Volume Title: Doctors and Their Workshops: Economic Models of Physician Behavior

Volume Author/Editor: Mark Pauly

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-65044-8

Volume URL: http://www.nber.org/books/paul80-1

Publication Date: 1980

Chapter Title: Physicians as Agents

Chapter Author: Mark Pauly

Chapter URL: http://www.nber.org/chapters/c11521

Chapter pages in book: (p. 1 - 16)

1   Physicians as Agents

Total medical care expenditures in the United States in 1978 were $192 billion. Of that amount, $35.3 billion, or only 18.3%, represented payments for physicians' services. In the manufacturing sector, profits before taxes were about 13% of total expenditures. Looking at this first set of figures, one might be tempted to conclude that physicians are of relatively minor importance in the health care system, and that the prices, fees, or other incentives which they experience are of little consequence. In a sense this is true; a substantial reduction in physician fees (say, 25%) would have little effect on total medical spending.[1] By the same token, however, a 25% reduction in profits would have little effect on most product prices. The major theme of this study is that physician-returns play a role in the medical care sector which is analogous to that played by profits in the for-profit sector. Just as the incentive provided by the relatively small fraction of total spending that is profits determines the form and use of all inputs and outputs in conventional markets, similarly the relatively small amount that goes to physicians provides the financial incentive which determines the bulk of resource use, output quantities and characteristics, and total costs in the health sector.

There has been fairly extensive study of the market for physicians' services, and of physician time as an input into these services, especially with regard to physician provision of ambulatory care. The general focus of the analysis here is different: this analysis will be concerned with the effects physicians have on the use of medical inputs other than the physician's own input or those inputs for which he pays directly. There will be a similar consideration of the effects of physicians and physician incentives on the markets for kinds of medical care other than physician's services. Physician behavior will be examined not just in the office, but also in the hospital and in all of the physician's "workshops." The

focus will be not just on the physician's own actions, but also on the effects of those actions on the environment in which he works and on the resources with which he works.

The parallel which has been drawn between physician expenditures and profits, like all analogies, is not perfect. There are some additional aspects of physician behavior to be emphasized. The most important of these is the physician's role as provider of information to demanders. The physician not only coordinates the output production process, as does the classical entrepreneur, but he also provides information to potential demanders on the need for and value of services that he or others will provide. This represents a kind of relationship between buyer and seller which falls outside the scope of the standard neoclassical models, whether competitive or monopolistic. While there are some similarities between physician advice and advertising, the integration of information and services provision will be the most distinctive feature of the models that will be discussed.

A second difference concerns what might be reasonable to postulate as the physician's maximand. Maximizing profits is an acceptable assumption for most businesses, but maximization solely of physician money income is not a plausible assumption for physicians. Instead, much of what is to be explained empirically will require broadening the set of arguments that concern the physician's utility function. While models of a utility-maximizing entrepreneur are not unknown, the economy of theory and the empirical relevance of the money income maximization approach will not be available in this study.

The physician, like the classical entrepreneur, really performs two functions: (1) he organizes and directs the production process, and (2) he provides some productive input. In the case of the entrepreneur, his equity capital investment provides some of the firm's capital stock. In the case of the physician, his own time is a useful and often essential productive unit. While the importance of capital to the production of output has always been recognized, and while the role of physician time in the production of office-based ambulatory care has also been subject to scrutiny, an important part of total physician time has been ignored or treated haphazardly. This is the time the physician spends caring for patients in hospitals. The fraction of total working time spent at the hospital varies according to the physician's specialization, of course, but what sketchy information we have suggests that it averages about 30% of physician time, and may be more than half of total time for the hospital-oriented surgical specialties.[2] In order to round out our understanding of how physicians affect the use of all inputs, including those not paid for by the physician, it is necessary to incorporate physicians' time into the analysis of hospital production.

In principle, of course, physician time input also affects the productivity of other ways of treating patients. More time spent diagnosing or monitoring may make drug therapy more effective and nursing home care more productive. But it will not be possible to treat those subjects here.

A Taxonomy of Physician Pricing Models

In order to explain how physicians affect the use of other inputs, we need to understand how physician prices are set. This is ultimately an empirical question, on which there is presently no consensus. It will be helpful to try to classify alternative theories which have been suggested. Figure 1, an extension of a scheme developed by Uwe Reinhardt,[3] provides one such classification.

