Microeconomic Foundations I:
Choice and Competitive Markets

Student's Guide

Chapter 6: Utility for Money

Summary of the Chapter

Chapter 5 discussed the basic models of choice under uncertainty in some generality. In this chapter, we specialize to the expected-utility model (where probabilities are given "objectively") and we further specialize to the case where the prizes are one dimensional and interpreted as "money." We always have a consumer who is an expected-utility maximizer for simple lotteries with money prizes and whose utility function is generally denoted by U. Within this context, three general tasks are undertaken:

- In the spirit of Chapter 2, we look in Section 6.1 for "natural" properties of the consumer's preferences over lotteries and for corresponding properties of the utility function U. Monotonicity and continuity of U are quickly disposed of (some discussion of the continuity of U requires some advanced mathematics; it helps if you know about the topology of weak convergence on probability measures). The heart of this section, though, involves risk aversion and the comparison of the levels of risk aversion of two different consumers or one consumer but at different levels of wealth. Risk aversion corresponds to concavity of U; measures of the level of risk aversion (for smooth enough utility functions) concern the so-called coefficient of risk aversion, $-U''/U'$. We conclude with a brief discussion (all proofs left as exercises) of two partial orders that are commonly applied to lotteries, first- and second-order stochastic dominance.
- Presumably, our consumer desires money because it will enable her to buy consumption goods. So preferences over lotteries with money prizes are *derived* preferences, derived from more primitive preferences over lotteries of consumption bundles. Section 6.2 discusses this and goes on to answer the question: What properties on preferences over lotteries of consumption bundles yield commonly as-

sumed properties for preferences on lotteries over income? (In setting up this discussion, some commentary is provided about how economists model "derived" or "induced" preferences. This is the sort of thing that is easy to skim through; this is one place where you might want to slow down. In particular, if you think the points made about this are all obvious, try Problem 6.7.)

- The most basic common applications of these theories concern consumer demand for insurance and for risky assets. Section 6.3 presents the most basic result in the theory of insurance; you are asked to tackle the basics of demand for risky assets in the problems.

The material in this chapter is a bare introduction to aspects of *The Economics of Uncertainty*, about which whole courses are sometimes offered. (And this doesn't touch at all on *The Economics of Information*, which is a natural partner with the economics of uncertainty, but which requires tools and concepts that we only will reach in Volume II.) As for applications, applied choice under uncertainty is the basis of much of the theory of financial markets, about which other entire books have been written and courses are taught. So your instructor is likely to have his or her own ideas about which material in this chapter and what sorts of supplements are worth your while.

Solutions to Starred Problems

■ 6.1. (a) Suppose U is strictly increasing. Then $\delta_x \succ \delta_y$ if and only if $U(x) > U(y)$ if and only if $x > y$. Note in this chain of implications that the expected utility of δ_x is $U(x)$, and that of δ_y is $U(y)$. Conversely, suppose $\delta_x \succ \delta_y$ if and only if $x > y$. By the representation, we know that $\delta_x \succ \delta_y$ if and only if $U(x) > U(y)$, so we have $x > y$ if and only if $U(x) > U(y)$, and U is strictly increasing.

(b) Suppose U is concave. Then for all x_1, x_2, and $a \in [0,1]$, $U(ax_1 + (1-a)x_2) \geq aU(x_1) + (1-a)U(x_2)$. Consider the induction hypothesis, that for all $n \leq N$, lists of real numbers $x_1, \ldots, x_n$ in the domain of U, and lists of nonnegative numbers $a_1, \ldots a_n$ such that $\sum_{i=1}^{n} a_i = 1$, $U(a_1x_1 + \ldots + a_nx_n) \geq a_1U(x_1) + \ldots + a_nU(x_n)$. The induction hypothesis is trivially true for $N = 1$, and it is true for $N = 2$ by the definition of a concave function. So now assume it is true for some N; we will demonstrate it is true for $N+1$ and therefore show it is true for all positive integer N.

To show it is true for $N+1$, take a list of real numbers $x_1, \ldots, x_n$ and a list of nonnegative real numbers $a_1, \ldots, a_n$, where $n \leq N+1$, and such that $\sum a_i = 1$. If $n \leq N$, we know that $U(a_1x_1 + \ldots + a_nx_n) \geq a_1U(x_1) + \ldots + a_nU(x_n)$ by the induction hypothesis, so we need only consider the case where $n = N+1$. And if any of the a_i are zero, then we essentially are in the case $n \leq N$, so we can assume that $a_{N+1} > 0$, at least. For $i = 1, \ldots, N$, let $b_i = a_i/(1-a_{N+1})$. We know that $\sum_{i=1}^{N} b_i = [\sum_{i=1}^{N} a_i]/(1-a_{N+1}) = 1$, and so by the induction hypothesis,

$$U(b_1x_1 + \ldots + b_Nx_n) \geq \sum_{i=1}^{N} b_iU(x_i).$$

But by concavity of U,

$$\begin{aligned} U\big(a_1x_1 + \ldots + a_{N+1}x_{N+1}\big) &= U\big((1 - a_{N+1})(b_1x_1 + \ldots + b_Nx_N) + a_{N+1}x_{N+1}\big) \\ &\geq (1 - a_{N+1})U\big(b_1x_1 + \ldots b_Nx_N\big) + a_{N+1}U\big(x_{N+1}\big) \\ &\geq (1 - a_{N+1})\left[\sum_{i=1}^{N} b_iU(x_i)\right] + a_{N+1}U(x_{N+1}) \\ &= a_1U(x_1) + \ldots + a_NU(x_N) + a_{N+1}U(x_{N+1}). \end{aligned}$$

This establishes the induction hypothesis. (Please note that nowhere in this proof did we use the fact that the x_is are real numbers. This works for any concave function on any convex domain.)

