Chapter 9: The Economics of Information

I. Information and Knowledge

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 allowed us to use expected utility and expected profit models to analyze 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 were wrong rather than rooted in observation, or simply because they were unlucky in what emerged from a well-understood stochastic setting—as when 00 turns up on a roulette wheel.

It is important to realize that in nearly all of the previous choice settings analyzed, the participants 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, but not the specific sequence of outcomes that will be generated by such processes. Similarly, if a process is truly chaotic, little systematic can be known about it beyond the fact that it falls into the category of Knightian uncertainty.

Chapter 8 was the first chapter in which investments partially determined what was known, although the knowledge-increasing aspects of investments in R&D were not emphasized. This chapter brings the extent of knowledge further into the analysis. It explores the demand and supply of various types of information and a subset of the implications that knowledge limitations and learning have for markets. Several other implications are worked out in the remaining chapters.

Knowledge as Internalized Information

Learning can be regarded as a form of household production. It takes place in the minds of individuals one at a time. Transforming information into knowledge is a time-consuming process. Because there are more types of knowledge than time available for accumulating them, all of us are less informed than we wish to be. A college student who majors in a particular subject knows his or her area of specialty much better—more completely—than fields to which he or she has given little or no attention.

Although the final steps that transform information into knowledge (or beliefs) are the product of an individual's own analysis and synthesis, the information that we use comes in many forms. There is the raw data of our senses, and there is "preprocessed" information provided by our friends, family, teachers, and publications of various kinds. That information is easier to transform into knowledge than raw data we have individually collected and analyzed, because it distills conclusions reached by their teachers, and by their teachers' teachers, going back past the dawn of history. Of course, not all of what is passed on to us is true, nor are our own conclusions from what we observe and learn from others—but preprocessed information makes learning far easier than it would have been without it. And much of it is sufficiently useful and accurate that it improves our awareness of opportunities and methods for ranking alternatives.

At the moment our choices are made, it is the information already acquired that provides the basis for our understanding of "what is," "what can be altered by our actions," and many of our ideas about how to rank such possible changes.

This chapter largely focuses on choice settings in which individual stocks of knowledge are sufficiently similar that they can be regarded as identical to simplify our models and discussion. The implications of various asymmetries in knowledge were taken up in the previous chapter, where entrepreneurs were assumed to know things that others did not. Others are taken up in the chapters on game theory (Chapter 13), crime (Chapter 14) and norms (Chapter 16) and in a few other places in the other chapters that follow.

In some cases, additional information helps us to improve our understanding of phenomena that we are already familiar with. 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 accurate estimates of parameters of relevant functions reduce errors in our solutions to the many optimization problems that we face.[1]

In other cases, new information may reveal possibilities or causal chains that we had not previously imagined or taken into account. At some point in our lives we learn that the length of the sunlit portion of the day varies systematically during the year—falling from the summer solstice until the winter solstice and rising again until the summer solstice is reached once again. We may learn that stock markets exist, which provide an alternative to savings accounts. Or we may invent entirely new ideas, as Aristotle, Newton, Smith, Einstein, and Arrow did in their lives.

In some cases, such breakthroughs may lead to new products and new markets that induce significant alterations in our pattern of life. For example, integrated circuits emerged in the late 1950s out of the laboratories of Texas Instruments and Fairchild Semiconductor. As their production costs fell and quality improved, a long series of products based on such

[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 probability of x (here, rain), f is the conditional probability function, and z is the conditioning variable. The 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.

miniaturized circuits were introduced. Examples include radically new products such as handheld calculators, personal computers, computer games, cell phones, and AI-bots. All these induced significant changes in contemporary lifestyles as we adapted our lives to take greater advantage of such products. Few people have landlines at home these days, a choice that was almost unimaginable at the beginning of the twenty-first century.

In addition, there are cases in which a better understanding of causality is used to improve our understanding of familiar phenomena. 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 adding previously unknown causal factors that influence patterns of weather.[2]

Living involves a good deal of learning. This is as true of economic life as it is of most other human activities.

II. Unique Features of the Demand and Supply of Information

The usefulness of much accumulated knowledge—and of other knowledge that we may imagine—creates a demand for information. As true of the demand for ordinary goods and services, the demand for information implies a willingness to pay for information (with time, attention, and money). And as true of the supply of ordinary goods and services, the existence of a demand for particular types of information implies that there are profits that can be realized by organizations that can supply the information demanded.

