

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

I. Information Problems and Rationality

With the exception of the Knightian analysis of entrepreneurship, all of the analysis and models to this point in the course have assumed that information is “perfect” in the sense that consumers knew enough to maximize their net benefits from market transactions and that firms knew enough to maximize their profits from market transactions. As a place to start the analysis of markets, this is a completely reasonable place to begin.

If trades are largely repeated, consumers will know or have accurate estimates of the average quality of the goods purchased. Firms will have accurate estimates of the costs required to produce their products and bring them to market. And, input providers will have accurate estimates of the prices that their goods and services will command in their various labor, raw, material, and intermediate goods markets.

In an evenly rotating economy, experience provides one with essentially all of the data that one needs to be “locally” informed—which is to say, to know or have good estimates of the qualities and prices of the goods and services on offer that one is most likely to purchase or produce. In such settings, market participants can maximize their net benefits from trade—even if they occasionally make mistakes.

In such cases, models that assume that consumers and firms have very good or perfect information are reasonable, and the conclusions drawn from them will be correct on average. Although those same models would

not be able to explain the existence of some types of markets or mistakes, the models would be useful, reliable, ways to think about most market transactions.

However, in settings where such information is not available, is too costly, or prices move more or less randomly because of variations in weather, input prices, innovation, or public policies, neither firms nor consumers will know exactly what prices or the quality of goods and services are, and mistakes—in a sense worked out in this chapter—will be more common. Such fluctuations need to be taken into account by both consumers and firms.

What this chapter does is to (1) analyze cases in which some aspects of a good’s quality, prices, or costs are randomly distributed, and so exact knowledge is not possible. The cases that are easiest to handle are ones where the fluctuations can be characterized with a probability distribution (as argued by Frank Knight). Economists generally assume that in such cases, consumers and firms maximize average or “expected” net benefits rather than net benefits per se. (2) It also discusses what some term “rational ignorance,” which differs from risk, but also may affect consumer choices, planning, and their demand for services.

Chapter 9 (the next chapter) takes up choice settings in which choices that take significant time to execute—which is to say, it examines choices among plans or long term commitments rather than immediate actions.

In the cases explored in most of this chapter, market participants do

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

not know exactly the quality of the goods on offer, nor the lowest price at which they can be purchased, but do know something about the probability function that generates market prices. So, some of their expectations about prices are mistaken, but if the probability function is known reasonably well, they will, on average, maximize their net benefits in such choice settings. This is really as well as they can do, if prices or other choice relevant variables are generated via a random process.

Randomness generates unavoidable risk and, in some cases, also uncertainty. The use of “expected values” for making decisions in those settings is the most systematic way to model rational decision making by consumers and firms. It is the effect of risky settings on decision makers and markets that is the focus of this chapter. The same method can be used to analyze about choices under Knightian uncertainty, although in such cases the probability function can only be guessed or intuited rather than actually known.

Economic and other social systems are complex, and always a bit chaotic at the margin, because a variety of random factors including weather, disease, innovation, and geopolitical risks constantly jostle market decisionmakers and thus induce equilibrium prices to move about.

Thus, the choice settings covered in this and the next chapter are more realistic in that they explicitly take account of risks (and to some extent uncertainties) that occur in the ordinary course of life in commercial societies. On the other hand, the assumption that individuals maximize expected net

benefits is a bit less general than simply assuming that individuals maximize net benefits as they themselves understand them. It takes a bit more sophistication to use expected values or median values than to maximize net benefits in familiar predictable circumstances.

Before moving on to these more realistic settings, it should be kept in mind that if people are on average correct in their expectations, then the models studied previously in the course also will be correct on average—although they cannot explain the existence of errors, risk, or uncertainty, nor the existence markets for goods and services that help to ameliorate risk and uncertainty or to increase knowledge. For example, insurance markets make sense only in risky environments. Libraries and universities exist only because people are less than perfectly informed. Self-help books make sense only if people are at least a little uncertain about their true preferences. Accountants are unnecessary every firm owner knew exactly what their marginal costs and marginal revenues were for every output and market condition. Planning advice is useful only if someone does not know all of the circumstances that will have to be addressed in the future.

In a contemporary market system, Schumpeterian innovation is one of the main causes of both risks and uncertainties. Innovations often affect many prices. They may reduce the value of one’s human capital (skills) or by change other relative prices in ways that were not anticipated (or anticipatable). Innovations, such as libraries, universities, and the internet, can reduce ignorance and makes one’s expectations more reliable predictors of

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

the future, but also reduce income from “expert” advisors who may help others make better decisions. Even the beneficial effects of innovations may undermine long-standing patterns of life. Thus, dealing with random events, risk, and uncertainty are all natural parts of life in a contemporary commercial society. They are part of the reason that so many people change jobs and careers in their adult lives.

Both risks and uncertainties tend to be greater when one plans over relatively long periods of time, because the future can never be estimated exactly. Chapter 9 explores some of the consequences of or rational decisions that take account of both risk and time when making such long-term commitments. Examples include efforts to obtain a college degree, building marriage, a new factory, and undertaking research to create new products such as new medicines, self-driving cars, AI, or new apps.