(c) We suppose that X is an interval of R and U is continuous. Take any simple lottery π. Since π has finite support, there are $\overline{x}$ and $\underline{x}$ in the support of π such that $U(\overline{x}) \geq U(x) \geq U(\underline{x})$ for all x in the support of π. Therefore, $U(\overline{x}) \geq \sum \pi(x)U(x) \geq U(\underline{x})$, where the sum is over the support of π. But then, by the intermediate value theorem, there is some x^* in the interval from $\overline{x}$ to $\underline{x}$ (we don't know which is greater as a real number, so I can't write the interval as $[\underline{x}, \overline{x}]$, for instance) such that $U(x^*) = \sum \pi(x)U(x)$, which makes x^* a certainty equivalent of π. (Note that this would work for arbitrary X, as long as it is simply connected.)

Suppose U is strictly increasing. Suppose by way of contradiction that some lottery π had more than one certainty equivalent; say, x_1 and x_2 are both certainty equivalents, with $x_1 \neq x_2$. Since the real line is simply ordered, if $x_1 \neq x_2$, then either $x_1 > x_2$ or $x_2 > x_1$. But then, in the first case, $U(x_1) > U(x_2)$, and in the second, $U(x_2) > U(x_1)$. In either case, they are unequal, so they both cannot equal expected utility under π, so they cannot both be certainty equivalents of π. (Note that this depends very much on the fact that the domain of X is the real line or, at least, is simply ordered with respect to the order under which U is strictly increasing.)

■ 6.3. The expected utility for a lottery with (net) prizes $x^0 + \epsilon$ and $x^0 - \epsilon$, each with probability 1/2, is $[U(x^0 + \epsilon) + U(x^0 - \epsilon)]/2$. Let δ be the risk premium for this lottery; in other words, the certainty equivalent of the lottery is its expected value, x^0, less δ. Therefore, δ is defined by the equation

$$U(x^0 - \delta) = \frac{U(x^0 + \epsilon) + U(x^0 - \epsilon)}{2}.$$

The Taylor's series expansion of the quantity on the right to the 2nd derivative is

$$\begin{aligned} & 0.5[U(x^0) + \epsilon U'(x^0) + (1/2)\epsilon^2 U''(x^0) + o(\epsilon^2)] + \\ & \qquad 0.5[U(x^0) - \epsilon U'(x^0) + (1/2)(-\epsilon)^2 U''(x^0) + o(\epsilon^2)] \\ = & \quad U(x^0) + (1/2)\epsilon^2 U''(x^0) + o(\epsilon^2). \end{aligned}$$

And the exact Taylor's series expansion of the term on the left, to the first derivative, is

$$U(x^0) - \delta U'(x^0) + \delta^2 U''(x^0 - \alpha\delta)/2,$$

where $\alpha \in [0,1]$. So we have

$$U(x^0) - \delta U'(x^0) + \delta^2 U''(x^0 - \alpha\delta)/2 = U(x^0) + \frac{1}{2}\epsilon^2 U''(x^0) + o(\epsilon^2), \text{ or}$$

$$\delta U'(x^0) + \delta^2 U''(x^0 - \alpha\delta)/2 = \frac{1}{2}\epsilon^2 U''(x^0) + o(\epsilon^2).$$

Divide both sides of this equation by ϵ. As $\epsilon \to 0$, the term on the right has limit zero. And on the left, since $\delta > -\epsilon$ by the monotonicity of U, the second term goes to zero. Therefore, $\delta/\epsilon \to 0$ as $\epsilon \to 0$; that is, δ is $o(\epsilon)$. But then the remainder term is $o(\epsilon^2)$, so the equation is

$$-\delta\, U'(x^0) = \frac{1}{2}\epsilon^2 U''(x^0) + o(\epsilon^2),$$

which, if you divide both sides by $-U'(x^0)$, is the conclusion of the proposition.

Concerning Problems 6.5 and 6.6. The solutions to Problems 6.5 and 6.6 that I give here require a bit of sophistication concerning the language of random variables, their expectations, and their conditional expectations. And in the solution to Problem 6.6, to avoid the use of algebra that hides some relatively simple ideas, I resort to "picture proofs" and "proof by example" when it is expositionally convenient. (It is worth noting that for the hard part in all this—part c of Problem 6.6—I essentially sketch the clever "direct construction" due to Machina and Pratt, 1997.)

■ 6.5. (a) We assume that π and ρ are simple lotteries such that for every nondecreasing utility function U, the expected utility under π is at least as large as under ρ. If this is so, it is so in particular for the utility function U_x defined by

$$U_r(x) = \begin{cases} 1, & \text{if } x > r, \text{ and} \\ 0, & \text{if } x \le r. \end{cases}$$

But for this utility function U_r, the expected utility under π is the probability that π gives a prize strictly greater than r, and similarly for ρ. Therefore,

$$\sum_{x>r} \pi(x) \quad \geq \quad \sum_{x>r} \rho(x), \tag{$\star$}$$

where the value of $\pi(x)$ is implicitly zero if x is not in the support of π, and similarly for ρ. Since the sum of all probabilities for any lottery is 1, this means that

$$F_\pi(r) = \sum_{x \leq r} \pi(x) \quad \leq \quad \sum_{x \leq r} \rho(x) = F_\rho(r),$$

and this is moreover true for all r. This is what we are asked to show.

You may be unhappy that the utility function U_r used here is discontinuous. If so, redo the analysis with $U_{r,\epsilon}$ defined by $U_{r,\epsilon}(x) = 0$ if $x \leq r$, $= 1$ if $x \geq r + \epsilon$, and $= (x - r)/\epsilon$ for $x \in (r, r + \epsilon)$. If you apply the fact that $\pi \geq^1 \rho$ for this utility function and then, for fixed r, take the limit of the resulting inequalities as $\epsilon \to 0$, you will get the inequality $\star$, from which the result follows.