Thus, supply and demand operate in information markets as well as in markets for potato and computer chips. It is largely for that reason that the services and products such as universities, books, and other fee-for-service sources of information exist. Universities, for example, would not exist without students who are willing to pay for the preprocessed information provided by their professors, libraries, and online resources.

[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.

It bears keeping in mind, however, that information is not a single thing any more than goods and services are. It is at least as varied as more tangible goods and services are.

However, in contrast to ordinary goods and services, a good deal of information is freely available both in its raw "data" sense and in more or less "preprocessed" forms. There are many sources of information that are free, which is to say, available without paying money for them. Much of that information has been preprocessed by others to make it more easily transformed into knowledge. For example, there are millions of books in public libraries and on websites.

We may be said to be satiated in our demands for such sources of information because we "consume" only a small proper subsample of free sources of information. However, satiation does not exhaust our demand for information.

Preprocessed Information as an Intermediate Good

Although we are in a sense satiated, we retain a demand for additional information. This is partly because transforming information into knowledge requires the investment of time and attention by individual learners. It is also partly because new information is often very useful. We have only 24 hours in every day. That ultimate constraint, together with the usefulness of so many types of information, implies that we do not have sufficient time and attention to read and master everything that is known. Even though many free sources of information exist, the time cost of transforming information into knowledge is non-trivial.

Individuals are thus willing to pay for curated information and also for information that is even more preprocessed than the books in public libraries and online resources. They are also willing to pay for informed advice in cases in which the required period of study to master a subject of interest is too long to make becoming informed oneself a sensible use of one's time. Thus, there is a demand for "experts"—such as doctors, lawyers, accountants, and economists—who have mastered particular subsets of information, among many others.

Such curated and preprocessed sources of information reduce the time and attention necessary for turning information into knowledge, and knowledge into plans and actions that advance one's other interests in health and wealth. And they do so without substantially reducing the supply of information available for others. Although there are only a finite number of seats in the classrooms of universities and appointment times for meeting with experts, the universities, experts, and books continue to exist even after we have been informed by them. Admittedly, like other capital goods, their useful lives are somewhat shortened by our use.

It is this subset of knowledge demands that provides the basis for information markets.

Some of that information is valued in its own right—i.e., it is a direct source of amusement or satisfaction. Other types of information are analogous to the intermediate goods of input markets. They help advance other interests for which more accurate or complete knowledge allows us to achieve better outcomes.

This chapter focuses on a subset of the full range of information types. It analyzes choice settings in which information has relatively clear effects on market-relevant decisions.

As true of many of the goods that we considered in Chapter 7, the quality and usefulness of information is rarely perfectly clear before we acquire it. We don't know what we will eventually know until we acquire the data and internalize it. Similarly, the authors of books such as this one rarely know exactly what will be included in their books until the final drafts of their manuscripts are completed. The author's internal processing of "data" already possessed, and other information that is discovered in the process of writing, implies that an author normally has more knowledge when they finish than when they begin.

Hayek (1945) 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 information is acquired and how it affects consumer and firm decisions is consequently necessary if we are to understand how extended market networks 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-alls," while comforting, tend to make such decision makers more error-prone than they could be. They also tend to make them more confident than they should be, and thus more likely to be risk

Some informational choices resemble choices in settings of risk. Others are imbued with Knightian uncertainty.

III. Information and the Distribution of Market Prices

Stigler's "Search" Model of the Demand for Information

We begin with one of the simplest and most general models of the demand for market-relevant information and its associated effects on markets. Stigler introduced a model of the demand for some types of information that is referred to as the "search" model of the demand for information. That model was developed to account for the dispersion of prices in real markets.[4] Outside of organized stock and futures markets, there is usually more than one price for goods sold in markets at every moment in time. Prices do not converge to a single price as implied by the core neoclassical model of price theory, but tend to cluster around some average price. (That average 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 and waiting times at different merchants—but the informational explanation is important and emerges from a straightforward extension of the rational choice model.

According to Stigler's theory, the prices that one pays for goods are partly a matter of choice rather than entirely determined by market forces beyond an individual's control. Given perfect information, the logic of "perfect competition" and its law of one price is clear—but what if information about prices is not freely available to all consumers?[5]

---

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 overconfident entrepreneurs.

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

[5] 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. Once the lowest sustainable selling price emerges, a price lower than

Stigler (1961) argues that in such cases consumers sample prices and purchase goods at the lowest price in their sample. They search for the lowest price. However, searching is a costly activity, and thus different consumers choose different sample sizes, and pay different prices, because the lowest prices in their samples differ.

Stigler goes on to argue that the easier it is to search for prices (i.e., the less costly sampling is) the narrower is the range of prices that tends to emerge in markets for a homogeneous product. (He also provides a bit of evidence supporting that hypothesis.)

