Chapter 7: Risky Choices and Market Outcomes

I. Deterministic and Stochastic Choice Settings

In the choice settings modeled in the first 6 chapters of the book, all the relevant decision makers have been assumed to have clear overarching goals (utility, net benefits, or profit) and complete knowledge about their possibilities (opportunity set, production functions, prices) and the consequences of their choices. In such settings, no mistakes are ever made by rational decisionmakers. Thus, consumer choices always truly maximize their utility, and firms always exactly maximize their profits.

Differences among markets and individuals may exist, but they are not due to differences in knowledge or risk aversion. There are no risks, because there is nothing stochastic in the choice settings modelled. Utility is maximized or it isn't. Costs are minimized or they are not. Both firms and consumers fully understand all the consequences of their decisions and actions. There is no such thing as "bad" or "good" luck in a deterministic world.

In choice settings that include stochastic elements, this is no longer true. There are often obviously random aspects of some choice settings, as when gamblers place bets on the roll of a pair of dice or purchase lottery tickets. There are others in which random elements are less obvious, as with the weather or day-to-day stock market price changes. There are still others in which random elements may or may not be present, but probabilistic ideas are used to understand unknown factors or complex systems. Everything in a system may actually be mechanically caused, as postulated by some physicists, but may appear to be random to persons with less than perfect knowledge. Computer generated random numbers are one of many instances of such phenomena. In all three cases, decision makers have to make choices without knowing the precise consequences of their actions.

In the first and second cases, one may truly understand the probability functions that lie behind a random event. In the third, probabilistic ideas provide useful explanations that reflect the absence of sufficiently complete knowledge to entirely understand the phenomenon of interest.

Either perspective implies that our previous models cannot be used to analyze all choice settings, because there are decisions for which the outcome or consequence is not known with certainty.

As with intertemporal choices, there are reasonable methods for making sensible choices in settings that include random aspects, although it is unlikely that all individuals use such methods. Nonetheless, a sufficient number may use them that simply analyzing how such individuals make choices sheds light on how markets operate in settings where time or randomness are important considerations. (The persons not making use of reasonable methods for taking account of time and randomness may make decisions that are randomly distributed about the choices of the more careful and informed decisionmakers.)

And as with intertemporal choice settings, the existence of probabilistic phenomena provides opportunities for mutual gains from trade that are neglected by models that ignore randomness. Whole industries have emerged in response to randomness. For example, insurance markets would be unlikely to exist were it not for the importance of random phenomena in day-to-day decision making and economics.

Two Different Types of Incomplete Knowledge: Risk and Knightian Uncertainty

This chapter focuses on settings in which individuals can (and do) know the probability distributions that lie behind economically relevant random phenomena. In other circumstances, they may not have such knowledge. Frank Knight (1921) refers to the first types of settings as "risk" and the second as "uncertainty." One can make reasonable choices in the first case—one can, for example, maximize "expected" value adjusted for risk—but cannot undertake such calculations in the latter case, because there is just too little information to do so.

Knight argued that risky choices can be completely modelled and have clear implications for market outcomes—as clear as models rooted in deterministic settings. For example, markets arise for insurance and other services that redistribute risks among persons and organizations. Many of those markets are competitive, while others are less so and thus the full range of price theory can be applied to them. It is these circumstances that are the main focus of this chapter.

Knight also argued that this was less true of settings of uncertainty. He argued that uncertainty is the best explanation for both higher-than-normal and lower-than-normal profits in competitive markets. Without truly knowing the odds, some "uncertainty takers" win one of life's lotteries and realize enormous profits, while others (the majority) earn lower rates of return on their investments than is typical in more familiar markets where phenomena are either mechanically caused or generated by well-understood probability functions. That line of reasoning and Knightian uncertainty itself are given more attention in the next chapter where entrepreneurship and innovation are analyzed.

The textbook is somewhat less careful in its vocabulary in this chapter than Knight was. It uses the terms risk and uncertainty more or less interchangeably, with minor exceptions, and uses the term Knightian uncertainty to describe choice settings in which the probability and density functions are less than fully known or unknown.

II. Probability Functions, Expected Values, and Expected Utility

A modest extension of the rational choice model can be used to characterize decisions in choice settings in which a well-understood probability function either generates the outcomes that follow from a decision or generates factors that determine the conditions in which choices are made. Rather than assume that individuals maximize utility, it is assumed that consumers maximize "expected" utility. Similarly, rather than assuming that firm owners maximize profits, it is assumed that firm owners maximize "expected" profits. In random conditions, truly maximizing utility or profits is beyond the control of individuals—although they can maximize average results if the relevant probability functions are understood.[1]

---

[1] Here, it should be acknowledged that the assumption that individuals maximize expected utility is less logically and empirically supported than the assumption that individuals maximize utility. If individuals have an over-arching goal, they will necessarily try to advance that goal with all of their "rational" decisions. Or, if

The notion of "**expected value**," itself, is an idea taken from statistics and means the average result associated with a large series of "draws" from a stable random process of some kind.

> DEF: Every **probability function** assigns probabilities to discrete events (here events 1, 2, … N) such that the sum of the probabilities is 1.0 and the numbers assigned to particular outcomes characterize the relative frequency or likelihood that that possibility occurs. (The probability that something will actually happen is 1, is completely certain, thus one of the possibilities will always occur.)

$$\sum_{i=1}^{N} P_i = 1 \qquad with\ P_i > 0 \tag{1}$$

All possibilities, i, have positive probabilities of occurrence $1 \geq Pi > 0$. All impossibilities, j, have a zero probability of occurring and are not considered parts of a probability function.

The mathematical expected value is the sum of the values of those possibilities (here V_1, V_2 … V_N) times their particular probabilities of occurrence (here P_1, P_2, … P_N). It characterizes the large-sample *average* value of the distribution of the possible values in such samples.

