

**The Interaction of Competition and Regulation in Affecting Product Quality
in Medical Services**

Mark V. Pauly, Ph.D.
Department of Health Care Systems
The Wharton School
The University of Pennsylvania
Room 208 Colonial Penn Center
3641 Locust Walk
Philadelphia, PA 19104
215.898.2838 (Tel.)
215.573.7025 (Fax)
pauly@wharton.upenn.edu

April 8, 2004

This is a draft of a paper prepared for the Federal Trade Commission's Conference on Health Care Information and Competition, to be held at the FTC Conference Center, Washington, D.C.,
on April 16, 2004.

Introduction

Both public policy and economic theory have traditionally been concerned with the effect of competition on prices of goods and services. The benchmark model predicts that competitive equilibrium will lead to a price equal to (long run and short run) marginal cost. That price, and the quantity demanded associated with it, can be said to be efficient.

In recent discussions of the functioning of a medical market, issues of quality have drawn attention. In particular, the emphasis (traditionally normative) on price or spending is being supplemented and in some cases supplanted by concerns for quality. The controversy stirred up by the IOM committees that recently studied medical errors has led to more discussion of the determinants of this dimension of quality, while the development of report cards of various types has called attention to the dimension of quality they measure.

In this paper, I review conceptual and empirical evidence on the relationship between competition and quality in both price-regulated and unregulated medical markets. Some of this work is quite old—the pre-managed care “arms race” model of quality determination—but all of it bears on current concerns.

The most fundamental positive question is how competition affects quality. The normative question is more complex: it asks whether more competition (compared to less) leads to a price-quality combination that is better for consumer welfare. In making this judgment, one cannot evaluate quality independently of the price being charged and the cost being incurred. Apart

from making the obvious points that movements which both lower cost and increase quality will be efficiency improving and therefore desirable, and those which raise cost and lower quality will be undesirable, we will not be able to offer definitive normative statements when or if competition should cause quality and price both to increase or both to fall.

Consumer knowledge and competition over quality

If consumers were perfectly informed, they could decide what level of quality they prefer, given some (marginal) price for quality. They would also choose a higher quality product or seller, given similar prices. These choices are what drive outcomes of competitive markets toward ideal quality, although the actual outcome depends on the supply side as well.

If consumers are not well informed, by contrast, it is very unlikely that markets will perform well, even if the other competitive preconditions are present. It is not even guaranteed that competition will represent an improvement over markets with fewer sellers and less competition. It is also not necessarily true that changes which add to but do not perfect consumer information, given some level of competition, will represent improvement. All bets are off.

These ideas have discouraging implications for the performance of medical markets and the development of anti-trust policy for them. Even without extensive empirical research, we know that for many medical products and services consumer information is imperfect, often woefully so. Indeed, many medical services are themselves nothing but the supply of information, whose quality cannot be known until production has already occurred.

Some Benchmark Models

The simplest way to add quality to the standard economic model of competition or monopoly constrained by regulation is to assume that all consumers prefer the same level of quality, no matter what their demand for quantity is. This optimal level of quality presumably reflects a choice based on a comparison of the marginal cost of yet-higher quality and its marginal benefit; there is no implication that (in a technical sense) quality is maximized. In an efficient outcome, the number of defects would still be positive; the best technology would remain unadopted; there could well be a “chasm” between the level of quality supplied in this ideal market and the maximum possible quality. If the market is competitive and all consumers well informed, the level of quality that will be supplied is the level consumers prefer, which is the (constrained) optimal level at which marginal benefit equals marginal cost. The optimal quality will fall if marginal cost rises or marginal value or benefit falls; if costs or benefits differ, quality should differ.

The more general and probably more realistic benchmark model is the classic version proposed by Sherwin Rosen. If there are (even modest) firm-level economies of scale in specializing in one level of quality rather than in supplying multiple levels, but no large scale economies in output, in equilibrium there will be many firms offering many different levels of quality. Consumers with different demands for different levels of quality will patronize different firms. Again there will be “gaps” between the lowest and highest levels of quality in the market, matched by gaps in unit prices, but the equilibrium will be efficient.

I explore the impacts on this equilibrium of three changes in the basic model: (1) Imposed monopoly; (2) economies of scale; (3) imperfect information.

