
The Monopoly Nonproblem

Taking Price Discrimination Seriously

— ◆ —

STEPHEN SHMANSKE

One of the most important distinctions that economists teach is that between competition and monopoly, or, more accurately, that between price takers and price searchers. The distinction is pervasive. It shows up in the simple diagrammatics of profit maximization, in the welfare economics of the deadweight loss to monopoly, in the monopolistic competition model, in classic duopoly, in the arithmetic of the value of the marginal product and the marginal revenue product, in partly disguised form in monopsony models, in basic theoretical prescriptions for industrial policy and antitrust, not to mention on examinations at every level from principles to the Ph.D. comprehensives. An understanding of the distinction is an absolute requirement for students at all levels.

Economists themselves sometimes tend to be smug in their understanding and application of the distinction. They always present the simpler model of price-taking behavior first. By comparison, they regard the more complex model of monopoly, or the price searcher, as an improvement. When studying firm behavior, economists make the leap from models with exogenous prices to models with endogenous price setting, and they look with jaded eyes upon many of the models in the management

Stephen Shmanske is a professor of economics at California State University, East Bay.

The Independent Review, v. X, n. 3, Winter 2005, ISSN 1086-1653, Copyright © 2005, pp. 337–350.

sciences because those models treat the quantity demanded or the price as constant in order to examine some other dimension of decision making.¹ Allowing price to vary along a downward-sloping demand curve represents a huge improvement over fixed-price models.

Notwithstanding the improvement that a consideration of price searching brings to the analysis, economists should not smugly accept the standard presentation of monopoly because the price-searching model itself is saddled with a hidden, simplifying assumption that removes all relevance from the model: the assumption that all units must be sold at the same price. This uniform-price assumption is untenable on theoretical grounds and unrealistic on empirical grounds. More important, the uniform-price assumption is not benign with respect to its effect on the conclusions drawn from the model. We have the worst of all possible worlds: a model with an unrealistic hidden or simplifying assumption that yields implications for policy that are not robust when the assumption is relaxed. The best that can be said for the monopoly model is that it is better than the price-taker models because it considers a downward-sloping demand curve. Still, the standard price-searcher model is woefully inadequate.

The remedy for this inadequacy is to treat nonlinear pricing and price discrimination much more centrally and seriously from the outset. Textbooks generally do not do so currently. Most textbook treatments suggest that the market structures worthy of consideration include price-taking competition and uniform-price monopoly; monopolistic competition and classic oligopoly receive less attention. Price discrimination, when it comes up, is most often treated in an offhand manner as an extension of the monopoly model. Indeed, in my own casual survey of twelve principles texts and ten intermediate theory texts, I found that the proportion of the coverage of the monopoly model that is devoted to price discrimination averages only 23 percent.²

The current emphasis is all wrong. In the next section, I review the current treatment, arguing that economists spend too much time juxtaposing price taking and price searching as if they were the only relevant forms of seller behavior. Economists even go so far as to present policy prescriptions based on market structure as depicted in the textbooks and the simple application of these models. In this approach, price discrimination is an afterthought, an interesting but computationally and diagrammatically complex extension that in some instances distorts the market equilibrium even more than does the nondiscriminating monopolist. Price discrimination may even seem to be interesting only occasionally because of the list of preconditions that must be met for the monopolist to be able to charge different prices for different units.

1. For example, queuing theory uses an “arrival rate,” which is akin to a demand, but the theory typically specifies the arrival rate as a constant and never considers whether it is influenced by price. In inventory-control theory’s models of economic-order quantity and economic-production quantity, both price and quantity are held constant. In the “stepping stone” transportation model, demand quantities, production quantities, and prices are held constant. In virtually every linear-programming example, prices are held constant. Input-output models also hold prices constant.