(b) The random variable $X = x_i$ if $U \in (F_\pi(x_{i-1}), F_\pi(x_i)]$, which has probability $F_\pi(x_i) - F_\pi(x_{i-1}) = \sum_{x \leq x_i} \pi(x) - \sum_{x \leq x_{i-1}} \pi(x) = \pi(x_i)$. The only exception to this is for the case $i = 1$, and by the rule given, the probability that $X = x_1$ is the probability that $U \leq F_\pi(x_1)$, which is $\pi(x_1)$.

(c) On a probability space on which a uniformly distributed random variable U is defined, let $X_F = F^{-1}(U)$ and let $X_G = G^{-1}(U)$. By part b, X_F has distribution F and X_G has distribution G. And since $F(r) \leq G(r)$ for all r, $F^{-1}(u) \geq G^{-1}(u)$ for all $u \in [0, 1]$, hence $X_F \geq X_G$ for all values of U, and $X_F - X_G \geq 0$ with probability one.

(d) If X_π is a random variable with distribution π, then the expected utility under π is the expectation of the random variable $U(x_\pi)$, for which we write $E[U(X_\pi)]$. And if X_π and X_ρ are (joint) random variables on some probability space such that $X_\pi - X_\rho \geq 0$ with probability one, then $X_\pi \geq X_\rho$ with probability one. Therefore, if U is a nondecreasing function, $U(X_\pi) \geq U(X_\rho)$ with probability one, which implies that $E[U(X_\pi)] \geq E[U(X_\rho)]$. Since this is true for any nondecreasing function, we have satisfied the conditions for the definition of $\pi \geq^1 \rho$ to hold.

▪ 6.6. Rather than write in full "cumulative distribution function," the abbreviation c.d.f. will be used.

(a) We suppose that $\pi \geq^2 \rho$, which means that expected utility computed for π is at least as large as expected utility computed for ρ, for any utility function U that is nondecreasing and concave. As suggested in the problem, we look at the expected utilities of π and ρ for the specific utility function $U_{r^0,r^1}(x)$ which equals $x - r^0$ for $x \leq r^1$ and equals $r^1 - r^0$ for $x \geq r^1$, where r^0 is any number less than all values in the supports of both π and ρ and r^1 is any number greater than r^0.

The expected utility of this utility function for π, which I'll denote by $E^\pi[U_{r^0,r^1}]$, is then

$$\sum_{x \in \text{supp}(\pi), x \leq r^1} (x - r^0)\pi(x) \quad + \quad (r^1 - r^0)(1 - F_\pi(r^1)),$$

and similarly for ρ. Because $\pi \geq^2 \rho$, we know that $E^\pi[U_{r^0,r^1}] \geq E^\rho[U_{r^0,r^1}]$, which implies that

$$(r^1 - r^0) - E^\pi[U_{r^0,r^1}] \leq (r^1 - r^0) - E^\rho[U_{r^0,r^1}].$$

But I assert that

$$(r^1 - r^0) - E^\pi[U_{r^0,r^1}] = \int_{r^0}^{r^1} F_\pi(r)dr, \qquad (\star\star)$$

and similarly for ρ, which (once I show this) gives part a. Why is $(\star\star)$ true? This is an application of integration by parts (which is how one would show this formally), most easily seen in a picture. See Figure G6.1. Panel a depicts (with dashed lines) the function F_π for a typical simple probability distribution π, writing the elements of the support of π as $\{x_1, x_2, x_3, x_4\}$ and with r^1 between x_2 and x_3. (You should mentally construct the picture for the cases where r^1 is in the support and where it lies above the elements in the support of π.) Panel b then shows, for this example, the terms that go into $E^\pi[U_{r^0,r^1}]$. There are two terms in the summation, the areas of the two more heavily shaded rectangles; the final term is the area of the more lightly shaded rectangle. Turning finally to panel c, you see a representation of $r^1 - r^0$, which is the area of the heavily bordered rectangle in both panels b and c less then shaded areas in panel b, or the shaded area in panel c. This is manifestly the integral on the right-hand side of $(\star\star)$.

(b) In this part of the problem, we are given two (joint) random variables X_π and X_ρ, where X_π has distribution π, X_ρ has distribution ρ, and $E[X_\rho - X_\pi | X_\pi] \leq 0$. The objective is to show that $\pi \geq^2 \rho$, which means that for all nondecreasing and concave utility functions U, the expected utility of U under π is greater or equal to the expected utility of U under ρ, or $E[U(X_\pi)] \geq E[U(X_\rho)]$. We verify this by showing that the conditional expectations, conditional on X_π bear this relation; that is

$$E[U(X_\rho)|X_\pi] \leq E[U(X_\pi)|X_\pi] = U(X_\pi);$$

once this is shown, the law of iterated conditional expectations gives the desired result. But since U is concave, we know that $E[U(X_\rho)|X_\pi] \leq U(E[X_\rho|X_\pi])$, and since U is

