Chapter 9: The Economics of Information

I. Information, Knowledge, and Markets

The last two chapters have explored implications of choice settings in which individuals and organizations cannot precisely know the consequences of their actions. The consequences of an action or purchase may be probabilistic—as when one purchases a lottery ticket or a product whose quality is characterized by a probability function such as the uniform or normal distribution. Or the choice setting may be one in which the consequences of an action or purchase are (or may be) results of such complex or chaotic processes that no probabilistic representation is possible. In the latter case, individuals may nonetheless use probabilistic ideas to try to assess the relative merits of alternatives using subjective (Bayesian) rather than objective frequentist characterizations of probability functions.

Such subjective assessments, for example, allowed us to use expected utility and expected profit models to characterize the research and development activities of firms. Probabilistic thinking—even when no specific probability function is known beforehand—allows individuals to be systematic in making choices in settings of uncertainty given what is known beforehand, although it does not eliminate mistakes. Individuals and organizations may later regret such decisions, because the probabilities used are wrong, rather than rooted in observation.

It is important to realize that in both choice settings, individuals have been assumed to know as much as it is possible to know. If a process is truly probabilistic, one can only know the probability function, not the outcome of such processes. Similarly, if a process is truly chaotic, little systematic can be known about it beyond the fact that it falls into that category of phenomena.

Essentially all of the models developed to this point have assumed that knowledge is exogenous, beyond the models. In the first seven chapters, everyone knew as much as could be known. In the last chapter, ignorance of various kinds existed, but there was little (other than R&D) that could be done to change it. This chapter brings the extent of knowledge further into the analysis. This chapter explores the market for information. A good deal of information is produced by the individuals themselves, and may be regarded to be instances of household production. Other information is produced by information specialists—as with doctors, lawyers, teachers, researchers, newspapers, and the marketing campaigns of firms. It is various combinations of information—that already acquired and internalized and that being acquired that allow individuals to assess the relative merits of the possibilities before them and make choices.

In some cases, additional information improves one's understanding of phenomena that one is already familiar with. In the previous models, these include estimates of prevailing market prices, income, efficient production methods, the preferences of consumers—and to some extent the nature of one's own utility function. More complete information improves one's decisions by improving one's estimates of parameters of relevant functions that affect choices within familiar settings.[1]

In other cases, new information may reveal possibilities or causal chains that had not been previously imagined or taken into account. For example, the new information provided by satellites has made short-term weather forecasts much more reliable than they were a century ago—not simply by improving our estimates of well-understood probability functions but by

---

[1] The idea of a conditioning variable for a probability function is not obscure—although we have not used that idea before. An example of a conditioning variable is cloud cover. When the sun is shining it is less likely to rain than when dense (dark) clouds are overhead. Conditioning variables are often written as $p(x) = f(x|z)$ where $p(x)$ is the conditional probability function, f is the conditional probability, and z is the conditioning variable. The conditional probability of rain $p(x)$ varies with the degree and darkness of the clouds, $f(x, z)$ in the example just used, where z is the cloud density.

adding previously unknown causal factors that tend to produce one pattern of weather rather than another.[2]

In both types of choice settings, additional data (information) is useful, which creates a demand for information. As a consequence, individuals and organizations are constantly gathering and interpreting information. Living involves a good deal of learning—which is true of economic and other areas of life.

Even when we walk across a well-marked pedestrian crossing of a paved road, rather than decide where each foot should be placed before crossing the street and then closing our eyes and walking as planned, we take a step or two and then consider where to place the next few steps based on fresh information about the road surface (potholes or trash), the apparent path of other pedestrians, and our balance after we have taken the previous step. As a consequence, most of us continually adjust our planned path as we walk across a street in light of new information. The same is true of many aspects of one's personal life—including many that are relevant for market prices, outputs, and patterns of exchange.

Time Allocation, Preprocessed Data, and the Demand for Information

Producing knowledge requires time and attention (and in some cases special tools for observing and measuring the phenomena of interest). There are a variety of ways that one can learn new things and people make use of all of them. For example, one can collect all data for oneself, or make use of "preprocessed data" to increase one's knowledge. Direct personal experience is an instance of the former and "book learning" of the latter.

The latter is important, because there is so much that one could usefully know and insufficient time for an individual to reinvent fire, the wheel, language, domestic animals, or skipping forward: the automobile, cell phone, and satellite. Thus, much of what we know is

[2] An example of this is the effect of El Niño conditions in the Pacific Ocean west of South America on weather patterns in North America.

learned indirectly from others because it allows much more to be known about the world than we could figure out on our own.

As a consequence, there is a demand for "preprocessed data" or, put more simply, a demand for information.

Demand implies a willingness to pay for information (as with time and money spent on a college education, training program, or informative books, magazines, websites, and AI programs). That demand implies that there are potential profits to be realized by supply. It is largely for that reason that the services and products just mentioned exist. There are markets for many types of information—even though there are many sources of information that are free. (Which is to say, available to those willing to devote part of their time and attention to acquiring it.)

That free information exists, however, does not allow one's demand for information to be fully informed or fully satisfied. For example, there are many thousands of books in university libraries and websites on the world-wide-web that an individual can access without paying anything in money. However, we have only 24 hours in every day, and much of that is allocated to sleep. We simply do not have the time or attention to read or master them all. Consequently, we all make choices about which books to read, which sections within books to focus our attention on, and thereby what will be ultimately learned by reading what we do.

Exactly what one will learn by reading a book or taking a course remains unknown until the book is read or course is taken. And, similarly, the authors of those books rarely knew exactly what will be included in their books until the final drafts of their manuscripts are completed. The demand for knowledge is not and cannot be a fully informed demand.

Relevance of Learning for Markets

A large subset of our informational choices and our associated accumulation of knowledge is relevant for understanding how markets operate. The most commonplace of these effects are the main focus of this chapter.

For example, no consumer knows all the prices of all the goods sold in markets, as assumed in Part I. Consumer knowledge and ignorance affects the kinds of goods and services demanded by individuals, and, indirectly, the kinds of goods and services produced and brought to market for sale. It does so because information about products and their uses is bounded by the time and energy one has devoted to learning about them. There are many products that we are unfamiliar with that may-well advance our interests (contribute to our utility) more effectively than the products that we purchase. Discovering such products is a good deal of the fun (useful information) that shoppers get from shopping. (How firms develop specific products is taken up in Chapter 10).

Specialization among humans is mostly an informational phenomenon. We know some things and how to do some things better than others (our specialities) and are largely ignorant of many others. Our inability to master all fields simultaneously is part of the reason that individuals and organizations specialize. The economic returns to various forms of specialization, in turn, influence most individual and firm-level decisions about what to specialize in—at least insofar as careers and products are concerned. Informational differences, thereby, affect the manner in which the things brought to market are produced.