2. Details available from the author upon request at stephen.shmanske@csueastbay.edu.

I then describe what needs to be done to change this emphasis. Economists should give more attention—in both research and textbook presentations—to the costs and benefits of price discrimination, the empirical measurement of price discrimination, and the strategic theory of price discrimination with multiple sellers. The professional literature in these areas is still in its infancy. Perhaps a humbler approach in textbooks, one that details the deficiencies of the distinction between price taker and price searcher and that points the student in the direction of learning more about what we do not currently know, will encourage future generations of students and economists to tackle these issues. However, this changed emphasis will not come about until economists repudiate for the most part the uniform-price monopoly model.

Critique of the Dominant Paradigm

Competition/Monopoly or Price Taker/Price Searcher

The authors of virtually every economics textbook at the principles and intermediate levels expend a great deal of effort in introducing and distinguishing between monopoly and competition. What most students probably understand intuitively becomes needlessly tortured as the literal “single seller” notion of monopoly brings up quibbles about whether brand-name differentiation confers monopoly or whether the local geographical monopoly of the corner grocery store should be the target of interventionist policy. The conclusion from these quibbles is that every seller is either a monopolist or not, depending on how narrowly or broadly one defines the relevant market. No doubt a clear difference exists between a single-seller public utility and the multiple sellers in any of a number of obviously atomistic markets. But what is the essence of the distinction?

This question has two answers, and, in my opinion, textbooks traditionally focus too much on the wrong one. By and large, textbooks concentrate on differences between price takers and price searchers. The right answer, however, has to do with the distinction between closed markets and open markets, not simply with the number of sellers.³ Monopoly status is a privilege granted by governments that allows privileged firms to be free from competition. The actual number of sellers in a closed market may be one, as in the case of a local cable franchise, or many, as in the case of taxicab medallion holders, but the restriction on entry is what consumers naturally recognize as restricting their choices and causing the “monopoly” problem. When markets are open, consumers naturally gain from the lower prices or better qualities that entrants are free to offer. In this view, governments are the cause of the monopoly problem and laissez-faire government

3. George Stigler (1968) points out the distinction between competition among large numbers of sellers, which he calls “market competition,” and competition owing to open markets, which he calls “industrial competition.”

policy is the appropriate prescription. Quite simply, closed markets restrict consumer choice, and open markets allow innovations that lower prices, raise quality, or introduce new products. This issue, however, is not my theme in the present article.

The open-market/closed-market approach to the monopoly problem usually gets only a fraction of the attention that the standard paradigm gives to a distinction between competition and monopoly that focuses on the slope of the demand curve—namely, the price-taker/price-searcher distinction. According to this view, in competitive markets the market sets the price, and the sellers take the price as a given in their decision making. The individual firm is then painstakingly analyzed as the language of marginal or incremental reasoning is developed and the diagram of the profit-maximizing, price-taking firm is presented and explained. Following this instruction, the relationship of price takers to the market is described and theorems are developed with respect to profits and losses, entry and exit, and Pareto optimality.

Alternatively, monopolists face the whole market demand curve, which by the first law of demand must be downward sloping. Such a seller is a price searcher and must work with a trade-off between price and quantity. After doing a little arithmetic and developing the familiar monopoly diagram, the economist derives the price-searching equilibrium conditions and derides the monopolist for restricting output in order to raise price and profit. The wielder of the standard model then delights in claiming that neither the high price nor the high profits are the real problem (they are simply transfers from consumers to producers), which is that the restricted quantity distorts resource allocation into a non-Pareto-optimal configuration. Only a little reflection is required to realize that neither monopoly nor competition drives this result; rather, the result hinges on whether the demand curve facing the firm is horizontal or downward sloping. At this point, students have moved beyond the definitionally ambiguous “competition versus monopoly” distinction to the language of price searching and price taking.

Price Taker versus Price Searcher

The teaching agenda just outlined has good and bad aspects. On the good side is the lesson involving the effect of profits on entry in the competitive models. This result seems to be extremely robust over a variety of settings in addition to price-taking pure competition. Another important lesson resides in the marginal or incremental reasoning used to describe the price taker’s optimizing decisions. The lesson of incremental reasoning pays off again and again in a variety of economic models. The illustration of what Pareto-optimal resource allocation amounts to is also a valuable lesson. The establishment of the price-taking equilibrium, with its individual optimization by buyers and sellers intermediated by market prices à la Adam Smith’s invisible hand, is an important lesson in political economy.