Imposed Monopoly, For-Profit and Non-profit

Now suppose that, even though it would be technically feasible to have a large number of sellers at each level of quality, in fact a market is forced by some cause to have a small number of sellers. What will happen to quality and price? The general idea is that, the fewer the firms, the more closely a typical firm's demand curve approximates the market demand curve. Let us begin with the single-quality model. The key parameters here are the price elasticity of demand—the percentage change in quantity, given quality, for a given percentage change in price—and the analogous quality elasticity of demand—the percentage change in quantity, given price, for a given percentage change in quality. If quality can be measured directly, or if we measure it by cost, a well-known result (the Dorfman-Steiner Theorem) says that the relative importance of price and quality in characterizing profit-maximizing firm level equilibrium will depend on the relative size of these two elasticities. That is, if demand from a firm is relatively less responsive to quality than to price, the firm will choose to supply low quality products at low prices. Conversely, if demand responds more to quality, quality and price will be high. In the initial competitive equilibrium, both of these elasticities will be very large: a firm will lose almost all its customers if its price is above the competitive level or its quality below. If we reduce the number of firms, we know both elasticities will fall in absolute value, but it is the relative change in the two that will be important for understanding which way quality will move

or, indeed, whether it will increase or decrease. For example, if the quality elasticity should rise substantially relative to the price elasticity (even while both are falling), it is possible that quality will be higher in the less competitive setting. If it is easier to convey information to buyers about price than about quality, competition would lead to lower prices and lower quality than monopoly.

It is often hypothesized that the movement away from indemnity insurance toward managed care led to an increase in insurer-level price elasticity of demand relative to quality elasticity. If this were so, note that quality would fall either under competition or monopoly.

These conclusions will continue to hold for a non-profit firm as long as that firm makes pricing and quality choices in the same way as a for-profit firm would. For example, the firm may choose to maximize profits and then devote those profits to some part of the non-profit's mission, such as assisting the poor or fostering research. The analysis will change if the non-profit attaches utility to quality per se.

Adding imperfect information to this model can be done in a simple way just by assuming that imperfect information tends to make firm level demand less elastic. Then its impact on quality (versus price) depends on whether it reduces one type of elasticity relative to another.

Departures from Cost Minimization

Some observers assert that quality could be improved and cost lowered in the medical care sector today if only providers could be offered enough of a reward or incentive for doing so.

Sometimes that initiative takes the form of payments (though usually small) for efficient performance. Alternatively, sometimes the suggestion is made that greater competition per se would pressure providers to take actions with nothing but good effects that they have somehow failed to undertake.

The problem with such assertions is that, if the firm was maximizing profits, it should not have been behaving inefficiently in the first place. So we have no basis for concluding that competition or incentives will make it more efficient in the technical sense. The most we can say is that incentives can cause it to shift to a higher quality, higher cost position, whether or not that position is more efficient. In a more practical sense, a competitive firm would be pressed toward improving quality if that would reduce costs by two incentives: by the desire to maximize profits, and by the desire to survive as a business. In contrast, a monopoly firm does not necessarily face the second incentive. If the firm's manager gets more benefit from avoiding cost reduction—a peaceful life, the avoidance of interpersonal conflict, or fewer enemies—he may permit operation at a point where higher quality could be associated with lower cost. But the monopolist that single-mindedly pursues profits will also never be technically inefficient in producing whatever quantity and quality level it does produce. The other possibility is that the firm might not *know* how to reduce cost and raise quality, even though external critics claim to know. But the problem with this story is why didn't the firm make the effort to find out?

The easiest explanation is that doing the things critics propose as reducing waste really would require effort, effort which management or labor is reluctant to supply. This is not pure inefficiency: reducing cost must save less than the effort to do so would cost. For example, the latest discussed technique to improve quality and increase efficiency is the greater use of computerization in clinical care in hospitals, but this has been slow to progress. It appears to be the case that computerized order entry requires more physician time than writing a short note in the patient's chart. If so, this cost should be taken into account and, without measure of it, we cannot be sure whether the new technology really should be spreading faster.

That so, there should be a positively sloped supply curve of higher quality, once the cost of effort is taken into account. The only positive issue is whether more of that effort would be supplied under an alternative arrangement, whether it be greater competition or pay for performance. The normative question is whether the payment needed to lower cost exceeds the cost reduction itself. Generally, we could expect the increase in effort to be sufficient, but some discrimination or some better targeting might help. But it is bound to be a difficult struggle.

Imperfect Information

It is obviously unrealistic to assume that every consumer can tell what the actual quality level is that is supplied by every seller. While this assumption would be sufficient (given appropriate supply-side conditions) to secure an optimal competitive equilibrium, it may not be necessary

and it may especially not be necessary for the equilibrium to be at a moderately good level (if not at an absolutely perfect one).