II. Statistical Methods for Calculating Averages: Expected Values and Expected Net Benefits

The choice settings of interest in this chapter are ones in which perfect predictions are impossible because of random factors that have to be taken into account or because the scope of one’s knowledge is insufficient to fully understand the processes that generate choice-relevant factors.

Some ideas from statistics help us (and individual decision makers) to think sensibly about possible outcomes and their average values when the

probability of particular outcomes are known or knowable. The four ideas from statistics that we’ll be using to model choices in “risky” settings are **probability distributions, expected (or average) values, variance, and sample averages.**¹

A **probability function** lists all the possible outcomes of a random process and assigns probabilities to those outcomes that characterize the frequency that the random process of interest generates those values. The sum of the probabilities always adds up to 1. This is partly by assumption (as a defining feature of a probability distribution) and partly as a matter of logic—the probability that something will happen, where something is one of the known possibilities is exactly 1. So, the sum of the probabilities of each possible event or value must also equal 1.

A probability function maps possible events into probabilities. For example, the rolling of a fair dice of the conventional cube-shaped form with numbers 1 through 6 on its sides has a relatively simple probability function—namely $P(1)=1/6$, $P(2)=1/6$, $P(3)=1/6$, $P(4) = 1/6$, $P(5)=1/6$, and $P(6)= 1/6$. Such distributions are said to be “uniform” because the probability of each possibility is the same. “Normal” distributions, in contrast, are bell-shaped and individual possibilities (or range of possibilities) often have different probabilities. In this chapter, we’ll mainly use probability functions where there are just two possibilities, because such probability functions (sometime called Bernoulli probability functions) are sufficient to

probabilities were worked out by games of chance that are much older than that..

¹ Probability theory began about five hundred years ago, although some ideas about

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

illustrate some of the main implications of risky settings.

The two statistical ideas that we'll use most are the ideas of a “**true average**” or expected value, and a **sample average**. If you roll a dice 10 times, write down the individual numbers, add them up and divide by 10, you have calculated a particular **sample average**. In most cases, you'll have calculated a number that in most cases is around 3.5. 3.5 is the true average value generated by rolling a dice in very large samples or calculated using the **expected value** formula. The larger the number of rolls that you tabulate the closer the typical sample average tends to be to 3.5.

That is to say, that the larger the sample is the smaller is the variation in the sample averages that you calculate around the true average value. The “average” deviation of randomly generated series of numbers around the process's true average is called its “standard deviation”. The standard deviation squared is the **variance of a distribution**, which is often easier to calculate than its standard deviation. The variance of a sample average tends to fall as the sample size increases.

Expected Values

In cases in which the probability function is known, one can also use mathematics or arithmetic to figure the average value that you would tend to observe from very large samples. For example, in cases where the probability of every possible random event ($P = P_1, P_2, P_3, \dots, P_N$) and the values are known ($V = V_1, V_2, V_3, \dots, V_N$), or can be accurately estimated, the average value or “**expected value**” can be calculated with the following

formula.

$$V^e = \sum_{i=1}^N p_i v_i$$

Economists for the most part use the term “expected values” for that calculation, although the term average value would more accurately describe the value calculated, V^e .

The expected value of a single roll of a dice is thus:

$$V^e = \sum_{i=1}^N p_i v_i =$$

$$(1/6)(1) + (1/6)(2) + (1/6)(3) + (1/6)(4) + (1/6)(5) + (1/6)(6) = \mathbf{3.5}$$

Notice that the “expected value” is actually impossible in this case, but 3.5 does predict the average value observed for very large samples of dice rolls. Thus, expected values are not always the “typical” value.

In continuous distributions like the normal distribution that are symmetric, the expected value tends to be in the middle of the distribution and is usually among the most frequently observed values. For the normal distribution, the average value is a typical value, rather than an impossible one. Indeed, it is the value that is most likely to occur.

III. Using Expected Values to Make Decisions in Settings where Outcome Are at Least Partly Random (e.g. Risky Circumstances)

This section illustrates how one can use expected values to help make

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

decisions in settings where the results are random and generated by a reasonably well-understood probability function.

For example, consider a career choice after graduation. Suppose that you are offered a job at a new firm (a startup). If the company survives for more than 10 years you'll make 750,000 in wages and have stock options worth another 500,000 in stock options. If the startup fails, you'll have to switch careers which may not be easy and realize no stock options. BLS statistics suggest that only about 35% of firms survive 10 years. You also have a good job offer with another offer from an established firm that will pay out 1000000 in wages over the ten-year period (counting average raises).

To simplify, we'll ignore taxes and time discounting (which is taken up later in the next chapter). The expected value of the startup—which is the risky option—can be written as $V^E = (P^f)(V^f) + (P^s)(V^s)$, where P^f is the probability of failure and P^s is the probability of success. There are only two possibilities so $P^f + P^s = 1$. (The superscripts are just “notation” [a way to distinguish the probability of success (P^s) from the probability of failure (P^f). They are *not exponents* in this case.)