Searching (sampling) limits the range of prices paid for goods and services because every consumer purchases goods from the merchant with the lowest price in their samples. Thus, firms with well-above-average prices sell relatively little of the good or service of interest—i.e., only to consumers who do very little or no searching.

If most people do a bit of shopping, the result is reduced sales and profits for the higher-priced merchants, and greater sales and profits for lower-priced firms. This tends to reduce the variance of the distribution of prices—as the high-priced firms exit from the industry or lower their prices. Thus, Stigler's theory accounts 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 do.

Modeling an Individual's Price-Sampling Decision

Suppose that Al knows his or 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, he or she knows from statistics that as more sampling is undertaken, the lowest price in his or her sample tends to fall. In this case, sampling might involve driving to numerous shops or visiting various websites at which a specific type of good can be purchased. In both cases, sampling has an

that price would attract all the buyers in the market, but firms charging lower than sustainable prices 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 for identical products than the lowest price available.

opportunity cost—namely the value of the time spent searching plus any other expenses such as fuel for one's vehicle or electricity for one's computer or cell phone.

The simplest way to model this choice is as an effort to maximize Al's expected net benefits from purchasing the good of interest. Al's expected net benefits (N^e) are the total benefit from purchases of the good (B=b(Q, Y)) less the expected price of the good purchased times the quantity purchased and search costs, $C^e = P^e Q + c(S)$, with $P^e = p(S)$. This can be written as:

$$N^e = \mathrm{B} - \mathrm{C}^e = \mathrm{b}(\mathrm{Q}, \mathrm{Y}) - p(S)Q - c(S) \tag{9.1}$$

where S is the number of shops or web-sources sampled, b is Al's total benefit function and c(S) is the opportunity cost of the searches undertaken. The expected price, p(S), is a function that declines with sample size at a decreasing rate ($P_S < 0$ and $P_{SS} > 0$). The sampling cost function, $c(S)$, increases with the sample size at an increasing rate ($C_S > 0$ and $C_{SS} > 0$) as increasingly valued bits of time are devoted to search. Differentiating equation 9.1 with respect to Q and S characterizes Al's ideal purchases and optimal sample size for his or her searching efforts—i.e., efforts to learn what the lowest available price is for the good of interest.

There are two first order conditions, one of which is similar to that used in Chapter 2 to characterize a demand function and the other of which takes account of the productivity of sampling prices at different shops.

$$B_Q - P = 0 \tag{9.2a}$$

$$-P_S[Q] - C_S = 0 \tag{9.2b}$$

The first term in equation 9.2a is the marginal benefit from consuming the good and the second is its marginal cost. The first term in equation 9.2b is the expected marginal benefit (total cost reduction) associated with searching more and the second is the marginal cost of search. Al's expected marginal benefit from searching varies with the extent to which prices tend to decline with search (P_S) and with the quantity that will be purchased (Q). The more units Al expects to purchase and the broader the distribution of prices the higher is the marginal expected benefit from searching.

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

An alternative modeling approach is to first characterize Al's purchase of the good at a given price as in Chapter 2, which yields a first order condition identical to equation 9.2a. This implies that $Q^* = q(S, Y)$. Substituting that ideal purchase into Al's expected net benefit function yields: $N^* = \mathrm{b}(Q^*, \mathrm{Y}) - p(S)Q^* - c(S)$. Differentiating that upper envelope of the net benefit function with respect to S and applying the envelope theorem yields:

$$-P_S[Q^*] - C_S = 0 \equiv \mathrm{H} \tag{9.3}$$

The first term, $-P_S[Q^*]$, is again the expected marginal benefit of searching and the second its marginal cost—now evaluated at ideal purchases of the good of interest. (Equation 9.3 can be interpreted as the first order condition that characterizes Al's expected cost-minimizing expenditure for Al's ideal quantity of the good of interest.) This first order condition is illustrated below. (Recall that $MB = -P_S[Q^*] > 0$.)