Market situation Physician Pricing	Market clears, demand not shifted	Market doesn't clear; demand not shifted	Market clears; demand shifted
Price-taker	1 Competitive equilibrium	3 Price ceiling	5 Oligopolistic demand creation; target income
Price-setter	2 Monopoly equilibrium	4 Chronic excess demand	6 Shift of entire demand curve

The rows refer to the extent of individual physician control over price. The physician may, as in the first row, take a going price as given, either because competition forces him to do so, or because insurance reimbursement policy, administrative difficulties of changing prices, or government price controls compel him to accept a given price for any output he produces. In the other row, the physician is assumed to have some control over the price he can charge and still obtain business. The columns reflect alternative assumptions, one dealing with whether or not the price attains a level at which supply and demand are equated, and the market clears; the other dealing with whether or not the individual physician is able to shift the demand for his own services by varying

what he suggests to patients. (There could also be a fourth column, in which demand is shifted and markets do not clear, but that would only represent a combination of the second and third columns.)

The cells provide a way to classify possible models. Cell 1 would be the usual competitive market. If the physician services' market were of this nature, both demand and supply curves could be estimated. If the market were the standard textbook monopoly in cell 2, the market, individual, or firm-level demand curves could be estimated. It would not, however, be possible to estimate a firm or physician supply curve of output. The output the monopolist physician would choose would depend upon the configuration of the entire demand curve, not just the price he happened to charge. There are some assumptions that might be made which would make it possible to estimate a quasi-supply curve or offer curve, but it is not clear that these assumptions would be legitimate.

So far, however, these are standard problems in estimating models of any market. The peculiar problems of the medical care industry begin to emerge in the other cells. A price which physicians take as given may be set below the level which permits markets to clear, as in cell 3. This is most likely to happen if, either because of insurance reimbursement limits or because of price controls, the price is not permitted to rise. It may also represent a state of disequilibrium, in which prices are gradually rising to their equilibrium levels. In this case, no demand curve for physician services can be estimated; the only points observed are those on the supply curve.

In cell 4 the physician controls his own price. If he is an income maximizer, he would set the price at the monopoly level, and satisfy all the demand at that price. But he may have goals other than income, and pursuit of these goals may prompt him to set price below the monopoly level, and refuse to produce enough to satisfy the demand at that price.

The distinction between cells 3 and 4 arises from the source of excess demand. In cell 3, price is exogenous to any individual physician. In such a case, it is the constraint on his willingness to supply that causes excess demand to occur. Excess demand may arise for the usual reason: price may equal the physician's marginal cost (including the opportunity cost of his own time) at a quantity which is less than that demanded from him. If he has preferences as to the kinds of cases he will treat, it is possible to have excess demand for some kinds, while other, more desirable kinds are selected to be treated. The trade-off between leisure and work may depend on the kind of work to be done. If there is an "insufficient" amount of desirable cases, he may not be willing to supply as much labor time at a given price as when the mix of outputs is more to his liking.

In cell 4, however, the physician sets the price. He may not be willing to produce all of the output demanded at this price, for the reasons to

be discussed later. In this case we obviously cannot observe the demand curve. It is also true that we cannot observe a supply curve either, since the monopolist by definition is not able to sell all he wants at given prices. In cell 3 excess demand need not require physicians to have more in their utility functions than income and leisure, if prices are fixed by law or are sticky. But in cell 4 excess demand can only exist if physicians for some reason receive utility from its existence.

In all of these cases, the individual physician takes the demand curve which confronts him as given. We will now consider models in which the physician can create or shift the demand curve which confronts him. In cell 5, he would obviously only have an incentive to do so if price exceeds marginal opportunity cost; at prices below that point, the situation is the same as that in cell 3. If price is above marginal cost, it is possible to show how demand will be created (this will be done in a later chapter).

Finally, in cell 6 the physician chooses the price he wishes to charge on a demand curve which he manipulates. Precise analysis is difficult here because behavior depends on the way in which the demand curve shifts as information is changed, and there are no comfortable a priori conjectures on this. We can, however, use this model to provide some conditional statements about market response.