Such instruction is time well spent because this relatively simple material is a necessary first step toward understanding markets, maximization, and equilibrium. It soon becomes obvious, however, that no person or entity called the “market” actually “sets” the price. Sellers have to do so for themselves. Furthermore, a little reflection on the real world indicates that no seller would lose all his customers if he raised his price by a penny. Logically, sellers do not take price as given. They have to set the price themselves, and they have a choice—that is, all sellers are price searchers.

Unfortunately for fans of markets and private enterprise, along with the added realism of the price-searcher equilibrium comes the unhappy news that the equilibrium is no longer socially optimal. Value is greater than cost on the margin, and not enough resources are devoted to making the price searcher’s product. If all sellers are price searchers, the monopoly “problem” is rampant throughout the economy. The door is then open for government regulation, and the interventionists are happy to enter. Economists obligingly show how the monopoly problem is related to the elasticity of demand, and they develop models that link elasticity to the number of firms or the distribution of market shares. Current U.S. antitrust policy, which focuses on the definitions of the relevant market and the distribution of market shares within that market, seems motivated by this view.

Price Discrimination to the Rescue, or Not?

The deadweight loss to monopoly arises because the marginal-revenue curve lies below the demand curve: when the seller expands the quantity sold by moving down the demand curve, he must also lower the price on all the preceding units, in accordance with the hidden assumption that all units must be sold at the same price. But why *must* he lower the price on the preceding units? Does anyone sell every unit at the same price? It is difficult to think of anyone who does so.⁴ The uniform-price assumption is neither realistic nor rational. After all, the seller would maximize profit by selling each successive unit at the highest price possible without lowering the price on the preceding units. The clever diagram of the simple monopoly problem notwithstanding, it represents neither what real-world price searchers do nor what they are attempting to do nor what they should be attempting to do. The price searcher should be attempting to practice “perfect price discrimination”—a goal that, if reached, results in a Pareto-optimal allocation of resources.

Textbooks, however, do not present the perfect-discrimination model until after exhausting the uniform-price model completely. This approach has both a conscious and a subconscious effect of making students suppose that something is more natu-

4. A student of mine suggested that the California State Lottery Commission sells every ticket for one dollar. If so, the commission is marketing the lottery poorly. It is trivial to show that quantity discounts should be used. Even my church will sell you six raffle tickets for the price of five.

ral about the uniform-price monopolist and the associated problems than about the discriminating monopolist, notwithstanding the Pareto optimality associated with the latter's successful operation.⁵ By this time the students (not to mention the professor) have too much invested in the model to scrap it completely. When price discrimination does come up, it seems to come up as an afterthought (which historically it was) and something to be discounted, shoved aside, or treated as an interesting but unimportant special case. Indeed, the standard treatment often starts with a definition that tries to marginalize the relevance of price discrimination. The standard treatment further circumscribes the applicability of price discrimination by presenting a set of (difficult?) conditions that must be met if price discrimination is to be practiced.

In defining price discrimination, many texts make a point of arguing that not all price differences are price discrimination. As Lott and Roberts (1991) argue, serving different customers may entail different costs, and even in atomistic markets long-run competition will allow—indeed, may even require—the prices to differ in such cases. Therefore, only price differences disproportionate to underlying cost differences truly exemplify price discrimination.⁶ The implication to be drawn from this definition is that price discrimination is less important than it seems in casual observation. Why would textbook writers want to give this impression? At stake is not really whether price discrimination occurs, but whether there exists the underlying monopoly power evidenced by price-searching behavior along a downward-sloping demand curve. Indeed, even in the face of the many real-world examples of price differentials, the world can be seen as close (enough?) to the competitive ideal if many of those price differentials are not indicative of true price discrimination.⁷

The distinction between cost-based price differences and price discrimination has always struck me as a rhetorical device meant to diminish the ubiquity of the monopolist/price-searcher effect with its concomitant implication of resource misal-

5. Students are not the only ones affected. Chiang and Spatt (1982) analyze imperfect price discrimination by a monopolist and conclude that the welfare effects are ambiguous when compared with uniform pricing. Implicit in their approach is that uniform pricing is taken as the baseline case. Therefore, in their view, the jury is still out on whether to allow or to prohibit price discrimination. They might just as easily (and in my view more properly) have said that compared to a baseline case of an unregulated world with realistic price discrimination, a law mandating uniform pricing has ambiguous welfare effects.

