Uncertainty and Information: An Expository Essay

Singh, Nirvikar

University of California, Santa Cruz

22 March 2010

Uncertainty and Information: An Expository Essay

Nirvikar Singh
Department of Economics
University of California, Santa Cruz

March 22, 2010

Abstract

This essay provides an elementary, unified introduction to resource allocation under uncertainty in competitive markets. The coverage includes decision-making under uncertainty, measuring risk and risk aversion, insurance and asset markets, and asymmetric information.

Keywords: Uncertainty, risk, risk aversion, insurance, asset markets, asymmetric information

Introduction

Uncertainty is pervasive in the world, and has implications for all kinds of economic decision-making. This essay provides a unified, elementary exposition of some aspects of how uncertainty and asymmetric information affect resource allocation. The motivation for this exercise is that the topics covered here are typically treated separately in textbooks (e.g., Mas-Colell et al., 1995; Varian, , making the conceptual unity of approach less than obvious.

The first step is to develop a framework for understanding how decisions are made in the face of uncertainty. This framework can then be used to examine how the presence of uncertainty affects the working of markets. One key idea that is related to uncertainty is that of risk. Risk is different from uncertainty. While there are several aspects of this difference (and there is not always total agreement on what the two terms should mean), an illustration of this difference can be given from thinking about insurance. Suppose I face some probability of being in an accident. I am uncertain about whether I will have an accident or not. The accident involves the risk of a loss. However, if I can purchase complete insurance coverage, there is still uncertainty about the accident (and the uncertainty is exactly the same if my actions do not affect the probability), but my risk has gone away. The risk is now borne by the insurer, but it may be pooling many such risks, so its overall or average risk may be lower. This essay will provide the analytical machinery to explore such market transactions.

Information is also related to uncertainty. In the case of an accident, there may be no information I can gather that affects my uncertainty – the probabilities I assign to having an accident or not. But suppose I am a farmer and my uncertainty is about the rainfall during the growing season. If I can get information about the likelihood of various levels of rainfall, I can modify my assessment of the uncertainty, and this may affect my planting and other farming decisions. In some cases, different economic actors have different information. I may know my exact probability of having an accident (I know how good a driver I am), but the insurance company may not be able to tell whether I am a good or bad driver. This asymmetry of information can exist even if there is no inherent uncertainty. I may know perfectly well whether my car is sound or is a poor quality "lemon." However, a prospective buyer does not have this information: she is uncertain about the quality of my car, even though I know for sure. Asymmetries of information can have serious implications for how well markets work, and we explain in this essay how such situations can be analyzed.

Decision Making under Uncertainty

There are two ways of building the framework for analyzing decision making under uncertainty. They lead to basically the same approach, but since they are both used, and have some subtle differences, it is useful to be aware of both starting points. Roughly, they correspond to two different views of uncertainty, familiar from basic statistics. The fundamental description of uncertainty is in terms of an outcome space. For example, if I toss a coin once, there are two possible outcomes, heads (H) and tails (T). If I toss the coin twice, there are four possible outcomes: HH, HT, TH, and TT. The outcome space is the set of all possible outcomes. In the case of two coin tosses, it is {HH, HT, TH, TT}.

Often, in using statistics, we work with random variables. Technically, a random variable is defined by a mapping from the outcome space to the space of real numbers (or a higher dimensional real space, if the random variable is multi-dimensional). Suppose that I win $ 1 every time a head comes up in a coin toss, and win or lose nothing otherwise. Then the random variable that represents my monetary winnings with a single coin toss maps H to 1 and T to 0. For the case of two tosses of the coin, HH is mapped to 2, HT and TH to 1, and TT to 0.

Continuing with the case of two coin tosses, we can describe the random variable by the three values it takes, and the associated probabilities, (0, ¼; 1, ½; 2, ¼), or, listing outcomes and probabilities separately, (0, 1, 2; ¼, ½, ¼), where the order obviously has to be preserved to match the numerical values and probabilities. We will use this second representation. We will refer to this vector of numerical outcomes and probabilities as a *lottery*. We can generalize the concept to include multi-dimensional outcomes, and any number of possibilities.[1] Thus if a consumer is faced with getting any one of N different consumption bundles, the lottery can be written ($x^1,\dots, x^N; \pi^1,\dots, \pi^N$), where the x^n are different consumption bundles, and the p^n are the associated probabilities. Of course, the probabilities sum up to 1, $\sum_1^N \pi^n = 1$.

For developing a theory of individual decision making under uncertainty, working with lotteries (essentially, random variables) is convenient. However, in conceptualizing how markets work in situations of uncertainty, it is useful to start with underlying outcome spaces. The fundamental building block of markets in the face of uncertainty is a contingent contract, where the contingencies are exogenously determined, and not influenced by economic actors' decisions. We call these exogenous contingencies *states of the world* or *states of nature*. Hence, if all the uncertainty in the world is the outcome of two coin tosses, there are four possible states of the world, the four components of the outcome space. There can be different consumption vectors associated with each member of the set of states of the world, but the set remains the same. Thus, the payoff structure above associates an amount 2 with HH, 1 with HT, 1 with TH, and 0 with TT, with each outcome having a probability of ¼. If I receive a dollar only for the first head, and nothing otherwise, the amounts are 1, 1, 0 and 0, but the probabilities are unchanged. In contrast, with the lottery representation, the vector is (0, 1; ½, ½) – this is a more compact representation, but varies with the payoffs. Of course, it could be also expressed as (0, 1, 2; ½, ½, 0).

Von Neumann-Morgenstern Expected Utility

If we begin with lotteries, we assume that an individual has preferences over these lotteries. We will call the numerical outcomes associated with a lottery *prizes*. Three basic assumptions about these preferences are:

A1 Getting a prize with probability one is the same as getting that prize with complete certainty.

