Cw) = (Ax + Bv)^\top \hat{P}(Ax + Bv),$$

where

$$\hat{P} = P + \theta^{-1} PC(I - \theta^{-1} C^\top PC)^{-1} C^\top P. \quad (32)$$

Comparing (32) with (29), we see that risk-sensitive and robust control lead to similar objective functions and produce identical decision rules if $\theta^{-1} = 2\alpha$.

A different representation of the problem leads to a solution that fits exactly into the traditional optimal control framework and is therefore amenable to traditional computational methods. The min-max theorem suggests that we can compute the solutions for v and w simultaneously. With this in mind, define

$$\hat{v} = \begin{bmatrix} v_t \\ w_t \end{bmatrix}, \quad \hat{Q} = \begin{bmatrix} Q & 0 \\ 0 & -\beta\theta I \end{bmatrix}, \quad \hat{S} = \begin{bmatrix} S & 0 \end{bmatrix}, \quad \hat{B} = \begin{bmatrix} B & C \end{bmatrix}.$$

Then the problem is one of optimal control, and can be solved using the Riccati equation (27) applied to $(\hat{Q}, R, \hat{S}, A, \hat{B})$. The optimal controls are $v = -F_1x$ and $w = -F_2x$, where the F_i come from partitioning F . A doubling algorithm applied to this problem provides an efficient computational technique for robust and risk-sensitive control problems.

Entropy constraints. As in the static case, dynamic robust control problems can be derived using an entropy constraint. Hansen and Sargent (2004, ch 6) suggest

$$\sum_{t=0}^{\infty} \beta^t w_t^\top w_t / 2 \leq I_0.$$

Discounting is convenient here, but is not a direct outcome of a multiperiod entropy calculation. They argue that discounting allows distortions to continue to play a role in the solution; without it, the problem tends to drive I_t and w_t to zero with time. A recursive version of the constraint is

$$I_{t+1} = \beta^{-1}(w_t^\top w_t - I_t).$$

A recursive robust constraint problem is based on an expanded state vector, (x, I) , and the law of motion for I above. As in the static case, the result is a theory of the Lagrange multiplier θ . Conversely, the solution to a traditional robust control problem with given θ can be used to compute the implied value of I_0 . The recursive version highlights an interesting feature of this problem: nature not only minimizes at a point in time, but allocates entropy over time in the way that has the greatest adverse impact on the agent.

Example 16 (robust precautionary saving). Consider a linear-quadratic version of the precautionary saving problem. A theoretical agent has quadratic utility, $u(c_t) = (c_t - \gamma)^2$, and maximizes the expected discounted sum of utility subject to a budget constraint and an autoregressive income process:

$$\begin{aligned} a_{t+1} &= r(a_t - c_t) + y_{t+1} \\ y_{t+1} &= (1 - \varphi)\bar{y} + \varphi y_t + \sigma \varepsilon_{t+1}, \end{aligned}$$

where $\{\varepsilon_t\} \sim \text{NID}(0, 1)$. We express this as a linear-quadratic control problem using c_t as

the control and $(1, a_t, y_t)$ as the state. The relevant matrices are

$$\begin{bmatrix} Q & S^\top \\ S & R \end{bmatrix} = \begin{bmatrix} 1 & -\gamma & 0 & 0 \\ -\gamma & \gamma^2 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

$$A = \begin{bmatrix} 1 & 0 & 0 \\ (1-\varphi)\bar{y} & r & \varphi \\ (1-\varphi)\bar{y} & 0 & \varphi \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ -r \\ 0 \end{bmatrix}, \quad C = \begin{bmatrix} 0 \\ \sigma \\ \sigma \end{bmatrix}.$$

We set $\beta = 0.95$, $r = 1/\beta$, $\gamma = 2$, $\bar{y} = 1$, $\varphi = 0.8$, and $\sigma = 0.25$. For the optimal control problem, the decision rule is

$$c_t = 0.7917 + 0.0500a_t + 0.1583y_t.$$

For the robust control problem with $\theta = 2$ (or the risk-sensitive control problem with $\alpha = 1/2\theta = 0.25$), the decision rule is

$$c_t = 0.7547 + 0.0515a_t + 0.1632y_t.$$

The impact of robustness is to reduce the intercept (precautionary saving) and increase the responsiveness to a and y . Why? The anticipated distortion is

$$w_t = -0.1557 + 0.0064a_t + 0.0204y_t,$$

making the actual and distorted dynamics

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0.2 & 1.0526 & 0.8 \\ 0.2 & 0 & 0.8 \end{bmatrix}, \quad A - CF_2 = \begin{bmatrix} 1 & 0 & 0 \\ -0.1247 & 1.0661 & 0.8425 \\ -0.1247 & 0.0134 & 0.8425 \end{bmatrix}.$$

The distorted dynamics are pessimistic (the income intercept changes from 0.2 to -0.1247) and more persistent (the maximal eigenvalue increases from 1.0526 to 1.1086). The latter calls for more aggressive responses to movements in a and y . [Adapted from Hansen, Sargent, and Tallarini (1999) and Hansen, Sargent, and Wang (2002).]

Beyond LQG

You might conjecture (as we did) that the equivalence of risk-sensitive and robust control hinges critically on the linear-quadratic-gaussian structure. It doesn't. The critical functional forms are the exponential certainty equivalent and the entropy constraint. With these two ingredients, the objective functions of risk-sensitive and robust control are the same.

We demonstrate the equivalence of risk-sensitive and robust control objective functions in a finite-state setting where the math is relatively simple. Consider an environment with conditional probabilities $p(z'|z)$. Since z is immaterial in what follows, we ignore it from here on. In a typical dynamic programming problem, the Bellman equation includes the term $EJ = \sum_{z'} p(z')J(z')$. A robust control problem has a similar term based on transformed probabilities $\hat{p}(z')$ whose values are limited by an entropy penalty:

$$\hat{J} = \min_{\{\hat{p}(z')\}} \sum_{z'} \hat{p}(z')J(z') + \theta \left\{ \sum_{z'} \hat{p}(z') \log[\hat{p}(z')/p(z')] \right\} + \lambda \left(\sum_{z'} \hat{p}(z') - 1 \right).$$

If $\hat{p}(z') = p(z')$, this is simply EJ . The new elements are the minimization with respect to \hat{p} (the defining feature of robust control), the entropy penalty on the choice of \hat{p} (the standard functional form), and the constraint that the transformed probabilities sum to one. For each $\hat{p}(z')$, the first-order condition for the minimization is

$$J(z') + \theta \{ \log[\hat{p}(z')/p(z')] + 1 \} + \lambda = 0. \quad (33)$$

If we multiply by $\hat{p}(z')$ and sum over z' , we get $\hat{J} + \theta + \lambda = 0$, which we use to eliminate λ below. Each first-order condition implies

$$\hat{p}(z') = p(z')e^{-[J(z')+\theta+\lambda]/\theta} = p(z')e^{-J(z')/\theta+\hat{J}/\theta}.$$

If we sum over z' and take logs, we get

$$\hat{J} = -\theta \log \left(\sum_{z'} p(z') \exp[-J(z')/\theta] \right),$$

our old friend the exponential certainty equivalent with risk aversion parameter $\alpha = \theta^{-1}$. If we place \hat{J} in its Bellman equation context, we've shown that robust control is equivalent (even outside the LQG class) to maximizing Kreps-Porteus utility with an exponential certainty equivalent. The log in the entropy constraint of robust control reappears in the exponential certainty equivalent. An open question is whether there's a similar relationship between Kreps-Porteus preferences with (say) a power certainty equivalent and a power-like alternative to the entropy constraint.

Discussion: interpreting parameters

Risk-sensitive and robust control raise a number of estimation issues, some we've seen, and some we haven't. Risk-sensitive control is based on a special case of Kreps-Porteus preferences and therefore leads to the same identification issues we faced in the previous section: we need variation over time in the conditional distribution of next period's state to distinguish time and risk parameters.

Robust control raises new issues. Risk-sensitive and robust control lead to the same decision rules, so we might regard them as equivalent. But they're based on different preferences and therefore lead to different interpretations of parameters. While risk-sensitive control suggests a risk averse agent, robust control suggests an agent who is uncertain about the model that generated the data. In practice, the two can be quite different. One difference is plausibility: We may find an agent with substantial model uncertainty (small θ) more plausible than one with enormous risk aversion (large α). Similarly, if we find that a model estimated for Argentina suggests greater risk aversion than one estimated for the US, we might prefer to attribute the difference to model uncertainty. Hansen and Sargent (2004, ch 8) have developed a methodology for calibrating model uncertainty ("error detection probabilities") that gives the robust-control interpretation some depth. Another difference crops up in comparisons across policy regimes: the two models can differ substantially if we consider policy experiments that change the amount of model uncertainty.

6 Ambiguity

In Sections 3 and 4, agents know the probabilities they face, and with enough regularity and repetition an econometrician can estimate them. Here we consider preferences when the consequences of our choices are uncertain or ambiguous. It's not difficult to think of such situations: what are the odds that China revalues this year by more than 10%, that the equity premium is less than 3%, or that productivity shocks account for more than half of the variance of US output growth? We might infer probabilities from history or market prices, but it's a stretch to say that we know (or can find out) these probabilities, even though they may affect some of our decisions. One line of attack on this issue was suggested by Savage (1954): that people maximize expected utility using "personal" or "subjective" probabilities. In this case, we retain the analytical tractability of expected utility but lose the empirical convenience of preferences based on the same probabilities that generate outcomes (rational expectations). Another line of attack generalizes Savage: preferences are characterized by multiple probability distributions or "priors." We refer to such preferences as capturing ambiguity and ambiguity aversion, and explore two examples: Gilboa and Schmeidler's (1989) "max-min" expected utility for static environments and Epstein and Schneider's (2003) "recursive multiple priors" extension to dynamic environments. The central issues are dynamic consistency (something we need to address in dynamic settings) and identification (how do we distinguish agents with ambiguous preferences from those with expected utility?).