Figure 9.1: Net Benefit Maximizing Search

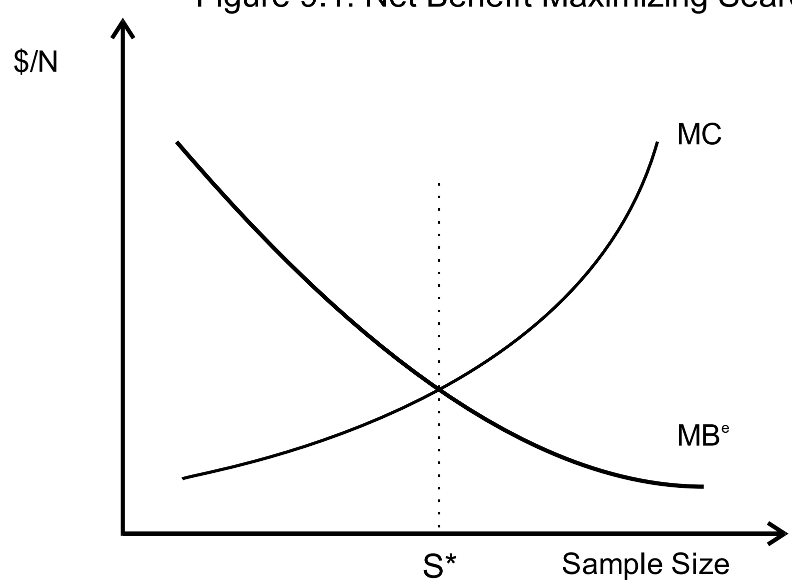

Searching for the lowest price puts downward pressure on the firms selling the good of interest at above average prices. It is the searching efforts of consumers that induce prices to cluster within a relatively small band. However, prices do not collapse to a single price because the marginal cost of searching is greater than zero and differs among consumers. Thus, some consumers pay more than the average price, because searching is a costly activity and they expect the cost of further search to be greater than its benefits.

Comparative statics of the demand for this form of information 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,

$$S_Y^* = \frac{H_Y}{-H_S} = \frac{-P_S[Q_Y^*]}{-({N^*}_{SS})} > 0 \tag{9.4}$$

Derived in this way, the numerator has an unambiguous sign. If the good of interest is a normal or superior good, then $Q_Y^* > 0$, and since $-P_S > 0$, the numerator is greater than zero. If the net benefit function is strictly concave, then so is any subsection of it, e.g. the upper envelope in this case, and thus the second derivative of expected net benefits is negative and $-{N^*}_{SS} > 0$. Intuitively, as Al's income increases, more of the good tends to be purchased at every price, and therefore searching is more worthwhile.

The model implicitly assumes that an increase in income does not raise the opportunity cost of searching. If it did, the marginal cost of search would rise with income, the term $-C_{SY}$ would be included in the numerator, and the overall sign of equation 9.4 would be ambiguous without further assumptions.

Stigler's model demonstrates how a rational choice model can be used to improve our understanding of market prices. Even in equilibrium, prices may vary a bit—how much varies with search costs and the variation in prices. Additional information allows better (less error-prone or more fruitful) decisions to be made, and actions taken. This together with variation in prices encourages price searching—which in turn tends to reduce the variance of market prices, although not to zero.

Note that Stigler's modest extension of the neoclassical model explains why observed distributions of prices do not converge to a single price. According to his theory, this is unlikely to happen unless knowledge of prices is costless.

The model also implies that informational differences are likely to exist among consumers and such differences are likely to be economically significant. Those who search less, on average, pay higher prices for the things that they purchase than those who search more.

The same logic directly applies to firms with respect to their input prices. Firms have an interest in purchasing their inputs at the lowest price possible, and so they engage in a search process similar to that described above, with their demand for those inputs based on the marginal revenue product of those factors of production.

IV. The Value of Better Estimates for Firms

Related but somewhat different informational problems are associated with choosing the price at which firms sell their goods and services. The highest price at which a price-making firm's output can be sold is determined by its demand function, as developed in Chapters 3 and 5.

As a consequence, price-setting firms profit from better estimates of their current and future demand functions. More accurate estimates of future demand functions enable firms to more accurately 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 demand function. More data and more sophisticated models are required for more accurate estimates and these improvements are costly. Thus, more accurate estimates are more costly than less accurate ones.

The Profits and Costs of More Precisely Estimated Demand Functions

The deterministic models of firm decision making developed in Chapter 3 ignored randomness and uncertainty. Demand functions, in principle, may be mechanically causal (as in the models of Chapter 2), or stochastically causal (as in the models of Chapter 7). They may also be affected by uncertainty, insofar as both firms and their consumers make decisions in settings in which there are chaotic or poorly understood phenomena.[6]

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

A Model of the Optimal Precision of a Firm's Estimated Demand Functions

One way to incorporate the effect of 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 they are 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 actual average income of its consumers and Q is its output, as in Chapter 2. Suppose that consumer income, 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. To simplify, we'll assume that unbiased methods of data collection and analysis are used, so that Y^e is fairly accurately estimated, but that the variance of the estimator employed is greater than the actual variance in income,

---

[6] 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.