6. In fact, the exact definition remains a contentious issue; see Clerides 2004.

7. Levine (2002) argues that price differences may occur as competitive (in the sense of open markets) firms seek to cover costs jointly incurred across units of the same product or across product lines. He recognizes that whether this phenomenon is price discrimination by some formal definition is partly a semantic matter, and he, like Lott and Roberts, is concerned for the most part with debunking the view that mere price differences evince underlying monopoly power. But Levine sees the problem of the underlying market power as the ability to earn monopoly rents by manipulating price and quantity along a downward-sloping demand curve, and he argues that the price differences can exist without the rents. He calls (erroneously in my view) such firms price takers because there is only one price (or set of prices) consistent with earning nonnegative profits. His firms, however, are no more price takers than are the firms in the textbook zero-profit (uniform price), monopolistic competition model. My argument is that the price differences are obvious indications that the demand curve is downward sloping and that the firm has market power. However, the monopoly "problem" of Pareto-relevant underproduction comes from the uniform-price assumption, which is demonstrated to be false by the very existence of the price differences.

location. The insinuation that monopoly power is not behind price differences perhaps makes some people more comfortable with free-market policy. Rhetoric aside, not even the logic of this insinuation is solid because it opens the door for instances of uniform pricing actually to be price discrimination, a point that not all texts are careful to make. Ultimately, however, all of this is beside the point. Consider, for example, perfect price discrimination, which has little or nothing to do with the cost structure. Perfect price discrimination comes from charging prices that trace the actual demand curve out to the point at which it meets the marginal cost curve. Whatever the structure of inframarginal costs and their relationship to demand prices, one would still call it perfect price discrimination if the seller had worked his way down the demand curve point by point.⁸

Closely following the definition in most textbooks comes a list of necessary conditions for the practice of price discrimination. Again, the rhetoric seems designed to limit the scope or applicability of price discrimination because price discrimination is impossible if the necessary conditions are not all satisfied. Usually three conditions are stated. First, the seller has to have market power—a downward-sloping demand curve. Indeed, in a perfect competition, price-taking setting, the seller has no power to raise the price above the given level and therefore cannot price discriminate.⁹ Second, the seller has to be aware of differences in demand prices in order to charge different prices on different units. Presumably, knowing that one buyer will pay \$12 and another only \$10 does the seller no good if he does not know which buyer is which. Third, the seller must be able to prevent the resale of low-priced units to other customers who are willing to pay higher prices. If the seller does not prevent such resale, he will quickly lose all of his higher-price market segments and make sales only at the lowest price if such arbitrage is feasible and costless.

One has to wonder why these conditions are even enumerated considering that they are met with trivial ease. Every seller is a price searcher. Some face more elastic demands than others, but this difference means only that some sellers' price discrimination will result in a lower-variance structure of prices relative to the prices that other price discriminators set. Sellers do not have to know the identity of each consumer because the consumers will separate themselves into submarkets based on numerous clever hurdles and restrictions that sellers employ, from trading stamps, coupons, and contests to advance reservations, tie-in sales, and time-of-day or day-of-week pricing

8. At the risk of belaboring the point, I ask readers to consider just one more angle. Suppose the marginal-cost curve exactly mimicked the demand-price schedule. If the seller charged prices equal to the demand prices (and to the incremental costs) for each unit and therefore just broke even, would not this case be characterized as one in which perfect price discrimination was necessary for the firm to exist? Would anyone other than Lott and Roberts call this a case with no price discrimination because prices were proportional to costs? Would anyone other than Levine call this a case of price-taking behavior because any other set of prices would lead to negative profits?