Static ambiguity

Ambiguity has a long history and an equally long list of terminology. Different varieties have been referred to as Knightian uncertainty, Choquet expected utility, and expected utility with non-additive (subjective) probabilities. Each of these terms refers to essentially the same preferences. Gilboa and Schmeidler (1989) provide a simple representation and an axiomatic basis for a preference model in which an agent entertains multiple probability models or priors. If the set of priors is Π , preferences are represented by the utility function

$$U(\{c(z)\}) = \min_{\pi \in \Pi} \sum_z \pi(z) u[c(z)] = \min_{\pi \in \Pi} E_{\pi} u(c). \quad (34)$$

Gilboa and Schmeidler refer to such preferences as “max-min” because agents maximize a utility function that has been minimized with respect to the probabilities π . We denote probabilities by π , rather than p , as a reminder that they are preference parameters. The defining feature is the set Π , which characterizes both ambiguity and ambiguity aversion. If Π has a single element, (34) reduces to Savage’s subjective expected utility.

Gilboa and Schmeidler’s max-min preferences incorporate aversion to ambiguity: agents dislike consequences with unknown odds. Consider an agent choosing among mutually exclusive assets in a three-state world. State 1 is pure risk: it occurs with probability $1/3$. State 2 is ambiguous: it occurs with probability $1/3 - \gamma$, with $-1/6 \leq \gamma \leq 1/6$. State 3 is also ambiguous and occurs with probability $1/3 + \gamma$. The agent’s probability distributions over γ define the set Π . We use the distributions $\pi_{\gamma}(\gamma = g) = 1$ for $-1/6 \leq g \leq 1/6$, which imply $(1/3, 1/3 - g, 1/3 + g)$ as elements of Π . These distributions over γ are *dogmatic* in the sense that each places probability one on a particular value. The approach also allows non-dogmatic priors, such as $\pi_{\gamma}(\gamma = -1/6) = \pi_{\gamma}(\gamma = 1/6) = 1/2$. In this setting, consider the agent’s valuation of three assets: A pays one in state 1, nothing otherwise; B pays one in state 2; and C pays one in state 3. How much is each asset worth on its own to a max-min agent? To emphasize the difference between risk and ambiguity, let $u(c) = c$. Using (34), we find that asset A is worth $1/3$ and assets B and C are each worth $1/6$. The agent is apparently averse to ambiguity in the sense that the ambiguous assets, B and C, are worth less than the unambiguous asset, A. In contrast, an expected utility agent would never value both B and C less than A.

Example 17 (portfolio choice and non-participation). We illustrate the impact of ambiguity on behavior with an ambiguous version of Example 6. An agent has max-min preferences with $u(c) = c^{\alpha}/\alpha$ and $\alpha = 0.5$. She invests fraction w of initial wealth a_0 in a risky asset with returns $(r(1) = \kappa_1 - \sigma, r(2) = \kappa_1 + \sigma)$, with $\sigma = 0.17$ and fraction $1 - w$ in a risk-free asset with return $r_0 = 1.01$ in both states. Previously we assumed the states were equally likely: $\pi(1) = \pi(2) = 1/2$. Here we let $\pi(1)$ take on any value in the interval $[0.4, 0.6]$ and set $\pi(2) = 1 - \pi(1)$. Two versions of this example illustrate different features of max-min preferences.

- Version 1: First-order risk aversion generates a non-participation result. With expected utility, agents are approximately neutral to fair bets. In a portfolio context, this means they'll buy a positive amount of an asset whose expected return is higher than the risk-free rate, and sell it short if the expected return is lower. They choose $w = 0$ only if the expected return is the same. With multiple priors, the agent chooses $w = 0$ for a range of values of κ_1 around the risk-free rate (the non-participation result). If we buy, state 1 is the worst state and the min sets $\pi(1) = 0.6$. To buy a positive amount of the risky asset, the first-order condition must be increasing at $w = 0$:

$$0.6(r_0)^{\alpha-1}(\kappa_1 - \sigma - r_0) + 0.4(r_0)^{\alpha-1}(\kappa_1 + \sigma - r_0) \geq 0,$$

which implies $\kappa_1 - r_0 \geq 0.2\sigma$ or $\kappa_1 \geq 1.01 + 0.2(0.17) = 1.044$. If we sell, state 2 is worst and the min sets $\pi(2) = 0.6$. The analogous first-order condition must be decreasing:

$$0.4(r_0)^{\alpha-1}(\kappa_1 - \sigma - r_0) + 0.6(r_0)^{\alpha-1}(\kappa_1 + \sigma - r_0) \leq 0,$$

which implies $\kappa_1 \leq r_0 - 0.2\sigma = 0.976$. For $0.976 \leq \kappa_1 \leq 1.044$ the agent neither buys nor sells.

- Version 2: Let $\kappa_1 = 1.07$. Then the mean return is high enough to induce the agent to buy the risky asset and state 1 is worst. The optimal portfolio is $w = 2.147$. In this two-state example, the result is identical to disappointment aversion with $\delta = 0.5$. With more states, this need not be the case.

[Adapted from Dow and Werlang (1993) and Routledge and Zin (2001).]

Dynamic ambiguity

Epstein and Schneider (2003) extend max-min preferences to dynamic settings, providing an axiomatic basis for

$$U_t = u_t + \beta \min_{\pi \in \Pi_t} E_\pi U_{t+1}, \tag{35}$$

where U_t is short-hand for utility starting at some date- t history z^t , u_t is utility at z^t , U_{t+1} refers to utilities starting with histories $z^{t+1} = (z^t, z_{t+1})$ stemming from z^t , Π_t is a set of one-period conditional probabilities $\pi(z_{t+1}|z^t)$, and E_π denotes the expectation computed from the prior π . Hayashi (2003) generalizes (35) to non-linear time aggregators: $U_t = V(u_t, \min_{\pi \in \Pi_t} E_\pi U_{t+1})$.

As in Section 4, the combination of time and risk raises a question of dynamic consistency: Can (35) be reconciled with some reasonable specification of date-zero max-min preferences? The answer is yes, but the argument is subtle. Consider the “dilation” example suggested by Seidenfeld and Wasserman (1993). The starting point is the event tree in Figure 1, to

which we add ambiguous probabilities. (We suggest you write them on the tree.) Date-one probabilities are $\pi(z_1 = 1) = \pi(z_1 = 2) = 1/2$; they are not ambiguous. Date-two (conditional) probabilities depend on z_1 and an autocorrelation parameter ρ , for which the agent has dogmatic priors on the values $+1$ and -1 . Listed from top to bottom in the figure, the conditional probabilities of the four date-two branches are $\pi(z_2 = 1|z_1 = 1) = \pi(z_2 = 2|z_1 = 2) = (1 + \rho)/2$ and $\pi(z_2 = 2|z_1 = 1) = \pi(z_2 = 1|z_1 = 2) = (1 - \rho)/2$. In words: the probabilities depend on whether z_1 and z_2 are the same or different and whether ρ is $+1$ or -1 .

With these probabilities, consider the value of an asset that pays one if $z_2 = 1$, zero otherwise. For convenience, let $u(c) = c$ and set $\beta = 1$. If the recursive and date-zero valuations of the asset differ, preferences are dynamically inconsistent. Consider recursive valuation. At node (A) in Figure 1, the value is $(1 + \rho)/2$. Minimizing with respect to ρ as suggested by (35) implies $\rho = -1$ and a value of zero. Similarly, the value at node (B) is also zero, this time based on $\rho = 1$. The value at date zero is therefore zero as well; there is no ambiguity, so the value is $(1/2)(0) + (1/2)(0) = 0$. Now consider a (naive) date-zero problem based on the two-period probabilities of the four possible two-period paths: $(1 + \rho)/4$, $(1 - \rho)/4$, $(1 - \rho)/4$, and $(1 + \rho)/4$. Ambiguity in these probabilities is again represented by ρ . Since the asset pays one if the first or third path occurs, its date-zero value is $(1 + \rho)/4 + (1 - \rho)/4 = 1/2$, which is not ambiguous. The date-zero value $(1/2)$ is clearly greater than the recursive value (0) , so preferences are dynamically inconsistent. The computational point: our recursive valuation allows ρ to differ across date-one nodes, while our date-zero valuation does not. The conceptual point: giving the agent access to date-one information increases the amount of information but also increases the amount of ambiguity, which reduces the value of the asset.

Any resolution of this dynamic inconsistency problem must modify either recursive or date-zero preferences. Epstein and Schneider propose the latter. They show that if we expand the set of date-zero probabilities in the right way, they lead to the same preferences as (35). In general, preferences depend on probabilities over complete paths, which in our example you might associate with the four terminal nodes in Figure 1. Epstein and Schneider’s “rectangularity condition” tells us to compute the set of probabilities recursively, one period at a time, starting at the end. At each step, we compute a set of probabilities for paths given our current history. In our example, the main effect of this approach is to eliminate any connection between the values of ρ at the two date-one nodes. The resulting date-zero probabilities take the form $(1 + \rho_1)/4$, $(1 - \rho_1)/4$, $(1 - \rho_2)/4$, and $(1 + \rho_2)/4$. The value of the asset is therefore $(1 + \rho_1)/4 + (1 - \rho_2)/4 = 1/2 + (\rho_1 - \rho_2)/4$. If we minimize with respect to both ρ_1 and ρ_2 , we set $\rho_1 = -1$ and $\rho_2 = +1$ and the value is zero, the same value we computed recursively. In short, expanding the date-zero set of probabilities in this way reconciles date-zero and recursive valuations and resolves the dynamic inconsistency problem.