$s^2 > \sigma^2$. The variance of estimates of average consumer income, s^2, falls 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 or analysis. For example, consumer income may still vary because of business cycles. Estimation errors are reduced through additional data and analysis, D, albeit with diminishing marginal returns and a lower bound (the true σ^2 of consumer income).

The risk premium for the investment of interest (here the production of goods for sale in the future) rises as the overall variance (estimation errors plus underlying randomness) increases. 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. s^2 tends to fall as more data 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 the risk (i.e., the 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).

As in the model of a consumer's ideal investment in search, we can gain some additional insights by using the envelope theorem and undertaking a two-stage optimization process. First, we can characterize Q* for given levels of data analysis by differentiating equation 9.5 with respect to Q and setting the result equal to zero (which will notationally be the same as equation 9.6a, although it will have a somewhat different meaning). We then use that first order condition to characterize Q* as a function of other parameters of the choice setting, $Q^* = q(D, Y^e, \sigma^2)$. Substituting that relationship into the profit function (Equation 9.5) to form the upper envelope of the profit function with respect to output, and then differentiating with respect to D and applying the envelope theorem, yields:

$$\Pi_D^* = -c_D - r_D Q^* = 0 \equiv \mathrm{H} \qquad (9.7)$$

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 the effect on risk-adjusted profits that accounts for most of the non-academic private-sector job market for economists.

The effect of an increase in average consumer income is similar to that of the search model:

$$D_Y^* = \frac{H_Y}{-H_D} = \frac{-r_D[Q_Y^*]}{-(\Pi^*{}_{DD})} > 0 \qquad (9.8)$$

An increase in average consumer income induces firms to engage in additional data analysis and more sophisticated methods of estimation (or greater expenditures on consultants who can undertake such estimates).[7]

[7] Although a profit-maximizing objective function is used above, a similar analysis can be undertaken using 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

V. Ignorance

The sampling approach to imperfect information is a plausible characterization of a particular class of informational problems—namely, those in which one has a clear understanding of the phenomena of interest (as with prices or demand functions), but a less perfect understanding of the actual properties of those phenomena. In such cases, much is initially 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 (more accurate estimates of the lowest price, demand functions, 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.

Examples of universal ignorance include the possibility of telegraphs and telephones for long-distance communication prior to 1800, low-cost methods for refining aluminum before

---

ideas from Chapters 2 and 7 allows the expected utility maximizing choice for a two-good utility function with uncertain income (here Π) to be written as $U^e = \int u(X_1, (\Pi - P_1X_1 - P_D D)/P_2) f(\Pi|D) d\Pi$. This formulation uses the expected-value formula for a variable with a continuous frequency distribution. Π is the individual owner's or entrepreneur's 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.

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

$$U^e{}_D = \int [-u_{X_2}\left(\frac{P_D}{P_2}\right) f(\Pi|D) + u(X_1, (\Pi - P_1X_1 - P_D D)/P_2) f_D] d\Pi = 0 \text{ at } D^*$$

Notice that the first first-order condition is satisfied where the expected value of $u_{X_1} - u_{X_2}(P_1/P_2)$ equals zero, as in the non-stochastic version of this choice problem—although in this case that choice is affected by variation in the firm owner's profits from producing another good or service.

The second first-order condition demonstrates that data and analysis are also 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.

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 personal 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—i.e., 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 learning about them. We fail to consult lists of languages and learn a few words in each one. We don't bother to fill gaps in our knowledge of science or engineering by reading a few books and journal articles on each of them, or by taking courses in such fields.

There are also cases in which personal ignorance is not the direct consequence of decisions, 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 of, because our personal experience has not included them—although they are known to others. Examples, for most of us, include the existence of coronaviruses (until 2019), important persons and their titles in national governments, 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. However, it is the natural ignorance of humanity that 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. We would all be true know-it-alls in that case—rather than pretenders to such claims.

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 (i.e., consequential innovations). A few entrepreneurs believe that particular latent demands exist, and they attempt to profitably elicit and serve those demands. The automobile is one such product, along with the long sequence of minor innovations that refined it over the decades that followed. The smart (cell) phone is another.

However, profitable latent demands do not always exist for the new products imagined by entrepreneurs. Nonetheless, successes are sufficiently numerous that most products are improved in multiple ways through time (albeit perhaps not as rapidly as their marketing teams suggest).