Hayek (1945) once argued that it is differences in the knowledge possessed by individuals that make market prices and specialization critical determinants of prosperity. Understanding how market-relevant decisions affect the distribution of information among the individuals in an economy is thus necessary if we are to understand how markets operate.[3]

---

[3] Some individuals may not be aware of the scope of their ignorance—in effect, they are ignorant of their own ignorance—and so act with what Adam Smith once termed "over-weening conceit." However, the attitudes of "know it all's," while comforting, tends to make such decision makers more error prone than they could be. It tends to make them more confident than they should be, and thus more likely to be risk takers than they should be. The latter leads to many bankruptcies and failures of new firms. However, it should be noted that their occasional surprising successes often benefit the more sensible prudent men and women who more accurately assess the riskiness of such ventures, and so would never have attempted to develop the unlikely—but sometimes very successful—products, production methods, and economic organizations created by such over-confident entrepreneurs.

II. Modeling the Demand for Information

The demand for information differs from the demand for most goods and services because it is a demand for something that cannot be known beforehand. If one already knows a price, probability, or other fact, then it is unnecessary to spend one's scarce time and attention acquiring such information. It is only in areas in which one is less informed than one believes would be useful that it can be worthwhile to devote part of one's scarce time and attention to acquiring additional information or to undertaking additional processing of information that one already has.

Many choices about information are instances of normal risk taking in that one has reasonable expectations about distribution of possible results associated with gathering more information. For example, the current price of a good or investment may be relevant for some other choice and being familiar with the past values of the price of interest, one knows quite a bit about the probability that the price will take various values through time. Knowing the current price, however, may be particularly relevant for a buy or sell decision of a stock or for the purchase of a particular good—is it "on sale" or not? In other cases, the expectations may be less concrete, but still reasonable. Students that have taken several economics courses may confidently predict that "If I take more economics courses, I'll know more about economics afterwards than when I started" without knowing exactly what will be learned.

Other informational choices resemble Knightian uncertainty. In cases of complete ignorance or in which one has only a bit of information about the phenomena of interest, one lacks sufficient knowledge to construct an observation-based probability function. One may make guesses about what one will know after such an investment in information is made, but prior to the investment, one does not really know what will be learned. A student that takes his or her very first economics courses may have no idea what will actually be learned—and only realize well after the course was taken whether it was worthwhile or not. The information acquired in such cases may be a complete surprise!

In both settings, choices are made to acquire additional information and in both cases, the information obtained may improve subsequent decisions—but the two types of decisions are quite different from one another. In the first case, more or less rational decisions can be made that can be characterized with the rational choice models developed in this textbook (and in economics and game theory more generally), in the second, such calculations cannot really be undertaken. Although as mentioned before, one can use the expected utility and expected profit approach to gain some insight into the factors that affect such decisions, those models cannot fully capture Knightian uncertainty or much about the individual calculations that induce them to seek such information.[4]

In many cases, the collection and processing of information reduces uncertainty and thereby makes decisionmaking less error prone. However, there are also cases in which more information reveals uncertainties that one did not know about beforehand, which increases uncertainty. Such learning can induce better planning not by reducing uncertainty, but by better accounting for risks that previously were unknown.[5]

---

[4] Again, this chapter is not attempting to characterize the demand for information as entertainment—which is an important rationale for collecting information characterized by Knightian uncertainty. Tourists and scholars, for example, often put themselves into unfamiliar contexts so that they will be surprised (hopefully more often pleasantly than unpleasantly) by the information (experiences) acquired. It is the discovery of the new that motivates the collection of such information, rather than any clear idea of improving one's future decision or earning profits by selling it.

[5] Accumulating more information does not always reduce uncertainties, as some statistical measures of information imply, but nonetheless may improve decisions by providing a more realistic understanding of a choice setting. This is a flaw in many statistical measures of information, which were initially developed as models of communication. Improved communication occurs when errors in interpretating the messages sent diminish. Such communication is said to be more informative. See, for example, Shannon's classic book (1948) or Shannon and Weaver (1998). This usage of the term information has by now produced a very large and technical field of research. Most of that field, however, is beyond the scope of this chapter, which attempts to model some of the essential features of more or less informed decisions, rather than communication, signaling, or artificial intelligence.

The "Search" Model of the Demand for Information

We begin with one of the simplest models of the demand for information and its associated learning. It is a model that is based on the "search" model introduced by George Stigler (1966) which attempted to analyze the extent to which informational costs could account for the dispersion of prices in real markets.[6] Outside of organized stock and futures markets, there is usually more than one price for goods sold in markets. Prices do not converge to a single price but to a cluster of prices around some average price—an average which may be regarded as the market clearing price of the models developed in the first part of this book. There are other explanations for price variation—for example differences in the "free" services or waiting times offered by different merchants—but the informational explanation is important and emerges from a straightforward extension of the rational choice model. It turns out that the prices that one pays for goods are partly a matter of choice rather than entirely determined by market forces beyond one's control.

In the leanest neoclassical models, all knowledge problems and transactions costs are assumed away, in which case the 'law of one price' holds. If consumers know the prices of every good on offer, they will always purchase goods at the lowest available price. Thus, only one price can exist, namely the lowest one sufficient to keep suppliers in business. The lowest one that can be sustained is the one that equals the average cost of production (for the marginal firm) as pointed out by Marshall in the late nineteenth century among many others. Once the lowest sustainable selling price emerges, a price lower than that price would attract all the buyers in the market, but firms charging that lower price will lose money and go out of business. Any price higher than the minimum on offer will attract no customers—since there is no reason to spend more on an identical product than the lowest price available. (Differences in services and/or locations in neoclassical models imply that products are different, rather than being identical products. It is the full package of goods,

[6] This paper was mentioned in Stigler's 1982 Nobel Prize announcement. https://www.nobelprize.org/prizes/economic-sciences/1982/press-release/

services, and location that identify a "product.") It is this type of logic that was implicitly relied on when the neoclassical theory of equilibrium prices was reviewed in chapter 5. Unique equilibrium prices emerge in such models.

The logic of "perfect competition" is clear---but what if information about prices is not freely available to all consumers? Stigler (1966) argues that in such cases consumers would sample prices and purchase goods at the lowest price in their sample. That is to say, they would search for the lowest price. However, different consumers would choose different samples, and so would purchase at different prices because the lowest price in their samples would differ. The cost of searching for the lowest price and the expected savings on purchases would determine how much searching an individual undertakes.