A related example illustrates the Epstein-Schneider approach in a somewhat more complex environment that allows comparison to an alternative based on entropy constraints. The setting remains the event tree in Figure 1. Date-one probabilities are $\pi(z_1 = 1) = (1 + \delta)/2$ and $\pi(z_1 = 2) = (1 - \delta)/2$, with a dogmatic prior for any δ in the interval $[-\bar{\delta}, \bar{\delta}]$ and $0 \leq \bar{\delta} < 1$. Date-two probabilities remain $(1 + \rho)/2$, $(1 - \rho)/2$, $(1 - \rho)/2$, and $(1 + \rho)/2$, but we restrict ρ to the interval $[-\bar{\rho}, \bar{\rho}]$ for $0 \leq \bar{\rho} \leq 1$. An asset has date-two payoffs of (from top to bottom in the tree) $1 + \varepsilon$, ε , 1 , and 0 , where $\varepsilon \geq 0$. The Seidenfeld-Wasserman example is a special case with $\bar{\delta} = \varepsilon = 0$ and $\bar{\rho} = 1$. The addition of ε to the payoffs introduces a concern for first-period ambiguity. Consider four approaches to the problem of valuing the asset:

- *Naive date-zero approach.* The four branches have date-zero probabilities of $[(1 + \delta)(1 + \rho)/4, (1 + \delta)(1 - \rho)/4, (1 - \delta)(1 - \rho)/4, (1 - \delta)(1 + \rho)/4]$. Conditional on δ and ρ , the asset is worth $(1 + \varepsilon + \delta + \delta\rho)/2$. If we minimize with respect to both δ and ρ , the value is $(1 + \varepsilon + \bar{\delta}\varepsilon - \bar{\delta}\bar{\rho})/2$. If $\bar{\delta} = \bar{\rho} = 1/2$ and $\varepsilon = 1$, the value is $9/8$.
- *Recursive approach.* We work our way through the tree, starting at the end, applying (35) as we go. At node (A), the value of the asset is $(1 - \rho)/2$. If we minimize with respect to ρ , the value is $\min_{\rho} (1 + \rho)/2 + \varepsilon = (1 - \bar{\rho})/2 + \varepsilon$ (set $\rho = -\bar{\rho}$). At (B), the value is $\min_{\rho} (1 - \rho)/2 = (1 - \bar{\rho})/2$ (set $\rho = \bar{\rho}$). At the initial node, the value is $(1 - \bar{\rho})/2 + (1 + \delta)\varepsilon/2$. Minimizing with respect to δ gives us $(1 - \bar{\rho})/2 + (1 - \bar{\delta})\varepsilon/2 = [1 - \bar{\rho} + (1 - \bar{\delta})\varepsilon]/2$ (set $\delta = -\bar{\delta}$). This is smaller than the date-zero valuation, which implicitly forced us to choose the same value of ρ at (A) and (B). If $\bar{\delta} = \bar{\rho} = 1/2$ and $\varepsilon = 1$, the value is $1/2$.
- *Rectangular approach (“sophisticated date-zero”).* As in the dilation example, we allow ρ to differ between the two date-one nodes, giving us two-period probabilities of $[(1 + \delta)(1 + \rho_1)/4, (1 + \delta)(1 - \rho_1)/4, (1 - \delta)(1 - \rho_2)/4, (1 - \delta)(1 + \rho_2)/4]$. Conditional on δ , ρ_1 , and ρ_2 , the asset is worth $(1 + \delta)(1 + \rho_1)(1 + \varepsilon)/4 + (1 + \delta)(1 - \rho_1)\varepsilon/4 + (1 - \delta)(1 - \rho_2)/4$. Minimizing with respect to the parameters gives us the same value as the recursive approach.
- *Entropy approach.* In this context, entropy is simply a way of describing the set Π : an entropy constraint places limits on (δ, ρ_1, ρ_2) that correspond to limits on the conditional probabilities at each node. We compute entropy at each node from equation (23) using $(1/2, 1/2)$ as the reference probabilities. The date-one entropy of probabilities following the initial node is

$$I_1(\delta) = (1/2)[(1 + \delta) \log(1 + \delta) + (1 - \delta) \log(1 - \delta)].$$

Note that $I_1(0) = 0$, $I_1(\delta) = I_1(-\delta) \geq 0$, and $dI_1/d\delta = (1/2) \log[(1 + \delta)/(1 - \delta)]$.

Similarly, the date-two entropy for the node following $z_1 = i$ is

$$I_{2i}(\rho_i) = (1/2)[(1 + \rho_i) \log(1 + \rho_i) + (1 - \rho_i) \log(1 - \rho_i)],$$

which has the same functional form as I_1 . The overall two-period entropy constraint is

$$I_1(\delta) + [(1 + \delta)/2]I_{21}(\rho_1) + [(1 - \delta)/2]I_{22}(\rho_2) \leq \bar{I} \quad (36)$$

for some number $\bar{I} > 0$ (a preference parameter). Our problem is to choose (δ, ρ_1, ρ_2) to minimize the value of the asset subject to the entropy constraint. What's new is the ability to shift ambiguity across periods implicit in the tradeoff between first- and second-period entropy.

We solve this problem recursively. To do this, it's helpful to break the constraint into pieces: $I_1(\delta) \leq \bar{I}_1$ and, for each i , $I_{2i}(\rho_i) \leq \bar{I}_2 \equiv \bar{I} - \bar{I}_1$. they are equivalent to the single entropy constraint (36) if the multipliers on the individual constraints are equal. In the first period, we choose not only the value of δ that satisfies the date-one entropy constraint but how much entropy to use now (\bar{I}_1) and how much to save for the second period ($\bar{I}_2 = \bar{I} - \bar{I}_1$). The solution is the allocation of entropy that equates the multipliers. Given a choice \bar{I}_1 , we solve the date-two problems. At node A, the entropy-constrained valuation problem is to choose ρ_1 to minimize $\varepsilon + (1 + \rho_1)/2$ subject to the entropy constraint $I_{21}(\rho_1) \leq \bar{I}_2$. If θ is the multiplier on the constraint, the first-order condition is

$$1/2 + (\theta/2) \log[(1 + \rho_1)/(1 - \rho_1)] = 0.$$

As with rectangularity, we set $\rho_1 < 0$ to reduce the probability of the good state ($z_2 = 1$). We're going to reverse engineer this and determine the constraint associated with setting $\rho_1 = -1/2$, the number we used earlier. With this value, entropy is $\bar{I}_2 = 0.1308$ and the first-order condition implies $\theta = 0.9102$. The value of the asset at this node is therefore $\varepsilon + 1/4$. At node B, if $\bar{I}_2 = 0.1308$ a similar calculation implies $\rho_2 = 1/2$, $\theta = 0.9102$, and an asset value of $1/4$. Note, in particular, that ρ is set differently at the two nodes, just as it is under rectangularity. At the initial node, we now have the problem of choosing δ to minimize $[(1 + \delta)/2](\varepsilon + 1/4) + [(1 - \delta)/2](1/4) = 1/4 + [(1 + \delta)/2]\varepsilon$ subject to the entropy constraint $I_1(\delta) \leq \bar{I}_1$. The first-order condition is

$$\varepsilon/2 + (\theta/2) \log[(1 + \delta)/(1 - \delta)] = 0.$$

If $\varepsilon = 1$ and $\bar{I} = 0.2616$, the solution includes $\delta = 1/2$, $\bar{I}_1 = 0.1308$, and $\theta = 0.9102$. As with rectangularity, the value is $1/2$. However, for other values of ε entropy will be reallocated between the two periods in the way that has the largest adverse impact on utility. If $0 \leq \varepsilon < 1$, the risk between nodes (A) and (B) is relatively small and entropy will be shifted from period one to period two, increasing $|\rho_i|$ and decreasing

$|\delta|$. If $\varepsilon > 1$, first-period risk is more important and entropy will be shifted from period two to period one, with the opposite effect. This reallocation of ambiguity has no counterpart with rectangularity, where the range of probabilities (and associated parameters) is unrelated to other aspects of the problem (the payoffs, for example, represented here by ε).

We have, then, four approaches to the same problem, each of which has arguments in its favor. The naive date-zero approach, which is in the spirit of Chamberlain's (2000) econometric application, allows less impact of ambiguity than the other approaches, but does so in a way that remains consistent with a version of date-zero max-min preferences. It does, however, place some importance on the choice of date zero: if we reoptimize in the future, we would typically compute different decisions. The recursive approach, without rectangularity, might be justified as a game among agents at different dates. The same idea has been widely used in other contexts (the next section, for example). The rectangular approach is a clever way to reconcile date-zero and recursive approaches and leads to a natural recursive extension of Gilboa-Schmeidler. One puzzling consequence is that it can induce ambiguity in events that have none to begin with. (Recall the joint probability of the first and third paths in the dilation example, which is $1/2$ regardless of ρ .) The apparent puzzle is resolved if we realize that the date-zero rectangular set does not represent date-zero ambiguity, it represents the date-zero probabilities needed to anticipate preferences over future ambiguity. Epstein and Schneider (2003, p 16) put it this way: "there is an important conceptual distinction between the set of probability laws that the decision maker views as possible ... and the set of priors that is part of the representation of preference." Finally, the entropy approach allows the "min" to operate not only within a period but across periods, as entropy and ambiguity are allocated over time to have the greatest impact. This violates conditional independence for reasons similar to habits (Example 13), but seems consistent with the spirit of pessimism captured by the "min" in (34).