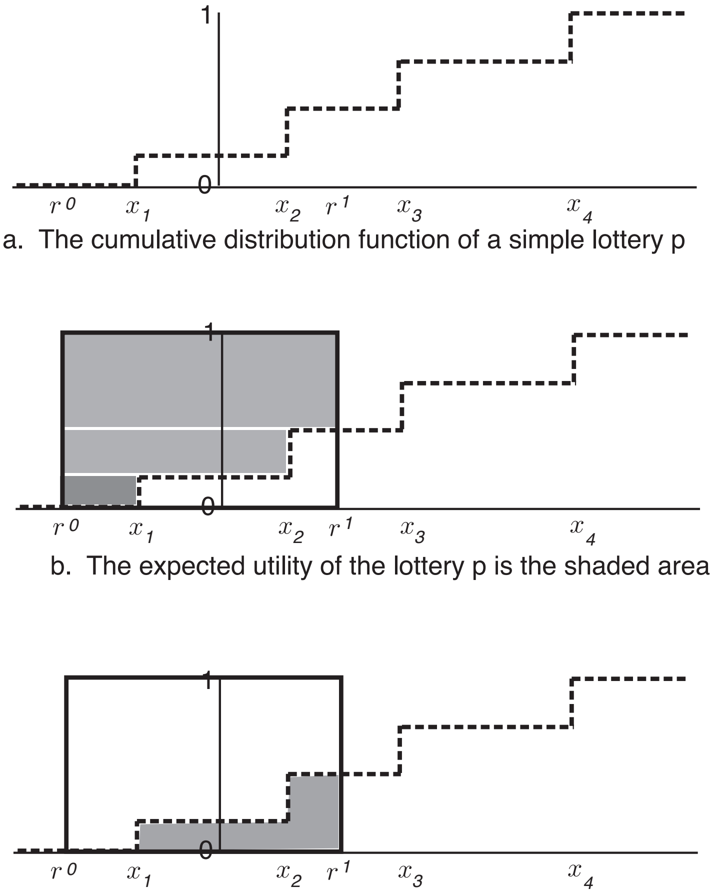

c. The shaded area is the heavy rectangle less the shaded area in panel b, which is the integral under the cumulative distribution function

Figure G6.1. Proving Part a of Problem 6.6. The key step in the proof is an application of integration by parts, to show equation ($\star\star$). Panel a graphs the cumulative distribution function of a simple lottery with a dashed line. In panel b, $E^{\pi}[U_{r^0,r^1}]$ is computed "graphically"; for points in the support of π that are less than r^1, we get the probability of the point (the height of respective rectangle) times $x - r^0$ (the base); plus we get the term $(r^1 - r^0)(1 - F_{\pi}(r^1))$, the area of the third rectangle. We subtract this expected utility from $r^1 - r^0$, the area of the heavily outlined rectangle in both panels b and c, and we are left with the shaded area in panel c, which is clearly the integral of π's c.d.f. from r^0 up to r^1.

nondecreasing and $E[X_\rho | X_\pi] \leq X_\pi$, $U(E[X_\rho | X_\pi]) \leq U(X_\pi)$. Combining these two inequalities does it.

This "proof" is likely to seem like hand-waving to readers who are not familiar with conditional expectations at a somewhat advanced level (where a conditional expectation is itself a random variable). In fact, it is not hand-waving at all. But for readers who find it hard to follow, let me explain it in terms of the sort of compound probability trees you see in Figure 6.2 of the text. When computing the expected utility under ρ, you can convert each endpoint to utility and then compute expected utility in two rounds, first computing the expected utility for the "second-stage" uncertainty and then overall. For any concave utility function U, the expected utility computed for the second-stage will, for each first-stage result, be less than or equal to the utility if there is no second stage. Look, for instance, at Figure 6.2(a) and (b) and, in particular, at the middle branches: In panel a, we have the outcome 5 as an outcome under π. But for ρ, we have in the middle the outcomes $5-2=3$ with probability 0.75 and $5+3=8$ with probability 0.25. Because U is concave, the expected utility of 3 with probability 0.75 and 8 with probability 0.25 can be no greater than the utility of the expected value, which is $(3)(0.75)+(8)(0.25)=4.25$. Since U is nondecreasing, this can be no greater than the utility of 5. In general, this always happens, because the "supplement" in the second stage to what happens in the first stage is always random (invoking the concavity of U) and has expected value less than zero (which invokes the assumption that U is nondecreasing.) That, in somewhat painful detail, is what the two inequalities with the conditional expectations are saying.

(c) This part of the problem is by far the hardest piece of these two problems. Because we are dealing with simple probability distributions, it can be done rather neatly and slickly—it takes a bit more math to do it in general. But it is not something that one would expect most students to come up with on their own. And even for simple probabilities, to write out all the details of a formal proof is a formidable task. So, instead, I will indicate how this is proved though an example.

We are trying to show that the "integral of the c.d.fs." condition implies the "$X_\rho = X_\pi +$ conditionally (on X_π) nonpositive-mean 'noise'" characterization. So the first thing to do is to understand the structure of of the integral-of-the-c.d.f. functions. Fix an r^0 less than every member of the support of π, and let

$$I_\pi(r) = \int_{r^0}^{r} F_\pi(r')dr', \quad \text{for } r \geq r^0.$$

This function is piece-wise linear and convex. Until the first (smallest) element of the support of π, it is zero (and in this respect, the choice of r^0 is irrelevant as long as it is smaller than this first element in the support); after the last (largest) element of the support of π, it has slope 1. Figure G6.2(a) provides an example; this is for the lottery π which gives prize 0 with probability 0.4, 5 with probability 0.2, and 10 with probability 0.4, the lottery in panel a of Figure 6.2 in the text. And in Figure G6.2(b) we superimpose (with dashed line segments) the function $I_\rho(r)$ for ρ that gives -3 with probability 0.2, 3 with probability 0.35, 8 with probability 0.25, and 12 with probability 0.2. (It is hard to see the kinks in I_ρ at the values 8 and 13, but they are there.) Note that $I_\pi \leq I_\rho$; we know this must be true because we know from parts a and b of this problem that ρ can be realized as π plus "noise" with nonpositive conditional means—that's what Figure 6.2 demonstrates—hence $\pi \geq^2 \rho$, and hence $I_\pi \leq I_\rho$.