[1] The outcomes of a lottery could also be further lotteries – this possibility will be addressed when we consider preferences over lotteries.

A2 Preferences are not dependent on the ordering of a lottery's components.

A3 If some or all of the prizes in a lottery are lotteries, preferences depend only on the net probabilities of receiving the different prizes.

The first two assumptions are almost trivial. A1 simply connects preferences over lotteries to standard preferences over certain consumption bundles. A2 and A3 involve perceptions, and there is some evidence that the last assumption can be violated in actual behavior. Since it involves an aberration, we will not question it further. It is important for reducing the prizes to be numerical outcomes such as consumption bundles.

As in the case of preferences over consumption bundles, we can assume that some rationality is built into the consumer's preference ordering over lotteries. Thus, we assume that these preferences are complete, reflexive and transitive. Further assumptions of continuity and monotonicity, similar to those imposed in the case of certainty, ensure that there is a continuous utility function that represents preferences over lotteries. Formally, if L_j denotes a lottery we have:

$$L_j \succ L_k \quad iff \quad V(L_j) > V(L_k) .$$

Representing preferences over lotteries by a utility function does not really simplify analysis, nor is the concept of preferences or utility with respect to lotteries very intuitive. Luckily, there is a further simplification that provides a much more intuitive analysis of decision-making in the face of uncertainty. Another assumption is needed, however, in addition to all those we have implicitly or explicitly imposed.

A4 (The independence axiom)

A preference relation $\succsim$ on the space of lotteries satisfies the independence axiom if for all L_j, L_k, L_m and $\alpha \in (0, 1)$ we have

$L_j \succsim L_k$ *iff* $\alpha L_j + (1\text{-}\alpha)\, L_m \succsim \alpha L_k + (1\text{-}\alpha)\, L_m$

This says that mixing in any third lottery does not affect the original preference ordering of the original two lotteries, or that the original preference ordering is independent of adding in the new lottery.

With the independence axiom, we have the result that the utility function has a particular form, namely that it is linear in the probabilities that are contained in the lottery. Suppose that the lottery L can be written out fully as $(x^1,\dots, x^N; \pi^1,\dots, \pi^N)$. The result states that

$$V(x^1,\dots, x^N; \pi^1,\dots, \pi^N) = \sum_1^N \pi^n U(x^n)\,,$$

for some function U.

Note that this new utility function U is defined over the space of certain outcomes. This makes the simplification a very intuitive one for analyzing decision-making in the face of uncertainty. The function U does not represent preferences over the different uncertain situations (the lotteries) faced by the decision maker. That role is played by the original utility function, which is V.

The last equation says that the original preferences over lotteries can be represented by the functional form on the right hand side of the equality. Sometimes the entire right hand side is referred to as the von Neumann-Morgenstern expected utility form. In other usage, just the function U is termed the von Neumann-Morgenstern utility function. U is also called the Bernoulli utility function.

Savage Expected Utility

An alternative derivation of an expected utility representation of preferences begins with the underlying outcome space rather than numerical values associated with the outcomes (e.g., {H,T} for a coin toss, rather than possible payoffs (1, 0)). In general, we can assume a set of states of nature {1, …, S}, and suppose that there is a preference relation

$\succsim$ defined over all random variables that map from this set to the space of real numbers, or vectors of real numbers. In the latter case, if the vectors are restricted to be nonnegative, they are again interpretable as consumption bundles.

The Savage derivation of expected utility starts from this framework. It generalizes the earlier framework, because it allows probabilities to be subjective, and therefore to differ across individuals. These probabilities are derived as part of the expected utility derivation, rather than being given. Furthermore, the Bernoulli utility function can differ across states of the world. Getting to the expected utility representation of preferences requires a generalized version of the independence axiom (in addition to completeness, transitivity and continuity). Since the end result is very similar in form to that of the earlier approach, we omit any further mathematical details or discussion.[2] In the following, we will put aside the possibility of the utility function varying with the state of the world.

Uniqueness of Bernoulli Utility Function

In the case of representing certain consumption bundles by a utility function, we saw that the utility function was not an exact numerical scale. Thus, any positive, strictly monotonic transformation of the utility function would represent the same preferences. This implied that diminishing marginal utility could not be a core concept for the theory of consumer choice. Instead, we relied on the marginal rate of substitution and its properties in analyzing consumer behavior.

In the case of the Bernoulli utility function, which forms part of the expected utility representation of preferences in situations of uncertainty, the only freedom we have in transforming the utility function is applying a positive linear (affine) transformation. Thus, if U is a Bernoulli utility function that is part of an expected utility function that represents preferences over lotteries or random variables, then any transformation $\beta U + \gamma$, $\beta > 0$, represents the same preferences. Furthermore, only such transformations are allowed. Hence the utility function is unique up to a positive linear transformation. An

[2] See Kreps (1990) or Mas-Colell et al. (1995) for mathematical details.

important implication of this is that the sign of the marginal utility is preserved under admissible transformations: in the case of Bernoulli utility functions, marginal utility has an economic meaning and implications.

Risk and Risk Attitudes

Having introduced a convenient framework for analyzing decision making in uncertain situations, we can now begin applying it. To proceed with economic interpretation and applications, it is useful to simplify to the case where there is a one-dimensional numerical outcome. This can be interpreted as money payoffs, as in the coin toss example. We will assume that the payoff is a continuous variable, so that we can continue to use calculus tools. In a similar spirit to our analysis of choice under certainty, we will also assume that the utility function is increasing and differentiable (as many times as we need).

While there are exceptions, much human behavior is guided by a dislike for uncertainty, in the following sense. A decision maker who dislikes uncertainty will prefer to receive the expected value of a lottery for sure to facing the uncertain outcome of a lottery. Mathematically,

$$U\left(\sum_1^N \pi^n x^n\right) \quad > \quad \sum_1^N \pi^n U(x^n) \tag{1}$$

A decision maker whose Bernoulli utility function satisfies this property is called *risk averse*. Risk aversion is equivalent to always preferring the expected value of a lottery to facing the lottery itself. In this sense, risk and uncertainty are essentially identical in this theoretical framework.