> DEF: The **mathematical expected value** of a set of possible outcomes, 1, 2, ... N with values V_1, V_2, ... V_N and probabilities of occurrence P_1, P_2 , ... P_N is:

$$E(V) = \sum_{i=1}^{N} P_i V_i \tag{2}$$

Expected utility is a special case of expected values, namely it characterizes the average utility realized when "value" is measured in terms of utility (as utils). The *expected* utility associated with a probabilistic setting is thus calculated in a similar manner:

---

they have an internally consistent preference ordering that ranks all relevant possibilities, they will behave in a manner consistent with models that assume that utility is maximized.

In contrast, maximizing expected utility is only one of many ways to deal with risky situations. Instead, one might, for example, attempt to minimize risk (e.g. attempt to maximize the worst possible outcome [maxi-min] rather than maximize the average outcome). They might focus on the most likely outcome rather than the average outcome, as for example when crossing a street, where catastrophic possibilities are often ignored. How one deals with risky settings is itself a choice. The use of average (expected) utility as metric for assessing the relative merits of risky choices is simply one of many reasonable strategies that individuals may employ.

$$E(U) = \sum_{i=1}^{N} P_i U(v_i) \qquad (3)$$

where the *N* possible outcomes ($v_1, v_2, \ldots v_N$) are associated with utility levels through an individual's utility function. To use this formula for expected utility calculations, one has to assume that the number of outcomes is finite and countable, that the values are finite, and that each outcome has a positive probability associated with it.

Expected values can also be calculated for random phenomena with a continuous domain. In those cases, a probability density function such as $f(x)$ is used for the calculations, rather than a probability function. A probability density function is constructed so that the area under that function equals 1 and the probability that x takes a value between x' and x'' is the area under that function between x' and x''. Expected utility in such cases is determined using an integral, rather than a summation, as with

$$E(U) = \int_{-\infty}^{\infty} \mathrm{f(x)U(x)dx} \qquad (4)$$

That a probability function or probability density function is known is not an unreasonable assumption in many circumstances, and it is a reasonable first approximation of many others. The probabilities assigned may be the result of careful empirical work (frequentist probabilities) or based on theoretical reasoning (many natural phenomena are normally distributed, so "this one" is likely to be as well). Or, it may reflect cumulative learning about the likelihood of particular events that are continually updated as more evidence is gathered (Bayesian updating).

Such probabilities are educated guesses rather than necessarily accurate. This last case is one way to use probabilistic choice models to think about choice settings that Knight would regard to be uncertain. Such models, for example, are used in chapter 8 to characterize a subset of entrepreneurial choices.

Most economists focus on circumstances in which all the possible outcomes are known, and reliable probabilities can be assigned to them. The assumption that probability or probability density functions are either known or accurately estimated allows models to be constructed that provide useful insights into how probabilities affect the choices made—conclusions that can be tested (at last to some degree). The latter is less true of models where probabilities are entirely subjective and not necessarily accurate or complete.

Illustration of the Difference Between Expected Values and Expected Utility

To illustrate the difference between expected values and expected utility consider the expected roll of a die. Suppose that a single die is to be rolled. The face that turns up on top is a random event. Suppose that you will be paid a dollar amount equal to the number on the face that winds up on top. Since the probability of a particular face winding up on top is 1/6 and the value of the outcomes is 1, 2, 3, 4, 5, 6, arithmetic implies that the expected value of this game in money terms is \$3.50 = (1)(1/6) + (2)(1/6) + (3)(1/6) + ...(6)(1/6). If you played the game dozens of times, your average payoff per roll would be approximately\$3.50.

Note that the expected value of a single roll of a die is 3.5, a number that actually is impossible, rather than "expected" in the usual sense in ordinary English. The values are all whole numbers. This is not always the case, but this example illustrates that the meaning of "expected value" is a technical one: namely the large sample average result, rather than the result that you would most commonly observe. There are many probability distributions in which the average value is also the mode, as with the normal distribution, in which case the expected value is both the average result and also the most likely value to be observed,

Next, we'll repeat the exercise for a concrete utility function, namely $u(V) \ = V^{0.5}$, where V is the value (and winnings) from a particular roll of the die. In this case

$$E(U(V)) = \sum_{i=1}^{N} (1/6)\ (V_i)^{.5} \tag{5}$$

Or,

$$U^e = (0.1667) + \ (0.2357) + (0.2887) + (0.3333) + (0.3726) + (0.4082) = 1.8053$$

If the intermediate cases (2, 3, 4, and 5) were for some reason impossible—or "rounded off" as far as prizes are concerned, there would be just two possibilities, each with a probability of .5, with a prize of V=1 for values three and lower and a prize of V=6 for values greater than three. In that case the expected utility is:

$$U^e = .5\,(1)^{.5} + .5(6)^{.5} = (.5) + (1.2247) = 1.7247 < 1.8053 \qquad (6)$$

Expected utility falls because the stochastic event becomes "riskier" when only the extreme outcomes are possible.

Notice that expected utility is not the same as the utility of the expected value, e.g. $((3.5)^{.5} = 1.8708)$. The expected value of a single roll of a die is 3.5 in each case. Al, as we shall see later in the chapter, is quite risk averse, because $u(V^e) > u^e(V)$.

Utility functions that can be used to calculate expected utility values that consistently rank alternative outcomes (according to expected utility) are called **Von Neumann–Morgenstern utility functions**. Von Neumann–Morgenstern utility functions are all complete, transitive, continuous, and exhibit monotonicity. In addition, they have the property of what is sometimes called substitutability which is a form of internal consistency with respect to the calculation of expected values. If one is indifferent between outcomes x and y, then one is also indifferent between px and $(1-p)z$ and py and $(1-p)z$, where p is the probability of event x. And, if z is regarded to be better than x, than $pu(x) + (1-p)u(y) < pu(x) + (1-p)u(z)$.

Experiments have been undertaken to use various gambles to create Von Neumann–Morgenstern Utility functions—which, as it turns out, do not perfectly explain individual behavior under uncertainty in laboratories, but do so reasonably well. Von Neumann–Morgenstern utility functions for particular individuals are also "unique" up to a linear transformation (and considered by some to be a form of cardinal utility), because one can do arithmetic with them.