Ignorance, Knightian uncertainty, and Knightian entrepreneurship make mistakes possible—indeed likely. Many innovations fail to elicit sufficient latent demand to be profitable. New firms very often fail. Nonetheless, changes in the product mix brought to market through time clearly demonstrate that latent demands have existed in the past and likely continue to do so in the present.[8]

---

[8] 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

Modeling Latent Demand

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

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 his or her utility function is $U = A^a B^b$—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 derive Al's demand for both goods by forming a Lagrangian, taking three partial derivatives, and undertaking a bit of algebra on the system of equations that emerges:

$$\mathcal{L} = A^a B^b - \lambda(W - P_a A - P_b B) \qquad (9.9)$$

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 Al's two demand functions.

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

$$B^* = bW/[\mathrm{P}_b(a+b)] \qquad (9.10b)$$

If, for example, a = b, Al would spend half of his or 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:

$$\mathcal{L} = A^a B^b C^c - \lambda(W - P_a A - P_b B - P_c C) \qquad (9.11)$$

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

---

college students, I developed many of my long-standing hobbies while in school. Other subsequent latent demands revealed for me include laptops, jet travel, the internet, dancing, and cell phones.

$$A^* = aW/[\mathrm{P}_a(a+b+c)] \tag{9.12}$$

$$B^* = bW/[\mathrm{P}_b(a+b+c)] \tag{9.13}$$

$$C^* = cW/[\mathrm{P}_c(a+b+c)] \tag{9.14}$$

Notice that the denominator is larger than it was in the previous case, so the amounts spent on goods A and B fall, 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.[9]

In extreme cases—although not ones associated with this specific 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 creative destruction can be modeled. The successful invention and marketing of a new product upsets the preexisting equilibrium—here because of new demands for latent goods elicited by innovations. This, in turn, reduces demands for at least a subset of preexisting goods in order to actualize demands for goods that previously had been only latent.

Rational Ignorance and Specialization

Adam Smith (1776), George Stigler (1951), and Buchanan and Yoon (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

[9] Here it should be noted that reductions in spending on all other goods do not necessarily occur with utility functions other 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.

and individuals relative to that which could be produced by "well-rounded" organizations and individual "jacks of all trades." Hayek (1937, 1945) and Machlup (1981/2014), in turn, note that specialization at the level of individual workers—although motivated by market incentives—is largely informational in nature. It arises because of differences in the productive knowledge acquired and internalized by individuals over the course of their lifetimes (i.e., their knowledge and skill sets).

Individuals know quite different things—partly because of differences in their choices about the kinds of information acquired through experience, reading, and education. It is such choices that allow one to say that a person has chosen to specialize in X rather than Y. It is a subset of these differences that creates comparative and absolute advantages in production, and it is the demand for the things and services that can be produced from those skills that accounts 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 (i.e., they are associated with different wages and/or job satisfaction) generates demands for the information sets required to internalize the knowledge behind the skills employers demand—which is to say, it creates individual 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 to the subset of information relevant for their chosen specialty, rather than encouraging them to divide their time and attention among a broad variety of types of information. They become doctors rather than lawyers, or economists rather than plumbers, or accountants rather than managers.

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 specialty, 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. As a consequence of such choices, 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.[10]

A Model of Career Choice

A 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 families and communities in which they grew up. Obviously, one cannot choose a career that one is unaware of.[11]

Our first model of career choice (i.e., specialization) involves a decision between two careers about which quite 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

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

[11] In my own case, with a blue-collar middle-class background, I had no idea what economics as a field entailed when I went off to college. 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—a field that turned out to hold great interest for me and occupied much of my future.

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

The first model characterizes choices when specializing is the only possibility. The second characterizes choice when one can choose between complete specialization and various combinations of those specialties (as with double majors). 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 specialty 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 or two.

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

Choosing among Specialties

Suppose there are just two specialties, 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 specialty. 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^{AH} for career A and P^{BH} for career B. The training periods differ, with T^A and T^B being the periods of study necessary to be certified, and S^A and S^B being the time actually invested in mastering the skill set during the period of training. After certification, the high-income group in specialty A earns wage rate A^H and the lower-income A specialists earn wage rate A^L. Similarly, the high-income group of occupation B earns wage rate B^H and the low-income group, wage rate B^L. 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 the 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. Since the probability of earning a high wage, A^H, is P^{AH}, the probability of earning a low wage, A^L, is $1 - P^{AH}$.

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 (from $t = 0$ until $t = T^A$). In the work period (from $t = T^A$ until $t = F$), Al earns either a high or a low income in his or her chosen profession, which cannot be determined beforehand. Job satisfaction increases with hours worked, but at a diminishing rate. Differences in job satisfaction are represented by parameters J^A and J^B, so that job satisfaction in career A is $J^A W^A$ and in career B is $J^B W^B$.