A general mathematical result (Jensen's inequality) tells us that the inequality holds if and only if the function U is strictly concave. If U is twice differentiable, and x is a real number, then this is equivalent to the following inequality for the second derivative: $U''(x) < 0$.

If the inequality is reversed, then U is strictly convex, and the individual is *risk loving*. If the relationship holds with equality, then the individual is *risk neutral*, and the utility function is linear.

It is entirely possible that a decision maker could display different risk attitudes over different ranges of the outcome variable. This might explain why people gamble as well as buy insurance. However, the true explanation for such combinations of behavior may lie in aspects of preferences in uncertain situations that do not fit the expected utility model at all. In practice, we will work almost exclusively with examples where individuals are risk averse or risk neutral, on the grounds that this covers the vast bulk of economically important situations (portfolio choice in asset allocation, smoothing consumption over time, insurance, and so on).

Next we provide some conceptual tools for analyzing risk averse behavior in the face of uncertainty. These concepts will also help to further understand the relationship between risk and uncertainty, in the context of the standard theoretical framework.

Markets with Uncertainty

First, consider the case where there are two states of the world, 1 and 2. An individual's expected utility with money as the payoff is $\pi^1 U(x^1) + (1 - \pi^1)U(x^2)$. We can also think of this as a utility function $V(x^1, x^2; \pi^1)$. If we draw the indifference curves of V in x^1-x^2 space, they must be downward sloping as long as U is strictly increasing in x (and hence V is strictly increasing in each x). In fact, we can show that these indifference curves are strictly convex precisely when the individual is risk averse.

The behavior of risk averse consumers is therefore amenable to the mathematical tools used for choice under certainty. In particular, a consumer's (expected) utility-maximizing choice is determined by the tangency of the budget line to the highest affordable indifference curve. However, the interpretation of the goods and prices is different from the certainty case. In the case of two states of the world, with only one good (money) the two goods in the expected utility function are contingent amounts of money. Hence x^1 is the amount of money that the consumer will have *if and only if* state 1 comes to pass. If state 2 occurs, the consumer will have an amount of money x^2 to consume.

How will the individual's prospective consumption be financed? Let us first suppose that the consumer has some money now to allocate toward consumption once the uncertainty is resolved. This could be denoted by I as before. The money cannot be stored, so must be spent now, or used to get commitments to money that will be received once the uncertainty is resolved. Also, suppose that she can spend x^0 now, which gives her utility $U(x^0)$. The price of consumption now is just 1, i.e., good 0 is the numeraire good. Suppose that the price of consumption in state s is p^s, $s = 1, 2$. Her choice problem is now

$$\text{Max} \quad U(x^0) + \pi^1 U(x^1) + (1-\pi^1)U(x^2) \quad \text{subject to} \quad x^0 + p^1x^1 + p^2x^2 = I \qquad (2)$$

What is the nature of the purchase of "money in state 1"? The consumer pays \$ p^1 now, sacrificing that much in current consumption, and in return will receive 1 unit of money (\$ 1) if state 1 actually occurs. The fulfillment of the other side of the contract is contingent on state 1 occurring. If state 2 is what actually comes about, then this contract pays nothing. A similar description can be given for contracts that pay off in state 2.

What is exchanged when the consumer pays money now? She receives a promise to pay if the relevant state occurs. Hence, if she pays \$ p^1 now, she can get a claim that guarantees her \$ 1 if state 1 actually occurs – a *state-1 contingent claim*. Similarly, she can buy state-2 contingent claims at \$ p^2 each. She can also guarantee a dollar after the resolution of the uncertainty, by purchasing one unit of each type of contingent claim, which costs her \$ $(p^1 + p^2)$.

Now suppose that the consumer has no money now, but knows that she will have some money in each of the states. If state s occurs, her endowment of money will be ω^s. Thus, her contingent endowment vector is (ω^1, ω^2). To keep matters simple, suppose that she does not consume anything or get any utility until the uncertainty is resolved. Now her choice problem is

$$\text{Max} \quad \pi^1 U(x^1) + (1-\pi^1)U(x^2) \quad \text{subject to} \quad p^1x^1 + p^2x^2 = p^1\omega^1 + p^2\omega^2 \tag{3}$$

Now the nature of the consumer's market transactions is as follows. Unless she is best off with her contingent endowment vector, she will sell contingent claims for one of the states, and buy contingent claims for the other state. Her motivation could be as follows. Suppose she is risk averse and $\omega^1 \gg \omega^2$. Thus, she will be much worse off if state 2 occurs than if state 1 occurs. Depending on the probabilities and prices, she may plausibly choose to sell state 1 contingent claims and buy state 2 contingent claims, to reduce the difference across her consumption in the two states – her exposure to risk.

Measuring Risk and Risk Aversion

Inequality (1) notes that for a risk averse decision-maker, the expected utility of a lottery is less than the certain utility of the expected value of the lottery. Therefore, by continuity, there is some amount $C(L;U) \equiv C\left(x^1,\ldots,x^{\mathrm{N}};\pi^1,\ldots,\pi^{\mathrm{N}};U\right)$ that is smaller than $\sum_1^N \pi^n x^n$, such that

$$U(C(x^1,\ldots,x^{\mathrm{N}};\pi^1,\ldots,\pi^{\mathrm{N}};U)) = \sum_1^N \pi^n U(x^n) \tag{4}$$

The amount $C\left(L;U\right)$ is called the *certainty equivalent* of the lottery. It is the amount that will make the decision maker indifferent between taking a certain amount and facing the lottery. For a risk lover, $C(L;U) > \sum_1^N \pi^n x^n$, while for a risk neutral person, $C(L;U) = \sum_1^N \pi^n x^n$.