Expected Utility with Continuous Probability Functions

Many economic choice settings concern variables that exist in a continuum, rather than being discrete. It is such choices that lend themselves to analysis using calculus-based models. Similar models can be developed to characterize choices where the outcomes are at least a bit uncertain, as the quality of an individual piece of fruit, bottle of wine, or automobile may not be known beforehand, because quality is itself a random variable. To see how uncertainty about quality affects consumer choices, consider the following choice setting.

Suppose that Al has a two-good strictly concave utility function, $U = u(A, B)$ where the prices of goods A and B are P_A and P_B respectively. Al has W dollars to spend in the period of interest. The quality of good A is not known at the moment of purchase, whereas that of good B is known with certainty. Given $f(q)$, the density distribution of the quality of some good that an individual may purchase, the expected utility for Q units of good A can be written as an integral of the following sort. The density function is distributed between the lowest quality, L, and the highest quality, H, possible for the good of interest. (Note that the substitution method has been used to characterize good B as a function of purchases of good A.)

$$\mathrm{U}^e = \int_L^H \mathrm{U}(\mathrm{Q}(\mathrm{q}), (W - P_A Q)/P_B)\mathrm{f}(\mathrm{q})\ d\mathrm{q} \qquad (7)$$

The integral written above is the expected (or average) utility associated with purchase of Q units of the good with stochastic quality distributed according to density function $f(q)$. Note that the quality is not entirely unknown, because it always lies between L and H. How it is distributed is also known. Its distribution can be characterized with a well-understood density function, $f(q)$, that can be integrated.

The quantity of this product that maximizes Al's expected utility can be found by differentiating expected utility with respect to Q and setting the result equal to zero. The first order condition in this case takes the form:

$$\mathrm{U}_Q^e = \int_L^H f(q)[U_A - U_B(\frac{P_A}{P_B})]d\mathrm{q} = 0 \text{ at Q*} \qquad (8)$$

where A and B subscripts indicate derivatives with respect to the variables subscripted. Notice that the partial derivatives are results obtained by differentiating the integrand. The integral domains are carried forward and the function being integrated (the integrand) is replaced with its relevant first derivatives. The first-order condition includes terms for the expected marginal benefit (the integral of the first term, in terms of utils) and for the expected marginal cost of units of A (the integral of the second term, again in terms of utils lost from the implied reduction in purchases of B).[2]

The implicit function theorem implies that Al's demand for good A, here A^*, can be written as $A^* = g(P_A, P_B, W, H, L)$. The density function of quality uncertainty affects the shape of this function. However, as written, the demand function does not include a variable that characterizes that effect, but does include the end points of the domain (L and H) of A's quality. If the utility function or the probability density function had included a "conditioning" variable, such as weather, that variable would have been included in the demand function.

The domain of the integral is determined by the probability density function. In the case used above, there is presumed to be a lowest (L) and a highest (H) quality. In other cases, such as the normal distribution, the limits would be minus infinity and plus infinity. Some density functions are "full domain." Anything may happen, but some events are more likely than others and some are extremely unlikely.

One uses the term "probability density function" (pdf) rather than "probability function" here, because probabilities are associated with integrals of (areas under) the density function, rather than by the function itself. Thus, the total area under both a conditional and unconditional probability function is 1 (by definition, the sum of all possible probabilities is still 1).

III. Risk Aversion and the Demand for Insurance

> **DEF**: An individual is said to be *risk averse* if the expected utility of some gamble or risk is less than the utility that would be generated at the expected value (mean) of the variable that determines utility.

---

[2] If one characterizes a utility function with a concrete function, then these integrals can often be evaluated. This is also possible for single-variable abstract functions—integrals of which simply return the initial integrand, which are then evaluated at the high end of the range of possibilities (here H) and then the low end of the possibilities (here L), which is subtracted from the high end value, which in turn gives one the (net) marginal utility of a change in Q.

A **risk-averse** person is one for whom the expected utility of a gamble (risky situation) is less than the utility of the expected (mean) outcome when obtained with certainty. In mathematical terms, a person is risk averse if and only if $U^e(x) < U(x^e)$ where X is a binary random event, with one possibility, x', occurring with probability P and the other occurring with probability (1-P). $x^e = Px' + (1-P)x''$. This property is true of every possible pair of possible outcomes for a risk averse person.

This property implies that any net benefit or utility function that is strictly concave with respect to income, exhibits risk aversion with respect to income or wealth types of variables. Why? Because expected utilities are convex combinations of utilities. Recall that a function is strictly concave only if af(x') + (1-a)f(x") < f(ax' + (1-a)x") for any x' and x" and any value of a with 0<a<1. If one substitutes a probability for the term a, you can see that the two definitions are essentially identical.

A ***risk neutral*** individual is one for whom the expected utility of a gamble (risky situation) and utility of the expected (mean) outcome are the same. $U(x)^e = U(x^e).$ A ***risk preferring*** individual is one for whom the expected utility of a gamble is greater than the utility of the expected (mean) outcome. $U(x)^e > U(x^e).$

The degree of risk aversion is often measured using the *Arrow-Pratt* measure of (absolute) risk aversion: $r(Y) = -(d^2U/dY^2)/(\frac{dU}{dY})$ which is a measure of how steeply downward sloping the marginal utility of income is at a particular point. In general, this implies that the more steeply downward sloping the marginal utility of income curve is, the more risk averse an individual is.

In the illustrating example above where $U = V^{.5}$, the marginal utility function is quite steeply downward sloping $\frac{dU}{dY} = .5V^{-0.5}$ and $\frac{d^2U}{dY^2} = -.25V^{-1.5}$, so

$$-\frac{[-.25V^{-1.5}]}{.5V^{-0.5}} = (.25V^{-1.5})(.5V^{0.5}) = .125/V \qquad (9)$$

Note that the degree of risk aversion for this function varies with V, decreasing as V increases.