In the first period Al chooses the amount of time to devote to studying, S^A, and in the second period Al chooses the amount of time to devote to work, W^A. Given all this, Al's expected utility from the pursuit of career A during the planning horizon, F, can be written as:

$$U^A = \int_0^{T^A} u(D - S^{Amin}, 0, 0, t)dt + P^{AH} \int_{T^A}^{F} u(D - W^A, J^A W^A, A^H W^A, t)dt + (1 - P^{AH}) \int_{T^A}^{F} u(D - W^A, J^A W^A, A^L W^A, t)dt \qquad (9.15)$$

Similarly, expected utility for career B can be written as:

$$U^B = \int_0^{T^B} u(D - S^{Bmin}, 0, 0, t)dt + P^{BH} \int_{T^B}^{F} u(D - W^B, J^B W^B, B^H W^B, t)dt + (1 - P^{BH}) \int_{T^B}^{F} u(D - W^B, J^B W^B, B^L W^B, t)dt \qquad (9.16)$$

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

The career choice in this case is determined by the higher of the maximums of the two expected utility functions. Al simply chooses the profession (specialty) that generates the highest expected utility, given how hard he or she expects to study in the training and work in the post-training periods and the future income and job satisfaction anticipated for the two careers under consideration.

During the training period, Al's studying effort may largely be prescribed by the requirements of certification. Insofar as studying is assumed not to be pleasurable, Al will invest in the minimum level required for certification, S^{Amin} or S^{Bmin} (i.e., the minimum necessary to pass all required examinations). In such cases, Al can be thought of as having only two "real" choices: his or her choice of major and how hard he or she plans to work after certification.[12]

The first-order conditions that characterize 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 career A is characterized by:

$$P^{AH} \int_{T^A}^{F} [-u_L + J^A u_J + u_Y A^H] dt + (1 - P^{AH}) \int_{T^A}^{F} [-u_L + J^A u_J + u_Y A^L] dt = 0 \qquad (9.17)$$

And that associated with career B is:

$$P^{BH} \int_{T^B}^{F} [-u_L + J^B u_J + u_Y B^H] dt + (1 - P^{BH}) \int_{T^B}^{F} [-u_L + J^B u_J + u_Y B^L] dt = 0 \qquad (9.18)$$

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

To choose one career or the other requires substituting the ideal daily work effort and the minimum effort level and time necessary to complete the training programs for each career path into Al's objective function, and then comparing their maximal utilities over the planning 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 envelope theorem).

[12] 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 to persons in the two groups. Also, it bears noting, if one of the two courses of study is regarded to be more enjoyable than the other, this will affect the perceived relative cost of the two career paths.

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 specialties, the objective functions provide many insights into factors that would affect choices. First, there is the obvious one: if career A has higher high and low wages 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 the 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 is 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 higher low salary rather than the higher average salary, other things being equal.

Fourth, in cases in which differences in job satisfaction and income are evident (but not in study periods), an individual will favor one or the other based on job satisfaction if that is a major concern (i.e., it has a higher marginal utility than income) or vice versa, if one is mainly interested in income (because the marginal utility of income is much greater than that associated with 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 and study times are all the same.

Lastly, it bears noting that although the choice is largely an informational one—i.e., 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 specialties was chosen might induce Al to reconsider his or her choice, and return to school.

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 modeled simply by changing the time constraint so that time is allocated among two types of work and two types of study, and then combining the two expected utility functions after deciding which training period (major or advanced degree) is completed first if they cannot be completed simultaneously. 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 specialty A takes place first, followed immediately by training for specialty B, and that the high- or low-income outcomes in the two fields are independent. Given the probabilistic and temporal aspects of the model, the new optimization problem looks quite complicated. Maximize:

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

$$P^{AH}P^{BH}\int_{T^B+T^A}^{F} u(D - W^A - W^B, J^AW^A + J^BW^B, A^HW^A + B^HW^B, t)dt +$$

$$(1 - P^{AH})P^{BH}\int_{T^B+T^A}^{F} u(D - W^A - W^B, J^AW^A + J^BW^B, A^LW^A + B^HW^B, t)dt +$$

$$(1 - P^{BH})P^{AH}\int_{T^B+T^A}^{F} u(D - W^A - W^B, J^AW^A + J^BW^B, A^HW^A + B^LW^B, t)dt +$$

$$(1 - P^{AH})(1 - P^{BH}) \int_{T^B+T^A}^{F} u(D - W^A - W^B, J^AW^A + J^BW^B, A^LW^A + B^LW^B, t)dt \quad (9.19)$$

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, which again is assumed to be unentertaining and so done at the minimum level that assures certification (e.g. graduation from the program of interest).