Who Knows What High Quality Can Do?

It is beyond commonplace to remark that defining and measuring the quality of medical care is difficult. The next remark is usually on an offsetting optimistic note, that some definitions are possible and some measures feasible. The whole edifice of quality assurance is built on the foundation of knowledge of some dimensions of quality, and true optimists even imagine that such measures can be constructed that will be useful to providers of medical care or health insurance linked to specific delivery systems. But is a little knowledge a good thing, or can it be dangerous?

It is easy to slide to either extreme, unreasonable optimism or unalloyed pessimism, when it comes to measuring the quality of care provided by different medical providers. Here I will simply assume that:

- a.) some indicators or signals of quality are possible to design and feasible to measure;
- b.) some agents have better levels of information about these measures than others;
- c.) No one has good enough measures to offer guarantees.

Given these assumptions, the fundamental behavioral implication is that better information can lead to better odds on the bets that all patients take when they subject themselves to medical care.

I will assume here that:

- a.) consumers as a group are not totally confused; they do have or can have valid information about quality;
- b.) some consumers will be better informed than others;
- c.) There will be some “experts” who will be better informed than the best-informed consumers.

We now ask how the equilibrium level of quality is affected by consumer information in competitive and non-competitive markets, and whether the information experts might have could be beneficially used in either equilibrium.

In competitive markets with some but not all consumers well-informed, the equilibrium can be close to that with perfect information as long as there are “enough” consumers well-informed about the prices and quality levels of different sellers. Competition forces price for a given level of quality closer to near-breakeven levels, and equivalently, competition leads to providers who nearly maximize quality for a given price. “Arms race” provision of excessive quality is deterred, since suppliers are pressured by the informed buyers not to offer low-value quality, and even poorly informed buyers need only to pick out one provider who supplies the quality they prefer, given price, confident that other providers will not offer much superior levels.

At the other extreme, if there are almost no informed buyers, competition will still lead to zero profits, but firm-level demand curves can still have negative slopes. Providers can survive by

offering quality worth less than the cost. If there are fixed costs at the firm level, there can be too many, too small firms.

But what does not come out of this model is an equilibrium in which quality is low and cheaply preventable errors survive, as long as buyers are attracted at all by lower prices or higher quality. What can happen, however, is that errors whose cost of remediation (though positive) is small may not get fixed if the new business that would be attracted is small relative to the cost of correction. The effect, as noted above, in a world of free entry is not high profits for the slothful providers, but either small size and too many such providers.

Not all information has to be given by some third party to buyers. Sellers also tell buyers about their quality and that of their competitors. But, anyone would say, sellers have an incentive to be selective in how they present this information to make themselves look good. Might that not do more harm than good if buyers are not super-experts?

Of course the answer is yes, but there are some reasons not to be discouraged. Assume away fraudulent information, but suppose that sellers provide incomplete if truthful information. My closest hospital tells me it is number 1 in cancer care but does not tell me it is number 18 in heart disease care. However, if I as a consumer am aware of the games sellers play, I should (unless I was born yesterday) be “skeptical” of this information. I should adjust my interpretation of the data being presented by my own suspicions about how hospitals are choosing what to present. In this example, I should interpret the absence of information from my closest hospital about heart disease care to indicate that it is probably not number 1 (or it would have said so). For the

moment I can only give it the average value of all hospitals except number 1. But then the true number 2 hospital will have an incentive to say so—being number 2 is not that great, but it is at least better than all the rest of the pack. In theory the remaining “conspiracy of silence” will unravel, until the only hospital in town left that is silent will be judged by skeptical consumers to be the worst.

Even efforts to manipulate outcome indicators by selecting for treatment only those patients who are not as sick as they look—behavior for which there is some strong evidence—will disappear if consumers are able to tell which providers are being selective (i.e., which are turning away patients). In choosing between two providers with good report card values, the skeptical consumer will want to avoid the one who has turned away many referrals in preference to the one who takes all comers.

Is there any way to have an equilibrium in a monopolistically competitive market if there are no fixed costs? If there is uniform quality, the answer is clearly no. There is no way to bring price into equality with average cost. With quality variable, this result would still seem to hold. Thus an indirect test for inefficient quality would be to look at the number of separate providers. For example, more consumer information should be associated with fewer firms, though still enough to have competition.