The utility Functions that imply risk-averse behavior are all strictly concave, as illustrated below. If the above individual, Al had been risk neutral, the Arrow-Pratt measure would have been zero. If $U = V$, then $\frac{dU}{dY} = 1$ and $\frac{d^2U}{dY^2} = 0$, which implies that $r(V) = \frac{0}{1} = 0.$

The Geometry of Risk Aversion and Risk Premia

The figure below illustrates a choice setting in which an individual is risk averse and facing a risky environment in which either an outcome with the value V_1 or another outcome with the value V_2 will occur. The individual cannot influence which outcome will occur, but knows that the probability of V_1 is P, which implies that the probability that V_2 occurs is (1-P) (Recall that the probabilities for the only two possible events have to add up to one.)

Let's refer to the individual as Al. Al's utility function is strictly concave, which means that a chord connecting any two points on it lies below the utility function (except for the two points used as end points—which, by definition, are not part of the chord). Assume that Al confronts an uncertain environment in which V_1 occurs with probability P and V_2 occurs with probability (1-P) Al's expected utility in that case is:

$$U^e = PU(V_1) + (1-P)U(V_2) \qquad (10)$$

As P increases from 0 to 1, the expected utilities trace out the chord between $U(V_1)$ and $U(V_2)$ and so will be below the utility function if it is strictly concave.[3]

Expected Utility, Risk Aversion, and Risk Premiums

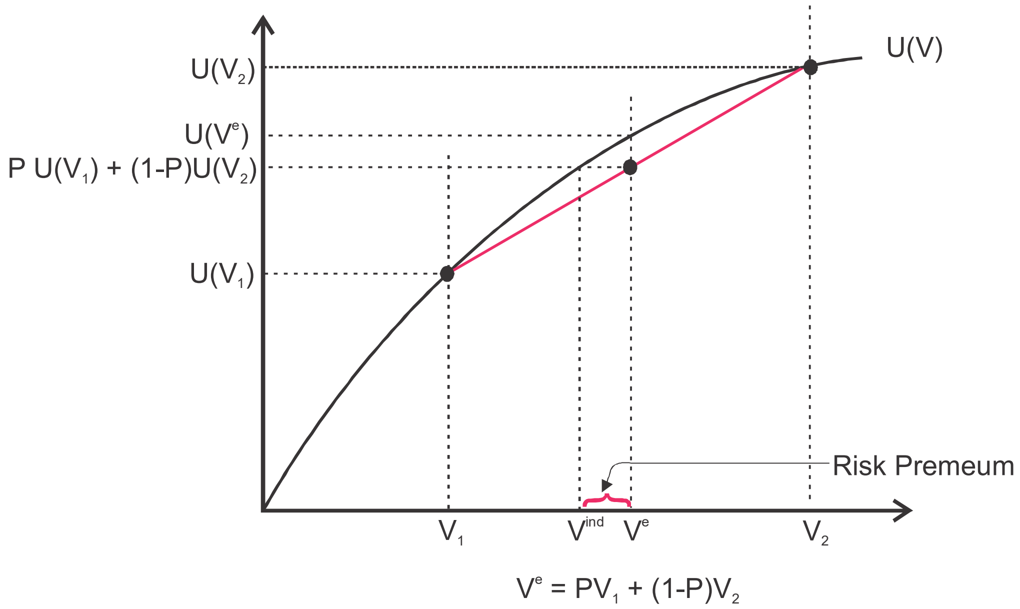

This geometry is illustrated in the diagram above for a probability, P, that is approximately equal to 0.5, but it would be true for all probabilities $0< P < 1$ and all strictly concave utility functions.

This diagram can also be used to determine how much an individual would be willing to pay to have a certain payoff rather than face a risky or uncertain future. This is done by looking at the certain outcome that a person would be equivalent in their mind to the risky event. If we go to the left from the expected utility associated with the two probabilistic outcomes

---

[3] We have until this point used a "sufficient condition" for strict concavity, namely that a utility function is strictly concave if it has a positive first derivative for V and a negative second derivative for V. In other words, Al's utility function is concave if it exhibits diminishing marginal returns from V. However, at this point the formal definition becomes a nice bridge between risk aversion and strict concavity. The expected value of V is $V^e = PV_1 + (1-P)V_2$. Note that if U is strictly concave then $U(V^e) > PU(V_1) + (1-P)U(V_2)$. As mentioned above, this satisfies the formal definition of concavity. We've simply substituted "P" for "α".

over to the utility function and then down to the horizontal axis, we find the value (labeled V^{ind}) that Al would find equivalent to the risky one faced. (V^{ind} is the certain outcome that generates the same expected utility as the risky one faced. "ind", stands for indifferent.) The difference in values, $V^e - V^{ind}$, (assuming that the values along the horizontal axis are in money terms) is the highest price that Al would pay to avoid the risk.

It also characterizes the lowest expected value that Al would accept to bear the risky environment shown rather than have outcome V^{ind} with certainty. That difference is called Al's **risk premium** for this choice setting or "gamble." Note that Al would accept the gamble (risky environment) rather than V^{ind} only if the expected value of the risky payoff is greater than V^{ind}. How much greater would vary with Al's degree of risk aversion. The more risk averse Al is, the greater the risk premium would have to be.

The latter has implications for businesses in risky circumstances. Risk averse firm owners will demand a risk premium to bear the risks associated with their businesses. In such cases, Marshallian competition would generate different equilibrium profit-rates (returns) in different industries, ones that vary according to the riskiness of the business environment and the risk aversion of firm owners. The Marshallian equal returns idea in that setting is revised to equal "risk adjusted" rates of return in a setting where risks vary among industries or firms.

Risk premia also have implications for an individual's demand for insurance. An individual's risk premium characterizes the highest amount that Al is willing to pay for insurance that eliminates the risk confronted. Note that the expected loss can be represented as $P(V_2\text{-}V_1)$ which is the distance from V_2 to $V^e=PV_1+(1\text{-}P)V_2 = P(V_1 - V_2) + V_2 = V_2 - P(V_2\text{-}V_1)$. This last expression characterizes the expected value of the risky setting in terms of the expected loss that occurs when the unfortunate event occurs—possibly a fire, accident, or a disease. The risk premium is the amount above the objective expected value of the risk that Al is willing to pay to avoid the risk.