Stigler's theory has often been referred to as the search theory of information, but it is better thought of as the sampling theory of prices. He argues that the easier it is to search for prices (e.g. the less costly sampling is) the narrower is the range of prices that tends to emerge in markets for a homogeneous product. Searching (sampling) limits the range of prices paid for goods and services because consumers purchase the good at the lowest price in their sample. Thus, firms with well above average prices sell relatively little of the good or service of interest—e.g. only to those who do very little or no searching. If most people do a bit of shopping around, the result is reduced sales and profits for the higher priced merchants, and greater sales and profits for lower priced firms. Thus, Stigler's theory can account for both the dispersion of prices observed in the real world and for the extent of that dispersion—something the law of one price cannot.

Modelling an Individual's Price Sampling Decision

Suppose that Al knows her own demand function $Q^D = q(P, Y)$ but doesn't know the price at which the good of interest will be purchased. However, she knows from statistics that as more sampling is undertaken, the lowest price in one's sample tends to fall. In this case, sampling might involve driving to numerous shops and inquiring about the price of the good she is thinking about purchasing or simply spending time on the computer looking at alternative web-sellers for the good of interest. In each case, sampling has an opportunity cost—namely the value of the time spent searching plus any other expenses such as fuel for one's vehicle or electricity and internet service for one's computer or cell phone.

The simplest way to model this choice is as an effort to minimize Al's overall expenditure on the good of interest. Al's expected overall expenditure (E^e) on the good of interest is the amount spent to purchase Q units of the good (PQ) plus the opportunity cost of searching for the best price (C). This can be written as:

$$E^e = P(N)q(P(N), Y) + c(N) \qquad (9.1)$$

where N is the number of shops or web-sources sampled, q is Al's demand function and c(N) is the opportunity cost of search—other things being equal. The expected price, P(N), is a function that declines with sample size at a decreasing rate ($P_N<0$ and $P_{NN}>0$). The sampling cost function, c(N), increases with the sample size at an increasing rate ($C_N>0$ and $C_{NN}>0$). Differentiating equation 9.1 with respect to N characterizes Al's optimal sample size for his or her searching efforts—e.g. efforts to learn what the lowest available price is for the good of interest.

The first order condition looks similar to the usual ones, but in this case, Al is interested in minimizing her total cost for purchasing the good of interest and undertaking search.

$$P_N[Q + PQ_P] + C_N = 0 \equiv H \ at \ N^* \qquad (9.2)$$

Or alternatively as:

$$-P_N[Q + PQ_P] = C_N \ at \ N^* \qquad (9.3)$$

The first term in equation 9.2, $P_N[Q + PQ_P]$, is Al's expected marginal benefit from sampling, which is the marginal reduction in the expenditures on the good associated with lower prices generated by sampling ($P_N <0$). The second is the marginal cost of sampling. Al's expected marginal benefit from searching varies with the extent to which prices tend to decline with search (P_N) with the quantity likely to be purchased (Q), the price ultimately expected to be paid (P), and the price sensitivity of his or her demand function (Q_P).

Statistical theory implies that as the variance of the price distribution increases, P_N tends to increase, which tends to induce greater search because the marginal benefits of searching increases.

Searching for the lowest price thus tends to put downward pressure on the firms selling the good of interest at above average prices. It is the searching efforts of consumers, thus, that induces prices to cluster within a relatively minor band. Prices do not collapse to a single price because the marginal cost of searching is greater than zero.

Equation 9.3 expresses the first order condition in somewhat more conventional MB=MC terms. Al's overall expected cost for the good is minimized when the marginal benefits of search (the negative of expected marginal expenditure reduction) equals the marginal cost of searching for a lower price (e.g. the marginal cost of an increase in sampling). Figure 9.1 illustrates equation 9.3.

Figure 9.1: Expected Cost Minimizing Search

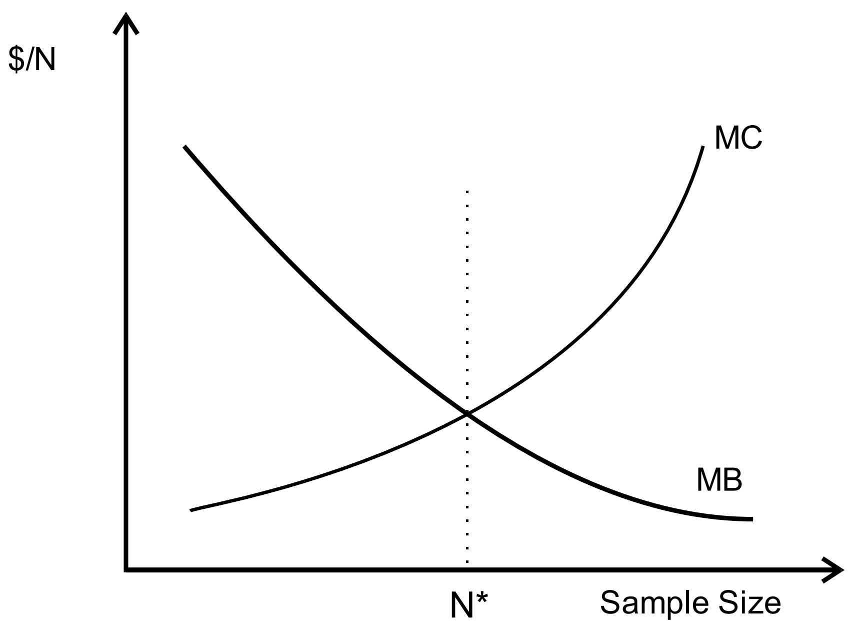

Comparative statics of this relationship can be characterized using the implicit function differentiation rule. As modeled, Al's income, Y, is the only exogenous variable and it affects demand rather than costs. Thus,

$$N_Y^* = \frac{H_Y}{-H_N} = \frac{P_N[Q_Y + PQ_{PY}]}{-(E^e{}_{NN})} > 0 \tag{9.4}$$

If an expenditure minimum exists for N>0, then the expenditure function is U-shaped, which implies that an increase in sampling at N* increases marginal expenditures, so $(E^e{}_{NN} > 0)$ at N*. Thus, the denominator is less than zero.

The numerator has an ambiguous sign. The first term of the numerator, $P_N Q_Y$, is negative. Purchase price falls with sample size ($P_N < 0$), and purchases of goods tend to rise with income ($Q_Y > 0$). The effect of income on the effect of price on demand tends to diminish (becomes less negative, e.g. shifts Q_P upwards, so $Q_{PY} > 0$). Thus, the overall sign of the numerator is negative, which given the negative sign of the denominator implies that the overall effect of an increase in income is positive—assuming that income does not affect the marginal cost of searching.

The model implicitly assumes that an increase in income does not increase the opportunity cost of searching, which it might, in which case the overall effect of an increase in income on searching is ambiguous. If income were included in the opportunity cost function—which is plausible in many cases, the marginal cost of searching would tend to increase with income. In that case, the term $-C_{NY}$ would be included in the numerator, and the overall sign of equation 9.4 would be undetermined without further assumptions. It would depend on whether the purchase effect, which is the effect characterized in the model as developed, dominates or is dominated by the marginal cost effect of higher income. It is probably the case that millionaires search less regarding the prices of ordinary goods and services than middle class persons do, because for them the opportunity cost effect dominates the purchase effect.