Example 18 (precautionary saving). Ambiguity generates precautionary saving through pessimism: pessimistic forecasts of future income reduce current consumption and raise current saving. The magnitude depends on the degree of ambiguity. We illustrate the result with a two-period example that shares several features with its robust control counterpart (Example 16). The endowment is y_0 at date zero and $y_1 \sim N(\kappa_1 + \gamma, \kappa_2)$ at date one. The parameter γ governs ambiguity: $\gamma^2 \leq g^2$ for some positive number g . An agent has utility function

$$U = u(c_0) + \beta \min_{\gamma} E u(c_1)$$

with $u(c) = -\exp(-\alpha c)$. The budget constraint is $c_1 = y_1 + r(y_0 - c_0)$. If we substitute this into the objective function and compute the expectation, we find

$$U = -\exp(-\alpha c_0) - \beta \min_{\gamma} \exp[-\alpha r(y_0 - c_0) - \alpha(\kappa_1 + \gamma) + \alpha^2 \kappa_2 / 2].$$

The minimization implies $\gamma = -g$ (pessimism). The first-order condition for c_0 then implies

$$c_0 = \log(\beta r)/[\alpha(r-1)] + (ry_0 + \kappa_1)/(1+r) - \alpha\kappa_2/[2(1+r)] - g/(1+r).$$

Here the second term is permanent income, the third is risk and risk aversion, and the fourth the impact of ambiguity. [Adapted from Miao (2003).]

Example 19 (sharing ambiguity). If agents have identical homothetic preferences, optimal allocations are proportional: the ratio of date-state consumption by one agent is proportional to that of every other agent. In stationary settings, we often say (with some abuse of the language) that consumptions are perfectly correlated. Observations of individuals and countries, however, exhibit lower correlations, suggesting a risk-sharing puzzle. One line attack on this puzzle is to let agents have different preferences. In international economics, for example, we might let the two countries consume different goods. A variation on this theme is to let preferences differ in their degree of ambiguity. In particular, suppose agents have less ambiguity over their own endowment than over other agents' endowments. A symmetric two-period, two-agent example shows how this might work. Agent i has utility function

$$U^i = \log c_0^i + \beta \min_{\pi^i \in \Pi^i} \sum_z \pi^i(z) \log c_1^i(z),$$

for $i = 1, 2$. In period zero, each is endowed with one unit of the common good. In period one, there are four states (z) with the following endowments (y^i) and probabilities (π^i):

z	y^1	y^2	π^1	π^2	c^1	c^2
1	2	2	$1/4 - \gamma_1$	$1/4 - \gamma_2$	2	2
2	2	1	$1/4 + \gamma_1$	$1/4 - \gamma_2$	9/4	3/4
3	1	2	$1/4 - \gamma_1$	$1/4 + \gamma_2$	3/4	9/4
4	1	1	$1/4 + \gamma_1$	$1/4 + \gamma_2$	1	1

Each set Π^i is constructed from dogmatic priors over values for γ_i between $-1/8$ and $1/8$. Note that each agent is ambiguous about the other agent's endowment, but not her own. Without ambiguity ($\gamma_i = 0$), the symmetric optimal allocation consists of one-half the aggregate endowment in all states: perfect correlation across the date-one states. With ambiguity, agent i chooses the value of γ_i that minimizes her utility, $\gamma_i = 1/8$. Since agent 1 applies a lower probability ($1/8$) to state 3 than agent 2 ($3/8$), she gets a proportionally smaller share of the aggregate endowment in that state. The resulting allocations are listed in the table and show imperfect correlation across agents. The amount of ambiguity in this case is so large that in states 2 and 3 the agent with the larger endowment consumes even more than her endowment. A simple decentralization makes the same point. Suppose agents at date zero trade claims to the endowments of the two countries. How much would each invest in her own endowment, and how much in the other agent's endowment? If w is agent 1's investment in her own endowment, it satisfies

$$wy^1(z) + (1-w)y^2(z) = c^1(z)$$

for all states z . The solution in this case is $w = 5/4$: agent 1 exhibits extreme home bias in her portfolio. [Adapted from Alonso (2004) and Epstein (2001).]

Discussion: detecting ambiguity

Preferences based on subjective probabilities capture interesting features of behavior that other preferences cannot, but they raise challenging issues for quantitative applications. Consider subjective expected utility. If we allow the probabilities that enter preferences (π) to differ from those that generate the data (p), we can “explain” many things that are otherwise puzzling. The equity premium, for example, could result from agents placing lower probability on high-return states than the data generating process. It is precisely the lack of predictive content in such explanations that led us to rational expectations ($\pi = p$) in Sections 3 and 4.

Ambiguity provides a justification for systematically pessimistic probabilities — they’re the minimizing choice from a larger set — but raises two new issues. One is how to specify the larger set of probabilities or models. Hansen and Sargent (2004) propose choosing models that have have similar log-likelihood functions, much as we do in hypothesis tests. Differences between such models are presumably difficult to detect in finite data sets. Epstein and Schneider (2004) suggest nonstationary ambiguous models that are indistinguishable from a reference model, even in infinite samples. The other issue is observational equivalence: robust control and recursive multiple priors generate behavior that could have been generated by an expected utility agent, and possibly by a Kreps-Porteus agent as well. In some cases, the agent seems implausible, but in others not. Distinguishing between ambiguous and expected utility agents remains an active area of current research. The most ambitious example to date is Epstein and Schneider (2004), who note that ambiguous news has an unusual asymmetric affect on asset prices, since bad news affects the minimizing probability distribution but good news does not.

7 Inconsistency and temptation

Economists often tell stories about the hazards of temptation and the benefits of reducing our choice sets to avoid it. We eat too much junk food, we over-consume addictive substances, and we save too little. To counter these tendencies, we may put ourselves in situations where the range of choices limits our ability to make bad decisions. We go to restaurants that serve only healthy food, support laws that discourage or prohibit addictive substances, and sequester our wealth in housing and 401(k) accounts that are less easily used to finance current consumption. The outstanding questions are why we make such choices, what the relevant welfare criterion should be, and how we might detect the impact of temptation on

observed decisions.

Inconsistent preferences

The traditional approach was outlined in Example 4: dynamically inconsistent preferences. This line of research is motivated by experimental studies, which suggest that subjects discount the immediate future more rapidly than the distant future. Common practice is to approximate this pattern of discounting with the “quasi-geometric” or “quasi-hyperbolic” scheme: $1, \delta\beta, \delta\beta^2, \delta\beta^3$, and so on, with $0 < \beta < 1$ and $0 < \delta \leq 1$. The critical parameter is δ : if $\delta < 1$, the discount factor between dates $t = 0$ and $t = 1$ (namely, $\delta\beta$) is less than the discount factor between dates $t = 1$ and $t = 2$ (β).

Let us say, then, that an agent’s utility from date t on is

$$\begin{aligned} U_t &= E_t \left[u(c_t) + \delta\beta u(c_{t+1}) + \delta\beta^2 u(c_{t+2}) + \delta\beta^3 u(c_{t+3}) + \dots \right] \\ &= u(c_t) + \delta\beta E_t \sum_{j=0}^{\infty} \beta^j u(c_{t+j+1}). \end{aligned}$$

The only difference from Example 4 is the introduction of uncertainty implicit in the conditional expectation E_t . The dynamic inconsistency of these preferences suggests two questions: With competing preferences across dates, what does such an agent do? And what preferences should we use for welfare analysis? We need an answer to the first question to derive the behavioral implications of inconsistent preferences, and an answer to the second to evaluate the benefits of policies that limit choice.

The consensus answer to the first question has become: treat the problem as a game with the agent at each date acting as a separate player. Each such player makes choices that maximize her utility, given the actions of other players (herself at other dates). There are many games like this, corresponding to different strategy spaces. We look at stationary Markov perfect equilibria, in which agents’ decisions are stationary functions of the current state for some natural definition of the state. Consider the classical consumption problem with budget constraint $a_{t+1} = r_{t+1}(a_t - c_t) + y_{t+1}$, where y and r are iid positive random variables, and a borrowing constraint $a \geq \underline{a}$ that we will ignore. Our objective is a stationary decision rule $c_t = h(a_t)$ that solves the game. With constant discounting ($\delta = 1$), the problem is the solution to the dynamic programming problem summarized by the Bellman equation,

$$J(a) = \max_c u(c) + \beta EJ[r'(a - c) + y'].$$

Under standard conditions, J exists, and is unique, continuous, concave, and differentiable. Given such a J , the maximization leads to a continuous stationary decision rule $c = h(a)$.

The equilibrium of a game can be qualitatively different. A stationary decision rule can

be derived with a “future value function,”

$$J(a) = u(c^*) + \beta EJ[r'(a - c^*) + y'], \quad (37)$$

where

$$c^* = \arg \max_c u(c) + \delta \beta EJ[r'(a - c) + y']. \quad (38)$$

Note the difference: when $\delta < 1$, the relation that generates J is different from that generating the choice of c . As a result, the decision rule need not be unique or continuous; see Harris and Laibson (2001), Krusell and Smith (2004), and Morris and Postlewaite (1997). For all of these reasons, there can be no general observational equivalence result between constant and quasi-geometric discounting. Nevertheless, the solutions are similar in some common examples.