Note that the demand for insurance is another instance of "income smoothing"—a method of shifting additional income to periods after a loss, where the marginal utility of income is higher than it is when no accident occurs. In this case, it is a stochastic event that induces the lower utility rather than some predictable effect of time such as season or differences in income associated with age.

IV. A Few Applications

Application (1): Selling Fire Insurance

The existence of risk premiums plus the effects of sample size on sample means implies that selling insurance can be profitable. In our example, fire insurance transfers risk from homeowners to insurance companies. However, if the probability function is well known and the insurance company has many customers, the insurance company has only a very small risk. The average payout from selling insurance converges to $P(V_2-V_1)$ per customer, per year, as the number of subscribers becomes large, because the payout in this case resembles a large sample average.[4]

The price for the insurance can be **up to $P(V_2-V_1)$ + the risk premium** from the above figure. This implies that selling insurance can be profitable as long as consumer risk premia are sufficiently high.

If a large number of purchasers for an insurance product exists, firms will have a quite predictable flow of expenses that are approximately equal to the expected value of the average loss, while customers are willing to pay more than that to avoid the risk of such losses. If the risk premium customers are willing to pay is more than sufficient to cover the cost of sales and administration of the insurance products sold, then insurance companies may be profitable investments.

However, competition among insurance providers, in turn, tends to bring profit rate down to the "risk-adjusted ordinary rate of return" that firm owners make from their other

---

[4] The simple formula for the variance of a sample average computed from a sample of size N is $var(X^{mean}) = \frac{\sigma^2}{N}$, where $\sigma^2$is the variance of random variable X, and N is the sample size used to compute the sample average X^{mean}. For insurance, variable X is the payout generated by the random variable "house fire." Note that sample variance converges to zero—e.g. complete certainty—as N becomes large.

investments. But the prices for insurance in equilibrium will be sufficient to keep them in business.

Note that both the buyer choices and the insurance provider choices involve the mathematics of expected values rather than utility or profit maximization under certainty. Insurance is a market that would not exist without measurable risks and the ability to shift risk from one group to another.

The ability to moderate risks through risk pooling—which is a property of the variance of sample averages, which tend to fall as sample size increases—is another key feature of the risks for which insurance is possible. (Recall that the variance of a sample mean is var(sample mean) = σ^2/n, where σ^2 is the variance of the variable being sampled and n is the sample size. The variance of a sample mean converges to zero as n becomes large.)

Application (2): Expected Benefits Maximization with Uncertain Product Quality and Money-Back Guarantees

Another case in which probabilistic thinking is useful is regarding products of uncertain quality at the point of sale. For example, suppose that Al is considering purchasing some produce from a farm and knows that some of the produce will be of high quality (H) and some will be of low quality (L) but simply can't tell the difference between the two types of produce at the time of purchase, as is true of many types of produce (corn, potatoes, tomatoes, squash, etc). However, suppose that Al has sufficient experience with the farm or produce shop to know what the probability of a defective product is.

Suppose that there are just two levels of quality that tend to turn up, high quality and low quality. Suppose also that the probability of high quality is F and that price per unit is simply P. Suppose that the benefits of high-quality units is B(Q, H) and the benefits from low-quality units is B(Q, L) where B(Q,H) > B(Q,L) for every Q.

How many units will Al purchase? Al's expected net benefit from purchasing produce is expected benefits less expected costs:

$$N^e = FB(Q,H) + (1-F)B(Q,L) - PQ \qquad (11)$$

To find Q*, differentiate N^e with respect to Q and set the result equal to zero.

$$F\left(\frac{dB^H}{dQ}\right) + (1-F)\left(\frac{dB^L}{dQ}\right) - P = 0 \qquad (12)$$

The first two terms are the expected marginal benefit of the produce and the last is its marginal cost.

To find a specific value we would need to use concrete functional forms for the two benefit functions, as with $B^H = HQ^{.5}$. and $B^L = LQ^{.5}$, with H>L, in which case our first order condition would be:

$$.5FH/Q^{.5} + .5(1-F)L/Q^{.5} = P$$

Multiplying both sides by $2Q^{.5}$ yields $HF + L(1-F) = 2PQ^{.5}$, which implies that

$$Q^* = [HF + L(1-F)]^2/4P^2 \qquad (13)$$

Equation 13 implies that the quantity Al purchases increases with F (the probability of the high-quality type) and with the benefit of the high-quality product, H (an indication of the quality of the high-quality type) and falls as the quality of the low-quality version of the produce decreases or the price of the good with uncertain quality increases. Demand curves again slope downward, but new “shift” terms are added to those developed in Chapter 2.

This quality risk is potentially insurable, but it may be too difficult (costly) for an independent insurance company to provide such an insurance policy, because it would require monitoring every transaction for the good whose quality is to be insured. Thus, such consumer risks, in practice, are not often insured by insurance companies.

An alternative way to reduce the risk that consumers face for products with probabilistic quality is a money-back guarantee. Such warranties shift the quality risk to the seller—who may then increase his or her price by the expected loss (average refund on low quality units) per customer. In effect, such warranties make the seller an insurance company. Sellers may charge their risk-averse customers for this service by raising prices a bit to take account of the average refund paid out on units of poor quality. Their risk-averse customers are willing to pay that premium as long as the surcharge is less than their risk premium.

The merchant's costs tend to be lower than that of an independent insurance company, because they are already involved in every sale of the good of interest.

Application (3): Quality Control—A Role for Management and Monitoring

Of course, quality variation is not only associated with agricultural products. All goods and services have some variation in quality. Within mechanized construction processes, wear and tear on both labor and capital can induce product variations that affect the usability of a firm's products. Such variation may affect the durability of the product as well as the immediate benefits that it provides to buyers in which case intertemporal aspects of quality control may also increase a firm's sales and profits.