We can also define the difference between the certainty equivalent and the expected value as the *risk premium*. Thus, the risk premium is defined by

$$\rho(L;U) = \sum_1^N \pi^n x^n - C(L;U) \tag{5}$$

or, equivalently, by

$$U\left(\sum_1^N \pi^n x^n - \rho(L;U)\right) \quad = \quad \sum_1^N \pi^n U(x^n) \tag{6}$$

In an analogy to an insurance premium, the risk premium is the maximum amount that a risk averse person will pay to avoid the uncertainty she faces. For a risk loving person, the risk premium can also be defined as above, but it is negative – she will pay to take the risk or face the uncertainty.

The certainty equivalent and risk premium depend on the uncertainty faced (the lottery), as well as the decision-maker's attitude toward risk, the latter being captured in the utility function. Next, we explore precisely how the attitude toward risk is measured by the shape of the utility function.

A simple first idea for measuring risk aversion is to use the second derivative of the Bernoulli utility function – this extends the idea that the sign of the utility function provides a dividing line between risk aversion and risk loving. For risk averse people, $U''(x) < 0$. We can multiply by -1 to make it a positive number, so that a higher value would indicate greater risk aversion. Hence, $-U''(x)$ would be the level of risk aversion at consumption x (this is a local measure, and can vary along the utility function). The problem is that the utility function is not uniquely defined. Anyone with $\beta U + \gamma$, $\beta > 0$, has exactly the same attitude toward risk at every point and in every situation. Hence, using the second derivative is misleading, because linear transformations change its magnitude.

Normalizing the measure of risk aversion takes care of the problem. The *coefficient of absolute risk aversion* is therefore defined as

$$A(x) = -\frac{U''(x)}{U'(x)} \tag{7}$$

This measure stays the same even if the utility function is subjected to a linear transformation, since the β in the numerator and denominator will cancel out.

The coefficient of absolute risk aversion is related to the risk premium as follows. In (6), denote the expected value of the lottery by $\bar{x}$, and the risk premium simply by ρ. So (6) can be written as

$$U(\bar{x}-\rho) \quad = \quad \sum_1^N \pi^n U(x^n) \tag{8}$$

Suppose that the risk is small, so that the risk premium is small. Now consider the Taylor series approximation (we go to second order because the term in the first derivative turns out to be zero):

$$U(x) \approx U(\bar{x}) + U'(\bar{x})(x-\bar{x}) + \frac{1}{2}U''(\bar{x})(x-\bar{x})^2 \tag{9}$$

Using this expression for each of the terms on the right hand side of (10.8), we get

$$\begin{aligned}\sum_1^N \pi^n U(x^n) &\approx \sum_1^N \pi^n U(\bar{x}) + U'(\bar{x})\sum_1^N \pi^n (x^n - \bar{x}) + \frac{1}{2}U''(\bar{x})\sum_1^N \pi^n (x^n - \bar{x})^2 \\ &= U(\bar{x}) + \frac{1}{2}U''(\bar{x})\sigma^2\end{aligned} \tag{10}$$

The second term is zero from the definition of $\bar{x}$, and σ^2 is the variance of x.

If we take the (first-order) Taylor series expansion of the left hand side of (10.8), we get $U(\bar{x}-\rho) = U(\bar{x}) - U'(\bar{x})\rho$. If we equate (allowing for the approximation) these last two right hand side expressions, and simplify, we finally get

$$\rho \approx -\frac{1}{2}\frac{U''(\bar{x})\sigma^2}{U'(\bar{x})} = \frac{1}{2}A(\bar{x})\sigma^2 \tag{11}$$

The risk premium is higher, the higher is the individual's coefficient of risk aversion, and the greater is the uncertainty as measured by the variance of the lottery.

An alternative measure of risk aversion modifies the coefficient of absolute risk aversion to take account of the level of consumption. The *coefficient of relative risk aversion* is therefore defined as

$$R(x) = -\frac{xU''(x)}{U'(x)} \tag{12}$$

The coefficient of relative risk aversion is also the elasticity of the marginal utility of consumption. Since we began our discussion of measurement with the connection between the level of consumption and the behavior of marginal utility, calculating this elasticity is a natural approach to measuring risk aversion.

It is easy to show that a utility function of the form $-e^{-Ax}$ (where A>0) implies a constant absolute risk aversion coefficient equal to A. Similarly, a utility function of the form x^{1-R} (where 1 > R > 0) has a constant relative risk aversion coefficient equal to R, while $\ln x$ has $R(x) = 1$.

Riskiness

The risk premium, as we saw in (11), depends on the decision-maker's attitude to risk and the riskiness of the situation she faces, where riskiness is approximately captured by the variance of the lottery. The idea of riskiness can be made more general. In doing so, it is convenient to work with distribution functions. A distribution function is the cumulative probability function, and has the virtue of being defined across discrete and continuous random variables. If the random variable is continuous, the distribution function is also continuous, and if the random variable is discrete, it is a step function. It is also convenient to scale outcomes so that they always lie within an interval [a,b], where a and b are finite. We shall suppress these limits of integration in the subsequent notation.

Now consider two lotteries, described by their distribution functions $F(x)$ and $G(x)$. The random variable represented by the outcome x could be the level of money, so that a decision maker views a higher x as a more favorable outcome. Suppose that $F(x) \le G(x)$ over the range of x. This means that for any given outcome, $F(x)$ has at least as great a chance as $G(x)$ of producing an outcome at least as good as the given outcome. Not surprisingly, we have that

$$\begin{aligned} &F(x) \le G(x) \quad \text{for every } x \\ &\Rightarrow \int U(x)dF(x) \ge \int U(x)dG(x) \quad \text{for every nondecreasing } U(x) \end{aligned} \tag{13}$$

Less obviously, the reverse implication also holds. Only if $F(x) \le G(x)$ is it true that the expected utility is higher with $F(x)$ than with $G(x)$ for *every* decision-maker who prefers more of x to less.