Stigler's model demonstrates how a rational choice model can be used to characterize individual investments in useful data for familiar things (here prices). Such information allows better (less error-prone or more fruitful) decisions to be made and actions taken.

Moreover, in this case, a modest extension of the neoclassical model can also explain why observed distributions of prices do not converge to a single price. Knowledge does not have to be perfect for markets to operate. Rather, imperfect information can be incorporated into models to provide new insights about prices.

Indeed, assuming that information is "given" is a mistake when analyzing choice settings and/or markets where information can be accumulated and informational differences exist that are likely to be economically significant. Those who search less, on average, pay higher prices for the things that they purchase.

The Value of Larger Samples for Firms

The logic of price searching also applies to firms with respect to their inputs and production processes. Firms have an interest in purchasing their inputs at the lowest price possible—and so they will engage in a search process similar to that described above for inputs—where their demand for those inputs is based on the marginal-revenue product of those factors of production. Similarly, sellers with price-making ability may search for the highest price that a good or service can be sold for. And, there are cases in which expert opinions are sought about production methods unfamiliar to firms, because of recent innovations.

Most firms that set their own prices sell products that are a bit unique, although other similar products may be available. The highest price at which their output can be sold is determined by their demand functions, as developed in chapters 3 and 5. As a consequence, such firms can often profit from better estimates of their future demand functions. Better estimates of their demand function helps a firm determine the optimal scale of their production facilities, the price at which their product or service is likely to be sold (on average), and also the extent to which the product is likely to be profitable.

This section models a single firm's effort to estimate its downward-sloping demand function when more accurate estimates are more costly than less accurate ones—partly because more data and more sophisticated models are required for more accurate estimates.

The Profits and Costs of More Precisely Estimated Demand Functions

The profitability of a firm is affected by the variance of sales, because production normally takes time, and inputs for production are often purchased or contracted for long before a product is brought to market. The higher the variance in estimated demand is, the more storage will be required during periods of lower-than-average demand when prices fall temporarily below production costs. Thus, the higher the variance in sales, the higher are production costs and the more likely a firm experiences excessive inventories of both outputs and inputs. Estimation errors are costly. Accurate demand estimates reduce costs and allow better estimates of profits to be undertaken.

The variance of profit flows—as well as their averages—are of interest to firm owners that are risk averse. Textbooks model firms often assume that firm owners are all risk neutral, and so maximize expected (average) profits without worrying about the variance of their profits.

However, that approach is clearly less realistic than it could be. Entrepreneurs and shareholders may be less risk averse than the average person, but relatively few investors in firms are indifferent to risk—because their own wealth and lines of credit are finite. Risk aversion implies that the variability in profits will matter to most owners, financiers, and shareholders.

The variance in the profits of a price-making firm is partly determined by the accuracy of its estimated demand and cost functions, and partly by irreducible random fluctuations in those functions and in their input prices. Better estimates of their demand functions cannot reduce such fluctuations to zero, but they can reduce profit variability by reducing forecast errors. It is for this reason that large firms sometimes have their own forecasting units or contract with consulting firms to provide accurate forecasts. A more accurate estimate (smaller variance) increases risk-adjusted profits by reducing planning errors, storage, and re-contracting costs.

The deterministic models of firm decisionmaking developed in chapter 3 ignored randomness and uncertainty. Demand, in principle, may be mechanically causal (as in the models in chapter 2), stochastically causal (as in the models of chapter 7), They may also be affected by uncertainty, insofar as consumers make decisions in settings in which there are chaotic or poorly understood phenomena.[7]

The issue addressed in the model developed below is how much data a profit-maximizing firm should collect and analyze. If firm owners have utility functions that are strictly concave (as assumed in previous chapters), they will be risk averse, and invest some of their wealth in data collection to reduce their risks from fluctuations in their income associated with profit variability.[8]

---

[7] Most of the statistical models that ground econometric studies of demand assume that some degree of randomness exists beyond the core deterministic aspects of demand. That is to say even if one knew exactly what consumer income was and the prices of other goods, the demand function would not be completely identified, because of random "shocks" of one kind or another beyond the variables of the demand function, as with $Q^D = a + bY - cP + e$, where a, b, and c are constants (parameters of the demand function), Y is consumer income, P is the selling price of the good, and e is a purely random factor with mean zero and variance σ^2.

[8] Although a profit maximizing objective function is used below. A similar analysis could have been undertaken within the utility maximizing framework. Suppose that f(Π) is the frequency distribution of profits (Π) and f(Π|D) is the conditional distribution of profits for given amounts of data collection and analysis (D). Combining ideas from chapter 2 and chapter 7 allows the expected utility maximizing choice for a two good utility function with uncertain income (here Π) to be written as $U^e = \int u(X_1, (\Pi - P_1 X_1 - P_D D)/P_2)) f(\Pi|D)$. This formulation uses the expected formula for a variable with a continuous frequency distribution. Π is the individual owner's or entrepreneurs' share of the firm's profits. X_1 is the quantity of good 1. P_1 is the price of good 1, and P_2 is the price of good 2. As, written, this firm owner/entrepreneur has two control variables, X_1 and D, assuming that the firm of interest produces a third good or service not purchased by the firm owner/entrepreneur.

Differentiating with respect to X_1 and D yields two first-order conditions that characterize the entrepreneur's consumption of goods 1 and 2, and his or her investment in data purchases and analysis.

$U^e{}_{X_1} = \int [u_{X_1} - u_{X_2}(P_1/P_2)] f(\Pi|D) = 0$ at X_1^*

The Optimal Precision of a Firm's Estimated Demand Functions

One way to incorporate the effect of reductions in risk on firm profits is to use an idea from Chapter 7, namely the idea of a risk premium. Risk-averse investors require a risk premium to undertake risky investments—a profit above that associated with the risk-free (low risk) options that he or she is aware of.

This idea can be used to characterize risk-adjusted expected profits by subtracting an entrepreneur/owner's risk premium from expected profits. The risk premium demanded by a risk-averse investor rises with the variance of the expected profits. That variance can be reduced by collecting and analyzing more data, although usually not to zero.