Example 20 (consumption and saving). Consider the classical saving problem with log utility ($u(c) = \log c$), budget constraint $a_{t+1} = r_{t+1}(a_t - c_t)$ (no labor income), and log-normal return ($\{\log r_t\} \sim \text{NID}(\mu, \sigma^2)$). With quasi-geometric discounting, we compute the stationary decision rule from

$$\begin{aligned} J(a) &= \log c^* + \beta EJ[r'(a - c^*)] \\ c^* &= \arg \max_c \log c + \delta \beta EJ[r'(a - c)]. \end{aligned}$$

We find the solution by guessing that the value function has the form $J(a) = A + B \log a$. The first-order condition from the maximization implies $c = (1 + \delta \beta B)^{-1} a$. Substituting into the recursion for J , we find $B = (1 - \beta)^{-1}$ and

$$c = \left(\frac{1 - \beta}{1 - \beta + \delta \beta} \right) a = h(a).$$

Compare this decision rule with two others:

- **Constant discounting.** The decision rule with constant discounting is $c = (1 - \beta)a$ (set $\delta = 1$). Note that with quasi-geometric discounting the agent consumes more, but not as much more as an agent with constant discount factor $\beta\delta$. The latter is the result of strategic interactions between agents. The data- t agent would like to save a fraction $\delta\beta$ of her assets at date t , and a larger fraction β at future dates $t + n > t$. She knows, however, that future agents will make the same calculation and choose saving rates less than β . To induce future agents to consume more (absolutely, not as a fraction of wealth), she saves more than $\delta\beta$ today. Note, too, that her consumption behavior is observationally equivalent to an agent with constant discount factor

$$\hat{\beta} = \frac{\delta\beta}{1 - \beta + \delta\beta} < \beta.$$

A similar result holds for power utility, and suggests that despite the difficulties noted earlier, constant and quasi-geometric discounting may be difficult to distinguish in practice.

- **Commitment.** Suppose the date- t agent can choose decision rules for future agents. Since the agent's discount factor between any future dates $t+n > t$ and $t+n+1$ is β , she chooses the decision rules $c_t = (1 - \delta\beta)a_t$ for date t and $c_{t+n} = (1 - \beta)a_{t+n}$ for all future dates $t+n > t$. This allocation maximizes the utility of the date- t agent, so in that sense “commitment” (limiting our future choice sets) is good. But it's not clear that date- t preferences are the appropriate welfare criterion.

[Adapted from Barro (1999), İmrohoğlu, İmrohoğlu, and Joines (2003), and Phelps and Pollack (1968).]

Example 21 (asset pricing). A similar example can be used to illustrate the role of quasi-geometric discounting on asset prices. The first step is to derive the appropriate Euler equation for (37,38). Define the “current value function” by

$$L(a) = \max_c u(c) + \delta\beta EJ[r'(a-c) + y']. \quad (39)$$

The first-order and envelope conditions are

$$\begin{aligned} u_1(c) &= \delta\beta E[J_1(a')r'] \\ L_1(a) &= \delta\beta E[J_1(a')r'], \end{aligned}$$

implying the familiar $L_1(a) = u_1(c)$. In the constant discounting case, $J(a) = L(a)$ and we're almost done. With quasi-geometric discounting, we need another method to express J_1 in terms of u_1 . Note that if we multiply (37) by δ and subtract from (39) we can relate J to L and u : $\delta J(a) = L(a) - (1 - \delta)u(c)$. Differentiating yields

$$\delta J_1(a) = L_1(a) - (1 - \delta)u_1(c)h_1(a).$$

If we multiply by β and substitute into the first-order condition, we get the Euler equation,

$$u(c_t) = E_t \{ \beta [1 - (1 - \delta)h_1(a_{t+1})] u_1(c_{t+1}) r_{t+1} \}.$$

This relation is a curious object: it depends not only on the current agent's decision problem, but (through h) on the strategic interactions among agents. The primary result is to decrease the effective discount factor, and raise mean asset returns, relative to the standard model. [Adapted from Harris and Laibson (2003), Krusell, Kuruşçu, and Smith (2002), and Luttmer and Mariotti (2003).]

Temptation

Many of us have been in situations in which we felt we had “too many choices.” (Zabar’s Delicatessen and Beer World have that effect on us.) In traditional decision theory this statement is nonsense: extra choices are at best neutral, because you can always decide not to use them. Gul and Pesendorfer (2001) give the phrase meaning: they develop preferences in which adding inferior choices (“temptations”) can leave you worse off. Among its features: utility can depend on the set of choices, as well as the action taken; temptation (in the sense of inferior choices) can reduce utility; and commitment (in the sense of restricting the choice set) can increase utility. We describe their theory in a static setting, then go on to explore dynamic extensions, including some that resemble quasi-geometric discounting.

Let us compare two sets of choices, A and B . In traditional decision theory, the utility of a set of possible choices is the utility of its best element. If the best element of A is at least as good as the best element of B , then we would say A is weakly preferred to B : $A \succeq B$ in standard notation. Suppose we allow choice over the potentially larger set $A \cup B$. The traditional approach would tell us that this cannot have an impact on our decision or utility: if $A \succeq B$, then we are indifferent between A and $A \cup B$. Gul and Pesendorfer suggest a *set betweenness* condition that allows inferior choices to affect our preference ordering:

$$A \succeq B \quad \text{implies} \quad A \succeq A \cup B \succeq B.$$

The traditional answer is one extreme (namely, $A \sim A \cup B$), but set betweenness also allows inferior choices B to reduce our utility ($A \succ A \cup B$). We say in such cases that B is a temptation.

Adding set betweenness to an otherwise traditional theory, Gul and Pesendorfer show that preferences can be represented by a utility function of the form:

$$u(A) = \max_{c \in A} [v(c) + w(c)] - \max_{c \in A} w(c). \quad (40)$$

Note that preferences are defined for the choice set A ; we have abandoned the traditional separation between preferences and opportunities. To see how this works, compare the choices $c^* = \arg \max_{c \in A} [v(c) + w(c)]$ and $c^{**} = \arg \max_{c \in A} w(c)$ for some choice set A . If $c^* = c^{**}$, then v and w agree on A and preferences are effectively governed by v (the w terms cancel). If not, then w acts as a temptation function.

Example 22 (consumption and saving). A clever use of temptations reproduces quasi-geometric discounting. Let

$$\begin{aligned} v(c_1, c_2) &= u(c_1) + \beta u(c_2) \\ w(c_1, c_2) &= \gamma [u(c_1) + \delta \beta u(c_2)], \end{aligned}$$

with $0 < \delta < 1$ and $\gamma \geq 0$ (intensity of temptation). The budget constraint has two parts: $c_1 + k_2 = rk_1$ and $c_2 = rk_2$, with k_1 given, which defines A . The agent solves

$$\max_{c_1, c_2 \in A} [(1 + \gamma)u(c_1) + (1 + \gamma\delta)\beta u(c_2)] - \max_{c_1, c_2 \in A} \gamma [u(c_1) + \delta\beta u(c_2)].$$

The first max delivers the first-order condition

$$1 = \left(\frac{1 + \gamma\delta}{1 + \gamma} \right) \beta \frac{u_1(c_2)}{u_1(c_1)} r.$$

The difference from the standard model lies in the first term. The two extremes are $\gamma = 0$ (which gives us the standard no-temptation model) and $\gamma = \infty$ (which gives us an irresistible temptation and the quasi-geometric discount factor $\delta\beta$). Since the term is decreasing in γ , greater temptation raises first-period consumption. [Adapted from Krusell, Kuruşçu, and Smith (2003).]

Gul and Pesendorfer (2002, 2004) and Krusell, Kuruşçu, and Smith (2003) have extended the temptation approach to quasi-geometric discounting to infinite-horizon settings. We illustrate the idea with a non-stochastic version of the consumption problem. Krusell, Kuruşçu, and Smith suggest an approach summarized by the “Bellman equation”

$$J(a) = \max_c \{u(c) + \beta J[r(a - c)] + L[r(a - c)]\} - \max_c L[r(a - c)],$$

where

$$L(a) = \gamma \{u(c^*) + \delta\beta L[r(a - c^*)]\}$$

serves as a temptation function and $c^* = \arg \max_c u(c) + \beta J[r(a - c)] + L[r(a - c)]$. Gul and Pesendorfer suggest the special case $\delta = 0$. The Krusell-Kuruşçu-Smith version reproduces the first-order conditions and decision rules generated by the Markov perfect equilibrium for quasi-geometric discounting. The Gul-Pesendorfer version avoids some of the mathematical oddities associated with the former. Each suggests an answer to the welfare question.

Discussion: detecting inconsistency and temptation

The difficulty of estimating the parameters of models based on quasi-geometric discounting is that the decision rules often look like those from traditional models with constant discounting. In some cases, they’re identical. One way to distinguish between them is to look for evidence of commitment. Agents with inconsistent preferences or temptations will typically be willing to pay something to restrict their future choice sets. In models with constant discounting there is no such incentive, so commitment devices provide a natural way to tell the two approaches apart. Laibson, Repetto, and Tobacman (1998, 2004) apply this logic and find that the combination of illiquid asset positions (pensions, 401(k) accounts) and high-interest liabilities (credit card debt) generates sharp differences between the two

models and precise estimates of the discount parameters ($\delta = 0.70$, $\beta = 0.96$, annual). With constant discounting, borrowing at high rates and investing at (on average) lower rates are incompatible.

The focus on commitment devices seems right to us, both for quasi-geometric discounting and for temptations more generally. There are, however, some outstanding questions, most of them noted by Kocherlakota (2001). One is whether tax-sheltered savings have other explanations (lower taxes, for example). If 401(k)'s were a pure commitment device, we might expect people to pay more for them and receive less, but this doesn't seem to be the case: sheltered and unsheltered investment vehicles have pretty much the same returns. Similarly, if commitment is valuable, why would an agent hold both liquid (uncommitted) and illiquid (committed) assets? The former would seem to undercut the bite of the latter. Finally, what is the likely market response to the conflicting demands of commitment and temptation? Will the market supply commitment devices or ways to avoid them? Is credit card debt designed to satisfy agents' desire to undo past commitments? Does it lower welfare? Perhaps future work will resolve these questions.