In this case, we say that $F(x)$ *first-order stochastically dominates* $G(x)$.

First order stochastic dominance gives a partial ordering of probability distributions that will be unanimously agreed on by all decision-makers. What about the subset of risk averse individuals? It turns out that there is another partial ordering of probability distributions that works in this case. First order stochastic dominance involves shifting probability weight from lower to higher outcomes, so that one distribution function is never above the other. If we weaken this requirement so that the distribution functions can cross, but there is a limit to how much probability weight can be moved around. Specifically, we will say that $F(x)$ *second-order stochastically dominates* $G(x)$ if the following holds:

$$\int_a^x F(t)dt \le \int_a^x G(t)dt \quad \text{for all } x \tag{14}$$

This is clearly a weaker condition than first-order stochastic dominance, since if $F(x)$ first-order stochastically dominates $G(x)$ then $F(x)$ also second-order stochastically dominates $G(x)$. Condition 910.14) compares the areas under the distribution functions, so that they may cross, but only if the cumulative area under $F(x)$ does not exceed that under $G(x)$.

We have the result that, if $F(x)$ second-order stochastically dominates $G(x)$, it is the case that

$$\int U(x)dF(x) \ge \int U(x)dG(x) \quad \text{for every nondecreasing concave } U(x) \tag{15}$$

In other words, second-order stochastic dominance gives a (partial) ordering of distributions that every risk averse or risk neutral decision maker will agree on.

Second order stochastic dominance ensures that the mean of the distribution is not any lower. It is useful to also impose the restriction that the means be the same for the two distributions (sometimes this condition is assumed in the stochastic dominance definition, but is strictly not part of it). This additional condition allows separation of the impact of differences in average return from that of higher risk. The condition of equal means is

$$\int_a^b F(t)dt = \int_a^b G(t)dt \tag{16}$$

If (16) also holds, then we say that $G(x)$ is a *mean-preserving spread* of $F(x)$. The result in (15) also justifies the assertion that $G(x)$ is *at least as risky as* $F(x)$, since all risk averse individuals will find the uncertain situation described by $F(x)$ at least as good as that described by $G(x)$. If the relevant inequality is strict, then we can say that $G(x)$ is *riskier* than $F(x)$. Note that this property is more general than that of a higher variance; though if a distribution is completely defined by its mean and variance, then the two concepts coincide: increasing risk is the same as higher variance.

Insurance and Asset Markets

We will now connect the consumer's behavior under uncertainty to some real world market institutions and phenomena. We begin with the choice problem as described in (3):

$$\text{Max} \quad \pi^1 U(x^1) + (1-\pi^1)U(x^2) \quad \text{subject to} \quad p^1x^1 + p^2x^2 = p^1\omega^1 + p^2\omega^2 \qquad (17)$$

We assume that the consumer is risk averse. The first order conditions for utility maximization, where λ is the Lagrange multiplier associated with the budget constraint, are:

$$\begin{aligned} \pi^1 U'(x^1) - \lambda p^1 &= 0 \\ (1-\pi^1)U'(x^2) - \lambda p^2 &= 0 \end{aligned} \qquad (18)$$

These conditions can be rearranged to give the usual condition for the marginal rate of substitution:

$$\frac{\pi^1 U'(x^1)}{(1-\pi^1)U'(x^2)} = \frac{p^1}{p^2} \qquad (19)$$

The ratio of probabilities on the left hand side is the odds of the two states, while the other part of the right hand side is the ratio of marginal utilities. If the odds equal the price ratio, then the ratio of marginal utilities is 1, and with risk aversion, it must be true that consumption in the two states of the world is equalized.

What is the significance of the equality of the odds and the price ratio? Suppose that $\omega^1 > \omega^2$, and that the consumer sells $\omega^1 - x^1$ state 1 contingent claims, while buying

$x^2-\omega^2$ state 2 contingent claims. From the budget constraint, $p^1(\omega^1-x^1)=p^2(x^2-\omega^2)$. Suppose that the counterparty to the trades is an insurance company. It also satisfies this equality, so has no net inflow or outflow before the uncertainty is realized. However, once the uncertainty is resolved, it either has to provide $x^2-\omega^2$ units of money if state 2 occurs, with probability $(1-\pi^1)$, or receives ω^1-x^1 units of money with probability π^1. Hence, its expected future profit is $\pi^1(\omega^1-x^1)-(1-\pi^1)(x^2-\omega^2)$. This is zero if and only if the odds equal the ratio of contingent claim prices. If there is competition, so that expected profits are zero, then that must be true. Hence, a competitive insurance market (with no operating costs) will lead to complete insurance.

We have interpreted the consumer's market transactions as a kind of insurance. Real world insurance contracts are not defined in terms of state contingent claims. However, the transaction above is equivalent to a standard insurance contract. To show this, we introduce some new definitions. Let $d=\omega^1-\omega^2$ be the damage that the individual suffers in state 2 (suppose this is a state in which she has an accident or illness or other adverse condition), against which she wishes to purchase insurance. Let P be the insurance premium she pays per unit of coverage, and I the indemnity or compensation she receives if state 2 occurs. Thus, I is the number of units of insurance, and her payment is PI. Now her choice problem is:

$$\text{Max} \quad \pi^1 U(\omega^1-PI)+(1-\pi^1)U(\omega^1-d-PI+I) \tag{20}$$

The only choice variable is I, and her first order condition is

$$-P\pi^1 U'(\omega^1-PI)+(1-\pi^1)(1-P)U'(\omega^1-d-PI+I)=0 \tag{21}$$

Rearranging and writing more compactly, this becomes

$$\frac{\pi^1 U'(x^1)}{(1-\pi^1)U'(x^2)}=\frac{1-P}{P} \tag{22}$$

Hence, if $\frac{1-P}{P} = \frac{p^1}{p^2}$, the first order conditions are identical for the contingent claims case and the traditional formulation of insurance. Also, the budget constraint can be written as

$$\begin{aligned} & x^1 = \omega^1 - \frac{p^2}{p^1}(x^2 - \omega^2) = \omega^1 - \frac{P}{1-P}(x^2 - \omega^2) \\ \text{or} \quad & (1-P)(x^1 - \omega^1) = -P(x^2 - \omega^2) \\ \text{or} \quad & -(1-P)PI = -P(1-P)I \end{aligned} \tag{23}$$