Suppose the firm knows that its inverse demand function is $P = p(Q, Y)$ where Y is the average income of its consumers and Q is its output, as in Chapter 2. Suppose that Y is randomly distributed with mean Y^e and variance σ^2, where both Y^e and variance σ^2 have to be estimated by the firm. The variance of both estimates is assumed to fall as more data are gathered and analyzed, as implied by standard statistical analysis. The real underlying variance of consumer income, of course, is not affected by data collection analysis, but the variance in a firm's estimate of average consumer income is. Other sources of variation, of course, still exist. For example, consumer income may still vary because of business cycles.

$$U^e{}_D = \int u_{X_2}\left(\frac{P_D}{P_2}\right) + u(X_1, (\Pi - P_1 X_1)/P_2))f_D = 0 \quad \text{at } D^*$$

Notice that the first first-order condition is satisfied where $u_{X_1} - u_{X_2}(P_1/P_2) = 0$, as in the non-stochastic version of this choice problem—although in this case that choice is affected by variation in the individual's profits from producing another good or service. The second first-order condition demonstrates that data and analysis are selected to maximize expected utility, in this case by setting the marginal opportunity cost of data collection and analysis (the first term) equal to the increase in expected utility generated by additional data and analysis. That increase in expected utility is generated by reducing the variance of the estimated distribution of profits. Both first-order conditions are simultaneously satisfied at the expected utility maximizing choice.

Estimation errors are reduced through additional data and analysis, D, (albeit with diminishing marginal returns).

The risk premium of the investment of interest (here the production of goods for sales in the future) rises as the overall variance (estimation errors plus underlying randomness) increase. Thus, the risk premium for choosing output Q can be written as $R = r(D, \sigma^2)Q.$ That premium diminishes subjective realized returns relative to a risk free return. It tends to fall as more analysis, D, is undertaken—although not to zero because the true underlying risk, σ^2, is not affected by the research and estimation process.

Given this characterization of the choice setting, the expected risk adjusted profit can be written as the expected profit minus the associated risk premium:

$$\Pi^e = p(Q, Y^e)Q - c(Q, D) - r(D, \sigma^2)Q \qquad (9.5)$$

Differentiating with respect to Q and D yields two first order conditions.

$$\Pi^e{}_Q = p_Q Q + p(Q, Y^e) - c_Q - r(D, \sigma^2) = 0 \qquad (9.6a)$$

$$\Pi^e{}_D = -c_D - r_D Q = 0 \qquad (9.6b)$$

Both first-order conditions have to be satisfied simultaneously, although some insights can be obtained by looking at them separately.

Equation 9.6a implies that the greater that risk (e.g. risk premium) the lower output tends to be. A risk-averse entrepreneur will produce somewhat less than the amount that maximizes expected profits. How much less will vary with his or her degree of risk aversion and the underlying variance of the phenomena (here average consumer income, Y^e) of interest.

Equation 9.6b implies that the extent of the research to estimate the probability distribution of the risky variable (here, average consumer income, Y^e) tends to fall as marginal cost increases, c_D, and to rise the more rapidly research reduces estimation errors (which reduces the marginal risk premium). There are benefits associated with greater data collection and analysis. Estimation errors are reduced. But also costs. The time and attention focused on information gathering and analysis implies that the marginal cost of data collection and analysis is greater than zero, $c_D > 0$.

Note that the term, $r_D Q$, is the source of a large firm's demand for economists and/or for its use of external consulting firms that employ economists. It is, thus, the effect on risk adjusted profits that accounts for most of the non-academic private-sector job market for economists.

III. Ignorance

The sampling approach to imperfect information is a plausible characterization of a variety of informational problems—namely, those in which one has an initial understanding of the phenomena of interest (as with prices or demand functions) but a less perfect understanding of the actual properties that phenomena—one that can be improved by data collection and analysis. In such cases, much is already known to the relevant decision maker—the domain of possibilities and often the type of probability distribution associated with the stochastic phenomenon of interest. More information allows one to better understand the phenomenon of interest (lowest price, demand function, etc.) but does not uncover entirely new phenomena or possibilities previously unknown to the decision maker of interest (consumer, firm, scientist).

Ignorance, in contrast, occurs when at least some relevant details about a phenomenon are entirely unknown. Indeed, the entire phenomenon may be unknown.

For example, of universal ignorance include the possibility of telegraphs and telephones for long distance communication prior to 1800, low cost methods for refining aluminum before the Hall process was put into practice in 1886, or such innovations as electric lightbulbs (1879), automobiles (1885), washing machines (1908), heavier than air passenger flights (1908), radio, television, solid state electronics, personal computers, cell phones, and satellite internet—to name just a few of the surprising innovations that created unexpected new markets for new products and services. Instances of contemporary ignorance include our limited knowledge of foreign languages (many of which we do not know exist), the manner in which computers are constructed and operate, and the distribution of talents and knowledge among persons that we have never met.

In some cases, ignorance can be said to be chosen—e.g. rational ignorance. There are things that we know exist (e.g. unknown languages or fields of science) yet take no steps to reduce our ignorance (by consulting lists of languages and learning a few words in each one), or by our failure to spend even a bit of time to fill gaps in our knowledge of science or engineering (by reading books and journals or by taking courses on them).

There are other cases in which ignorance is not the consequence of a decision, but of the human condition. There are relationships and things not yet known that will be known in the future. There are many others that we are individually unaware, because our personal experience has not included them--although they are known to others. Examples, for most of us included the existence of coronaviruses (until 2019), important figures and names of their offices in other national governments than our own, tax laws that do not affect our own tax returns, how the software behind contemporary AI programs works, whether there is a national flower of Japan or the United States, and so forth.

Both categories of ignorance are relevant for economics. In the cases where we choose our level of information, more or less information will allow us to make better or worse decisions, as in the case of the Stigler search model. The natural ignorance of humanity creates opportunities for significant innovations to occur, as developed in Chapter 8. Other aspects of natural ignorance imply that mistaken decisions are likely even if we use the information we have as well as it can be (Congleton 2001).

If we all knew everything that is possible—then books such as this one, as well as universities, would be unnecessary—indeed, a complete waste of time, and they would not exist.

Ignorance, Latent Demand, and Innovation

Ignorance is relevant for the theory of demand for several reasons. One is that ignorance creates the possibility of latent demands—demands that individuals have but are unaware of, and which become actualized when a new product is brought to market and seen for the first time or learned about through the mass media or one's social networks. Such products are often Schumpeterian innovations (e.g. consequential innovations). A few entrepreneurs believed that particular latent demands existed, and they attempted to profitably elicit and serve that demand. The automobile and all its refinements constitute one such product, or perhaps rather, line of products. The smart (cell) phone is another—but there are many dozens of such products brought to market every year. Not all new products are successful, nor disruptive—which demonstrates that sufficient latent demands do not always exist for every new product imagined by an entrepreneur.

Ignorance and Knightian uncertainty makes mistakes possible—indeed likely. Nonetheless, changes in the product mix brought to market through time clearly demonstrate that many instances of latent demands have existed in the past.[9]

Modeling Latent Demand

Such demands can be easily modelled, if we drop the assumption that individuals fully understand or know their own tastes and utility functions.