The BLS (Bureau of Labor and Statistics) statistics on startups suggest that $P^s = .35$, so P^f must equal .65. If we substitute these into the expected value formula, we get:

$$V^E = (P^f)(V^f) + (P^s)(V^s) =$$
$$(.65)(750,000) + (.35)(750,000 + 500,000) = 925,000.$$

In this case, the expected value of the start up job is less than the certain value of the conventional firm. So, if you want a job that maximizes your expected income, in this case, you should take the offer from the conventional firm (1,000,000 > 925,000).

Notice that if the odds of success were quite a bit higher, perhaps because of the talent of the entrepreneur or because the idea behind the new firm is especially good, you might reach the opposite conclusion. Suppose, for example, that the probability of success is $P^s = 0.60$ and of failure is $P^f = 0.40$. In that case, the expected value of joining the startup is:

$$V^E = (P^f)(V^f) + (P^s)(V^s) =$$
$$(.40)(750,000) + (.60)(750,000 + 500,000) = 1,050,000.$$

Since there is a greater probability of success, the stock options become more valuable and are now enough to compensate for the lower salary paid by the start up.

Risk Aversion and Risk Premiums

When individual use expected money values to make decisions, they are said to be **risk neutral**. They use expected income—for example—to choose among occupations that are otherwise very good substitutes for one another. There are no adjustments to take account of risk. Such a person would be indifferent between the job at the start up and the job with the job with the regular firm above, if the expected salaries including stock options were the same.

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

However, not everyone is risk neutral. Many folks would want to would prefer a “safe bet” to a “risky bet” if the expected values of the startup job and the conventional job are the same. Such a person would choose the conventional job. To take a job at the start up would require a “risk premium.” They would take the startup job only if it had a sufficiently greater expected value. That extra amount is the risk premium.

Whenever individuals “demand” a premium (a risk premium) to accept a risk, they are said to be **risk averse**. The more risk averse a person is, the larger the risk premium that he or she demands to accept a risky job, plan, or type of asset.

Folks that prefer a bit of risk if the high side reward is great enough are said to be **risk preferring**. Risk preferers are willing take on a lot of risk even if there is only a small chance of a large prize. (A lot of heroes on TV are risk preferers—but far more of them “win” on TV than actually would given the odds of success claimed by the script writers. Although to be fair to the script writers, uncertainty builds tension and the release of tension when the hero wins is enjoyed by their viewers—partly because it is unrealistic and so in a sense special—like winning a lottery. Knight’s entrepreneurs are also risk preferers.)

IV. Buying Products of Random Quality

Consider a consumer’s decision to purchase products when the probability of a defect is known (or can be estimated). A consumer’s willingness to pay for a product can be characterized with his or her total benefit curve for that product. The highest price that a consumer is willing to knowingly pay for Q units of a good is a just bit less than the total benefit generated by that good (his or her reservation price), measured in the currency used in his or her country—dollar, euro, yen, etc.—or in units of satisfaction or pleasure—utility. Successive units of goods and services normally produce additional benefits and so total benefits tend to rise with the number of units obtained.²

Recall that the additional benefit generated by one additional unit of a good or service is called its marginal benefit and can be interpreted as the highest price that a person will knowingly pay for the Q -th unit of a good. Economists normally assume that marginal benefits fall with quantity, so marginal benefit curves tend to slope downward, and total benefit curves tend to rise more slowly as the quantity acquired increases. The highest valued use receives the first unit, the second highest, the second, and so on.

Next we introduce some quality uncertainty. When the quality of a product varies, the marginal benefit received from a given unit cannot be known with certainty unless defects are obvious at the point of sale.

² This subsection is largely taken from Solving Social Dilemmas (2022), with some

revision to make the discussion developed there more relevant for this course.

Principles of Microeconomics: Chapter 8


Random Events, Expected Values, and Market Outcomes

To simplify a bit, we'll assume that only two degrees of quality exist: perfect and defective. Successive units of goods or services have one marginal benefit associated with them when they are perfect and another when they are defective. The marginal benefit curve associated with defective units lies below that of perfect units.

To further simplify, assume that consumers are risk neutral—which means that consumers use statistical averages when determining the “expected value” of a particular unit of a particular good. If probability of a “perfect” unit of a good is P and a “defective” unit of the good is $1-P$, the expected or average marginal benefit (MB) from a particular unit of the good is $P \cdot MB(Q)^+ + (1-P) \cdot MB(Q)^-$ where $MB(Q)^+$ is the marginal benefit of the Q -th unit of the good when it is perfect and $MB(Q)^-$ is its marginal benefit when defective.

Figure 8.1 illustrates a consumer's decision about the number of units of a good to purchase when its per unit price (marginal cost) is P . The expected marginal benefit curve (MB^e) lies between the marginal benefit curves of the perfect and defective units of the product. If $P=0.5$, then the expected MB is exactly midway between the MB^+ and MB^- curves. As the probability of perfect units increases, the MB^e curve moves closer and closer to the MB^+ curve.