9. A quibble might be added here. An inframarginal price-taking firm can price discriminate by charging some customers less than the price set by the market. Doing so does not raise profits, but it is a case of setting different prices on different units without corresponding differences in those units' costs.

or to quantity discounts, quality variations, and even random price dispersion. Finally, prevention of resale is automatic for any good with a service component and for most retail goods owing to the small potential profit from arbitrage and the magnitude of transactions costs. These conditions can always be met, at least in part, and therefore every seller practices price discrimination.

It would seem more appropriate to attempt to list the conditions under which no price discrimination occurs: perfectly open markets, perfect information about prices on the part of consumers, no transactions costs, and a constant-cost industry structure—that is, precisely the conditions required for perfect or atomistic competition. The conclusion should be obvious: these conditions are never met, profit seekers never use uniform pricing, and therefore all pricing entails some form of price discrimination.

After the definition of price discrimination and the discussion of its necessary conditions, the standard treatment moves from the theoretic to the purely descriptive by labeling three types or degrees of price discrimination. First-degree, or perfect, price discrimination is introduced and dismissed as irrelevant in almost one breath. Perfect price discrimination is consistent with Pareto optimality, but because it requires perfect information to be carried out perfectly, it is rejected as inapplicable to the real world. This quick rejection is a mistake and evinces a fundamental misunderstanding of the model's meaning—a point to which I return in the next subsection.

After dismissing perfect price discrimination, textbook writers present many realistic examples of imperfect price discrimination, both second and third degree. Again, however, the message is decidedly mixed. If imperfect price discrimination leads to a greater output than uniform pricing, then the equilibrium allocation of resources is improved, moving closer to that described as Pareto optimal. Such a change might often, or even usually, occur.¹⁰ Nevertheless, authors give more attention to a specific example of third-degree price discrimination in which an additional deviation from Pareto optimality occurs. Indeed, in some cases where one market with a linear demand curve in the relevant region can be broken into two submarkets, each with a linear demand curve, total output is not affected. Because the two submarkets have different prices, however, the consumers in those markets have different marginal values, so now in addition to a misallocation of resources to producing the product, we also have a Pareto-inefficient distribution of the product. Instead of making society better off, this type of price discrimination makes society even worse off than the uniform-price monopolist's pricing does.

On the whole, the textbooks create the following impression. Competitive price-taker markets are fine, but price-searching monopolists restrict output in order to raise price and profit. Perfect price discrimination would restore optimality, but it is irrelevant because of its unrealizable information requirement. Not all price differ-


10. Quantity discounts, declining block pricing, two-part pricing, and many instances of third-degree market segmentation all lead to greater output than does uniform pricing.

ences reflect price discrimination, but when a number of conditions are met, price discrimination is possible, and it may make society even worse off.

Taking Price Discrimination Seriously

A reexamination of the dismissal of perfect price discrimination can serve as a wedge into an improved understanding of price discrimination. Consider the following metaphor, supported by figure 1, which shows a market with constant marginal cost of production for the good. Point A is the competitive-equilibrium point if the costs are given by MCA . Point A also implies a Pareto-optimal allocation of resources to this industry. Suppose, however, that because of a hidden or simplifying assumption, MCA is drawn with the assumption that oil, a necessary input, is free. The real world, in which oil has a positive cost, can reach equilibrium only at point B on the demand curve, given that MCB includes the cost of oil. In this case, no one claims that a market failure exists. The real-world equilibrium cannot match the Pareto optimum calculated in a model that leaves out a real-world cost, nor should it.

Why then do economists claim market failure in the case of perfect price discrimination? Consider point A again. This time MCA includes all of the production costs but does not include the costs of the information required to price discriminate perfectly. The perfectly discriminating monopolist would work his way down the demand



curve, eventually reaching point A as an equilibrium and Pareto optimum. In the real world, however, no one can imagine reaching point A because of the information cost and the cost of price discrimination, so this potential solution is discarded. The market does not reach point A, and economists mistakenly suggest that a deadweight loss is incurred. In the case in which oil is mistakenly assumed to be free, no one is deluded into thinking that point A is the desired target, and no one claims that a market failure occurs because the market does not settle at that point. However, in the case of price discrimination, economists either implicitly or explicitly maintain the belief that A is the right target and that a market failure occurs because A is not reached.