---

[9] In my own case, I remember learning about a number of new products that I did not know existed when I attended college as a young man. For example, my hometown did not have any high-end stereo shops, and I learned of the possibility of high-quality reproduction of sound both from previously unknown fellow students who became friends and who had such systems and also from the local stereo shop. The latter also taught me about higher quality sound sources than I had known about. (Similarly, it was a surprise to me that not all albums were equally well recorded in those days.) Higher quality reproduction systems involved many higher quality products that jointly more faithfully reproduced the sounds from vinyl discs and CDs(phonographs, cartridges, CD-players, preamps, power amps, and speakers). Who knew? Obviously many others did or there would not be such shops, but not me. My latent demand was activated, and I began aspiring to purchase, and eventually purchasing, relatively high-end (accurate) stereo equipment. Like many college students, many of my long-standing hobbies were developed while a student. Other later latent demands revealed by innovation, for me at least, include laptops, jet travel, the internet, and cell phones.

The effect is easiest to demonstrate with a concrete three-good utility function. We'll assume that Al's true utility function is $U = A^a B^b C^c$ and that Al has W dollars to spend. The prices of the three goods are P_a, P_b, and P_c. However, initially Al does not know that good C exists and so acts as if her utility function is $U=A^aB^b$—which is what it would have been if c=0 reflected the actual marginal utility of good C.

In that case, as developed in Chapter 2, one can demand Al's demand for both goods by forming a Lagrange, taking three partial derivatives, and undertaking a bit of algebra on the system of equations that emerges:

$$\mathcal{L} = A^a B^b - \lambda(W - P^a A - P^b B) \qquad (9.7)$$

Differentiating with respect to A, B, and λ, setting the results equal to zero, and a bit of algebra (see that worked out in Chapter 2) generates the Al's two demand functions.

$$A^* = aW/\mathrm{P}^a(a+b) \qquad (9.8a)$$

$$B^* = bW/\mathrm{P}^b(a+b) \qquad (9.8b)$$

If for example, a=b, Al would spend half of her money on each good, with the actual quantities purchased varying with the price of the good or service.

Now consider Al's optimization problem after good C has been discovered and exponent c has been found to be different from zero. The Lagrange function becomes:

$$L = A^a B^b C^c - \lambda(W - P^a A - P^b B - P^c C) \qquad (9.9)$$

Differentiating with respect to A, B, and C, setting the results equal to zero and a bit of algebra now produces three demand functions—each with a similar functional form.

$$A^* = aW/\mathrm{P}^a(a+b+c) \qquad (9.10)$$

$$B^* = bW/\mathrm{P}^b(a+b+c) \qquad (9.11)$$

$$C^* = cW/\mathrm{P}^c(a+b+c) \qquad (9.12)$$

Notice that the denominator is larger than it was in the previous case, so the amounts spent on goods A and B falls, and that decrement is now spent on good C.

In this manner, the effect of activating a previously latent demand is to shift expenditures around in consumer budgets—reducing expenditures on many previously known goods and services so that some of one's wealth or income can be spent on the "new" good—at least goods that are new to Al.[10]

In extreme cases---although not ones associated with this concrete form of a utility function—expenditures on some of the preexisting goods may fall to zero.

Note that in either case, this model illustrates one way that a commercial society's impulse for creative destruction can be modeled. The successful invention and marketing of a new product upsets the pre-existing equilibrium—here through effects on demand for goods and services—generally reducing demands for pre-existing goods while increasing demand for goods that previously had only latent demand.

Rational Ignorance and Specialization

Adam Smith (1776), George Stigler (1951), and James Buchanan (2000) all have argued that much of the efficiency of market production is generated by specialization—by which they mean production by organizations and individuals that produce a narrow range of goods and services for sale in markets. That narrowness in some cases reduces the satisfaction associated with work itself, but it can substantially increase the outputs of firms and individuals relative to that which could be produced by "well rounded" organizations and individual "jacks of all trades." Hayek (1937, 1945), in turn, suggests that specialization at the level of individual workers—although motivated by market incentives—is largely

[10] Here it should be noted that reductions in spending on all other goods does not necessarily occur with different utility function than the multiplicative exponential function used here—but the overall expenditure on previously known goods has to fall if new expenditures on the good learned about are to be undertaken. Expenditures on some previously purchased goods may actually increase (e.g. complements to the new good) while most others fall.

informational in nature. It arises because of differences in the productive knowledge acquired and internalized by individuals over the course of their lifetimes (e.g. their "skill sets").

Individuals know quite different things—partly because of differences in their choices about which kinds of information to acquire through experience, reading, and education. It is such choices that allow one to say that a person has chosen to specialize in X and Y. It is a subset of these differences that create comparative and absolute advantages in production, and it is the demand for the things and services that can be produced from those skills that account for the pattern of market rewards for different kinds of market-relevant productive knowledge.

Market prices have effects on expected future income that encourage individuals to specialize. That rates of return among skill sets differ (e.g. are associated with different wages and/or job satisfaction) generate demands for the information sets associated with those skills—e.g. demands for training, experience, and education. Those rewards often induce "corner solutions" in individual demands for information—which is to say, they induce many individuals to devote essentially all of their time and attention on the subset of information relevant for their chosen speciality, rather than encouraging one to divide one's time and attention among a broad variety of types of information. They become doctors rather than lawyers, or economists rather than plumbers or accountants.

In that manner, decisions to pursue a career induce individuals (and organizations) to remain rationally ignorant about most types of information that are irrelevant for their speciality, but well-informed about the types of information that tend to increase their marginal productivity. For example, if one pursues a degree in economics, one usually will forgo learning about accounting, physics, electrical engineering, and so forth because the time and attention required to obtain the training necessary to qualify as an economist precludes such investments. Labor is not homogeneous, as often assumed in macroeconomic models.

In contemporary markets some forms of "narrowness" yield much higher salaries than a broader more general education normally does. The average doctor has a higher income than the average lawyer and the average lawyer a higher income than the average economist.[11]

A Model of Career Choice

The conscious interest in career-specific information requires the accumulation of a good deal of information over many years. For example, most students begin to think seriously about careers after they have reached high school or college. As young adults, they are no longer naturally ignorant about all topics. They all have some information about career opportunities—although their knowledge of possibilities is usually quite incomplete and based on the communities in which they lived. Obviously, one cannot choose a career that one is unaware of.[12]

Our first model of career choice (e.g. specialization) involves a decision between two careers about which a bit is known. The decision of interest is a time-allocation choice. How should one's time be divided up between acquiring skills (knowledge) of one kind or another. In the case of college education and other forms of technical education, personal expenditures may also be involved, but the time allocation problem is sufficient to illustrate many of the factors that influence decisions to specialize.