Mathematically, the first order conditions closely resemble those in the previous section except that the starting time of the employment period 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 9.19 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 the 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 best feasible 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. Such interior solutions are not always optimal for consumers or firms 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). Nonetheless, 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—i.e., 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 devote a little time to but do not. Graduate educational programs at a university are even more closely linked to job-market specialties than undergraduate training. So, with respect to specialization, corner solutions are important (i.e., choosing one of the single-career paths characterized by equations 9.15 and 9.16).

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

If the expected utility function in equation 9.19 is strictly concave in W^A and W^B, the usual geometry and mathematics of optimization imply that, after both certifications have been obtained, dividing one's work time between the two fields would be the best strategy for most individuals, including our representative individual, Al. That conclusion, however, neglects the time required to qualify for a second career.

A dual career requires completing both training programs, so Al's income-earning period begins at $T^A + T^B$ rather than at T^A or T^B. That additional training time is a fixed cost: forgone leisure, income, and job satisfaction that must be incurred before even a single hour can be devoted to the second field. Such fixed costs make the choice non-convex. In the case modelled, the dual-career path is not simply a blend of the two specialized paths, and a single-career "corner" can be optimal even when utility is strictly concave in hours worked. The longer the required training periods and the shorter the planning horizon, F, the more likely it is that one of the specialized careers dominates. Certification requirements thus provide a reason for specialization that does not depend on differences in wage rates, talent, or job satisfaction.

Specialization is further encouraged by significant differences in wage rates, talent, and/or job satisfaction that make incomes and/or job satisfaction much higher for specific specializations than for non-specialists. Such wage differences would arise if specialists were far more productive than jacks-of-all-trades in the positions that firms hire individuals to fill. Rational ignorance (specialization) in these cases is a profitable strategy.

VI. Conclusions: Information and the Extent of Markets

This chapter has explored a subset of the many ways in which the extent and distribution of information affect the extent and nature of markets. First, it has demonstrated that consumer information affects the distribution of prices observed. The more price information consumers assemble, the narrower the range of prices for particular goods and services, and the lower on average prices tend to be. Since demand curves slope downward, assembling information (search) tends to increase overall market sales beyond what they would have been without that information. Second, specialization among humans is largely informational in character. The more specialized (lumpier) is the distribution of productive knowledge and its associated skills, the greater is the potential output of an economy. Specialization increases the productivity of team production—just as it does for sports teams. The productivity of specialization tends to induce stronger demand for specialists (because of their higher marginal revenue products), which, in turn, induces more people to specialize. Third, the existence of latent demands implies that new types of products can be profitably introduced to markets in full equilibria. This type of ignorance, when recognized, motivates a variety of entrepreneurial and innovative activities—which expand the extent of markets as new products are successfully conceived and brought to market. In all three cases, informational factors influence the size and nature of markets.

Moreover, as emphasized by Hayek, the very limited knowledge that all of us have about how and where products are produced would, by itself, tend to make markets far less efficient as institutions for advancing the material interests of consumers than they are. However, when markets are at least roughly in equilibrium, prices provide useful information about the relative importance of products to consumers. Prices also provide useful information about the relative scarcities of inputs. Together these motivate individuals to specialize in the occupations deemed most important by consumers (through effects on marginal revenue products). The net effect is not chaos, but (adequately) informed individual decisions about what to purchase, what occupations to specialize in, what products to produce, and where innovations are likely to be most profitable. In this manner the price system can be said to coordinate the decisions of millions of relatively uninformed consumers, firms, and innovators—because so much information is implicitly included in those prices.

Market prices, when approximately at equilibrium levels, allow consumers, firms, and innovators to make far better-informed and more productive decisions than their own knowledge alone would allow.

References and Suggested Readings

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *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*(3), 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.

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

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

Machlup, F. (1981/2014). *Knowledge: Its creation, distribution and economic significance, Volume I: Knowledge and knowledge production*. Princeton University Press.

Mortensen, D. T., & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *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. *Quarterly Journal of Economics*, *90*(4), 629–649.

Schumpeter, J. A. (1942/2013). *Capitalism, socialism and democracy*. Routledge.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423.

Shannon, C. E., & Weaver, W. (1998). *The mathematical theory of communication*. University of Illinois Press.

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. *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.