Two facts about these functions should be noted:

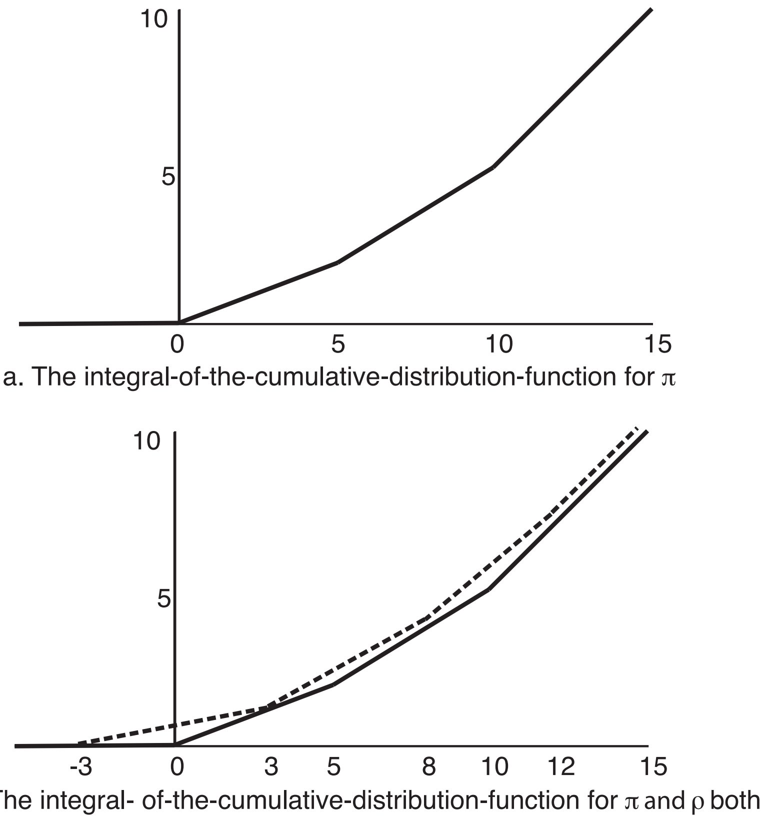

Figure G6.2. Graphing the integral-of-the-c.d.fs. In panel a, the function $I_\pi(r) = \int_{r^0}^{r} F_\pi(r')dr'$ is graphed, for π the simple lottery in panel a of Figure 6.2 of the text; π gives 0 with probability 0.4, 5 with probability 0.2, and 10 with probability 0.4. Note that this function is zero up to the first (smallest) value in the support of π, and then piecewise linear and convex, with "final" slope of 1. Panel b superimposes the same function for the simple lottery that is depicted in Figure 6.2(c) of the text. This second lottery is 2nd-order stochastically dominated by the first, so its integral-of-the-c.d.f. function is "higher", as guaranteed by parts a and b of this problem.

1. Any function of this sort—that is, piecewise-linear and convex, with value (and slope) zero below some value and with "terminal" slope 1 after some value—can be "unwound" into a corresponding simple lottery π: The location of the kinks constitute the support of π, and the size of the change in slope at the kink gives the probability of that value.

2. If I_π and I_ρ are two of these functions, the "higher" of the two once both hit slope 1 corresponds to the lottery with the greater mean. Therefore, π and ρ have the same mean if and only if I_π and I_ρ coincide after the last (greatest) value in the

union of their supports. To prove this, you need to enlist the integration-by-parts argument alluded to (and shown graphically) in part a of this problem. This is left for you to do.

Therefore, we know that if $\pi \geq^2 \rho$, then π has mean greater or equal to the mean of ρ. (We know this directly by using the utility function $U(x) = x$, of course.)

So much for preliminaries. Now for the heart of the argument. Supposing $I_\pi \leq I_\rho$, we "convert" I_π into I_ρ one step at a time. Each step involves adjusting one linear segment at a time, moving from the right (greater values of x) to the left. We are looking at a picture like panel b of Figure G6.2: In the first step, we take the "terminal segment" of I_ρ—the segment with slope 1—and extend it down and to the left, until it hits I_π. And we replace I_π with the integral-of-c.d.f. function that is I_π up to this point of intersection and is the extended "terminal segment" thereafter, creating a π^1 whose integral-of-the-c.d.f. function this is. (If π and ρ have the same mean, this first step is unnecessary, as their "terminal segments" are already coincident.) This in essence will "wipe out" one or more of the largest values in the support of π, replacing them with a lower value. If the extension (down and to the left) of I_ρ's final segment hits I_π right in a kink of I_π, it will mean increasing the probability of the value at that kink.

So, for instance, if we extend the final segment of I_ρ in Figure G6.2 down and to the right, we hit I_π at the value 9.625. We are, in essense, supplementing the prize 10 in π by -0.375, so the result is a lottery with prizes 0, 5, and 9.625, with probabilities 0.4, 0.2, and 0.4, respectively. Note that the value -0.375 is computed once we know the target value of 9.625, and we know that because that is where the extended final segment hits I_π. See panel a Figure G6.3.

In panel b of Figure G6.3, we have I_{π^1} and I_ρ, where π^1 is π modified as in the first stage; that is, the old prize of 10 is replaced by 9.625. Note that ρ and π^1 coincide past the value 12 (by design), so ρ and π^1 have the same mean. Hereafter, any changes we make will be zero conditional-mean changes.

And to make the first of these (which is our second change, from π^1 to π^2), we extend the next segment of I_ρ, the segment from 8 to 12, down and to the left, until it hits I_{π^1}. This, you can compute, happens at the value 7.25, and the extended segment "cuts out" the value 9.625 in π^1. In place of the 0.4 mass at 9.625, it puts mass 0.2 at 12 and 0.2 at 7.25. As shown in the right-hand portion of Figure 6.3b, this amounts to a "supplement" to the prize 9.625 of 2.375 with probability 0.5 and -2.375 with probability 0.5, giving π^2 which is 0 with probability 0.4, 5 with probability 0.2, 7.25 with probability 0.2, and 12 with probability 0.2. The fact that the supplement gives ± 2.375, that is, the same size prize on the upside as on the downside, is specific to this example; this does not happen in general. But the supplement will always be a zero-conditional mean supplement, because the means of the two lotteries start the same and end the same.