Figure 8.1 Maximizing Expected Consumer Surplus



The risk-neutral consumer purchases all of the units for which the expected marginal benefits are greater than their marginal costs, which implies that Q^* units of the good are purchased by the consumer illustrated. Economists refer to the difference between the total benefit received from Q units of a good and the total cost of Q units of a good as its “net benefit” or “consumer surplus” [$CS(Q) = TB(Q) - TC(Q)$].

The total expected benefits of Q units of the good is the area under the MB^e curve from 0 to Q , and the total cost of Q units is the area under the marginal cost (MC) curve from 0 to Q , so the expected consumer surplus realized by purchasing Q^* units of the good is the triangular area between the MB^e and MC curves from 0 and Q^* . The area of this triangle is a measure of a consumer's expected net gains from purchasing Q^* units of the good. However, the amount realized is randomly distributed around

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

the expected value, because the exact proportion of good and bad units varies with the sample taken (the actual units bought). In this case, this is just a matter of good or bad luck, not ignorance on the part of the consumer if the probabilities of good and bad units have been accurately estimated or is known with certainty.

The quantity characterized by the intersection of the MB^e and MC curves, Q^* , is the quantity of the good that maximizes this buyer's expected consumer surplus—which is to say the average CS realized by a long series of such purchases.³

Notice that as the average quality increases, the expected marginal benefit curve shifts toward the higher curve (in blue)—the one characterizing perfect units, and the quantity purchased increases. If the average quality falls, then the expected marginal benefit curve shifts back toward the lowest MB curve (in red) and the consumer's purchases fall.

Notice, we are assuming that different degrees of quality occur randomly. They are not efforts by the seller to defraud their consumers—although such cases do exist. They are simply consequences of variations in

quality associated with their production methods. An agricultural product may receive more or less sun, be growing on a more or less fertile bit of land, be genetically a bit better or worse tasting instance of the fruit or vegetable being purchased. The equipment used to produce the good of interest may be more or less precise and so the items produced may vary in quality.

However, if the typical firm in an industry increases the average quality of the goods produced, the demand for their products would increase because the expected marginal benefit increase for their customers. Contrariwise, if the average quality declines, their market demand curve would shift back to the left (fall) because consumers realize smaller expected marginal benefits from their purchases of the good or service of interest.

This provides firms with an incentive to invest in quality control whenever the additional revenue produced by increased quality and sales is more than sufficient to pay for the cost of the quality control. Perfection, however, is normally too expensive to be worth pursuing.

³ In the case usually assumed by economists, consumer surplus is maximized at the quantity where marginal benefit equals marginal cost. This follows from elementary calculus. If $CS(Q)=TB(Q)-TC(Q)$ then CS is maximized when the first derivative of CS with respect to Q is zero or when $CS'=TB'-TC'=0$. TB' is marginal benefit and MC' is

marginal cost. The mathematics are not important for this chapter but are the simplest way to explain why finding the quantity that sets $MB=MC$ maximizes consumer surplus—the net gains from trade for a consumer.

Principles of Microeconomics: Chapter 8

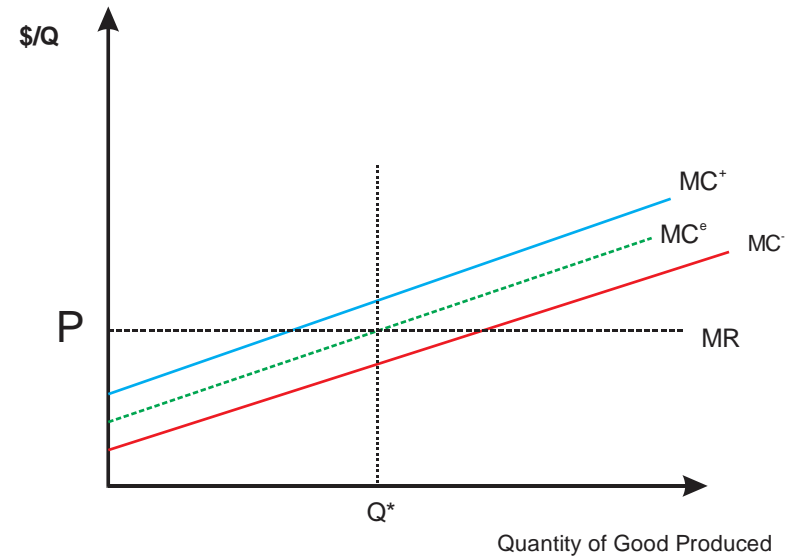
Random Events, Expected Values, and Market Outcomes

V. Producing Products with Random Costs

The choice faced by firms with random costs is very similar to that of consumers facing random quality. If there are just two states of the world that influence a firm's marginal cost function, and the probability of being in the high-cost state is known (or can be estimated with reasonable accuracy), a risk neutral firm will produce at the output level that maximizes **expected profits** rather than certain profits as in the models in prior chapters.