Hence the two budget constraints are also identical, so the solutions to the two consumer choice problems are identical.

We can think of state contingent claims as special kinds of assets, ones that payoff in specific states of the world. Next we show how the contingent claim formulation also covers more general kinds of assets.

A state 1 contingent claim can be described by the vector of returns that it promises in each state of the world, which is (1,0). Similarly, a state 2 contingent claim is described by its return vector (0,1). Since holding a state contingent claim now involves a specification of payments once some uncertainty is resolved, such claims are assets, or securities. An asset with return vector (1,1) is a safe, or fixed-income asset (the income in this case being zero). In general, any asset is equivalent to a return vector $r = (r^1, r^2)$. These returns could be negative, positive or zero. Let us denote a security with return vector *r* by S_r. Let S_1, S_2 denote state 1 and state 2 contingent claims, respectively. Then, based on the relationship between the return vectors, we can see that holding one unit of the security S_r is equivalent to holding r^1 units of S_1 and r^2 units of S_2. We can express this as

$$S_r = r^1 S_1 + r^2 S_2 \tag{24}$$

We can also assert a relationship between the price of the security S_r and the prices of the two state contingent claims. If markets for all three securities are available, then an individual could make infinite profits by appropriate trades (called *arbitrage*) of the securities. In the presence of the possibility of arbitrage, therefore, it must be true that

$$p^r = r^1 p^1 + r^2 p^2 \tag{25}$$

For example, if the left hand side of (25) is greater, then the individual could sell any number of units of S_r, while simultaneously buying the same number of units of S_1, S_2, and make a profit on that trade, without any risk or net outlay of funds.

Another important property of general securities is that if there are as many securities as states of the world, and the return vectors of all the securities are linearly independent (so that no return vector can be expressed as a linear combination of the others), then any pattern of returns that can be achieved with a full set of state contingent claims can also be achieved with this set of securities.

Essentially, this is a property of linear spaces. A linear space of dimension *N* can be spanned by *N* linearly independent vectors. Such a set of vectors is said to form a *basis* for the space, and to *span* the space. Any N linearly independent vectors will do. The full set of state contingent claims is a special case, of an *orthonormal basis*: the vectors are all orthogonal to each other, and have length 1.

We will illustrate with the case of two states, and choose two kinds of general securities, a safe asset S_0, with return vector $f = (1,1)$, and a risky asset, S_r, with return vector $r = (r^1, r^2), r^1 \neq r^2$. Consider any other security S_t, with return vector $t = (t^1, t^2)$. Then $S_t = t^1 S_1 + t^2 S_2$. The returns to security S_t can be obviously achieved by holding t^1 units of S_1 and t^2 units of S_2. Is there a holding of S_0 and S_r that achieves this pattern of returns?

Suppose the quantities of the two securities are q_0 and q_r. The pattern of returns that this holding implies is $(q_0 + q_r r^1, q_0 + q_r r^2)$. Hence, we require that the following two equations are satisfied:

$$\begin{aligned} q_0 + q_r r^1 &= t^1 \\ q_0 + q_r r^2 &= t^2 \end{aligned} \tag{26}$$

As long as $r^1 \neq r^2$, the two equations are linearly independent, and we can always solve them:

$$q_0 = \frac{(t^2 r^1 - t^1 r^2)}{(r^1 - r^2)},\ q_r = \frac{(t^1 - t^2)}{(r^1 - r^2)} \tag{27}$$

Hence, with the ability to purchase a safe and a risky asset, and two states of the world, the decision maker can achieve any desired pattern of returns, and hence consumption.

We can also note that if all four securities S_0, S_1, S_2, S_r are available, then the arbitrage argument implies that (25) holds as well as $p^0 = p^1 + p^2$.

Remark: While consumption levels must be nonnegative (and positive for an interior solution), there is nothing to prevent quantities of assets from being negative. Hence short sales are permitted in the above formulation. The constraint that the decision maker can always deliver on promises is inherent in the budget constraint.

Now consider the choice problem in (3) (there is no consumption before the uncertainty is resolved), but with a fixed money income rather than endowments (as in (2)). The choice problem for the purchase of contingent claims is:

$$\text{Max} \quad \pi^1 U(x^1) + (1-\pi^1)U(x^2) \quad \text{subject to} \quad p^1 x^1 + p^2 x^2 = I \tag{28}$$

Our argument about spanning the space of returns or consumption patterns tells us that this problem is equivalent to the following:

$$\begin{aligned} &\text{Max} \quad \pi^1 U(x^1) + (1-\pi^1)U(x^2) \\ &\text{subject to} \quad p^0 q^0 + p^r q^r = I, \quad x^1 = q_0 + q_r r^1, \quad x^2 = q_0 + q_r r^2 \end{aligned} \tag{29}$$

Substituting in the second and third constraints, and therefore making q_0 and q_r the choice variables, the first-order conditions are

$$\begin{aligned} &\pi^1 U'(x^1) + (1-\pi^1)U'(x^2) - \nu p^0 = 0 \\ &r^1 \pi^1 U'(x^1) + r^2 (1-\pi^1)U'(x^2) - \nu p^r = 0 \end{aligned} \tag{30}$$

It is easy to check, using the relationships among prices implied by arbitrage, that these conditions are equivalent to the conditions for the choice problem with two state contingent claims. This is just a reiteration of the conclusion that two assets with linearly independent return vectors will allow the achievement of any pattern of returns with two states of the world, just as would two state contingent claims.