Insofar as quality variability can be estimated by consumers, and purchase decisions are essentially independent of one another, we can use the net benefit maximizing model to characterize the demand for such products. $N^e = B^e(Q) - C(Q)$ which can be represented as $N^e = Fb^L(Q) + (1-F)b^H(Q) - PQ$ for the two-quality case, where F is the relative frequency of low-quality units, $(1-F)$ is the relative frequency of high value units, and P is the price of the units purchased. The benefit functions are both assumed to be strictly concave.

The quantity that a purchaser would acquire satisfies the first order condition:

$$dN^e/dQ = Fdb^L/dQ + (1-F)db^H/dQ - P = 0 \equiv H.$$

The implicit function theorem implies that a typical individual's demand for this product can be characterized as

$$Q^* = f(P,F).$$

The implicit function differentiation rule implies that:

$$dQ^*/dP = \frac{dH/dP}{-dH/dQ} = \frac{(-1)}{-[\frac{F(d^2b^L}{dQ^2}) + (1-F)(\frac{d^2b^H}{dQ^2})]} < 0$$

given the strict concavity of the benefit functions. The demand function is downward sloping in price. Similarly,

$$dQ^*/dF \;=\; \frac{dH/dF}{-dH/dQ} \;=\; \frac{\left(\frac{db^L}{dQ}\right) - (\frac{db^H}{dQ})}{-[\frac{F(d^2b^L}{dQ^2}) \;+\; (1-F)(\frac{d^2b^H}{dQ^2})]} \;<\; 0$$

As the probability of a defective unit increases, demand for the product of interest falls.

Notice that a firm selling this product and facing a downward sloping demand curve can influence the extent of demand through decisions that affect the frequency of low-quality units. For example, a firm's monitoring expenditures, M, may reduce F, with F = h(M). In effect firms have two controls in their efforts to maximize profits—monitoring (M) and output levels (Q), rather than simply one (quantity) as usually assumed. In this case, a firm's demand curve is partly determined by the firm's decision about quality control.

Let $P \;=\; g(Q,F) = g(Q,h(M))$ characterize its inverse demand function. The firm's profits in this case can be characterized as $\Pi \;=\; g(Q,h(M))Q \;-c(Q,M)$ (assuming that input prices are constant in the period of interest). There will be two first order conditions for its profit maximizing efforts:

$$(dg/dQ)(Q) \;+\; g \;\;-\; dc/dQ \;=\; 0 \tag{14}$$

$$(dg/dF)(dF/dM)Q \;-\; dc/dM \;=\; 0 \tag{15}$$

Both first-order conditions are simultaneously satisfied at the firms' profit-maximizing output.

This makes the firm's decision a bit more difficult to characterize in words. The monitoring level adopted affects the output decision through effects on the demand curve for the firm's product. And the output decision affects the optimal degree of monitoring by changing the extent to which prices respond to monitoring expenditures (dg/dF). In each case, the ideal levels occur where marginal revenue generated by changing Q or M equals the marginal cost of Q or M.

A market demand curve is not always beyond the influence of a price-making firm or industry. Perceived quality matters as does perceived average quality. Both are at least partly determined by a firm's monitoring efforts.

V. Some Additional Applications of expected utility and expected profit maximizing models to choice settings that indirectly affect market outcomes

The 1960s was a period in which rational choice models were applied to fields generally regarded by most economists to lie outside of economics. These new areas of research gradually gained sway inside economics and expanded the field to areas of the economics of regulation, socio-economics, law and economics, and political economy. Several of these areas made use of the expected utility maximizing model and expected net benefit maximizing model. Several Nobel prizes were awarded to the pioneers in these new areas of research, several of which are taken up in part III of this book.

Application (4): Expected Values and the Effects of Regulation

One can also use this type of model to model the effects of economic regulation. For example, in the area of environmental regulations, firms will take account of their overall net benefits from pollution including both cost savings and anticipated regulatory fines when choosing their production methods. In the absence of fines or fees for pollution and in the absence of enforcement of fines greater than 0, firms will choose their production methods to minimize their production costs—as in the models developed in the first part of the course (prior to the midterm).

(This does not necessarily mean that firms will pay no attention to air or water pollution, but they will do so only insofar as it affects the firm's expected profit through productivity and cost effects. Air or water quality that *affects the productivity of the firm's workforce* will be taken account of, but not spillovers on others outside the firm.)

In the real world, regulations are only imperfectly enforced, and firms know this. Consequently, it is not simply the magnitude of the fine or penalty schedule that affects a firm's decision to "pollute illegally or not," but also the probability that a person that violates the law will be caught, convicted and punished. Analyzing regulatory law and its enforcement on a firm's choice of production method and output level requires taking account of both the "expected cost" and "expected marginal cost" of any fines or penalties that might be associated with its production and output decisions.

(In addition, firms might face a loss of reputation and therefore reduced demand for their products if they are found guilty of violating regulatory law, but that effect will be ignored or assumed to be part of the fine.)

Consider a case in which production methods are fixed and output is regulated—which is the easiest case to model. In a regulatory environment with fines, a pragmatic firm's expected profits equal its total revenues less its production costs less its expected fines: $\Pi = R - C - F^e$ where $F^e = PF$. Suppose that Acme's output is sold in a competitive market, its cost function is $C = cQ^2wr$ and that its expected fine is the probability of being caught and convicted, which increases with output in excess of the regulatory limit, $p(Q-Q^R)$ and a fine schedule that increases with the extent of the violation $f(Q-Q^R)$ for $Q > Q^R$.

$$\Pi^e = PQ - cQ^2wr - p(Q-Q^R)f(Q-Q^R) \tag{16}$$

To make the functional form a bit more concrete, let us assume that $\mathrm{p(Q-QR)} = a(Q-Q^R)$ and $f(Q-Q^R) = b(Q-Q^R)$. In this case, Acme's expected profits are:

$$\Pi^e = PQ - cQ^2wr - a(Q-Q^R)\,b(Q-Q^R)$$

$$= PQ - cQ2wr - ab(Q-Q^R)^2 \tag{17}$$

Assume that the regulatory constraint is binding on Acme, and so it will take the expected fine schedule into account when making its output decision. Its expected profit maximizing output can be characterized by differentiating the above function with respect to Q, which is a bit more complex than usual because of the "$Q - Q^R$" terms.

$$\Pi^e_Q = P - 2cQwr - 2ab(Q-Q^R) = 0 \ at\ Q^* \tag{18}$$

This can be solved for Q^*. First, shift all the Q terms to the left side of the equal sign:

$$P = 2cQwr + 2ab(Q-Q^R) = Q(2cwr + 2ab) - 2abQR$$

Adding $2abQ^R$, dividing and reversing sides yields:

$$Q^* = (P + 2abQR)/(2cwr\ +\ 2ab) \qquad (19)$$

This is Acme’s supply function in the regulatory environment modeled.