The Physician as Agent

Many medical goods and services are not demanded directly by patients, but are requested on the patient's behalf by a physician. Indeed, in some cases, such as prescription drugs or hospital care, the patient is not even legally permitted to demand the items. And yet economists have persisted in estimating consumer demand curves for medical care in the same way as for other goods over whose purchase the consumer has direct control. How can the supposed response of use to consumer income and user price be rationalized when consumers are not, in large part, making the relevant decisions? Feldstein has suggested that an answer might be constructed on an assumption that the physician acts as the patient's agent.[4] In this role, the physician demands (or has the patient acquiesce in demanding) exactly those quantities of care of various types that the patient would choose if he had the information and knowledge that the physician has.[5] If the physician were to act as a pure or perfect agent, his ostensible maximand would be the patient's utility, and his choices would duplicate the choices the patient would make, if the patient had the same information as does the physician.

While much of medical ethics can be interpreted as an attempt to use moral suasion to persuade physicians to adopt the agent's role, there are reasons to expect that the physician may not act as a pure agent and

that his maximand need not be the fully informed patient's maximand. First, it is not obvious that any physician will have the information needed to permit him to act as a perfect agent, even if he wishes to do so. While the patient may be ignorant about the way to produce health or other attributes of medical care, he alone knows how much health he desires. The physician may be better informed about the technology for producing health, but he may be poorly informed about the patient's demand for health. On balance, the patient may be better informed than the physician. This assymetry of information has often been noted, but it has less often been recognized that it is an assymetry in type of information, not necessarily in amount. Indeed, Smallwood and Smith have even suggested that, when confronted with the difficult task of reading his patient's preferences, the physician may give up altogether and concentrate only on equal physically measurable outputs, such as changes in probability of death.[6] Whether or not physicians behave in this way, it is worth noting that physician ignorance can be as important as patient ignorance in causing deviations of actual from ideal outcomes.[7]

Second, even if he had the requisite information, the physician may not in fact act as a perfect agent. It is not usually assumed that any other economic actor chooses only with the interests of his customer in mind. From Adam Smith on, economists have argued that benevolence in producers of goods and services is neither to be expected nor necessarily desired. But if the physician manages the care of the patient with his own interests—income, leisure, interest of work—in mind, then he would not necessarily be acting as a perfect agent.

Thus there are reasons to expect that the agency relationship will not be perfect. But in order to get much further, two additional questions must be answered. How would the physician be expected to act if he is *not* acting as the patient's agent? And equally important, how is the extent of his departure from perfect agency determined? I intend to show that it will generally be both possible and likely that the physician will depart from the role of perfect agent, but that there are constraints on the extent of this departure. Perhaps more surprisingly, I shall show that the physician will choose medical care inputs so as to minimize the cost of producing a given level of health *even if he is not acting as the patient's agent, but rather as a selfish income maximizer.* I shall show that, given a suitably broad definition of cost, a similar proposition holds for a much wider class of physician utility functions as well.

Some Difficulties in the Notion of Agency

The concept of agency is most transparent for the situation in which all services are sold in competitive markets. The agent's only task then is to choose the quantities of those inputs that would be demanded by

the consumer at the given prices if the consumer had as much information as the agent has. That one of the services may be produced by the agent is of little consequence. As long as the agent accepts the price as given, and believes that he can sell as much as he wishes at that price, the agent has no incentive to alter the quantity of any service that any individual consumer demands, and no ability to alter the price he pays. Of course, if consumers are so poorly informed as to need agents, then a competitive market might not emerge. However, if some consumers are well informed, even if others are not, a competitive market may exist, and agents may behave optimally.

When the agent is also a monopolist in the provision of some productive input, there are two possible assumptions one might make. One might assume that he has the patient demand the quantities a fully informed patient would demand if price were set equal to marginal cost. Alternatively, the agent might try to determine the quantity that patients would demand at any given price if they were informed. Then, given this "true" demand curve, the physician would set the price at the monopoly level, and recommend the "fully informed" quantity corresponding to that price.

If the market is perfectly competitive, there is no reason for the physician not to act as a pure agent, while if the market is monopolistic, it appears to be inconsistent to assume that the physician is an income maximizer when he sets prices but not when he offers advice. The same motivation which leads him to set prices above marginal cost will lead him to distort the advice he gives.

A second difficulty with the agency notion arises in situations in which the quantity used is not the quantity that would be demanded by fully informed consumers when price equals social marginal cost. One possibility is that there is excess demand. If the physician were to act as the agent for a given patient, he would choose to satisfy that patient's fully informed demand. But if supply constraints make it impossible for him to do so for all patients, then the physician cannot feasibly act as a perfect agent for all patients.

