III. Uncertainty, Expected Utility, and Competitive Trade

Neil Wallace

June 9, 2008

1 Introduction

For some of the topics we want to discuss, we need models with uncertainty. We are all familiar with the general idea of uncertainty. We are uncertain about tomorrow's weather, about whether we will wake up with a headache tomorrow morning, and about whether someone's estimate of the labor required to repair our car is correct. Considerable effort is directed toward coping with uncertainty. Some farmers have costly irrigation systems in order to make output less dependent on variations in rainfall. And many of us buy insurance of various sorts to limit our exposure to some kinds of uncertainty. Moreover, there are government programs like disaster aid and unemployment insurance that are intended to offset some of the effects of uncertainty.

In order to proceed, we need four concepts from probability theory: the set of outcomes, the probability distribution over outcomes, random variables, and expected values. The set of outcomes describes all possible occurrences. An example is all possible levels of rainfall in a world in which rainfall is the only source of uncertainty. Throughout, we assume that the set of outcomes has a finite number of elements. A probability distribution is a function that describes the probability with which each element in the set of outcomes occurs. A random variable is any function whose domain is the set of outcomes and whose range is the real numbers. The expected value of a random variable is its average value computed using the probability distribution over outcomes.

2 Preferences under uncertainty

We begin with a general formulation called *the state-preference* approach to uncertainty. Then we specialize that approach to what is called *expected utility.* Both approaches begin from the notion of *contingent commodities.*

The notion of contingent commodities applies the idea of labeling commodities by dates to labeling them by outcome or *state (of the world)*. To explain this, let's begin by describing the sequence of actions that we will work with throughout this chapter. People start out and agree about the set of outcomes. Hence, if the only uncertainty is rainfall, then they agree about the finite set of possible rainfall levels. Prior to the realization of the actual rainfall, they may make various agreements or bets. Then the realization occurs and everyone sees the realization. Then, they honor their agreements. (Here, just as we did when we studied the inter-temporal version of the competitive model, we assume that people can commit to future actions.) Indeed, to focus on uncertainty, we assume here that there is one underlying good; say, rice. If there are two states of the world, high rainfall and low rainfall, then there are two contingent commodities, *rice-if-rainfall-is-high* and *rice-if-rainfall-is-low*. If there are K levels of rainfall, then there are K contingent commodities, each contingent on the level of rainfall.

Exercise 1 *Suppose that there are two states of the world, high rainfall and low rainfall, and that there are two groups of rice farmers. One group has irrigated fields that make the resulting crop independent of rainfall. The other has fields that depend on the rainfall; they get much lower output if rainfall is low than if it is high. What kind of deals might they enter into before the state is realized? Explain.*

A systematic approach to this question requires a specification of preferences and a notion of equilibrium. We start with preferences. Then we turn to the competitive equilibrium notion of equilibrium.

2.1 The state-preference approach

This approach starts from the notion of contingent commodities and treats the different contingent goods as different goods. Thus, if there are two states of world and one underlying good, there are two commodies and it is assumed that each person has an indifference curve map in the plane—where, for example, the horizontal axis is the quantity of rice-if-rainfall-is-low and the vertical axis is the quantity of rice-if-rainfall-is-high. This approach makes no explicit use of a probability distribution over outcomes. Instead, it makes assumptions directly about the shape of indifference curves. Standard assumptions are the same as those we made for two different goods in the static version of the competitive model without uncertainty.

2.2 Expected utility

One widely used hypothesis is that people evaluate distributions of consumption by what is called *expected utility.* Let $\{c_1, c_2, ..., c_K\}$ be a set of possible consumptions of a single underlying good for some person and let π_i denote the probability of that the outcome c_i occurs, where $\pi_i \geq 0$ and $\sum_{i=1}^{K} \pi_i = 1$. Let u denote the utility that this person assigns to sure or certain outcomes. (An example is the kind of u function that we have used in describing period utility in the discounted utility formulation.) The *expected utility hypothesis* is that the payoff that this person assigns to the above distribution is

$$\pi_1 u(c_1) + \pi_2 u(c_2) + ... + \pi_K u(c_K) = \sum_{i=1}^{K} \pi_i u(c_i). \tag{1}$$

In words, the payoff is the probability weighted sum of utilities or the expected value of the random variable $u(c_i)$.