8 Questions, answers, and final thoughts

We have described a wide range of exotic preferences and applied them to a number of classic macroeconomic problems. Are there any general lessons we might draw from this effort? We organize a discussion around specific questions.

Why model preferences rather than behavior?

Preferences play two critical roles in economic models. The first is that they 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 modelled behavior directly, we would also have to model how it adjusted to changing circumstances. The second role of preferences is to evaluate the welfare effects of changing policies or circumstances. Without preferences, it's not clear how we should distinguish good policies from bad. In our view, this is a major accomplishment of the "temptation" interpretation of quasi-geometric discounting: it suggests a clear welfare criterion.

Are exotic preferences simply an excuse for free parameters?

Theoretical economists think nothing of modifying the environments faced by their agents. Aggregate and individual risk, length of life, information structures, enforcement technologies, and productivity shocks are all fair game. However, many economists seem to believe that modifying preferences is cheating — that we will be able to explain anything (and hence nothing) if we allow ourselves enough freedom over preferences. We would argue instead that we have restricted ourselves to an extremely limited model of preferences for no better reasons than habit and convenience. Many of the weaknesses of expected utility, for example, have been obvious since the 1950s. We now have the tools to move beyond additive preferences in several directions, why not use them?

Equally important, the axiomatic foundations that underlie the preferences described above impose a great deal of discipline on their structure. We have let these foundations go largely without mention, but they nevertheless restrict the kinds of flexibility we've considered. Chew-Dekel risk preferences, for example, are more flexible than expected utility, but far less flexible than general preferences over state-contingent claims. One consequence: exotic preferences have led to some progress on the many empirical puzzles that plague macroeconomics and finance, but they have yet to resolve them.

Some exotic preferences make greater — or at least novel — demands on the data than we are used to. Kreps-Porteus and Epstein-Zin preferences, for example, require time-dependence of risk to identify separate time and risk preference parameters. Robust control comes with an entropy toolkit for setting plausible values of the robustness parameter, but comparisons across environments may be needed to distinguish robust from risk-sensitive control. Applications of temptation preferences to problems with quasi-geometric discounting rely heavily (entirely?) on observed implications of commitment devices, about which there is some difference of opinion. In short, exotic preferences raise new empirical issues that deserve open and honest debate. We see no reason, however, to rule out departures from additive utility before we start.

Are exotic preferences “behavioral”?

Many of the preferences we've described were motivated by discrepancies between observed behavior and the predictions of the additive preference model. In that sense, they have a “behaviorial” basis. They are also well-defined neoclassical preference orderings. For that reason, we think our approach falls more naturally into neoclassical economics than into the behavioral sciences.

We regard this as both a strength and a weakness. On the one hand, the strong theoretical foundations for exotic preferences allow us to use all the tools of neoclassical economics,

particularly optimization and welfare analysis. On the other hand, these tools ignore aspects of human behavior stressed in other social sciences, particularly sociology and social psychology. Kreps (2000) and (especially) Simon (1959) are among the many economists who have argued that something of this sort is needed to account for some aspects of behavior. We have some sympathy for this argument, but it's not what we've done in this paper.

Are there interesting preferences we've missed?

If you've gotten this far, you may feel that we can't possibly have left anything out. But it's not true. We barely scratched the surface of robust control, ambiguity, hyperbolic discounting, and temptation. If you'd like to know more, you might start with the papers listed in Appendix A. We also ignored some lines of work altogether. Among them are:

- **Incomplete preferences.** Some of the leading decision theorists suggest that the most troubling axiom underlying expected utility is not the infamous "independence axiom" but the more common assumption of completeness: that all possible choices can be compared. Schmeidler (1989), for example, argues that the critical role of the independence axiom is to extend preferences from choices that seem obvious to those that do not — that it delivers completeness. For this and other reasons, there is a long history of work on incomplete preferences. Notable applications in macroeconomics and finance include Bewley (1986) and Kraus and Sagi (2002, 2004).
- **Flexibility, commitment, and self-control.** Kreps (1979) describes environments in which agents prefer to maintain flexibility over future choices, just as agents with temptations prefer commitment. Amador, Werning, and Angeletos (2003) characterize optimal allocation rules when you put the two together. Ameriks, Caplin, Leahy, and Tyler (2004) quantify self-control with survey evidence and relate it to individual financial decisions. Benhabib and Bisin (2004) take a cognitive approach to a similar problem in which agents choose between automatic processes, which are subject to temptations, and control processes, which are not.
- **Social utility.** Experimental research suggests that preferences often depend on comparisons with others; see, for example, Blount (1995) and Rabin (1998). Abel (1990) and Galí (1994) are well-known applications to asset pricing.
- **Other psychological approaches.** Bénabou and Tirole (2002) model self-confidence. Bernheim and Rangel (2002) build a cognitive model and apply it to addiction. Brunnermeier and Parker (2003) propose a model of subjective beliefs in which agents balance the utility benefits of optimism and the utility cost of inferior decisions. Caplin and Leahy (2002) introduce anxiety into an otherwise standard dynamic choice framework and explore its implications for portfolio choice and the equity premium.

We find all of this work interesting, but leave a serious assessment of it to others.

Have we wasted your time (and ours)?

It's too late, of course, but you might ask yourself whether this has been worth the effort. To paraphrase Monty Python, "Have we deliberately wasted your time?" We hope not. We would guess that additive preferences will continue to be the industry standard in macroeconomics, finance, and other fields. Their tight structure leads to strong and clear predictions, which is generally a virtue. But we would also guess that exotic preferences will become more common, particularly in quantitative work. Who knows, they may even lose their claim to being "exotic."

We think several varieties of exotic preferences have already proved themselves. Applications of Kreps-Porteus and Epstein-Zin preferences to asset pricing, precautionary saving, and risk-sharing are good examples. While these preferences have not solved all of our problems, they have become a frequent source of insight. Their ease of use in econometric work is another mark in their favor.

The preferences described in the last three sections are closer to the current frontiers of research, but we are optimistic that they, too, will lead to deeper understanding of economic behavior. Certainly robust control, recursive multiple priors, and temptation are significant additions to our repertoire. They also raise new questions about identification and estimation. Multiple priors is a good example. When the probabilities affecting an agent's preferences are not characterized simply by the probabilities generating the data, we need to parameterize the agent's uncertainty and describe how it evolves through time. We also need to explore ways to distinguish such agents from those with expected utility or Kreps-Porteus preferences. Temptation is another. As a profession, we need to clarify the features of data that identify the parameters of temptation functions, as well as the kinds of temptations that are most useful in applied work. None of these tasks are simple, but we think the progress of the last decade gives us reason to hope for more.

Let's get to work!

A Reader's guide

We have intentionally favored application over theory, but if you'd like to know more about the theoretical underpinnings of exotic preferences, we recommend the following:

Section 2. Koopmans (1960) is the classic reference. Koopmans (1986) lays out the relevant theory of independent preferences. Lucas and Stokey (1984) approach the problem from what now seems like a more natural direction: they start with an aggregator function, while Koopmans derives one. Epstein and Hynes (1983) propose a convenient functional form and work through an extensive set of examples.

Section 3. Kreps (1988) is far and away the best reference we've seen for the theory underlying the various approaches to expected utility. Starmer (2000) gives a less technical overview of the theory and discusses both empirical anomalies and modifications of the theory designed to deal with them. Brandenburger (2002) describes some quite different approaches to probability assessments that have been used in game theory.

Section 4. Our two favorite theory references on dynamic choice in risky environments are Kreps and Porteus (1978) and Johnsen and Donaldson (1985). Epstein and Zin (1989) describe the technical issues involved in specifying stationary recursive preferences and explain the roles of the parameters of the constant elasticity version.

Section 5. Our primary reference is Hansen and Sargent's (2004) monograph on robust control; we recommend chapters 2 (overview), 5 (static robust control), 6 (dynamic robust control), and 9 and 17 (entropy constraints). Whittle (1990) is an introduction to linear-quadratic robust control for engineers. Hansen and Sargent (1997) introduce risk-sensitive control in chapters 9 and 15. Gianonni (2002), Maenhout (2004), Onatski and Williams (2003), and Van Nieuwerburgh (2001) are interesting applications.

Section 6. The essential references are Gilboa and Schmeidler (1989) and Epstein and Schneider (2003). Among the other papers we have found useful are Ahn (2003), Casadesu-Masanell, Klibanoff, and Ozdenoren (2000), Chamberlain (2000), Epstein and Schneider (2002, 2004), Gilboa and Schmeidler (1993), Hayashi (2003), Klibanoff, Marinacci, and Mukerji (2003), Sagi (2003), Schmeidler (1989), and Wang (2003).

Section 7. The relevant theory is summarized in Gul and Pesendorfer (2004), Harris and Laibson (2003), and Krusell, Kuruşçu, and Smith (2004). DeJong and Ripoll (2003), Esteban, Miyagawa, and Shum (2004), and Krusell, Kuruşçu, and Smith (2002) are interesting applications.

B Integral formulas

A number of our examples lead to normal-exponential integrals, most commonly as expectations of log-normal random variables or exponential certainty equivalents of normal random variables. The following definitions and formulas are used in the paper.