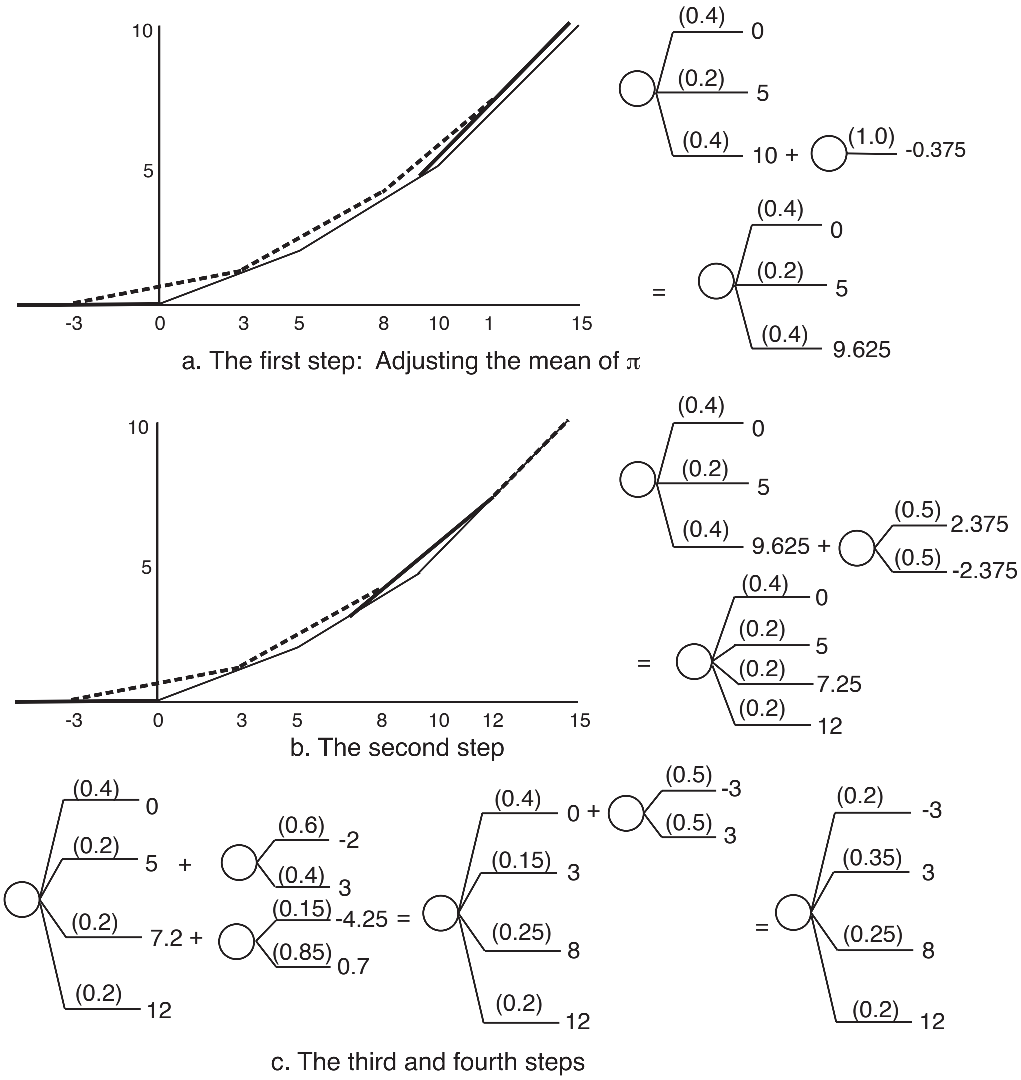

Figure G6.3. Converting π to ρ, one step at a time. See text for details.

In the next step (passing from π^2 to π^3), we "extend" the next segment of I_ρ down and to the left, until it hits I_{π^2}. In fact, since $I_\rho = I_\pi$ at the value 3, no extension is needed; we are going to replace in π^2 the *two* intermediate prizes 7.25 and 5 with prizes 8 and 3. For 7.25, this means a supplement of 0.75 on the upside and -4.25 on the downside which, to balance (to have zero conditional mean) requires probability 0.15 of -4.25 and 0.85 of 0.75. And for the supplement to 5, we need an addition of 3 and a subtraction of 2; to have zero conditional mean, this will take 0.4 probability for the supplement of 3 and 0.6 for the supplement -2. It takes a proof to show that, in general, these "balanced" or "zero conditional mean" supplements to the "passed over" kinks give the right probabilities in the end, but because the lotteries have to start and end with the same mean, this can be proved. In panel c of Figure 6.3, we show (without the graph) the transition from π^2 to π^3. And we add to panel c the final transition, from π^3 to ρ. This involves moving the 0.4 mass at 0 in π^3 to -3 and to $+3$.

Figure G6.4 summarizes the various steps in our transition from π to ρ and collapses (for each outcome of π) the changes made that led to ρ. This makes for an interesting comparison with Figure 6.2, which shows a different way to shift from π to ρ. Think in terms of the "supplementary random variable" Y that satisfies $X_\pi + Y = X_\rho$. In both Figure 6.2 in the text and here, the supplementary random variable has nonpositive conditional mean (conditional on X_π). But the supplement in Figure 6.2 is quite a different animal from the supplement here. In particular, they have different supports! The supplement in Figure 6.2 has support $\{-3, -2, 2, 3\}$, while the support of the supplement here includes -7 as well as the four other values. In general, there will be more than one way to move from a random-variable realization of one probability distribution π to the random-variable realization of a 2nd-degree stochastically dominated ρ, while respecting the rule that the supplement that accomplishes the move should have nonpositive conditional mean.