Exercise 2 *Suppose $K = 2$ in (1). Also suppose that u is strictly increasing and strictly concave. Consider the equation $\pi_1 u(c_1) + \pi_2 u(c_2) = b$ for given π_1, π_2, and b. There are many pairs (c_1, c_2) that satisfy this equation. (i) Sketch such pairs, an indifference curve, in a diagram. (ii) One such pair also satisfies $c_1 = c_2 = c$. What is the slope of the indifference curve at $(c_1, c_2) = (c, c)$?*

Exercise 3 *Let x satisfy the equation $\pi_1 u(x) + \pi_2 u(x) = b$. (i) In the diagram of the last exercise, plot all the pairs (x_1, x_2) that satisfy the equation $\pi_1 x_1 + \pi_2 x_2 = x$. (ii) If (x_1, x_2) satisfies the equation $\pi_1 x_1 + \pi_2 x_2 = x$ and $x_1 \neq x_2$, argue that $\pi_1 u(x_1) + \pi_2 u(x_2) < b$.*

From now on, we use the expected utility hypothesis.

3 An uncertainty interpretation of 2-good pure exchange

We here interpret the two-good pure exchange model as a model with uncertainty.

3.1 The environment and allocations

We start by supposing that there are only two elements in the set of possible outcomes and we label them outcome 1 and outcome 2. We assume that outcome 1 occurs with probability π_1 and outcome 2 with probability π_2, where $\pi_1 + \pi_2 = 1$.

For the inter-temporal interpretation with 2 dates, we assumed that there was one underlying good (rice) and distinguished between rice at date 1 and rice at date 2. As noted above, the labeling we want now is labeling by outcome, and the results of the relabeling are outcome-contingent commodities, or, more simply, contingent commodities. We assume that there is one underlying good, rice. We distinguish between rice-if-the outcome-is-outcome-1, which we call the outcome-1 good or good 1, and rice-if-the-outcome-is-outcome-2, which we call the outcome-2 good or good 2.

There is no production in this world and the total resources are total amounts of good 1, denoted W_1, and of good 2, denoted W_2. (You could think of W_1 as being the exogenous total crop of rice if outcome 1 occurs and W_2 as the exogenous total crop if outcome 2 occurs.) There are N people. If person n consumes (c_{n1}, c_{n2}), then person n gets utility $\pi_1 u^n(c_{n1}) + \pi_2 u^n(c_{n2})$, where the function u^n is strictly increasing and strictly concave. (In most applications, we assume that the everyone has the same utility function. When we do that, we drop the superscript.)

While this completes the environment, it is helpful to review the sequence of actions. As noted above, we have to distinguish between prior to and subsequent to the realization of the outcome. The preferences just described pertain to the situation prior to the realization of the outcome. These are the preferences that are relevant to decisions about the purchase of risky assets, the purchase of insurance, and betting.

An allocation is exactly what it was before, a list of consumption pairs, one pair for each person. Also, the definition of feasible allocations is unchanged, as is the definition of Pareto efficient allocations. All of these pertain to the situation prior to realization of the outcome.

3.2 Competitive equilibrium (CE)

As regards private ownership, we assume that each person starts out owning some of each good. We let w_{n1} denote the amount of good 1 owned by person n and let w_{n2} denote the amount of good 2 owned by person n. We will call the pair (w_{n1}, w_{n2}) person n's endowment. We also assume that everything is owned by someone. That is, $\sum_{n=1}^{N} w_{n1} = W_1$ and $\sum_{n=1}^{N} w_{n2} = W_2$.