I believe this to be a fundamental misunderstanding of the model. The perfect-information assumption is adopted to simplify the model. The model shows that in a world with perfect information, a monopolist or price searcher would attempt to sell each unit for a successively lower price, working down the demand curve to point A, the point of Pareto-optimal allocation. In the free-information setting, the model does show that the market's equilibrium has no defect. What the seller is trying to do is consistent with the optimum. Note that this statement cannot be made for the uniform-price monopolist. Once the added assumption is made that all units must sell for the same price (perhaps to simplify the model), then a market failure is alleged: a Pareto-relevant distortion of resource allocation. What the seller is trying to do is inconsistent with the optimum. This distortion is mistakenly blamed on the price-searching firm, and policies may be proposed and implemented to correct the alleged market failure. But this situation is not a market failure at all. The "failure" arises because of the unrealistic and constraining assumption of a uniform price, an assumption that has little, if any, real-world relevance.

It is a major misplacement of emphasis to claim over and over again that the price-searching monopolist distorts resource allocation and that the perfectly discriminating monopolist fails because of information costs. The correct emphasis is that the monopoly distortion as commonly understood is a mistake that springs from a hidden assumption that should be recognized and discarded. The model with perfect information works perfectly, just as perfect competition works perfectly. The main point of each of these claims is to demonstrate that no built-in failure comes from combining selfish, atomistic consumer behavior with selfish, profit-maximizing seller behavior, even when the seller is a monopolist or price searcher.

What remains is to take price discrimination seriously. Economists should be looking at the costs and benefits of price discrimination, just as they examine any other decision dimension they might analyze. Instead of rejecting the perfect-discrimination model as irrelevant and concluding that monopoly entails a deadweight loss, economists should realize that because perfect price discrimination is costly, its achievement is not desirable. This realization parallels the economists' common recognition that because perfect environmental purity is costly, its achievement is not optimal. The "economic" level of environmental purity is the amount beyond which

the marginal cost of additional purity exceeds the marginal benefit. So, too, the “economic” level of perfect price discrimination is that which equates costs and benefits of additions to price discrimination at the margin. Economists have completely neglected this task. Even the language that admits to a “level of perfect price discrimination” or to marginal “additions to price discrimination” is unfamiliar to the economist’s ear. This gauntlet is well worth picking up.

A Research Agenda

If price discrimination is to be taken seriously, several themes must be pursued vigorously at the research level and incorporated in the standard textbook treatments. I briefly discuss in this section a critical analysis of the uniform-price assumption, the pure theory of monopolistic price discrimination, the pure theory of price discrimination with oligopolistic interaction, and the applications of these theories in econometric work.

First, most textbooks need to be rewritten to be more critical of the uniform-price assumption. I am curious as to exactly what authors and students think of this assumption. Is it a simplifying one? After all, $TR = PQ$ (total revenue equals price times quantity) is simpler than $TR = \sum P_i$. It is simpler to write and simpler to analyze. Unfortunately, it is totally misleading. Does it seem natural somehow to assume that only one price can be charged? If so, how can it be so in the face of obvious non-uniform pricing? Is the assumption made just temporarily for the sake of argument? I suppose this option comes closest to the way I treat it when teaching why some textbooks draw the marginal revenue curve below the demand curve, but only a slim minority presents it this way. Some authors seem to quibble with phrases such as: “*if* or *usually* only one price can be charged.” But there is no explanation of the “if,” and stating that “usually” only one price can be charged has to be confusing to students when it is painfully obvious that almost never is only one price charged. By far the most common treatment, however, is simply to ignore the issue. Whether purposefully or not, most authors give the impression that it is normal or natural to be constrained to a single price—a presumption that is simply wrong.