Note that its output now varies with the regulatory standard (Q^R) its input costs (w and r) and parameters of the probability of being fined and fine schedules (a and b). Acme’s output declines as input prices and the expected fines increase (w, r, a, or b increase) and increases as the regulatory threshold (Q^R) increases.

Another possible output is simply Q^R, but this cannot be modeled with calculus because of a discontinuity in the expected cost function at that quantity. See below.

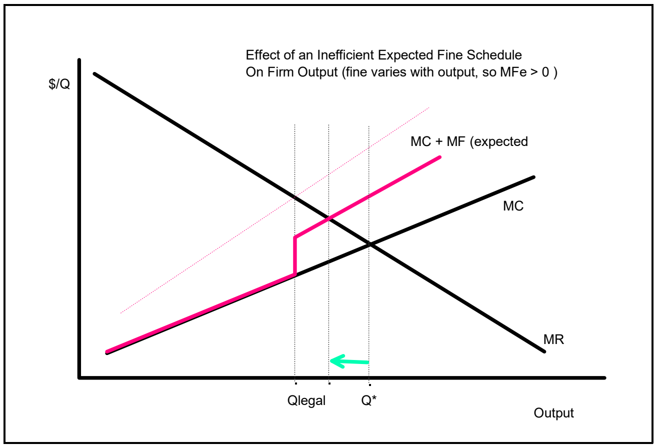

The diagram to the left illustrates Acme’s decision in this type of setting (with somewhat simpler probability and fine schedules). For students that have had public economics, note the similarities between Pigovian taxes and optimal enforcement with fines. If the regulation attempts to solve an externality problem and achieve Pareto efficiency, Q**, then the smallest fine sufficient to induce the target Q** has the **same expected value** as a Pigovian tax at Q** (with $Q^R \leq Q^{**}$). The expected fine should equal the expected marginal damages done by the Q**th unit of output.

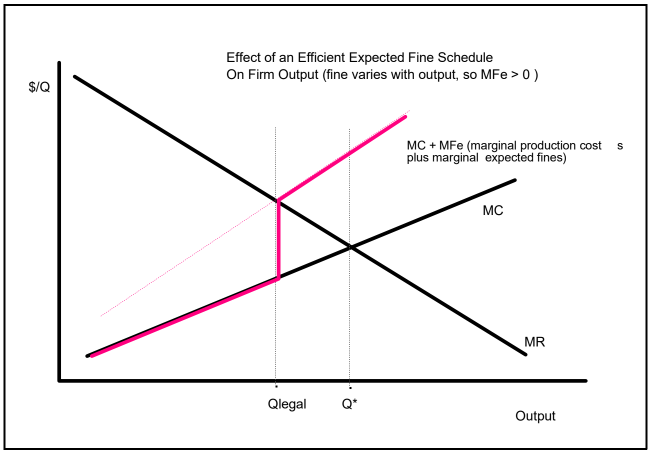

Note also that there is a policy-tradeoff between the probability of conviction and the level of punishment. In general, the larger the fine, the smaller the probability of capture can be to generate the same effect (expected marginal cost) on individuals. The larger is the probability of detecting a regulatory violation, the smaller the fine can be and still have the same effect.

The probability that an illegal activity is detected and punished varies with the resources used to enforce the law and the flagrancy of the violation. Thus, a more complete model would make the probability of detection an increasing function of law enforcement budgets in addition to the size of the violation.[5]

Application (5): Expected Values and the Logic of Crime and Punishment

The economic analysis of crime derives from a classic paper written by Gary Becker (1968), who subsequently won a Nobel prize in economics, partly for that contribution. In that paper, and in many others published since then, a criminal is modeled as a rational agent interested in maximizing his expected income or utility from criminal activities, given the probability of punishment and the punishment that he or she will be subject to if caught and convicted. Becker's model can be used to model theft and many other economically motivated violations of criminal and regulatory law.

In the real world, criminal laws are only imperfectly enforced, and both criminals and ordinary persons who occasionally think about violating a law or two know this. For example, a net income maximizing criminal would maximize an expected function like

$$\Pi^e = PQ - cQ^2 - p(Q)F \tag{20}$$

---

[5] A Few Law and Economic Puzzles for Students. Given this, how would you pick the appropriate punishment for speeding? for theft? For murder? etc.

How would the relative importance of the probability of detection and the expected fine be affected by the process of a jury trial and a long delay between being detected and being fined? (Some ideas for doing so are provided in the next chapter.)

Write down an expected profit function for a firm facing a fine schedule that is imperfectly enforced, but where the fine increases as Q exceeds Q legal. Find the first order conditions and compare them to the above diagrams.

Draw examples of a perfectly enforced and imperfectly enforced fine that does not vary with the degree of the offense. (Hint: such fines do **not** affect expected marginal costs.) Compare your graph with the mathematics of expected profit maximization in this case. Are such fines always irrelevant?

where Q is the number of crimes (thefts), price is the average price received by "fencing" the stolen goods, $p(Q)$ is a probability function describing the way that the probability of being caught and convicted varies with the number of crimes and F is the financial penalty assessed (or if jail time is spent, the opportunity cost of the time spent in jail and any subsequent losses in earnings).