Standard normal density and distribution functions. If $x \sim N(0, 1)$, its density is $f(x) = (2\pi)^{-1/2}e^{-x^2/2}$. Note that f is symmetric: $f(x) = f(-x)$. The distribution function is $\Phi(x) \equiv \int_{-\infty}^x f(u)du$. By symmetry, $\int_x^\infty f(u)du = 1 - \Phi(x) = \Phi(-x)$.

Integrals of “ $e^{a+bx}f(x)$.” We come across integrals of this form in Section 3, when we compute certainty equivalents for log-normal risks, and Section 4, when we consider the exponential certainty equivalent of a linear value function (Weil’s model of precautionary saving). Their evaluation follows from a change of variables. Consider the integral

$$\int_{-\infty}^{x^*} e^{a+bx}f(x)dx = (2\pi)^{-1/2} \int_{-\infty}^{x^*} e^{a+bx-x^2/2}dx.$$

We solve this by completing the square: expressing the exponent as $a+bx-x^2/2 = d-y^2/2$, where d is a scalar and $y = fx - g$ is a linear transformation of x . We find $y = x - b$ ($f = 1$, $g = b$) and $d = a + b^2/2$, so the integral is

$$(2\pi)^{-1/2} \int_{-\infty}^{x^*} e^{a+bx-x^2/2}dx = e^{a+b^2/2} \int_{-\infty}^{x^*-b} f(y)dy = e^{a+b^2/2} \Phi(x^* - b). \quad (41)$$

A common special case has an infinite upper limit of integration:

$$E(e^{a+bx}) = (2\pi)^{-1/2} \int_{-\infty}^{\infty} e^{a+bx-x^2/2}dx = e^{a+b^2/2}. \quad (42)$$

Example: Let $\log y = \mu + \sigma x$; then $Ey = E(e^{\log y}) = E(e^{\mu+\sigma x}) = e^{\mu+\sigma^2/2}$.

Integrals of “ $e^{a+bx+cx^2}f(x)$.” Integrals like this arise in Section 5 in risk-sensitive control with a quadratic objective. Consider

$$\int_{-\infty}^{\infty} e^{a+bx+cx^2}f(x)dx = (2\pi)^{-1/2} \int_{-\infty}^{\infty} e^{a+bx-(1-2c)x^2/2}dx.$$

We assume $1-2c > 0$; otherwise the integral diverges. We solve by the same method: express the exponent as $a+bx-(1-2c)x^2/2 = d-y^2/2$ for some $y = fx - g$. We find $f = (1-2c)^{1/2}$, $g = b/(1-2c)^{1/2}$, and $d = a + b^2/(1-2c)$, so that $y = (1-2c)^{1/2}x - b/(1-2c)^{1/2}$. The integral becomes

$$\begin{aligned} \int_{-\infty}^{\infty} e^{a+bx+cx^2}f(x)dx &= (1-2c)^{-1/2}e^{a+b^2/[2(1-2c)]} \int_{-\infty}^{\infty} f(y)dy \\ &= (1-2c)^{-1/2}e^{a+b^2/[2(1-2c)]}. \end{aligned} \quad (43)$$

References

- Abel, Andrew B., 1990, "Asset pricing under habit formation and keeping up with the Joneses," *American Economic Review (Papers and Proceedings)* 80, 38-42.
- Ahn, David, 2003, "Ambiguity without a state space," manuscript, September.
- Alonso, Irasema, 2004, "Ambiguity in a two-country world," manuscript, March.
- Amador, Manuel, Iván Werning, and George-Marios Angeletos, 2003, "Commitment vs. flexibility," NBER Working Paper 10151, December.
- Ameriks, John, Andrew Caplin, John Leahy, and Tom Tyler, 2004, "Measuring self control," NBER Working Paper 10514, May.
- Anderson, Evan W., 2004, "The dynamics of risk-sensitive allocations," manuscript, February.
- Anderson, Evan W., Lars Peter Hansen, Ellen R. McGrattan, and Thomas J. Sargent, 1996, "Mechanics of forming and estimating dynamic linear economies," in H. Amman, D.A. Kendrick, and J. Rust, eds., *Handbook of Computational Economics (Volume 1)*, Amsterdam: Elsevier.
- Bansal, Ravi, and Amir Yaron, 2004, "Risks for the long run: A potential resolution of asset pricing puzzles," *Journal of Finance*, forthcoming in the August issue.
- Barro, Robert J., 1999, "Ramsey meets Laibson in the neoclassical growth model," *Quarterly Journal of Economics* 114, 1125-1152.
- Bénabou, Roland, and Jean Tirole, 2002, "Self-confidence and personal motivation," *Quarterly Journal of Economics* 115, 871-915.
- Benhabib, Jess, and Alberto Bisin, 2004, "Modeling internal commitment mechanisms and self-control: A neuroeconomics approach to consumption-saving decisions," manuscript, February.
- Bernheim, B. Douglas, and Antonio Rangel, 2002, "Addiction and cue-conditioned cognitive processes," NBER Working Paper 9329, November.
- Bewley, Truman F., 1986, "Knightian decision theory: Part I," Cowles Foundation Discussion Paper No 807, November.
- Blount, Sally, 1995, "When social outcomes aren't fair," *Organizational Behavior and Human Decision Processes* 63 (August), 131-144.
- Brandenburger, Adam, 2002, "The power of paradox," manuscript, December.
- Brunnermeier, Markus K., and Jonathan A. Parker, 2003, "Optimal expectations," manuscript, June.
- Caplin, Andrew, and John Leahy, 2001, "Psychological expected utility theory and anticipatory feelings," *Quarterly Journal of Economics* 116, 55-79.

- Casadesus-Masanell, Ramon, Peter Klibanoff, and Emre Ozdenoren, 2000, "Maxmin expected utility over Savage acts with a set of priors," *Journal of Economic Theory* 92, 35-65.
- Chamberlain, Gary, 2000, "Econometric applications of maxmin expected utility," *Journal of Applied Econometrics* 15, 625-644.
- 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 under uncertainty: Weakening the independence axiom," *Journal of Economic Theory* 40, 304-318.
- Dolmas, Jim, and Mark A. Wynne, 1998, "Elastic capital supply and the effects of fiscal policy," *Economic Inquiry* 36, 553-574.
- Dow, James, and Sérgio Ribeiro da Costa Werlang, 1992, "Uncertainty aversion, risk aversion, and the optimal choice of portfolio," *Econometrica* 60, 197-204.
- Epstein, Larry G., 2001, "Sharing ambiguity," *American Economic Review (Papers and Proceedings)* 91, 45-50.
- 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, 2002, "Learning under ambiguity," manuscript, September.
- Epstein, Larry G., and Martin Schneider, 2003, "Recursive multiple-priors," *Journal of Economic Theory* 113, 1-31.
- Epstein, Larry G., and Martin Schneider, 2004, "Ambiguity, information quality, and asset pricing," manuscript, May.
- 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, 1990, "'First-order' risk aversion and the equity premium puzzle," *Journal of Monetary Economics* 26, 387-407.
- Epstein, Larry G., and Stanley E. Zin, 1991, "Substitution, risk aversion, and the temporal behavior of consumption and asset returns: An empirical analysis," *Journal of Political Economy* 99, 263-286.
- Epstein, Larry G., and Stanley E. Zin, 2001, "The independence axiom and asset returns," *Journal of Empirical Finance* 8, 537-572.
- Esteban, Susanna, Eiichi Miyagawa, and Matthew Shum, 2004, "Nonlinear pricing with


- self-control preferences,” manuscript, March.
- Farmer, Roger E.A., 1990, “Rince preferences,” *Quarterly Journal of Economics* 105, 43-60.
- Galí, Jordi, 1994, “Keeping up with the Joneses: Consumption externalities, portfolio choice, and asset prices,” *Journal of Money, Credit, and Banking* 26, 1-8.
- Gertler, Mark, 1999, “Government debt and social security in a life-cycle economy,” *Carnegie-Rochester Conference Series on Public Policy* 50, 61-110.
- Giannoni, Marc P., 2002, “Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model,” *Macroeconomic Dynamics* 6, 111-144.
- Gilboa, Itzhak, and David Schmeidler, 1989, “Maxmin expected utility with non-unique priors,” *Journal of Mathematical Economics* 18, 141-153.
- 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.
- Gul, Faruk, and Wolfgang Pesendorfer, 2001, “Temptation and self-control,” *Econometrica* 69, 1403-1435.
- Gul, Faruk, and Wolfgang Pesendorfer, 2002, “Self-control and the theory of consumption,” manuscript, November; forthcoming *Econometrica*.
- Gul, Faruk, and Wolfgang Pesendorfer, 2004, “Self-control, revealed preference, and consumption choice,” *Review of Economic Dynamics* 7, 243-264.
- Hansen, Lars Peter, 1982, “Large sample properties of generalized method of moments estimators,” *Econometrica* 50, 1029-1054.
- Hansen, Lars Peter, and Thomas J. Sargent, 1997, *Recursive Models of Dynamic Linear Economies* manuscript, December.
- Hansen, Lars Peter, and Thomas J. Sargent, 2004, *Misspecification in Recursive Macroeconomic Theory* manuscript, January.
- Hansen, Lars Peter, Thomas J. Sargent, and Thomas Tallarini, 1999, “Robust permanent income and pricing,” *Review of Economic Studies* 66, 873-907.
- Hansen, Lars Peter, Thomas J. Sargent, and Neng E. Wang, 2002, “Robust permanent income and pricing with filtering,” *Macroeconomic Dynamics* 6, 40-84.
- Harris, Christopher, and David Laibson, 2001, “Dynamic choices of hyperbolic consumers,” *Econometrica* 69, 935-957.
- Harris, Christopher, and David Laibson, 2003, “Hyperbolic discounting and consumption,” *Advances in Economics and Econometrics*, Cambridge UK: Cambridge University Press.
- Hayashi, Takashi, 2003, “Intertemporal substitution, risk aversion, and ambiguity aversion,” manuscript, December.