Suppose that Al is allocating time between the skill sets of two possible careers and leisure. The careers in turn are expected to produce various combinations of income and job

---

[11] That such income differences continue to exist is generated by various barriers to entry. These may include special talents or physical characteristics that are required or there may be legal barriers to entry. Once cannot be a doctor or a lawyer in the United States without a medical or law degree and certification from their respective associations.

[12] In my own case, with a blue-collar middle-class background, I had no idea what economics as a field was when I went off to college, and I took a college course in economics only after the strong recommendation of a close friend—although I still had little idea about what would be learned. Although a reasonably smart and well-informed student, I had remained naturally ignorant about economics as a field. I sampled it only because the encouragement of a good friend.

satisfaction. The exact relationship between investing in a career (studying particular subjects) and income and satisfaction are not known, but enough information is possessed that a statistical model of possible outcomes can be used for Al's choice of a speciality.

We'll first model choices when specializing is the only possibility, and then consider a choice where one can choose between complete specialization and various combinations of those specialities (as with double majors) are possible. Musicians, for example, often pursue two careers—a "day job" and a "night job" (the latter being their more enjoyable music performances).

In both cases, a forward-looking individual will imagine the expected utility associated with each of the known career paths, which is determined partly by expected income, partly by job satisfaction, and partly by the leisure sacrificed to qualify for positions in the speciality of interest. In many cases, decisions to specialize are once in a lifetime choices, but they may also be regarded as choices made with a shorter finite time horizon—a plan for the next several years or decade.

As is true of most models, the aim here is to isolate some essential features of career choice—e.g. decisions to specialize—that are likely to be factors in most such decisions, rather than all of them.

Choosing Among Specialities

Suppose there are just two specialities, A and B, that are considered by Al and that Al's planning horizon consists of two periods—a training period and a finite income-earning period of employment in his or her speciality. Assume that there is no prospect of failure during the training period as long as one invests the required time in the training. Suppose further that the distribution of income in each of the fields is bimodal, with the probability of high-income being P^A for career A and P^B for career B. The training periods differ, with T^A and T^B being the periods of study, and S_A and S_B, being the time required to master the skill set during the period in which specialized information is acquired and internalized (becomes part of one's personal knowledge). After certification, the high-income associated with speciality A earns A^H and the lower income A specialists earns A^L. The probability of being in the high-income group of occupation A is P^{AH}. Similarly, the high-income group of occupation B earns B^H and the low-income group, B^L, with the probability of being in the high-income group being P^{BH}. Each person allocates D hours each day between leisure and study or between leisure and work.

In the work period, individuals are assumed to be free to pick the length of their work years (although that is not, of course, always true), although not in the training period. Utility is assumed to be generated by Leisure (L), Job Satisfaction (J), and material comforts associated with income (Y). We'll use an intertemporal form of utility. Al's instantaneous utility is U=u(L, J, Y, t). Al's realized utility in the time period of interest is an integral of that function over the relevant time period. Argument "t" takes account of any time discounting that Al may engage in.

Both the lack of certainty about wage rates and the planning horizon affect the choice and will be taken into account by a forward-looking potential specialist. Al's planning horizon is assumed to be F years.

There are two relevant periods, the training period and the income-earning period to be taken account of. In the first period, the training period, Al has no income and no job satisfaction. In the work period, Al either earns a high or low income in his or her chosen profession, which cannot be determined beforehand. Job satisfaction increases with hours worked, but at a diminishing rate.

Given all this, Al's expected utility from the pursuit of career A during the planning period, F, can be written as:

$$U^A = \int_0^{T^A} u(D - S^A, 0, 0, t)dt + P^{AH}\int_{T^A}^{F} u(D - W^A, W^A, A^H W^A, t)dt +$$

$$(1 - P^{AH})\int_{T^A}^{F} u(D - W^A, W^A, A^L W^A, t)dt \qquad (9.13)$$

Expected utility for career B, can be written as:

$$U^B = \int_0^{T^B} u(D - S^B, 0, 0, t)dt + P^{BH} \int_{T^B}^{F} u(D - W^B, W^B, B^H W^B, t)dt +$$

$$(1 - P^{BH}) \int_{T^B}^{F} u(D - W^B, W^B, B^L W^B, t)dt \qquad (9.14)$$

Many aspects of career decisions are taken into account in this model. Nonetheless, there are just two control variables for Al. He or she can choose a speciality and decide how many hours to work in the post-training period.

The career choice in this case is determined by the higher of the two expected utility functions—which is partly a consequence of calculus (determining the ideal work rate in each career), although specialization is not a choice that calculus sheds direct light on. Al simply chooses the occupation (speciality) that generates the highest expected utility, given how hard he or she expects to work in the training and post-training periods and the income and job satisfaction anticipated for the two careers under consideration.[13]

The first-order condition that characterizes Al's post training work effort are notationally similar but differ because of the parameters of the post-training work periods, including differences in the wage rates. Al's utility maximizing work effort in the post-training period if career A is characterized by:

$$P^{AH} \int_{T^A}^{F} -u_L + u_J + u_Y(A^H)dt + (1 - P^{AH}) \int_{T^A}^{F} -u_L + u_J + u_Y(A^L)dt = 0 \qquad (9.15)$$

And that associated with career B is:

$$P^{BH} \int_{T^B}^{F} -u_L + u_J + u_Y(B^H)dt + (1 - P^{BH}) \int_{T^B}^{F} -u_L + u_J + u_Y(B^L)dt = 0 \qquad (9.16)$$

In both careers, the anticipated work effort varies with the expected present value of the marginal opportunity cost of working more (reductions in utility from diminished leisure)

[13] In a more detailed model, study time and work time might contribute to increased income either through effects on the probability of being in the high income group or through effects on marginal productivity and the wage rate paid persons in the two groups.

and the marginal returns from working (including both additional job satisfaction and material comfort).

To choose one or the other careers requires substituting the ideal work-year back into the objective function and then comparing maximal utilities over the training horizon.

The effects of the differences among careers (training period, required study time, odds of being among the high versus low-income graduates, job satisfaction, and the high and low wage rates) on expected (ideal) work effort can all be characterized using the implicit function differentiation rule. Their effects on expected total utility in the planning period can also be characterized (using the envelop theorem). However, only in cases in which Al is initially indifferent between the two career paths would small changes in any of the parameters of the choice setting affect his or her choice, given his or her planning horizon, wage rates, and odds of being in the higher of the two jobs that specialization qualifies one for.

In cases where the individual is not nearly indifferent between the two specialities, the objective functions provide many insights into factors that would affect choices. First, there is the obvious one: if career A has higher income in both cases than career B, equal or higher job satisfaction, and takes less time to acquire the skills, Al will choose career A over career B. Second, if opportunity cost of one of the careers is higher than the other (because the period of required study is longer) that career will not be chosen if the incomes and job satisfaction are the same. Longer periods of study will be undertaken only if income or job satisfaction are expected to be significantly higher than for professions requiring shorter periods of study. Third, if one is risk averse, one may choose the field with the highest low salary rather than the higher average salary, other things being equal. Fourth, in cases, in which differences in income and job satisfaction are evident (but not study periods), individual will favor one or the other based on job satisfaction if that is a major concern (e.g. has a higher marginal utility than that of income) or vice versa, if one is mainly interested in income (because the marginal utility of income is higher than that of job satisfaction).