The rational thief chooses Q^* such that $\Pi^e_Q = 0$, which in this case requires Q^* to satisfy $P - 2cQ - p_Q F = 0$ or $P = 2cQ + p_Q F$ (set the marginal revenue from theft equal to its **expected marginal cost**, which is not known with certainty).

In order to have a more concrete result, let's give the probability function a specific functional form as with: $p = aQ^2$ then $p_Q = 2aQ$. Given these additional assumptions, the above first order condition becomes $P - 2cQ - 2aQF = 0 \ or \ P = 2cQ + 2aQF$, which can be solved for Q.

$$P = Q(2c + 2aF) \rightarrow Q^* = P/(2c + 2aF) \tag{21}$$

Note that this implies that the rational criminal responds to incentives, his or her crime rate falls as the probability of being caught and convicted rises (e.g., with 2a), as the fine increases, and as the marginal cost of theft increases. Note also that there are tradeoffs between the size of the fine and the probability that a criminal is caught in terms of their overall effect on the criminal.

Many other examples from law and economics can also be similarly modelled. One does not have to be a more or less professional criminal for this logic to apply. One can think of choices to drive faster than the speed limit on a highway or to park without putting money in a parking meter, or to trespass on a neighbor's property, fail to report some income on one's taxes, and so on in much the same manner.

Chapter 14 analyzes how crime rates can affect the size and scope of markets.

VI. Some General Conclusions about Risky Choices

The main implication of this chapter is that neoclassical economics and its associated models can easily be extended to take account of risk—choice settings in which outcomes are uncertain, but the probability of various outcomes can be accurately estimated. In such cases, the logic of optimizing choice applies and the choices can be modelled using expected values rather than average values in the usual way. Both diagrams and calculus-based models can be developed to show how risky choices are made by reasonable persons.

However, taking account of risky choice settings alters many of the conclusions of the core model.

First, individuals with similar tastes may differ in their degree of risk aversion. Individuals may all prefer more income or wealth to less income or wealth, but their utility functions may differ in their curvature, in their degree of strict concavity. The Arrow-Pratt measure of risk aversion is one way to measure such differences. Some people are more risk averse than others, and so willing to pay more for insurance and other risk avoiding activities than others.

Second, risk aversion creates markets for insurance. If everyone were risk neutral, insurance-like products might still exist, but there would be no sales of such products, because consumers would not be willing to pay a premium for those products—and that premium is necessary to cover the cost of administering insurance products. Insurance is another instance of "consumption smoothing." When one is subject to a loss (here from a random event such as a hurricane), the insurance provides additional income to the person insured at a time when the marginal utility of income (or compensation) is relatively high—paying for that insurance (net costs greater than zero) occurs when the marginal utility of income (or compensation) is relatively low.

Third, the same logic implies that average returns on investment across markets tend to vary with risk, because investor-owners will demand a risk premium for investing in market activities that have higher risks. How much will vary with the risk aversion of such investors. Thus Marshall's equalization of profits across firms and industries now is modified by risk premia that were neglected in the risk-free environment of the core models.

Fourth, insurance products and similar products not always considered to be insurance imply that risks can be shifted from one individual to another and from one group to another. Moreover, risks can be pooled by insurance companies in a way that reduces overall risks, because of the statistical properties of sample means. As the sample size increases, the variance of the sample mean falls, and payouts tend to fall in a narrower range. Given their large "samples," insurance companies know almost precisely the amount that they will pay out in claims from year to year.

Competition among insurers—as noted by Frank Knight—tends to reduce insurance company profits (for honest companies) to "ordinary" rates of return—the so-called zero-profit equilibrium of Marshallian perfect competition. It may not do so entirely, because of differences in risk aversion, size, and organization. The relevance of Marshallian logic is limited to some degree by the fact that the variance of insurance payouts falls as the number of insurance subscribers increases. The latter, implies that economies of scale in risk pooling exist, and thus the number of insurance companies supported in the long run may be less than that required to assure Marshallian competition.

Fifth, there are many situations in which risks are not pooled and in which consumers, firms, or entrepreneurs have to make decisions in risky choice settings—shall I bring an umbrella when shopping today or not? Shall I save for a rainy day or not? In these cases, one can use expected utility or expected profit maximization to model the decisions reached and their comparative statics. However, the probabilities used for these calculations are often ungrounded in data analysis. Experimental evidence suggests that individuals do not always behave as predicted. Rational choice models of risky choices provide a good first approximation for the average behavior and for how changes in probabilities, benefits and cost affect behavior, but not a complete model of such choices.[6]

All this is not to say that the expected utility maximizing model of volition is the only model that makes sense; nor is it to say that observed behavior is entirely consistent with that model, but it is to say that this model sheds very useful light on choices made in circumstances where random phenomena are important.[7]

References

Arrow, K. J. (1958). Utilities, attitudes, choices: A review note. *Econometrica: Journal of the Econometric Society*, 1–23.

Arrow, K. J. (1984). *Individual choice under certainty and uncertainty* (Vol. 3). Harvard University Press.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, *76*(2), 169–217.

Camerer, C. F. (1995). Individual decision making. In J. H. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics* (pp. 587–703). Princeton University Press.

Edwards, K. D. (1996). Prospect theory: A literature review. *International Review of Financial Analysis*, *5*(1), 19–38.

Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of Political Economy*, *56*(4), 279–304.

Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, *92*(5), 1644–1655.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–292.

Knight, F. H. (1921). *Risk, uncertainty and profit*. Hart, Schaffner and Marx.

Shackle, G. L. S. (2010). *Uncertainty in economics and other reflections*. Cambridge University Press.

---

[6] For an overview of prospect theory, see Edwards (1996) or Tversky and Kahneman (1992)—who won a Nobel prize for their work in this area and others—research that shows some limits of the rational choice model of decision making under uncertainty (risk).

[7] There are many possible explanations for the departures from the prediction associated with maximizing expected utility in experimental settings. One is that there are many other plausible strategies that individuals might adopt for coping with risky choice settings as mentioned in footnote1. Another is that individuals are not very good at statistical theory, in which case, individual choices tend to be error prone.

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323.