Let p denote the price of good 2 in units of good 1 (p is measured in units of good 1 per unit of good 2). Throughout, we suppose that p is positive. We say that person n can *afford* to buy the pair (c_{n1}, c_{n2}) at the price p if there exists (q_1, q_2) such that

$$c_{n1} = w_{n1} - q_1,\; c_{n2} = w_{n2} - q_2, \text{ and } q_1 + pq_2 = 0, \tag{2}$$

where q_j is to be interpreted as sales of contingent commodity j (a negative q_j is a purchase of good j). As was demonstrated above, it is equivalent to say that person n can afford to buy the pair (c_{n1}, c_{n2}) at the price p if

$$c_{n1} + pc_{n2} = w_{n1} + pw_{n2}. \tag{3}$$

Given the above equivalence, we can again use satisfaction of (3) as our definition of affordability. The definition of CE is unchanged. That is, an allocation A and a price p is a CE if two conditions hold: (i) A is feasible, (ii) for each person n, the consumption pair assigned to n by A is liked by n as well as anything else that n can afford at the price p.

In the inter-temporal interpretation of two-good CE, we said that trade occurs at date 1 and involves promises. There, with one good per date, nothing happens at date 2 except that deliveries are made in accord with the promises made at date 1. Now we distinguish between *prior to* and *subsequent to* the realization of the outcome. Prior to the realization of the outcome, only promises are made. After the realization, deliveries are made and consumption occurs. In that respect, trade here is like betting with nothing changing hands at the time of the bet. After the outcome occurs, you pay if you lose and are paid if you win. Here, with two goods, trade is exactly like a bet on one outcome. If you sell outcome-1 good ($q_1 > 0$), then you are necessarily buying outcome-2 good ($q_2 < 0$). This is like a bet on outcome 2 and, therefore, against outcome 1. If outcome 1 occurs, then you lose and have to hand out q_1 units of the good. If outcome 2 occurs, then you win and collect q_2 units of the good. The price p gives the terms on which you bet.

Exercise 4 *Show the following. If $W_1 = W_2$, then there exists a CE with $p = \frac{\pi_2}{\pi_1}$. There is no other CE.*

Exercise 5 *Show the following. If $W_1 > W_2$, then any CE has $p > \frac{\pi_2}{\pi_1}$.*

Exercise 6 *Interpret the claims in the last two exercises in terms of the notions of fair and unfair insurance, where fair insurance means that the expected value of the premium is equal to the expected value of the payment.*

4 Pure exchange CE with many outcomes

The generalization to K outcomes is straightforward. It follows exactly what we earlier said about pure exchange with K commodities. In defining CE, we now need as many prices as there are outcomes. As before, we denote them as

$(p_1, p_2, ..., p_K)$, where we fix $p_1 = 1$ and where p_k is the price of outcome-k good in units of outcome-1 good.

Exercise 7 *Show the following. If* $W_1 = W_2 = ... = W_K$, *then there exists a CE with* $p_k = \frac{\pi_k}{\pi_1}$.

4.1 Fewer markets[1]

Many simple natural situations seem to give rise to many states, so many that trade contingent on all those states seems far from what we see. Here is one such situation. Suppose that there are N people and that each person will realize an endowment of w_i with probability π_i, where $i = 1, 2$, $w_2 > w_1 > 0$, and $\pi_1 + \pi_2 = 1$. The realizations for the different people are *independent* realizations. Hence, if we label the realization 2 a success (because $w_2 > w_1$), then the (total) number of successes follows a binomial distribution. In particular, the probability that r successes occur is $\binom{N}{r}\pi_1^{N-r}\pi_2^r$, where $\binom{N}{r} = \frac{N!}{(N-r)!r!}$ and $j! = j(j-1)...1$ and $0! = 1$. (You can think of outcome 2 as being healthy and outcome 1 as being sick or outcome 2 as not having an automobile accident and outcome 1 as having an accident.) Let's also suppose that each person maximizes expected utility, and has the same utility function for sure or certain consumption of the one underlying good.

This kind of setting gives rise to a large number of outcomes or states; namely, 2^N. However, because of the symmetry among people, we can group those outcomes and work with a much smaller number of *aggregate* states; namely, the different possible outcomes for the number of successes, which we denote r. Obviously, there are only $N + 1$ possibles outcomes for r.