In other cases, the magnitude of a variety of tradeoffs will influence the choice made. Note, that in none of the cases would a person choose solely on the basis of income unless job satisfaction has no effect on utility or study times are all the same.

Lastly, it bears noting that although the choice is largely an informational one—e.g. choosing to acquire one type of information rather than another, the model implies that the choice is one between two expected present value (of utility) calculations. Ignorance plays a role here insofar as career C, about which Al is completely uninformed, might have been chosen had he or she known that it existed. Learning about that career possibility after one of the two specialities was chosen, might induce Al to reconsider his or her choice.

Choosing among Degrees of Specialization

Modest extensions of the model can also be used to think about why individuals choose to specialize rather than to divide their time between the two fields under consideration. The joint choice can be modelled simply by changing the time constraint so that time is allocated between two types of work and combining the two expected utility functions after deciding which training period (major or advanced degree) is completed first. We'll assume that what might be called certification matters, and that to qualify for either specialty requires spending a given time in training—as in the first model. We'll also assume that training for speciality A takes place first, followed immediately by training for speciality B. Given the probabilistic and temporal aspects of the model, the new optimization problem looks quite complicated. Maximize:

$$U = \int_0^{T^A} u(D - S^A, 0, 0, t)dt + \int_{T^A}^{T^B+T^A} u(D - S^B, 0, 0, t)dt +$$

$$P^{AH} \int_{T^B+T^A}^{F} u(D - W^A - W^B, W^A, A^H W^A, t)dt +$$

$$(1 - P^{AH}) \int_{T^B+T^A}^{F} u(L - W^A - W^B, W^A, A^L W^A, t)dt +$$

$$P^{BH}\int_{T^B+T^A}^{F} u(D - W^A - W^B, W^B, B^H W^B, t)dt +$$

$$(1 - P^{BH})\int_{T^B}^{F} u(D - W^A - W^B, W^B, B^L W^B, t)dt \qquad (9.17)$$

There are two first order conditions—one for W^A and one for W^B, rather than one, because a dual career is contemplated—and certification mandates a particular level of study. Mathematically, the first order conditions closely resemble those of in the previous section except that the starting time of the employment period which is now $T^B + T^A$, rather than T^B or T^A because following both careers requires that Al meet the qualifications for each. Substituting the optimal values for W^A and W^B into equation 9.17 allows the expected utility of the two-career future to be characterized.

The choice about specializing requires comparing the expected utility associated with the idealized two-career future to two idealized specialized career paths of the previous section. Specialization is the rational choice whenever one of the single-career futures is expected to generate more satisfaction (utility) than the combined career.

Normally, in neoclassical economic models, "corner solutions" are neglected and "interior solutions" are the main focus, as with the focus of consumer choice models on the middle points of their budget constraints rather than their end points. That is not always true of consumer or production models for reasons that we have noted in passing in earlier chapters (the most common quantity purchased by a consumer in a grocery store is zero), but most people spend their incomes on many different goods and so interior solutions are the ones that characterize the most relevant aspects of the spending decision—e.g. the variety of goods and services actually purchased.

However, in the case of specialization, corner solutions are commonplace. Most college students earn only a single degree, and most people after college focus on one specialty (albeit not always one in which their college training is most important). There are a huge number of occupations that persons could have devoted a little time to, but do not. Graduate educational programs while at a university are even more closely linked to job-market specialities than undergraduate training. So, with respect to specialization, corner solutions are important (e.g. solutions in which either W^A or $W^B = 0$.

This implies that the combination of job satisfaction and wages associated with mastering a speciality are generally anticipated to be well above those associated with double or triple careers for most students.

If the expected utility functions after all the substitutions are strictly concave and there are no certification requirements for employment, then the usual geometry of optimization would imply that being a jack-of-all-trades would be the best strategy for most individuals, including our representative individual, Al.

To induce specialization, there must be some non-concave aspects to the decision problem. This tends to be generated by differences in wage rates and/or job satisfaction that make incomes and/or job satisfaction much higher for specialists than non-specialists. Rational ignorance can be a profitable strategy.

IV. Prices as Information

The previous models illustrate why prices affect the knowledge base of individuals. Variation in prices can induce individuals to invest time shopping. Variation in wage rates can induce individuals to acquire in particular types of information. In these cases, and many others, expected prices inform choices. When accurate, price forecasts provide useful information about the opportunities and tradeoffs associated with activities in all the markets that individuals are familiar with.

When markets are at least roughly in equilibrium, prices provide useful information about the relative importance of products (their marginal utilities) and occupations (their value added) as measured by willingness to pay for them, while simultaneously characterizing the boundaries of one's opportunity set.

The latter affects choices for reasons developed throughout this textbook and by the majority of the microeconomics literature. The closer markets are to equilibrium, the more reliable and stable that information tends to be. And, the closer markets are to equilibrium, the more productive is the distribution of skill-oriented information in the economy—given the distribution of talent and legal barriers to entry among the professions.

References

Akerlof, G. A. (1970). Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, *84*(3), 488–500.

Buchanan, J. M., & Yoon, Y. J. (2000). A Smithean perspective on increasing returns. *Journal of the History of Economic Thought*, *22*(1), 43–48.

Cai, J., Liu, Y., Qian, Y., & Yu, M. (2015). Information asymmetry and corporate governance. *Quarterly Journal of Finance*, *5*(03), 1550014.

Congleton, R. D. (2001). Rational ignorance, rational voter expectations, and public policy: A discrete informational foundation for fiscal illusion. *Public Choice* , *107* (1), 35-64.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory*, *3*(2), 156–168.

Downs, A. (1957). *An economic theory of democracy*. Harper & Row.

Mortensen, D. T., & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, *61*(3), 397–415.

Rothschild, M., & Stiglitz, J. E. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, *90*(4), 629–649.

Schumpeter, J. A. (1942). *Capitalism, socialism and democracy*. Harper & Brothers.

Smith, A. (1776). *An inquiry into the nature and causes of the wealth of nations*. W. Strahan and T. Cadell.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, *87*(3), 355–374.

Stigler, G. J. (1951). The division of labor is limited by the extent of the market. *Journal of Political Economy*, *59*(3), 185–193.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, *69*(3), 213–225.

Von Hayek, F. A. (1937). Economics and knowledge. *Economica*, *4*(13), 33–54.

Von Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review*, *35*(4), 519–530.