The first order conditions in (30) can also be written more compactly as:

$$\begin{aligned} &E_r U'(q^0 + q^r r) - \nu p^0 = 0 \\ &E_r [r U'(q^0 + q^r r)] - \nu p^r = 0 \end{aligned} \tag{31}$$

Here r is a random variable, and E_r denotes the expectation over r. If we normalize $p^0 = 1$ (as would be required if there were also initial pre-uncertainty consumption from the same income I), we can simplify (31) to:

$$
\begin{aligned}
&E_r[rU'(q^0+q^r r)] = p^r E_r U'(q^0+q^r r) \quad \text{or} \\
&E_r[(r-p^r)U'(I+q^r(r-p^r)] = 0
\end{aligned}
\tag{32}
$$

The last expression comes from substituting in the budget constraint, and is expressed in terms of the net return on the risky asset, $(r-p^r)$.

We started with the case of two states of the world, but the last expressions are quite general, and apply to the choice problem with any number of states, or even a continuum – so that the random variable is continuous – as long as there is a single safe and a single risky asset. In the case of two states of the world, this was equivalent to having complete contingent claim markets (one contingent claim for each state), but obviously in the more general case of many states, the equivalence does not hold.

Typically, in finance, choice problems are immediately framed in terms of continuous random variables, and conditions such as (32) are derived directly. We have taken the indirect route here, to show the connection of portfolio choice to the more fundamental idea of contingent claim markets. In turn, the contingent claim formulation is the conceptual generalization to the uncertainty case of the theory of competitive markets under certainty.

We close this section with another illustration of the connection between finance and the contingent claim-based formulation of assets and choice under uncertainty. Suppose first that there is a single asset, S_r, with return vector $r=(r^1,r^2), r^1 \neq r^2$. There is no safe asset, nor any contingent claims. Then a decision maker would not be able to achieve any arbitrary pattern of returns – markets are no complete.

Now introduce a new asset, which is an option to buy security S_r at a pre-specified price, say $\overline{p}$. This purchase can be made after the state is revealed, but before the returns are paid. It makes sense to exercise this option only if the return exceeds this price, i.e., $r^s > \overline{p}$. The return vector for this option is therefore given by

$r(\overline{p}) = (\max\{r^1 - \overline{p}, 0\}, \max\{r^2 - \overline{p}, 0\})$. For example, if $r^1 > \overline{p} > r^2$, then $r(\overline{p}) = (r^1 - \overline{p}, 0)$. This option therefore has a return vector that is linearly independent of the original return vector of S_r. The option is not a fundamental security, but instead a *derivative security*, or just a *derivative*. Nevertheless, S_r and the option to purchase S_r together allow spanning of the space of all possible returns: this is equivalent to a complete set of contingent claims. In general, derivatives can expand the set of achievable returns, even if they do not result in complete markets.

Information

We have assumed so far that the nature of the uncertainty faced by an individual is given and exogenous. It is possible, however, that some information becomes available before a decision has to be made. In the decision problem (10.3), the consumer chooses purchases and sales of contingent claims for the two possible states of the world. Suppose that before this decision is made, this consumer receives new information about the likelihood of the two states. In general, this can be in the nature of an imperfect signal. A signal is defined by the probabilities of observing signal values. If there are two states, a natural assumption is that the signal takes on two values, also indexed 1 and 2. However, the signal may be imperfect, and the value 1 may be observed even if the underlying state of the world is 2. The signal is therefore defined by the probability matrix

$$\sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{21} & \sigma_{22} \end{pmatrix}, \quad \text{where} \quad \sigma_{ts} \equiv \Pr(\text{signal} = t \,|\, \text{state} = s) \tag{33}$$

Hence, the posterior probability that the state of the world is *s*, when the observed signal is *s*, is given by

$$\pi_s \,|_s = \frac{\sigma_{ss}\pi_s}{\sigma_{ss}\pi_s + \sigma_{st}\pi_t} = \frac{\sigma_{ss}\pi_s}{(\sigma_{ss} - \sigma_{st})\pi_s + \sigma_{st}} \tag{34}$$

Note that if the signal probabilities are all ½, then the posterior probability is the same as the initial probability: the signal is completely uninformative, since it does not affect the decision maker's beliefs. On the other hand, if a correct signal is more likely than an incorrect signal, then $\sigma_{ss} > \frac{1}{2} > \sigma_{st}$, and the posterior probability of state *s* is higher when that state is signaled, $\pi_s |_s > \pi_s$. In the limit, if the signal is perfectly informative, then $\sigma_{ss} = 1, \sigma_{st} = 0$ and $\pi_s |_s = 1$.

Consider for simplicity the perfectly informative case, and assume that only this consumer receives the perfect signal of which state of the world will occur. The individual is too small relative to the market to have any impact on prices, so the market prices of contingent claims are unaffected by what this consumer does. With perfect knowledge, the consumer will buy the maximum number of state 1 contingent claims if the signal value is 1, and similarly the maximum number of state 2 contingent claims if the signal value is 2. The perfect foresight gives the consumer the opportunity to avoid any risk at all. The maximized expected utility (viewed before the signal is received) in this case becomes

$$\pi^1 U((p^1\omega^1 + p^2\omega^2)/p^1) + (1-\pi^1)U((p^1\omega^1 + p^2\omega^2 / p^2)) \tag{35}$$

Since consumption in each state is greater than the endowment in that state, this is clearly better for the consumer than the solution to (3), where the consumer will be worse off *ex post* in one of the states (but better off in expected utility terms).