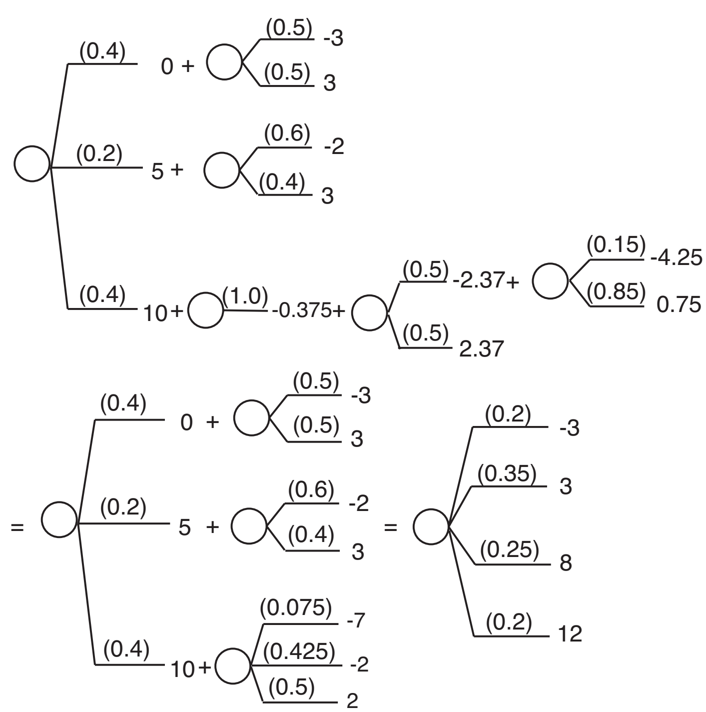

Figure G6.4. Summarizing the transition from π to ρ in Figure G6.3.

It hardly needs saying, I hope, that what just happened does not constitute a proof of what you were asked to prove in part c of Problem 6.6. The pictures we've drawn may convince you that this can be proved. But all we've done is to work through an illustrative example.

Writing out all the details of a proof for the case of simple probability distributions is not complex, just tedious. But since this result generalizes beyond simple probabilities, if you wish to pursue it further, you might want to see a proof for the more general case. If that is so, see Machina and Pratt (1997).

■ 6.8. (a) For these preferences, if $p = (1, 3)$, the consumer will consume good one only, or the bundle $(y, 0)$, while at the prices $p = (3, 1)$, she'll choose $(0, y)$, so her utility (for certain) will be $f(y)$. But if we take the average prices $(2, 2)$, she chooses (for instance) $(y/4, y/4)$ (or any bundle where $x_1 + x_2 = y/2$), for utility $f(y/2)$. Clearly she prefers the risky prices.

(b) At any prices (with this utility function), the consumer chooses to allocate her income to equalize the amounts of the goods she purchases. So at the prices $p = (\gamma, \gamma)$, she will consume $y/2\gamma$ of each good, for utility $f(y/(2\gamma))$, while at the prices $p = (1/\gamma, 1/\gamma)$, she'll get utility $f(\gamma y/2)$. Therefore, her average utility is

$$\frac{f(y/(2\gamma)) + f(y\gamma/2)}{2}.$$

But at the average prices, her utility is

$$f(y/(\gamma + 1/\gamma)).$$

Fix strictly increasing f. As γ increases toward infinity, the average utility is approximately $[f(0) + f(\gamma y/2)]/2$, while the utility at average prices approaches $f(0)$. So the former is greater. But fix γ. Since γ is greater than one,

$$\frac{y}{2\gamma} < \frac{y}{\gamma + 1/\gamma} < \frac{\gamma y}{2}.$$

No matter what are the values of $f(y/(2\gamma))$ and $f(y/(\gamma + 1/\gamma))$, by making f very concave, we can make $f(\gamma y/2)$ small enough so that the term in the middle is greater than the average of the two on the outside—the consumer prefers the sure prices.

■ 6.10. As in the proof of Proposition 6.17, the second half of Proposition A3.17(b) tells us that v as defined in (6.3) is concave. Moreover, the derivative of v is

$$v'(a) = \sum_{\theta \in \text{supp}(\pi)} V'(a(\theta - r) + rw)(\theta - r)\pi(\theta).$$

Therefore, $v'(0) = \sum(\theta - r)V'(rw)\pi(\theta) = V'(rw)(E\theta - r)$.

(a) If $E\theta < r$, then $v'(0) < 0$, and since v is concave, $v'(a) \le v'(0) < 0$ for all $a \ge 0$. Therefore, the solution, and the only solution, is $a = 0$.

(b) If $E\theta = r$, then $v'(0) = 0$, and $a = 0$ is a solution.

(c) If $E\theta > r$, then $v'(0) > 0$, hence $a = 0$ is not a solution.

(d) If $V(y) = Ay + B$ for $A > 0$, the first-order condition is $\sum_{\theta \in \text{supp}(\pi)}(\theta - r)A\pi(\theta) = A(E\theta - r) \le 0$.[1] Since a is entirely removed from v' in the first-order condition, there are three possibilities: If $E\theta < r$, the consumer chooses $a = 0$ (which we already knew). If $E\theta = r$, any a is a solution. And if $E\theta > r$, there is no solution.

(e) Since V is strictly increasing amd concave, $V' > 0$. If $\theta \ge r$ for all $\theta \in \text{supp}(\pi)$, then each term in $\sum_\theta(\theta - r)V'(a(\theta - r) + rw)\pi(\theta)$ is nonnegative. And if $\theta > r$ for at least one θ with positive probability, the corresponding term in the sum is always strictly positive. Hence the sum can never be nonpositive, and there is no solution to the problem.