A somewhat similar problem arises when insurance covers all or part of the cost of services the physician orders. From the viewpoint of any individual patient, his utility is maximized if he receives (approximately) the quantity he would have demanded if he were fully informed and faced a price equal to the net or user price. Even though the value to him of the last units he purchases is probably less than their social marginal cost, he ignores this "welfare loss" because the cost is spread over the usually very numerous other insureds. Even an individual physician's entire list of patients is likely to be a sufficiently small fraction of total insureds to permit the physician to ignore this welfare loss. But consumers would prefer that all physicians adopt a rationing policy in

which they receive the quantity they would have demanded at price equalling marginal cost, not the quantity at which price equals the marginal user price. However, one would not expect the individual physicians to adopt such a rationing policy. When insurance is present, we need to recognize that it may not maximize aggregate patient welfare if each physician acts as a perfect agent, in the sense of choosing the quantities his patients demand at the user price. This ambiguity in the notion of agency means that welfare evaluations will often be ambiguous as well.

The Tasks of the Agent

I first indicate the problem that the consumer would like the physician acting as agent to solve. The consumer is initially assumed to maximize a utility function in health H and other goods X.

$$(1) \qquad U = U(X, H)$$

Other goods may be purchased at a price of one dollar, but health is not purchased directly. Rather, its production is given by:

$$(2) \qquad H = H_o + g(M; H_o)$$

where H_o is initial endowment of health, and M is a vector $M = (M_1, M_2, \ldots, M_n)$ of medical services. Time cost is ignored in this exposition; it can easily be added without altering conclusions. The vector of prices for medical care is P, and the consumer's income is Y. The consumer's problem is to maximize the utility function (1) subject to the production function constraint (2) and the income constraint

$$(3) \qquad Y = X + PM$$

Optimality requires that levels of medical inputs be chosen so that the marginal health products of a dollar spent on each are equal.

$$(4) \qquad \frac{\partial H / \partial M_1}{P_1} = \frac{\partial H / \partial M_2}{P_2} = \ldots = \frac{\partial H / \partial M_n}{P_n} = 1/\Pi$$

where Π is the shadow price (equal to marginal cost) of an increment in health. Optimality requires that the marginal rate of substitution between health and other goods equal its shadow price:

$$(5) \qquad \frac{U_H}{U_X} = \Pi$$

The vector M^* which satisfies (4) and (5) is the quantity the physician would choose if he were able and willing to act as a perfect agent. Note that this choice involves two subproblems. First, for whatever level of health produced, costs should be minimized (equation [4]). Second, the level of health should be that quantity that is demanded at the schedule

of shadow prices (equation [5]). One may observe that the physician qua physician is more likely to have the knowledge needed to achieve (4) rather than (5), although even achievement of (4) requires that the physician know input prices as well as marginal effects on health.

Using the Model

From the utility function (1) and the constraints, it is obviously possible to derive a patient demand schedule for health as a function of the shadow price of health the patient faces. The shadow price in turn depends upon the price of health inputs and the patient's estimate of the marginal health products of those inputs. Given a homothetic production function, the mix of inputs chosen depends only on their relative marginal products and prices, not on the total amount of health demanded. The absolute level of marginal products determines the level of health the patient expects to achieve. Obviously, the information provided by the physician, or the choices that he makes, can reflect distorted relative *or* absolute health products.

Even if the physician is an income maximizer, he will act as the patient's agent in the sense of behaving in response to true relative marginal health products: he will minimize cost and will satisfy equation (4). But I also show that in his advice to the patient he would be expected to distort the absolute increase in health to be expected from a combination of medical inputs, and so he will not satisfy equation (5).

For simplicity, I begin by assuming that the physician maximizes his net money income. The health production function is altered to include explicitly two medical inputs. Input M_1 is produced by the physician (e.g., physician office visits) using his own time; M_1 is available at a constant marginal cost P_1. Input M_2 is sold at price P_2. No physician time is used in its production (e.g., prescription drugs or in-patient hospital care). In terms of the intermediate inputs M_1 and M_2, the health production function is:

$$H = H(M_1, M_2; H_o)$$

Suppose some increment to health $\Delta \bar{H}$ is to be produced. If the physician acts as a perfect agent, he will choose quantities ΔM_1 and ΔM_2 such that the ratios of their marginal products equal the ratios of their marginal costs. We may call these quantities $M_1{}^*$ and $M_2{}^*$. Since the production function is assumed to be homothetic, $M_1{}^*/M_2{}^* = \bar{k}$, a constant given any level of relative prices.