Figure 8.2 two illustrates such a choice for a “price taking” firm, which is a firm that has many rival producers for essentially the same good or service. It has a horizontal marginal revenue equal to the prevailing market price for the good or service that it produces and sells. Uncertainty of this variety can be the result of weather, which may disrupt the firm's production in various ways—but in a more or less predictable manner—as for example very cold weather may make it more difficult to harvest trees from a tree farm, or particularly stormy weather may interfere with fishing, or early frosts or droughts may reduce a farm's output level for given expenditures on fertilizer and tilling for particular crops.

Figure 8.2 Maximizing Expected Profit



Geometrically, expected marginal costs are a curve (or line) between those associated with the high marginal cost and low marginal cost states of the world. Producers maximize expected profits by producing outputs where **expected marginal costs** equal (expected) marginal revenues.

The higher the probability of a “high cost” setting, the closer the MC^e curve is to the higher of the two marginal cost curves (MC⁺) and the smaller is the firm's expected profit maximizing output. Changes in technology that makes crops more robust may have the opposite effect, moving expected costs toward the lower MC curve (by increasing the average output associated with expenditures on seed, tilling, and fertilizer etc.). Such innovations reduce rather than increase downside risks.

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

VI. Producing Products with Random Prices

Market demand or market supply may also be affected by random factors that together jointly determine the price at which a firm's output is sold. Prices in such cases vary in an unpredictable random manner. This, for example, is true of all agricultural products. The world-wide market supply of agricultural products is affected by weather patterns throughout the world. For many other products, variation in demand cause prices for various raw materials to vary widely, because their supply is relative inelastic (their supply curves are nearly vertical) in the short run.

In cases where products take time to be produced, firms cannot simply adjust their day-to-day output for the "spot" price of the service that they are selling. Their time-intensive production processes require firms to produce quantities and at costs that they expect to be most profitable, but they cannot really know for sure what the prevailing market price will be at the time that their products are ready to be sold.

The first case examined in this section is straight forward. If there are just two prices that might prevail in the market of interest, a high one and low one, then their expected marginal revenue is simply the expected or average price that prevails, as with $P^e = fP^- + (1-f)P^+$, where P^- is the low-price choice setting and P^+ is the high price choice setting⁴. The expected marginal revenue curve is just a horizontal line at the expected price, P^e .

⁴ Keep in mind that these models can easily be extended to settings where lots or outcomes are possible. The formula for calculating expected values is sufficient for any

Firms that are aware of the price variation will produce output levels that maximize profits at the expected price. (As a practice exercise, draw a diagram of this choice setting, illustrate the choice using expected marginal revenues and its effect on a firm's expected profits.)

The case in which a firm faces its own downward sloping demand curve is also fairly easy to characterize when the revenue effects are caused by random shifts in the demand curve rather than changes in its slope. If the demand curve shifts in a random way between high demand and low demand states, the expected (average) demand curve and their associated expected marginal revenue curves will be in between those in the high and low demand states.

For example, if each demand curve occurs with probability 0.5, then the expected demand curve will be halfway between the high and low demand curves and the expected marginal revenue curve will be halfway between the marginal revenue curves associated with their respective demand curves.


Figure 8.3 illustrates such a choice setting, but it includes only the expected marginal revenue curve to reduce the number of lines that readers need to keep track of. (Recall that with a straight-line demand curve, its associated marginal revenue curve is downward-sloping line halfway between the vertical axis and the demand curve of interest. In this case depicted, the

countable number of outcomes. Two outcome illustration are used to simplify the discussion and drawings—and without much loss of generality.

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

expected marginal revenue curve is halfway between the vertical axis and the expected demand curve.



The firm chooses the output that sets expected marginal revenue equal to expected marginal cost. It then sells its output at the price associated with the actual conditions, which will vary according to which demand curve it actually faces after its production has taken place. On average this will be the price associated with the expected (average) demand curve. The average price is the one that is depicted in the diagram. The actual selling prices will tend to be higher or lower than the average price at which the product is sold.

Effects of Risks on the Extent of Markets

Notice that in each of the cases analyzed so far in this chapter, the

effect of “downside” risk is to reduce sales and output, as consumers and firms “hedge” their “bets,” and buy less or produce less than they would have in circumstances where only the “good” outcome are possible. **Thus, these models and their associated diagrams can be used to understand why downside risks tend to reduce the size of markets and extent of market networks.**

The effects of risk on the size and extent of markets can be reduced in various ways. For example, decision makers may adopt more flexible consumption or production plans that allow one to rapidly adjust to new circumstances. Travelers may, for example, always pack an umbrella for a long trip. Firms may engage in short-term rather than long-term contracting. Both may accumulate rainy day funds to help cope with unexpected unpleasant surprises.