- İmrohoğlu, Ayşe, Selahattin İmrohoğlu, and Douglas H. Joines, 2003, "Time-inconsistent preferences and social security," *Quarterly Journal of Economics* 118, 745-784.
- Johnsen, Thore H., and John B. Donaldson, 1985, "The structure of intertemporal preferences under uncertainty and time consistent plans," *Econometrica* 53, 1451-1458.
- Kan, Rui, 1995, "Structure of Pareto optima when agents have stochastic recursive preferences," *Journal of Economic Theory* 64, 626-631.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji, 2003, "A smooth model of decision making under ambiguity," manuscript, April.
- Kocherlakota, Narayana R., 1990, "Disentangling the coefficient of relative risk aversion from the elasticity of intertemporal substitution," *Journal of Finance* 45, 175-190.
- Kocherlakota, Naryana R., 2001, "Looking for evidence of time-inconsistent preferences in asset-market data," *Federal Reserve Bank of Minneapolis Quarterly Review* 25 (Summer), 13-24.
- Koopmans, Tjalling C., 1960, "Stationary ordinal utility and impatience," *Econometrica* 28, 287-309.
- Koopmans, Tjalling C., 1986, "Representation of preference orderings with independent components of consumption," and "Representation of preference orderings over time," in C.B. McGuire and Roy Radner, eds., *Decision and Organization (Second Edition)*, Minneapolis, MN: University of Minnesota Press.
- Kraus, Alan, and Jacob S. Sagi, 2002, "Intertemporal preference for flexibility and risky choice," manuscript, November.
- Kraus, Alan, and Jacob S. Sagi, 2004, "Asset pricing with unforeseen contingencies," manuscript, March.
- Kreps, David M., 1979, "A representation theorem for 'preference for flexibility'," *Econometrica* 47, 565-578.
- Kreps, David M., 1988, *Notes on the Theory of Choice*, Boulder, CO: Westview Press.
- Kreps, David M., 2000, "Beliefs and tastes: Confessions of an economist," remarks made at the AAU Centennial Meeting, October.
- Kreps, David M., and Evan L. Porteus, 1978, "Temporal resolution of uncertainty and dynamic choice theory," *Econometrica* 46, 185-200.
- Krusell, Per, Bruhanettin Kuruşçu, and Anthony A. Smith, 2001, "Temptation and taxation," manuscript, June.
- Krusell, Per, Bruhanettin Kuruşçu, and Anthony A. Smith, 2002, "Time orientation and asset prices," *Journal of Monetary Economics* 49, 107-135.
- Krusell, Per, and Anthony A. Smith, 2004, "Consumption-saving decisions with quasi-geometric discounting," *Econometrica*, forthcoming.

- Kydland, Finn E., and Edward C. Prescott, 1977, "Rules rather than discretion: The inconsistency of optimal plans," *Journal of Political Economy* 85, 473-492.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman, 1998, "Self-control and saving for retirement," *Brookings Papers on Economic Activity* (1), 91-196.
- Laibson, David, Andrea Repetto, and Jeremy Tobacman, 2004, "Estimating discount functions from lifecycle consumption choices," manuscript, January.
- Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter, 2003, "The declining equity premium: What role does macroeconomic risk play?" manuscript, September.
- Ljungqvist, Lars, and Thomas J. Sargent, 2000, *Recursive Macroeconomic Theory*, Cambridge MA: MIT Press.
- Lucas, Robert E., and Nancy L. Stokey, 1984, "Optimal growth with many consumers," *Journal of Economic Theory* 32, 139-171.
- Luttmer, Erzo, and Thomas Mariotti, 2003, "Subjective discounting in an exchange economy," *Journal of Political Economy* 111, 959-989.
- Maenhout, Pascal J., 2004, "Robust portfolio rules and asset pricing," *Review of Financial Studies*, forthcoming.
- Mehra, Rajnish, and Edward C. Prescott, 1985, "The equity premium: A puzzle," *Journal of Monetary Economics* 15, 145-161.
- Miao, Jianjun, 2003, "Consumption and saving under Knightian uncertainty," manuscript, December.
- Morris, Stephen, and Andrew Postlewaite, 1997, "Observational implications of non-exponential discounting," manuscript, November.
- Onatski, Alexei, and Noah Williams, 2003, "Modeling model uncertainty," manuscript, April.
- Phelps, E.S., and R.A. Pollack, 1968, "On second-best national saving and game-equilibrium growth," *Review of Economic Studies* 35, 185-199.
- Rabin, Matthew, 1998, "Economics and psychology," *Journal of Economic Literature* 36, 11-46.
- Routledge, Bryan R., and Stanley E. Zin, 2001, "Model uncertainty and liquidity," NBER Working Paper 8683, December.
- Routledge, Bryan R., and Stanley E. Zin, 2003, "Generalized disappointment aversion and asset prices," NBER Working Paper 10107, November.
- Sagi, Jacob S., 2003, "Anchored preference relations," manuscript, August.
- Savage, L.J., 1954, *The Foundations of Statistics*, New York: Wiley.
- Schmeidler, David, 1989, "Subjective probability and expected utility without additivity," *Econometrica* 57, 571-587.


- Seidenfeld, Teddy, and Larry Wasserman, 1993, "Dilation for sets of probabilities," *Annals of Statistics* 21, 1139-1154.
- Shi, Shouyong, 1994, "Weakly nonseparable preferences and distortionary taxes in a small open economy," *International Economic Review* 35, 411-428.
- Simon, Herbert A., 1959, "Theories of decision-making in economic and behavioral science," *American Economic Review* 49, 253-283.
- Starmer, Chris, 2000, "Developments in non-expected utility theory," *Journal of Economic Literature* 38, 332-382.
- Tallarini, Thomas D., 2000, "Risk-sensitive real business cycles," *Journal of Monetary Economics* 45, 507-532.
- 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.
- Van Nieuwerburgh, Stijn, 2001, "Robustness and optimal contracts," manuscript, July.
- Wang, Susheng, 1993, "The local recoverability of risk aversion and intertemporal substitution," *Journal of Economic Theory* 59, 333-363.
- Wang, Tan, 2003, "Conditional preferences and updating," *Journal of Economic Theory* 108, 286-321.
- Weil, Philippe, 1989, "The equity premium puzzle and the risk-free rate puzzle," *Journal of Monetary Economics* 24, 401-421.
- Weil, Philippe, 1990, "Nonexpected utility in economics," *Quarterly Journal of Economics* 105, 29-42.
- Weil, Philippe, 1993, "Precautionary savings and the permanent income hypothesis," *Review of Economic Studies* 60, 367-383.
- Whittle, Peter, 1990, *Risk-Sensitive Optimal Control*, New York: Wiley.
- Yaari, Menachem E., 1987, "The dual theory of choice under risk," *Econometrica* 55, 95-115.

Figure 1
A representative event tree




This event tree illustrates how uncertainty might evolve through time. Time moves from left to right, starting at date $t = 0$. At each date t , an event z_t occurs. In this example, z_t is drawn from the two-element set $\mathcal{Z} = \{1, 2\}$. Each node is marked by a box and can be identified from the path of events that leads to it, which we refer to as a history and denote by $z^t \equiv (z_0, \dots, z_t)$, starting with an arbitrary initial node z_0 . Thus the upper right node follows two up branches, $z_1 = 1$ and $z_2 = 1$, and is denoted $z^2 = (z_0, 1, 1)$. The set \mathcal{Z}^2 of all possible 2-period histories is therefore $\{(z_0, 1, 1), (z_0, 1, 2), (z_0, 2, 1), (z_0, 2, 2)\}$, illustrated by the far right “column” of nodes.

Figure 2
State-space indifference curves with Chew-Dekel preferences




The figure contains indifference curves for three members of the Chew-Dekel class of risk preferences. In each case, the axes are consumption in state 1 and state 2 and states are equally likely. The risk preferences are expected utility (upper left, $\alpha = 0.5$), weighted utility (upper right, bold line, $\gamma = -0.25$), and disappointment aversion (lower left, bold line, $\delta = 0.5$). For weighted utility and disappointment aversion, expected utility is pictured with a lighter line for comparison. For disappointment aversion, the indifference curve is the upper envelope of two indifference curves, each based on a different set of transformed probabilities. The extensions of these two curves are shown as dashed lines. The lower right figure has all three together: expected utility (dashed line), weighted utility (solid line), and disappointment aversion (dash-dotted line). Note that disappointment aversion is more sharply convex than weighted utility near the 45-degree line (the effect of first-order risk aversion), but less convex far away from it.

Figure 3
Risk and risk premiums with Chew-Dekel preferences



The figure illustrates the relation between risk and risk premiums discussed in Example 5 for three members of the Chew-Dekel class of risk preferences. The preferences are: expected utility (dashed line), weighted utility (solid line), and disappointment aversion (dash-dotted line). The point is the nonlinearity of disappointment aversion: the ratio of the risk premium to risk is greater for small risks than large ones. Parameter values are the same as Figure 2.

Figure 4
Transformed probabilities: Entropy and disappointment aversion



The figure illustrates two sets of transformed probabilities described in Example 15: one set generated by an entropy constraint and the other by disappointment aversion. The bold triangle is the three-state probability simplex. The “+” in the middle represents the reference probabilities: $p(1) = p(2) = p(3) = 1/3$. The area inside the egg-shaped contour represents transformed probabilities with entropy less than 0.1. The dashed line represents probabilities implied by disappointment aversion with δ between 0 to 1.5.