Based solely on feasibility, we expect consumption to depend on r. And because people are identical in preferences and endowments, we also expect that each person will have the same allocation in a model in which people choose subject to a budget. Therefore, we let (c_{1r}, c_{2r}) be each person's consumption when there are r successes (in the economy), where c_{1r} is the person's consumption when this person has a failure and there are r successes and where c_{2r} is the person's consumption when this person has a success and there are r successes.

Let $\gamma(n, r)$ be the probability of j successes in n trials, where γ depends on n, r, and the π_i according to the binomial formula described above. We can

[1] This section is a special case of the model in "Individual Risk and Mutual Insurance," by David Cass, Graciela Chichilnisky, and Ho-Mou Wu, *Econometrica*, Vol. 64, No. 2, (Mar., 1996), pp. 333-341. They study pure exchnage just as we do here. An extension to a production economy appears in "A General Equilibrium Interpretation of Damage Contingent Securities," by R. Anton Braun, Richard M. Todd, and Neil Wallace, *Journal of Risk and Insurance*, December 1999.

express the expected utility for a person of an allocation of the form (c_{1r}, c_{2r}) as

$$\sum_{r=0}^{N-1} [\gamma(N-1,r)\pi_1 u(c_{1r}) + \sum_{r=1}^{N} \gamma(N-1,r-1)\pi_2 u(c_{2r})]. \tag{4}$$

This expression uses independence, which implies that the probability of a joint event is the product of the probabilities of its component events. Notice that if $r = 0$, then $u(c_{2r})$ does not appear (a success for this person cannot happen) and if $r = N$, then $u(c_{1r})$ does not appear (a failure for this person cannot happen).

Now, suppose that the person faces the following budget constraints. For each aggregate state r different from 0 or N, let s_r be the price of an additional unit of c_{1r} in units of c_{2r} surrendered. That is suppose the person faces the following separate constraints:

$$s_r c_{1r} + c_{2r} = s_r w_1 + w_2 \text{ for } r = 1, 2, ..., N-1. \tag{5}$$

The person chooses (c_{1r}, c_{2r}) for $r = 1, 2, ..., N-1$ taking s_r for $r = 1, 2, ..., N-1$ as given.

Definition 1 *An equilibrium is* $c_{10} = w_1$, $c_{21} = w_2$, (c_{1r}, c_{2r}) *and* s_r *for* $r = 1, 2, ..., N-1$ *such that each person's choice is best for the person given* s_r *and* (c_{1r}, c_{2r}) *for* $r = 1, 2, ..., N-1$ *is feasible.*

The following exercise contains a conjecture about the equilibrium.

Exercise 8 *Show the following. The unique equilibrium is* $c_{1r} = c_{2r} = c_r$, *where* $c_r = [(N-r)w_1 + rw_2]/N$ *and*

$$s_r = \frac{\gamma(N-1,r)\pi_1}{\gamma(N-1,r-1)\pi_2} \tag{6}$$

for $r = 1, 2, ..., N-1$.

Several comments are in order about this model. First, the equilibrium $c_{1r} = c_{2r} = c_r$ is an instance of the result found above: namely, that under exepcted utility maximization and no aggregate risk, the equilibrium has no individual risk. Here, conditional on an aggregate state (the number of successes), there is no individual risk. Second, the above allocation can be interpreted as a mutual insurance arrangement. To do this, think of $w_2 - c_r$ as the premium paid by those who experience a success in aggregate state r and $c_r - w_1$ as the net payment received by those who experience a failure in aggregate state r. Under mutual insurance, the company pays a dividend based on its experience.[2] The more successes, the better its experience (the fewer the claims) and the greater the dividend. That resembles how the above premiums and net payments vary with r. Third, the equilibrium is Pareto efficient, but only because of the symmetry in the environment. As we now explain, even if people differ, a combination of the above kind of insurance and markets contingent on the aggregate state suffice to make equilibrium Pareto efficient.