The key idea in the example above is that the market does not respond to this consumer's additional information, because it is not known that the consumer has this special knowledge, nor can anyone else infer anything about that information or its existence from the consumer's behavior.

What happens if the market knows that the consumer is informed? Then any attempt by the informed consumer to trade will signal what she knows, and other market participants could conceivably adjust their behavior to take account of this inferred knowledge. How precisely this inference takes place and is incorporated in market prices is a complicated matter, and a delicate one in the sense that changes in the precise assumptions can have significant impacts on the market equilibrium.

To pursue these ideas further, we will use the specific case of a competitive insurance market, and assume that there are two types of individuals, high risk and low risk. Each type of individual knows her risk level, as defined by the probability of having an accident. This is a bit different from the earlier perfect signal formulation, in the sense that individuals are perfectly informed about their risk level, but do not know for sure if they will have an accident or not. Their information is akin to an imperfect signal.

Insurance companies do not observe individual risk levels (accident probabilities), but know that there are these two types of individuals, what the two possible risk levels are, and what the probability is of any given individual being high risk or low risk. Consider again the insurance problem for a consumer

$$\text{Max} \quad \pi^1 U(\omega^1 - PI) + (1-\pi^1)U(\omega^1 - d - PI + I) \tag{36}$$

The insurance company's expected profit is

$$\pi^1 PI + (1-\pi^1)(PI - I) = PI - (1-\pi^1)I \tag{37}$$

This is zero (when there is competition) if $P = (1-\pi^1)$. In this case, $1-P = \pi^1$, and the price ratio faced by the consumer is equal to the odds, and she purchases full insurance, so that $x_i^1 = x_i^2$ or $I_i^* = d$. She faces no risk with the insurance, which covers the entire damage or loss.

If there are consumers with different levels of risk, measured by different levels of π^1, and the insurance company knows what the probability of not having an accident is for each consumer, the price can be adjusted accordingly. Thus, the premium per unit of insurance can be set as $P_i = (1-\pi_i^1)$ for each *I*, i.e., the premium price is the accident probability. On the other hand, if the insurance company does not have this information, it cannot tailor contracts in this precise way. Note that the implicit assumption is that identities can be used to determine the prices paid by individuals, so there is some institutional mechanism for enforcing this.

Instead, suppose that the company knows the proportion or probability of each type of risk level. Let the proportion of low-risk individuals be μ. The ideas can be illustrated with just two levels, π_H^1 and π_L^1, where *H* stands for high risk and *L* for low risk, so that $\pi_H^1 < \pi_L^1$. Thus *H*-types have a higher accident probability. The first order condition (21) defines the demand function for each risk-type, denoted

$$I_i^*(P, \pi_i^1), \quad i = H, L \tag{38}$$

If both types do purchase insurance, the insurance company's expected profit is therefore

$$\begin{aligned} &\mu[P-(1-\pi_L^1)]I_L^* + (1-\mu)[P-(1-\pi_H^1)]I_H^* \\ &\text{or} \\ &P[\mu I_L^* + (1-\mu)I_H^*] - [\mu(1-\pi_L^1)I_L^* + (1-\mu)(1-\pi_H^1)I_H^*] \end{aligned} \tag{39}$$

With competition, this expression is zero, and can be solved for the equilibrium price of insurance. In this case, it is easy to see that $(1-\pi_H^1) > P^* > (1-\pi_L^1)$. Hence, the high risk type gets better than fair odds, while the low risk type gets worse. With fair insurance, we would have $I_i^* = d$. Now the low-risk type buys less than full insurance.

It is possible, however, that the higher price of insurance deters low-risk types from purchasing insurance at all. As μ, the proportion of low-risk individuals, goes down, P^* increases, which reduces I_L^*. It is possible that there is a corner solution for low-risk individuals, in which case they drop out of the market (their proportion in the pool of customers goes to zero), and $P^* = (1 - \pi_H^1)$.

This phenomenon of low-risk types dropping out is called *adverse selection*. The term was originally applied to the insurance market, but can also be used in labor markets (where abilities differ), product markets (qualities differ), and so on. In the illustrative case of automobiles, this phenomenon has been termed the *lemons problem*, where a lemon is a low quality used car. In the original formulation of the lemons problem, the distribution of qualities was a continuous variable, and the argument we have sketched here led to every type dropping out of the market except the lowest qualities (which would have probability mass zero in the case of a continuous density function). In other words, the market basically would cease to function in that extreme case.

Two kinds of actions by insurance companies can limit the adverse selection problem. First, companies can gather information on an individual's risk type. Age, gender, location and employment can all be used, as can the history of previous accidents or claims. In the case of health insurance, a medical examination can be required by the insurer. All these information gathering methods are signals that improve the insurance company's knowledge of the true π_i^1 for an individual *i*. Second, insurance companies can offer contracts, which are not simply agreements to supply any quantity of insurance at a market determined price. Instead, an insurance contract can place limits on coverage, including deductibles, copayments and maximum payouts. A menu of contracts can be designed and offered that leads individuals with different risk-levels to *self-select* different contracts. For example, low-risk individuals may be willing to accept higher deductibles or lower payout limits, since their odds are better of avoiding damage or loss. These contracting approaches are considered in Singh (2010).

References

Kreps, David, 1990, *A Course in Microeconomic Theory*, Princeton, NJ: Princeton University.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, 1995, *Microeconomic Theory,* New York: Oxford University Press.

Singh, Nirvikar, 2010, “Market Institutions: An Expository Essay,” University of California, Santa Cruz, Department of Economics.

Varian, Hal R., 1992, *Microeconomic Analysis*, 3rd ed.**,** New York: Norton.