The question now is whether the physician will choose the combination of inputs $\bar{k}$ if he acts to maximize his own income. We suppose that the patient as patient has no preferences as to the mix of medical care inputs, and regards the production function as a "black box"; he

only cares about the health he will achieve and what he will have to spend to get it. The price he will have to pay for any batch of medical inputs is just the sum of their prices, but the health he expects to get depends in part on what the doctor tells him. However this expectation is formed, it does not, we assume, depend upon the mix or types of intermediate inputs. The patient views the physician as a prime contractor, who quotes him a total price and an expected final output. For a given expenditure, let the increment in health the physician tells the patient he will receive be $\Delta\hat{H}$; then the shadow price of a unit of health is approximately $\hat{\Pi} = \dfrac{P_1\Delta M_1 + P_2\Delta M_2}{\Delta\hat{H}}$. Obviously, the higher $\Delta\hat{H}$ the lower $\hat{\Pi}$ will be. The lower $\hat{\Pi}$, the more likely it is that the consumer will be willing to buy a larger number of units of health.

Consider first the situation in which the particular physician from whom the consumer purchases medical care is given, and the only question is how much the consumer will buy. If the physician is an income maximizer, he will set his price with regard to the patient's *total* demand curve for health. If he could get by with providing no inputs, he would set his price and choose what he would tell the patient so as to yield that shadow price for health which maximizes total revenue. If he must provide some inputs, for reasons to be suggested later, he would still maximize his income for any level of cost by representing $\Delta\hat{H}$ as greater than ΔH, i.e., by overstating the marginal health product of M. If, on the other hand, there is some competition among physicians and if the physician is a profit maximizer, his price will be set with regard to the *firm-level* demand curve for health. The physician will want to represent his performance of a particular procedure as likely to add as much or more health as performance of that procedure by any other physician.

Once the physician and patient agree on a total expenditure and an expected final level of health, we suppose that there must actually be some costs incurred and that there is some increment in health that must actually be delivered. In general, this actual level will differ from $\Delta\hat{H}$, but we suppose that it is determinate. It might be, for example, the minimum amount needed to forestall a malpractice suit, or undesirable competitive repercussions, or the consumer's prior knowledge might determine a maximum ratio of $\Delta\hat{H}$ to ΔH. With total expenditures already determined, and with the actual level of ΔH to be produced fixed by these minimal consumer expectations, physician income maximization will clearly require that the actual level of ΔH be produced in least-cost fashion. This is so even if some of the total expenditure goes for inputs which the physician himself does not employ, since the initial bargaining and information exchange only fixed the total expenditure for treatment, not its division among types of treatment. Were the physician to select

anything other than the cost-minimizing input combination for whatever level of H he finally produces, he would reduce the amount he himself could collect. In summary, deviations from perfect agency would be limited to misrepresenting the marginal health product of care in general, and hence the shadow price of health.

Choice of Input Combinations in Practice

While it is a simple point that the income-maximizing physician will be likely to choose the cost-minimizing combination of inputs, it seems not to be recognized in much of the discussion of physician behavior, and in comparisons between that behavior under fee-for-service and prepaid group practice. Consider two kinds of other inputs: hospital admissions and prescription drugs. Both of these are almost always purchased (at least initially) in conjunction with physicians' services. Given a demand constraint (whether demand is competitive or not), the discussion above indicates that the maximizing physician has an incentive to choose the least-cost combination and type of inputs.

Yet it has often been alleged that such incentives are absent. For example, it has often been argued that physicians will tend to overuse hospital inputs primarily because they are "free." It is true that the cost of these inputs is not billed to the physician, but they are far from free if insurance coverage is not full. Of course, full insurance coverage makes them appear to be free, but this incentive would also appear in a prepaid group if it purchased its hospital insurance as a member of a large pool.