Suppose people have different preferences, different u functions, but that endowments are as just described. Let us denote person n's utility function by u^n and his consumption in aggregate state r by (c^n_{1r}, c^n_{2r}). Then people might want to trade claims across aggregate states. To permit such trade, we write the choice problem of person n as follows: choose (c^n_{1r}, c^n_{2r}) for $r = 0, 1, 2, ..., N-1, N$ to maximize

$$\sum_{r=0}^{N-1}[\gamma(N-1,r)\pi_1 u^n(c^n_{1r}) + \sum_{r=1}^{N}\gamma(N-1,r-1)\pi_2 u^n(c^n_{2r})] \tag{7}$$

subject to

$$p_0(c^n_{01} - w_1) + \sum_{r=1}^{N-1} p_r[s_r c^n_{1r} + c^n_{2r} - (s_r w_1 + w_2)] + p_N(c^n_{02} - w_2) = 0. \tag{8}$$

Here, p_r is the price of a unit of the good in aggregate state r.

Definition 2 *An equilibrium is* (c^n_{1r}, c^n_{2r})*,* s_r *for* $r = 1, 2, ..., N-1$*,* p_r *for* $r = 0, 1, 2, ..., N$ *such that each person's choice is best for the person given* s_r *and* p_r *and* (c^n_{1r}, c^n_{2r}) *is feasible.*

Exercise 9 *(i) Show that an equilibrium satisfies* $c^n_{1r} = c^n_{2r} = c^n_r$ *for* $r = 1, 2, ..., N-1$*. (ii) Show that if* $u^n = u$ *for all* n*, then* $c^n_r = c_r = [(N-r)w_1 + rw_2]/N$ *is an equilibrium allocation.*

It is straightforward to show that an equilibrium is Pareto efficient.

4.2 Competitive asset pricing[3]

We use the framework of a K-outcome CE to price assets. We define assets by their payoffs. In particular, an asset is a vector of outcome-specific payoffs—say,

[2] Insurance companies organized as *mutuals* were quite common 40 years ago; they seem to be rare now.

[3] This kind of material is now standard. An early exposition is in "A Re-Examination of the Modigliani-Miller Theorem," by Joseph E. Stiglitz, *The American Economic Review*, Vol. 59, No. 5, (Dec., 1969), pp. 784-793.

$(a_1, a_2, ..., a_K)$ with the interpretation that the owner of this asset receives a_k units of the good if outcome k occurs. (If $a_k < 0$, then the owner must pay out a_k units of the good if outcome k occurs.) The price of the asset is defined to be the amount of outcome-1 good for which one can buy the asset prior to the realization of the outcome. Let us call this price $p(a_1, a_2, ..., a_K)$.

Proposition 3 *In a CE in which* $(p_1, p_2, ..., p_K) = (p_1^*, p_2^*, ..., p_K^*)$,

$$p(a_1, a_2, ..., a_K) = \sum_{k=1}^{K} p_k^* a_k. \tag{9}$$

Exercise 10 *Argue that proposition 3 is true. (Hint: Show that each person would profit by buying the asset if* $p(a_1, a_2, ..., a_K) < \sum_{k=1}^{K} p_k^* a_k$ *and that no-one wants to buy it and each person would profit by issuing it if* $p(a_1, a_2, ..., a_K) > \sum_{k=1}^{K} p_k^* a_k$*.)*

Exercise 11 *Consider two assets:* $(a_1, a_2, ..., a_K)$ *and* $(b_1, b_2, ..., b_K)$*, where* $b_k = za_k$ *for* $k = 1, 2, ..., K$*. Show that in a CE in which* $(q_1, q_2, ..., q_K) = (q_1^*, q_2^*, ..., q_K^*)$, $p(b_1, b_2, ..., b_K) = zp(a_1, a_2, ..., a_K)$*.*

The price vector $(p_1^*, p_2^*, ..., p_K^*)$ is sometimes called the "pricing kernel."

4.2.1 The pricing of limited-liability debt and equity in CE

We can also use proposition 3 to price the debt and equity of limited liability "corporations." We define the corporation as a vector of outcome-specific payoffs—say, $(x_1, x_2, ..., x_K)$. As part of "limited liability," we assume that $x_k \geq 0$ for each k.