VII. Risk Premiums and the Demand for Insurance

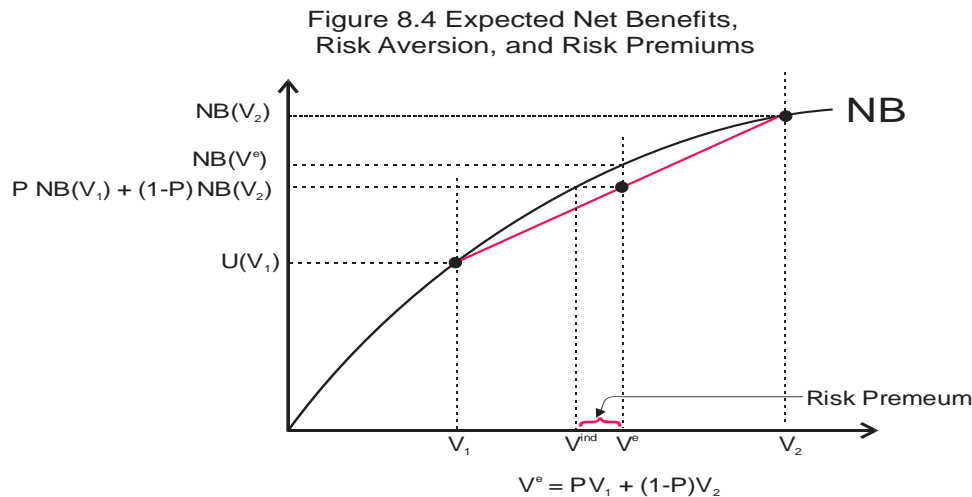
In cases in which a risk can be insured (and there are a variety of insurance and insurance like products that can do so other than the ones sold by insurance companies) it will often be worthwhile to purchase insurance and so reduce the risks associated with long term commitments, such as purchasing an automobile or house.

Risk-averse individuals are willing to pay a “safety premium” to realize a certain return rather than a risky return. That safety premium provides the basis for their **demands for insurance.**

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

Figure 8 provides the geometric logic and geometry behind the existence of risk premiums and thereby on the demand for insurance. The curve depicted is the **total net-benefit curve**. Diminishing marginal benefits and more or less constant marginal costs, imply that the total Net Benefit Curve has a shape similar to this one up to the point where maximum net benefits are realized with certainty. (When the payoffs are risky, a risk averse person will purchase fewer units than that best case NB-maximizing amount, so it is this portion of the NB curve that is relevant here.)



The net benefit curve of a risk-averse individual (or profit curve for a risk averse firm) is upward sloping and curved so that it becomes flatter as the choice variable (value, quality, or quantity) increases (e.g. as the values along the horizontal axis increase). Such curves are said to be **strictly concave**. It turns out the more rapidly NB curves flatten out as one moves to

the right, the larger their risk premium is, and so the more risk averse an individual is.

The diagram illustrates the case where a risky purchase or investment may have two values in the future (as, for example, the value of a house that might or might not catch fire or be affected by flood or landslide). It will either have value V_1 or value V_2 —which one is determined by some random process, with P being the probability of V_1 and $(1-P)$ being the probability of V_2 .

As P varies from one to zero, the expected net benefit traces out the straight line that lies below the NB curve (the net benefit curve). The expected net benefit is the highest price that a risk-averse consumer would pay for this risky asset. Notice that this amount is always below the NB line because of the curvature of the NB curve. If P is 0.5, the expected or average net benefit will be at exactly the midpoint of that line connecting the $NB(V_1)$ and the $NB(V_2)$.

Note also that the average value of the asset (V^e) lies between V_1 and V_2 on the horizontal axis. The net benefit of the various possible outcomes differs from $NB(V^e)$ because of the way the consumer surplus (or reservation prices) varies with different values of V for risk averse individuals. (If V^e always equaled NB^e , the individual would be **risk neutral**, and the NB line would be a straight line in the range depicted.)

A risk-averse person is only willing to pay an amount that is less than the average value of the risky asset. That is to say, he or her risk aversion

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

implies that the expected net benefit of a risky investment is less than a certain investment that returned the average value, V^e , with certainty. This is shown on the graph by the fact that at V^e , the expected net benefit is lower than the value on the NB curve at V^e .

The equivalent certain value to the risky investment can be found by going from NB^e over to the NB curve and then straight down to the V axis. That value is labeled V^{ind} , because the individual would be indifferent between that amount (with certainty) and the risky investment.

Notice that that value is less than V^e . **The difference between V^e and V^{ind} is the risk premium that an investor would require to accept the risky investment** (V_1 with probability P and V_2 with probability $1-P$) rather than a certain one that returns V^{ind} .

It is this risk premium that creates the demand for insurance. It is the highest price that an investor, consumer, or firm would be willing to pay to totally avoid the risk associated with the investment (asset) characterized in the diagram. Thus, it might also be called the individual's "safety premium," the highest amount that he, she, or it is willing to pay to avoid a risk rather than bear it. If full insurance can be purchased for less than that amount the investor would be better off buying it than bearing the risk personally.

Thus, the risk aversion and its associated risk premiums is the ultimate source of the demand for insurance.