As another example, Victor Fuchs has stated, in connection with prescription drugs, "If (contrary to fact) the physician had a financial stake in keeping down the cost of drugs he prescribes, as he would under a comprehensive prepayment plan, he might be motivated to examine more closely drug prices and alternative products and he undoubtedly would also be less susceptible to persuasive detail men and high pressure advertising."[8] In fact, the physician *does* have a financial stake under an ideal fee-for-service market. Given the above formulation of demand, higher patient payments for prescription drugs reduce what the physician can charge for his own services. If the patient were willing to pay a higher price for, say, the more expensive branded drug the physician prescribed, he would also have been willing to pay the physician a higher fee. Indeed, the incentive to cost minimize is likely to be stronger at the level of the individual physician under fee-for-service than in a prepaid group, since a dollar in extra drug or hospital user costs reduces the fee-for-service physician's income by a full dollar, while in a typical prepaid practice the reduction in an individual physician's share would ordinarily be somewhat less.

All of this does not, of course, imply that actual choices one observes will be efficient; it only says that money income is maximized if they are efficient. If we look at studies of physicians' office practices, and examine how physicians use inputs for which they *do* pay directly, one generally finds that physicians have not made efficient choices. While there has not been much evidence of technical inefficiency, there is evidence that physicians do not use the net income-maximizing (or cost minimizing) number of aides.[9]

If physicians do not cost minimize in the choice of inputs in their own offices, one should probably not expect them to cost minimize in their choice of inputs elsewhere. Thus it seems to be very misleading to view such alleged inefficiencies as arising from the lack of financial incentives, or to view the prepaid group's incentive structure as something which would lead to improvement.

The question then becomes: Why do they not cost minimize? There are three possible explanations offered for the results in office practices. The first is that physicians do not know what the cost-minimizing combination is; for example, they mistakenly underestimate aide productivity. Usually it is not alleged that physicians are unaware of studies which show high aide productivity. Rather, they are held to be ignorant of how to use aide time efficiently in their own practices; they may not know how to delegate.

But this only appears to push the argument back a step. It is not that physicians underpurchase aides; it is that they underuse information on how to use aides. Welfare evaluation of such arguments is difficult, because of the problem of specifying the value of information when the buyer is intrinsically uncertain about what the information may tell him. Nevertheless, there appear to be no barriers to the spread of information on the use of aides. Such information is not generally of a public goods nature, since the physician would need to know how to use aides in his own practice. The probable reason why a physician has not experimented with more aides, and so does not know how to use them, is that the cost of experimentation exceeds the expected gains from future use of the aides. Similarly, the reason why physicians do not know when generic drugs are therapeutically equivalent to branded drugs is because the cost of the information is too great. Here the "information" is not generally the facts themselves, for those are reported in the *Medical Letter*, available at a nominal price. It is the cost of perusing and digesting this information. In summary, while lack of information may be a reason for nonuse of the least cost technique, it is not obvious that, when the costs of obtaining and processing information are added, the total costs of adopting and using the "least-cost" technique really are the lowest.

The second explanation is that output is measured improperly: a task performed by an aide really is not as productive or valuable as a task performed by a physician. While questionnaire surveys indicate that those patients who use services do not always feel this way, this question has never been answered definitively. It is hard to believe that, for tasks which involve diagnosis or judgment (as opposed to mechanical skill), a patient would not prefer a physician. Recent work by Coate on optometry, in which outputs can be measured much more easily, indicates little or no aide underuse.[10]

The third reason is that there are psychological costs to the monitoring and supervision of aides. Although monitoring an aide's performance may not take the physician more time, monitoring may take more effort or require him to reduce patient contact in which he takes pleasure. He may simply not prefer to work in the style of practice associated with larger numbers of aides; there is little in the process of physician selection, training, or continuing education which is likely to result in high levels of skill in personnel management. This last reason may be a plausible explanation for the observed facts: physicians do not choose the cheapest drug *or* the ideal number of aides because they do not want to take the effort to seek the most efficient method of practice.

It is important to recognize the welfare implications of this explanation: if these costs do exist, it is *not* necessarily inefficient to do the inefficient thing. Put differently, the compensation that physicians would have to be paid to induce them to search out the cheapest drug or make the effort to supervise aides would exceed the cost savings to be realized. Put still differently, a rule which required physicians to use the optimal number of aides, or search for the cheapest drug, might indeed reduce costs. But it would make physicians worse off. The amount physicians might be willing to pay to remove the rule might exceed the amount of cost saving. Distributional questions aside, "efficiency" is not necessarily desirable; psychic costs are not imaginary; they are as real as the value of leisure time that the worker gives up or the enjoyment of goods and services that a consumer must sacrifice. The utility-maximizing behavior of physicians may make this case differ from the normal one of joint products. Utility maximization by physicians, as suppliers and managers, may indeed make medical services in general, and prescription drugs in particular, different from furniture or automobiles.