We next define what we mean by debt of the corporation. We consider debt to be an asset with a *promised* pay-off that is not contingent. The amount of debt is measured by the magnitude of that *promised* pay-off. Let $P \geq 0$ be the promised pay-off. The magnitude P is to be distinguished from what the owners of the debt will actually get. (After all, junk bonds are called junk because it is understood that holders of junk bonds may not receive the promised pay-off.)

Let $\min\{a, b\}$ denote the smaller of a and b. The owners of the debt of the corporation $(x_1, x_2, ..., x_K)$ with promised pay-off P receive the vector $(\min\{x_1, P\}$, $\min\{x_2, P\}$, ..., $\min\{x_K, P\})$. We let D denote the pre-outcome value of this debt in units of outcome-1 good. Using proposition 3, we have

$$D = \sum_{k=1}^{K} p_k^* \min\{x_k, P\}. \tag{10}$$

Now we define the equity of the above corporation. The owners of the equity of this corporation receive the vector $(\max\{x_1 - P, 0\}, \max\{x_2 - P, 0\}, ..., \max\{x_K - P, 0\})$, where $\max\{a, b\}$ means the larger of a and b. We let E denote the pre-outcome value of this equity in units of outcome-1 good. Using proposition 3, we have

$$E = \sum_{k=1}^{K} p_k^* \max\{x_k - P, 0\}. \tag{11}$$

A well-known result, called the Modigliani-Miller Theorem, is that the total value of a corporation, $D + E$, does not depend on how much debt it has. This conclusion follows from (10) and (11).

Exercise 12 *Show that* $\min(a, b) + \max(a - b, 0) = a$. *(Hint: Show that the conclusion holds in three mutually exclusive and exhaustive cases:* $a < b$, $a = b$, *and* $a > b$.*)*

Exercise 13 *(i) Use the result of the last exercise and (10) and (11) to compute* $D + E$. *(ii) Explain why your part (i) result is a Modigliani-Miller Theorem.*

Exercise 14 *We are describing a corporation by a vector of payoffs,* $(x_1, x_2, ..., x_K)$. *Suppose* $K = 2$. *(i) How would you describe the "riskiness" of a corporation? (ii) Are less risky corporations worth more than more risky corporations? Explain.*

Because D is the pre-outcome value of the debt and P is the promised pay-off, the ratio $\frac{P}{D} - 1$ is the promised yield or promised rate-of-return on the debt. As we next show, this yield is in general increasing in P, our measure of the amount of the debt. We draw this conclusion by studying the ratio $\frac{D}{P}$. If $b > 0$, then $\frac{\min\{a,b\}}{b} = \min\{\frac{a}{b}, 1\}$. Therefore, from (10), if $P > 0$, then

$$\frac{D}{P} = \sum_{k=1}^{K} p_k^* \min\{\frac{x_k}{P}, 1\}. \tag{12}$$

Exercise 15 *According to (12), higher debt as measured by* P *either does not affect the promised yield or increases the promised yield. Describe the circumstances in which each conclusion holds.*

Exercise 16 *At times, corporations are given loan guarantees. In a full analysis, we would have to describe how the guarantee is financed. Here we ignore that. If the guarantee is genuine, then the owners of the debt of the corporation* $(x_1, x_2, ..., x_K)$ *with promised pay-off* P *will actually receive* P *in every outcome. Suppose that is the case. (i) Give a formula for the pre-outcome value of such* *debt. (ii) Give a formula for $D + E$ for such a corporation. (iii) Does $D + E$ depend on P and in what way? (iv) Suppose the guarantor of the debt imposes some maximum amount of debt. Given that maximum, does the value of the equity of the corporation depend on the riskiness of the corporation?*

4.2.2 Hedge-fund contracts[4]

As has been widely reported, there seems to be a fairly standard hedge-fund contract for the manager of a hedge fund. It is called "2 and 20," where the manager gets 2% of the size of the fund per year independent of performance and gets 20% of the annual earnings of the fund that exceed some benchmark. In what follows, we ignore the 2% flat fee and focus on the part that depends on performance. (The results do not depend on the flat fee.) We let $\lambda > 0$ denote the share of the excess payoff that the manager gets. We first describe the investment strategy of a manager with such a contract. Then we discuss whether the model is appropriate for the analysis of hedge funds.