Individuals are not willing to pay more than their "safety premium" for

complete insurance but would be happy to pay less than that (because then he or she would gain expected net benefits from the insurance).

VIII. The Supply of Insurance

In principle, there are gains from trade between more risk averse people and less risk averse people, since less risk averse people are willing to pay a higher price for a risky asset (with a well-understood risk) than a more risk averse person. So, one market response to risk, is that risky assets tend to "move" from more risk-averse persons and firms to less risk-averse firms. For example, risk-averse persons would be out bid for ocean-front properties along the South Atlantic Coast of the United States because of the risk of hurricanes. The risk averse lovers of oceans would prefer property that is inland a bit, but perhaps in walking distance of the beach, or a really solidly built beach-side house over one that is a bit ramshackle.

Another effect of risk is the emergence of insurance markets.

We'll focus on the standard ones that consumers deal with rather than various futures contracts and similar investments that some investors use to reduce their risk. It turns out that insurance markets can be quite profitable because of the statistical properties of large samples. So, to understand the basic logic of an insurance company, first we must understand a bit about those statistical properties.

On the Properties of Sample Averages

Every random variable has a probability function associated with it

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

that assigns probabilities to the possible values of the random variable of interest.⁵ Every probability function, thus, has a range of values. A measure of the breadth of that range is called the **variance** of a probability function. Given the probability function over variable X, **the variance of that probability function can be calculated as:**

$$\text{Var}(X) = \sum_{i=1}^N p_i (X_i - X^e)^2$$

A related measure of the breadth of the outcomes associated with a probability function is its standard deviation, which is simply the square root of a probability function's variance. $S.D. (X) = (\text{Var}(X))^{0.5}$. (Greek letters are sometimes used for the mean or average value of a probability function as with mu, μ , and the standard deviation of a probability distribution is often written as sigma, σ , and its variance as sigma squared, σ^2 . But we do not need that notation for the purposes of understanding the supply of insurance.)

What is important about variance (and stand deviations) for insurance markets is that the smaller those values are, the narrow is the range in

which most values of the probability distribution fall and the more an asset resembles one without risk.

It turns out that the larger a sample size is, the more narrowly the values of a sample average are distributed. The variance of a sample mean, $\left(\frac{\sum X}{N}\right)$, falls as the sample size, N, increases.

$$\text{Var}\left(\frac{\sum X}{N}\right) = \text{Var}(X)/N = \left(\frac{1}{N}\right) \left(\sum_{i=1}^N p_i (X_i - X^e)^2\right)$$

Selling insurance is like sampling, because each insurance policy for a similar house or asset, simply another instance of the random process that is generating the risk (fire, flood, loan default, etc.). So, a company the sells lots of insurance for risks that are well understood, pays out, on average, the sample average of its collection of insurance policies. The insurance company's risk is basically the variance about its sample average, which becomes smaller and smaller as more and more insurance policies are sold.

So, for large insurance companies, the payout for insured losses is, on

⁵ Technically this is only true of probability functions that have only discrete values. For random variables that have continuous values, a probability density function is used to characterize the probabilities that values fall within a range. The "normal distribution" is an example of such a probability density function. However, the intuition that is associated with probability functions extends pretty well to probability density functions, so

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

average very predictable. It is approximately the true average value of the losses generated by the random phenomena being insured (as with fire or flood insurance). In other words, “pooling risks” essentially eliminates the risk (random nature) of the loss being insured for by the insurance company. (The initial risk sometimes said to be eliminated or greatly reduced by pooling.)

The Profitability of Insurance Companies

This sampling effect allows even risk averse firms to profit by selling insurance if they can sell insurance and settle claims at a reasonable cost. Unlike the owner of a house, whose may burn down or not in a given year, the insurance company essentially pays out the average or expected loss every year.

Since those demanding insurance are willing to pay up to the average loss plus the safety premium as long as its selling and settling costs are lower than the risk premiums that its customers are willing to pay, selling insurance can be profitable. In such cases, there are potential gains to trade between homeowners and insurance companies. Competition between insurance companies keeps their prices down to more or less ordinary rates of return for the products that they sell. Once up and running, insurance markets often resemble ordinary competitive markets.

The emergence of insurance markets is one of the implications of commonplace risks such as those associated with owning a house or car—or even offering loans (banks often insure against non-performance). Such

markets would be unnecessary in a world in which everything was certain and well understood by consumers and firms.

IX. Rational Ignorance, Markets, and Specialization

Prior to World War II, there was very little consideration of the role of information in choices by consumers and firms and thereby on the market equilibria. In 1961, George Stigler wrote an important paper on the economics of information in which he suggests that gathering information is analogous to sampling a random distribution. Sampling allows an individual to know with greater precision the random process generating market prices and many other things. The more one samples, the more precise is that understanding—the more precise is one’s estimates of the mean, variance, and trends in the phenomenon of interest.