There is, of course, a distributional effect, a question of property rights. Forcing physicians to do the efficient thing may permit the class of consumers to be better off, but it makes physicians worse off. More to the point, the amount by which physicians are made worse off exceeds the amount by which consumers are made better off. The class of physicians could offer a bribe to consumers *not* to impose "efficiency"

rules on them which would exceed the gains consumers might expect from increased efficiency. The fairytale character of this political fable suggests that we might expect such "efficiency" rules to be imposed even if they are inefficient, just because consumers are more numerous (though perhaps not always more politically potent) than physicians.

Unequal Information, or A Little Knowledge Can Be a Dangerous Thing

In the preceding discussion I assumed that the individual consumer determined how much medical care to buy in total, and how much to buy from each provider, on the basis of the total price (over all inputs) that he would have to pay for an increment of health. The consumer was assumed to pay no attention either to the mix or to the prices of individual intermediate inputs. There are two possible objections to this assumption. First, individuals may judge the health outcomes they expect to receive not just by what the physician tells them, but by the actual intermediate inputs provided. If consumers think they will be more likely to recover with a branded drug than a generic one, or with any prescription at all (even an unnecessary one), rather than just advice and hand-holding, then obviously the cost-minimizing mix of inputs may not maximize physician income. Income-maximizing physicians may find it worthwhile to cater to uninformed consumer desires. Second, and possibly more important, it is plausible to conjecture that the elasticity of demand for health from a given provider may not be the same for all input mixes. (Even if this elasticity were the same, of course, the elasticities of demand for individual inputs would generally differ. Doubling the surgeon's fee will generally have less effect on the demand for hospital admissions than would doubling the hospital price, even given equal insurance coverage.) Differences may arise because the consumer has different amounts of information about the prices and practices of various physicians.

Concretely, consumers may have a rough idea of what a physician's fee for an office visit ought to be. Suppose the consumer has some information about the going prices for routine physician services. He knows that an office visit typically costs between $10 and $20. He has very little information on the prices, types, and quantities of prescription drugs provided for various illnesses, or the usefulness of hospitalization or laboratory tests in connection with a particular illness. Suppose that physicians in his area typically order $30 worth of diagnostic tests for a particular condition, and charge $15 for an office visit. Suppose the tests are in fact worthless, or nearly so. By the argument in the section "Using the Model," each physician would have an incentive to stop ordering the test, thus saving the patient $30, and raising his fee to capture some

of the gains from not prescribing the test. Suppose a particular physician raised his office fee to $30, and prescribed no tests. Under full information this would attract patients to him, since his cost for treating the illness is $30, as opposed to $45 total cost for other physicians. But if consumers do not observe the total cost, but only the $30 versus $15 difference in office fees, this cost-minimizing physician may actually lose patients even while making them better off.

More formally, the problem is that, if information differs in this way, the firm-level demand elasticity may be smaller with respect to the prices of the "ordered" inputs than with respect to the physician's own input. While the physician would have an incentive to inform patients of his low-cost mode of treatment (and many do), this ability is limited by restrictions on advertising and the heterogeneity of treatment modes. An interesting empirical question would be whether the elasticities do differ in this way.

A second possibility is that, however desirable it might be, the physician may be unable to vary his own charge. This may happen because of third party reimbursement arrangements or because of price controls. However desirable price limitation may be for other reasons, it does have the effect of eliminating any cost-minimization incentive to the physician who does not have excess capacity, for he will be unable to capture any cost savings in his own fee. If he is willing to supply more output at the going price than is demanded of him, then he may be able to increase the quantity demanded from him by lowering the cost of other inputs.

Conclusion

That physicians only direct rather than pay for many other health inputs need not lead to departures from cost-minimizing behavior. This conclusion holds even though physicians may be able to induce consumers to obtain amounts of health, and consequently amounts of medical care services of all types, that differ from those that would have been provided were consumers not ignorant. How consumers decide how much to believe of what physicians tell them and how their beliefs affect their consumption decisions and constrain the actions of physicians will be treated in more detail in subsequent chapters.

.