Now we change the model to assume that barriers to entry of some (unspecified) sort cause there to be firm-level monopoly. Consider the case of a single firm. If it faces informed buyers who desire different levels of quality, and if it must choose a single level of quality, it may choose a

level high enough and costly enough to drive the low value buyers out of the market. (Perloff and Carlton, 2000). In contrast, with the same demand curve in the absence of monopoly, either the single level of quality will be lower, or multiple levels will be provided, possibly with a lower average level of quality.

In short, with reasonably well informed buyers, quality can be higher under monopoly than amongst competitors. The intuitive reason is that the monopolist may gain more by monopoly pricing high quality than by monopoly pricing low quality.

What if consumers are imperfectly informed and face a monopolist? The outcome here might then be lower (actual) quality (even if consumers perceive it has acceptable quality), and very high prices.

Incentives for quality in competitive markets

It has become fashionable to lay the blame for low quality in part on the payment system. The fee-for-service administered price system offers the same price to all providers regardless of their quality, and costly activities to improve quality either have no effect on reimbursements or, if the higher quality reduces volume, there may be even negative effects. An obvious response is that creation of incentives for quality regardless require efficient payment for quality. Brent James has given an example in a monopoly setting in southern Utah.

But imposed financial incentives may be wholly unnecessary in competitive markets. The reason is simple: if buyers can detect higher quality and prefer sellers with higher quality, they will reward such sellers with greater volume. If reimbursement is above marginal cost, this higher volume turns directly into higher profits. Higher quality, in a sense, generates its own reward. Of course, in unregulated markets, firms facing demands from consumers who place different values on quality will in equilibrium charge higher prices for higher quality: The Rosen Equilibrium. But the key point is that this is the equilibrium with competition; it need not be and usually is not imposed by any outside source. Moreover, in this equilibrium those lower-quality suppliers whose quality is not worth what its costs, savings will be disciplined by low quantity, not by lower prices.

To what extent can this efficient “invisible hand” equilibrium be reproduced by regulatory or quasi-regulatory intervention encouraged by governments or non-profit foundations, intended to “incentivize quality”? One key issue here as always is the number of sellers. For one thing, if there is competition, there would be no need for an external entity to try to solve the difficult question of how best to reward quality. So long as prices an insurer pays are initially correct, and patients can know about quality, the higher volumes produced by the migration of informed patients should be enough to reward quality. Intervention is not necessary in markets where providers are competitive.

So interventions might not be needed. But could they do any harm? If the market is reasonably competitive, errors in setting the reimbursement for quality will tend to be self-correcting. In contrast, if there is a monopolist or a few firms, errors will be self-reinforcing.

To see why, suppose that the market was in a Rosen Equilibrium with price variation across multiple quality levels reflective of the marginal cost of quality at the quantity informed buyers would choose. Suppose that a foundation-funded project exactly reproduced this set of relative prices at various quality levels. If there is competition, firms will be eager to supply the intended quality level at the prespecified price. The competitive equilibrium will be reproduced. The foundation funding will not have done any good, but it will not have caused harm, either.

Now let us modify this story in three different ways. We introduce monopoly; we retain large numbers of suppliers but get the relative prices wrong; or we leave correct prices and many sellers, but buyers are poorly informed. If we have small numbers of sellers, those sellers will want to earn profits by skimping on quality. It will be an eternal struggle to get them to hit the desired quality. There will be scandals, recriminations about under-payment, and perpetual instability. Next, suppose that there are many sellers but the improved price-quality schedule is “wrong” in the sense that the incremental price for improvements in quality understates the cost. Just as before, quality will be undersupplied. If the price is set too high, quality will be oversupplied. Or, if the marginal price of quality varies with quantity, there may be too many or too few sellers, and corresponding situations of shortage and surplus. Finally, let us suppose that consumers cannot tell what quality level is being provided. Then the target quality will not be sufficient to attract sellers. Firms will be tempted to modify the types of quality they provide to those types that consumers can detect and therefore value. Again there will be tension among patients, providers, and regulators.

To sum up: paying for performance will be a recipe for frustration if price, quality, and quantity are all specified in advance. The less aggressive strategy of determining how much some increment in quality is worth to the buyer, and announcing that to a set of competitive sellers, has some rationale. It cannot do any harm, and if some sellers accept the offer, buyers are probably better off. But this approach cannot optimize the situation because it may lead either to underpayment or overpayment of suppliers: it is likely to lead to too little supply or too much (or to unintended losses or excess profits). The decentralized competitive equilibrium is better in theory if it can be accomplished in practice.

References

Carlton, Dennis W., and Perloff, Jeffery M. (2000). *Modern Industrial Economics*, 3rd Edition. Addison Wesley: Harlow.