The extent of that information can affect market equilibria. For example, if a sufficient number of consumers search for the best price, they tend to induce market prices to converge on the lowest ones that they have been able to find. This is one of the mechanisms that tends to support the “one” price equilibria of the demand and supply models that we studied in the first half of the course. On the other hand, if consumers do not search for good prices, there will be less demand pressure on firms to sell at low prices and the range of prices among shops will be broader. Although this seems obvious now, Stigler won a Noble price, partly for that paper.

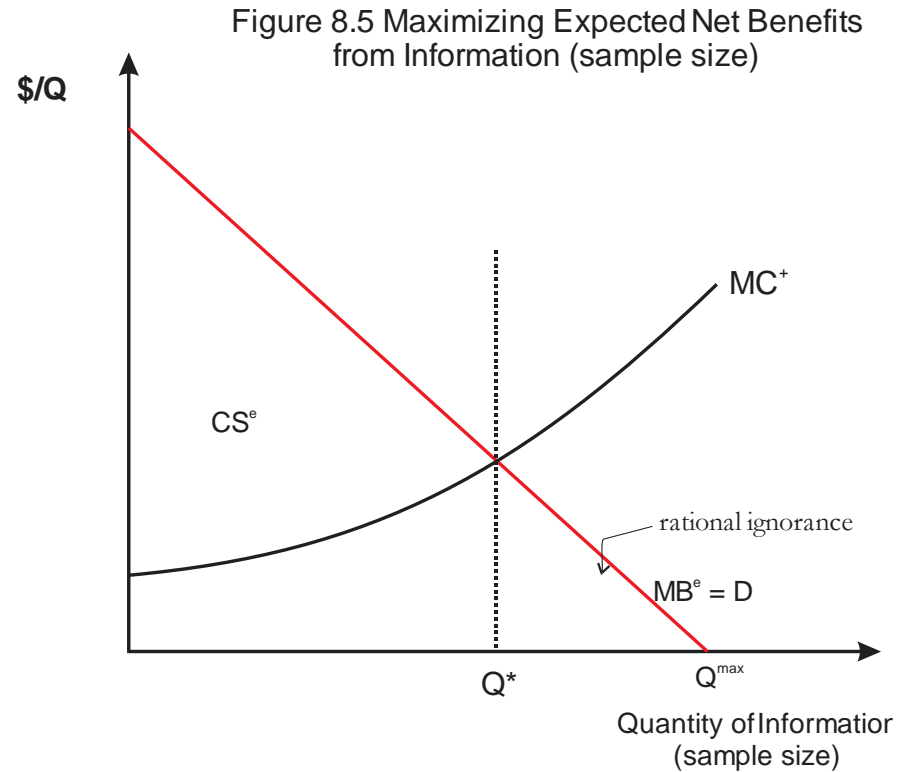
The basic idea is that seeking information has expected benefits and expected costs, so individual’s rationally economize on information. (A

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

very similar idea was worked out by Anthony Downs at about the same time with respect to voter information about candidates and public policies.) As a consequence, neither consumers nor producers nor voters nor students, etc. know as much as possible about most things. They remain rationally ignorant of many potentially useful things because they only have so much time to devote to collecting and analyzing data.

Figure 8.5 illustrates the logic of an individual's choice of the amount of information to gather from a randomly generated distribution that he or she is aware of. The marginal cost line represents the opportunity cost of using time to gather and process the information of interest (for example, about the price of cell phones or their capabilities). The marginal benefits are "expected" because one never really knows what one will learn until one does so.



Rational ignorance (and natural ignorance—ignorance that one has not chosen) is an important phenomenon in any setting in which information is not naturally very complete and reliable. In some cases, ignorance can induce systematic errors. In others, ignorance simply implies that one's estimates are less accurate (lower variance) than they could have been, although they are still unbiased (the expected values are correct rather than different from the actual value).

In some cases, the marginal opportunity cost is also informational—that is, the best alternative use of one's time may be learning about

Principles of Microeconomics: Chapter 8

Random Events, Expected Values, and Market Outcomes

something else. If you have T hours to divide up across informational investments, a rational student will allocate time so that the marginal benefit from time is the same for each use. This may imply very different amounts of time devoted to different types of study because the expected marginal benefits may differ substantially.

For example, the marginal benefits from learning more about one's expected major has relatively larger marginal benefits in terms of adding to skills that you are likely to find useful in your future career than in studies that have no direct entertainment value or affect on one's future earnings in one career.

Markets make use of the information that individuals have and what they lack. Knowledge and skills that are scarce relative to demand tend to have higher prices associated with them. This in turn encourages other students to specialize in acquiring such knowledge and skills. In the very long run, prices may thus cause the market value of skills that require equal investments to acquire to converge on roughly the same wage rates unless the capacity to acquire skills is itself scarce, because of variations in talent, diligence, or regulations that make it difficult for some persons to acquire the skills with the highest return.

Specialization, as an economic activity, often is associated with focusing most one's attention on specific types of information that are useful for one's career (or expected career). Thus, to a considerable extent, it tends to reflect choices about which kinds of knowledge to acquire and which

types one chooses to remain ignorant of.