The manager with a fund of size z chooses $y = (y_1, y_2, ..., y_K)$ subject to

$$z = \sum_{k=1}^{K} p_k^* y_k$$

and $y_k \geq 0$. (This last inequality is limited liability for the investors in the fund.) Given y, if outcome k occurs, then the manager gets $\max\{0, \lambda(y_k - x)\}$, where x is the benchmark. Then the pre-outcome value of what the manager gets is

$$\sum_{k=1}^{K} p_k^* \max\{0, \lambda(y_k - x)\}.$$

For reasons we will spell out below, let us suppose that the manager chooses y to maximize this last expression (subject, of course, to the above constraints). This is a linear programming problem and the solution is straightforward. Let $p_h^* = \min\{p_1^*, p_2^*, ..., p_K^*\}$. If $0 < x < 1/p_h^*$, then the solution is $y_h = 1/p_h^*$ and $y_k = 0$ for $k \neq h$; namely, bet everything on the longest long shot.

Is this really a good outcome for the manager? After all, he gets a payment only in outcome h, which can be interpreted as (one of) the least likely outcomes. However, that is not true if he can trade on his own private account and hedge(!) that risk. To hedge on private account, he sells claims on outcome h good (most of y_h) and uses the proceeds to buy other claims. He can even assure himself a safe outcome.

[4] This material is essentially identical to "Deposit Insurance and Bank Regulation: A Partial Equilibrium Exposition," by John Kareken and Neil Wallace, *Journal of Business*, July 1978.

The outcome we have found is not surprising. The contract has a "heads-*we*-win tails-*you*-lose" feature. Any contract of that sort gives the manager an inducement to take on risk.

Is the model appropriate for the analysis of hedge funds? In one sense, it is not. If there were prices at which anyone could buy any portfolio of outcome-dependent payoffs, then why would anyone share gains with a so-called hedge-fund manager? But even if the manager has access to some trades and some information that the investor does not have, the incentives of the manager described above are present. In particular, two aspects of hedge funds make that likely. First, the manager does not reveal his strategy; if he did, then investors could emulate it without paying fees and sharing gains. Also, hedge funds are unregulated. Therefore, any private-account hedging that the manager does is not made public. (Notice the contrast with insider-trading of publicly held firms.)

Why do hedge-fund contracts take the above form? Such contracts are similar to those that are common for managers of many corporations. (Stock options have payoffs that resemble sharing the gains, but not the losses.) Do such contracts also suffer from the risk-taking inducement set out above? To some extent they do. But, there are normally mitigating factors in place. Such contracts are intended to deal with a moral hazard problem of the following sort: effort by managers is imperfectly observable and the level of effort affects the distribution of payoffs. The models in which such contracts are shown to be good for the investor are ones in which the only choice by the manager is effort. That may be a good approximation in many situations. In most corporations, a manager's attempt to take on risk could not be hidden from its board of directors. And insider-trading rules tend to inhibit attempts by managers to privately hedge such risk. But such a model is almost certainly a poor approximation for hedge funds as they now exist. There is no board of directors to which the portfolio strategy is described and there are no insider trading rules. And, because the only role of the manager is to choose a portfolio, controlling risk and maintaining secrecy seem at odds with one another. Those features make it difficult to monitor and inhibit the kind of porfolio decisions shown above to be best for the manager of a hedge fund. The mystery, then, is why many investors are willing to enter into such contracts.

5 Concluding remarks

We have explored in some detail what is called the contingent claims approach to uncertainty. In what follows, we will be making assumptions inconsistent with that approach. We will sometimes assume that there is asymmetric information—that not everyone sees the entire outcome. And, we will sometimes assume that people cannot commit to future actions. Nevertheless, if only as a point of departure, students should be aware of the contingent claims approach to uncertainty.