(f) We show that v is strictly concave. (Once we know v is strictly concave, we know from Proposition A3.21 that (6.3) can have at most one solution over any convex set.) We do this by showing that v' is strictly decreasing. Note that v' is composed of a sum of terms. For those terms for which $\theta > r$, $\theta - r > 0$, and increasing a increases the argument of V', hence decreases V'. Hence increasing a decreases (strictly) all terms with $\theta > r$. While for terms for which $\theta < r$, increasing a decreases the argument of V', increasing V', but the coefficient $\theta - r$ is negative. So these terms get increasingly negative—that is, they decrease—as a is increased. (For $\theta = r$, changing a has no effect on the corresponding term in the sum.) Since the support of π contains at least two elements, one of them must be different from r, hence increasing a strictly decreases the sum, that is, $v'(a)$. ∎

■ 6.12. Suppose there are three assets available: A riskless asset, which returns $r = 5\%$ for sure, risky asset #1, which (say) returns 10% with probability 0.9 and 4% with probability 0.1, and:

- Risky asset #2, which returns 10% with probability 0.1 and 4% with probability 0.9, for an overall expected return of 4.6%. But this risky asset #2 returns 10% in precisely those states of the world where risky asset #1 returns 4%; the two are (perfectly) negatively correlated. Then consider any investor who, if given a portfolio choice of the safe asset and asset #1, would choose some of the safe asset. (Make the investor sufficiently risk averse, and this must happen.) If we add risky asset #2 to the possible investments, then this investor will necessarily choose some of risky asset #2: A portfolio of equal dollars invested in the two risky assets produces a riskless asset with overall (sure) return of 7%, dominating the riskless asset, so the optimal choice is necessarily a split between the two risky assets (only).

[1] Don't confuse A and a here!

- Risky asset #2, which returns 9% with probability 0.9 and 4% with probability 0.1, for an overall expected return of 8.5%. But this asset returns 9% precisely in those states of the world where asset #1 returns 10%. So this asset is dominated by risky asset #1, and no investor (who has an increasing utility function) would ever choose to invest in it.

The moral (well known in the theory of finance) is: When constructing an optimal portfolio out of various assets, how they correlate with one another is hugely important.

■ 6.14. Suppose that our consumer has constant absolute risk aversion; that is, if we denote her wealth by W, we have $u(W) = -e^{-\lambda W}$. Let $\tilde{X}_n$ be the (random) outcome of the nth gamble, so that if the consumer is given m gambles, her final wealth W will be

$$W = \sum_{i=1}^{m} \tilde{X}_i.$$

If all gambles are mutually independent, we have

$$E\left[u(W)\right] = E\left[-e^{-\lambda \sum_{i=1}^{m} \tilde{X}_i}\right] = E\left[-\prod_{i=1}^{m} e^{\lambda \tilde{X}_i}\right] = -\prod_{i=1}^{m} E\left[e^{-\lambda \tilde{X}_i}\right] = -\left(E\left[e^{-\lambda \tilde{X}_i}\right]\right)^m,$$

where the second-to-last equality follows by the independence assumption.

Therefore, independently of m, the consumer will take m copies of the gamble if and only if she is willing to take one, depending on whether $E[e^{-\lambda \tilde{X}_i}]$ is less than 1 or greater. This fits very well with the intuitive meaning of constant absolute risk aversion.

As for the challenge, (and using some results from the theories of random walks and martingales that you may not know): Whatever is the consumer's initial wealth (as long as it is not zero), as $N \to \infty$ the probability that the consumer will go bankrupt before the Nth gamble approaches some constant α (depending on the initial wealth) bounded away from one, and the probability that her wealth exceeds any finite level approaches $1 - \alpha$; that is, either she goes bankrupt eventually, or her wealth goes off to ∞. See, e.g., Chung (1974, Chapter 8).

Suppose u is unbounded above. Let Z be a wealth level sufficiently large so that $(1 - \alpha)u(Z) + \alpha u(0) > u(W_0)$, where W_0 is the consumer's initial level of wealth. (Since u is unbounded, such a Z exists.) Then for N sufficiently large, so that the probability of exceeding Z after N or more gambles is sufficiently close to $1 - \alpha$, the consumer will take the gambles. (Assume the consumer's initial wealth is a multiple of \$500, so we don't need to worry about her going bankrupt on a "part-loss." Also, throughout we are assuming that the consumer must take all the N gambles, limited by the bankruptcy constraint, or none at all, although you can work around that assumption for this part of the problem at least.)

When u is bounded above, things are a bit more delicate. Let $\hat{u} = \lim_{z\to\infty} u(z)$. Then as the number of gambles grows, and if the consumer chooses to gamble, her expected utility will approach $(1-\alpha)\hat{u} + \alpha u(0)$. To know whether this entices the consumer to take many gambles, we need to know whether $(1-\alpha)\hat{u} + \alpha u(0) > u(W_0)$. Assuming we know u, we know $u(0)$, $u(W_0)$, and $\hat{u}$. But what is α? What is the probability that the consumer's wealth will go off to infinity before she is bankrupt?

We can use martingale theory to find α.[2] Let W_n be the level of the consumer's wealth after n gambles (where we stop the process if the consumer is bankrupt, and we assume that the consumer takes the gambles), and let $X_n = .999610175^{W_n}$. The stochastic process $\{X_n\}$ is then a bounded martingale (or, rather, it will be if in place of the number .999610175 we have the root of the polynomial $.4\gamma^{1000} + .6\gamma^{-500} = 1$ that is near .999610175). This martingale either "absorbs" at the value of 1 (if the consumer goes bankrupt) or approaches 0 (if the consumer's wealth rises to infinity), and so by the martingale stopping theorem, we know that $1 \cdot \alpha + 0 \cdot (1-\alpha) = \alpha = .999610175^{W_0}$. This, then, gives us α; and—except for the knife-edge case where $(1-\alpha)\hat{u} + \alpha u(0) = u(W_0)$—given $\hat{u}$, $u(0)$, and W_0, we have everything we need to determine whether the consumer will take "many" gambles or not.

[2] Most advanced probability texts will cover this, as well as the random-walk theory we used previously. But for a definite reference, see Chung (1974), this time Chapter 9.