The pure theory of monopoly price discrimination must go beyond the categorization of types or degrees of price discrimination, with the associated examples pioneered by Pigou (1920) and Robinson (1933), to an examination of the costs and benefits of price discrimination. The starting point for this research might be two polar cases. In the first case (no price discrimination), we essentially argue that uniform pricing is the only feasible alternative because price discrimination is infinitely costly. At the opposite pole is the case in which the seller practices perfect price discrimination without cost. Neither of these poles usefully depicts the real world. Chiang and Spatt (1982) made significant strides by analyzing the monopolist’s ability to separate consumers into self-selected groups, each with a different price function. Yet

this analysis is still not a cost-benefit analysis. Without cost to himself, the monopolist knows some things about the structure of demand, but he has infinite costs in finding out other things in order to discriminate more perfectly. Analyzing this so-called incentive compatibility on the part of the consumers is a major step, but more work remains. Few economists, if any, focus on the costs of achieving more perfect price discrimination. I have yet to see a textbook whose author suggests that government policy might attempt to lower the costs of price discrimination. By contrast, writers give ample coverage to antitrust efforts aimed at prohibiting price discrimination, monopoly, conspiracy, monopolization, or incipency.

A monopoly model of the costs and benefits of price discrimination will be only a first step because all firms exist in markets with competition from sellers of substitutes. We need a model of equilibrium price discrimination with multiple sellers, as argued at least since Wright (1965) made some of the observations included in the present article but whose insights have still not entered the mainstream. Several attempts have been made, but much needs to be done.¹¹ Taken piecemeal, the various attempts to analyze price discrimination in multiple-firm settings deal with many of the important issues. Some examine the constraints that active consumer choice implies for a firm's pricing decision. Others consider the constraints implied by the existence of other firms, the timing of the resolution of demand uncertainty, and market entry. Whether the variety of results offered in these models will hold up in more general settings is an open question.

Empirical demand studies represent a final area in which the serious consideration of price discrimination holds the potential for fostering a great improvement in our knowledge. Consider the typical demand function, which states that quantity is a function of price and of numerous demand-shift variables that measure income, tastes, product quality, population, and the like. Even before we get to the usual problems of specifying the functional form and dealing with the simultaneity and identification issues, we face the basic problem of measuring price. Some studies use list prices, others modal prices, still others average revenue as the measure of price. The measurement problem is usually noted and dismissed as the research forges ahead. None of these measures, however, is the theoretically correct one, given the multiple-price strategies that firms choose. When complex pricing strategies are in use, the pricing function cannot be captured in one dimension. For example, two-part pricing requires two prices, as does third-degree price discrimination with two market segments. Meanwhile, two-block pricing requires two prices and a quantity threshold at which the lower price kicks in. Though improvements, each of these approaches remains inadequate to capture the complexity of pricing in the real world.

Until recently (see Shmanske 1998) no one attempted to tackle the problem. Price discrimination was considered in applied models at least as early as Oi's 1971

11. See Spulber 1979; Norman 1981; Oren, Smith, and Wilson 1983; Katz 1984; Borenstein 1985; Shmanske 1991a, 1991b; Dana 1998; and Levine 2002.

analysis of Disneyland pricing, but the early studies were descriptive, institutional case studies, with little in the way of econometric modeling or statistical hypothesis testing. My study of almost fifty individual golf courses put together a data set rich enough to explore the types and extent of price discrimination used in an industry of competing firms. As it turns out, econometric models that use only a single measure of price as an explanatory variable are underspecified. Their parameter estimates suffer from an omitted-variable bias. My 1998 article shows how the addition of a variable to capture price discrimination (several variables were explored) improves the results.

Much work unfortunately remains to be done, a great deal of it being the tedious collection of usable microdata on prices and quantities. Accomplishing this job will have a significant payoff, however, because nearly the whole received body of empirical demand studies currently suffers from an omitted-variable bias owing to its avoidance of the price-discrimination issue.

Several different problems in demand studies may be affected by the omission of a price-discrimination variable. Consider just one, from the sports economics literature. In attendance studies in professional sports, it is usual to include price as an explanatory variable in a regression equation with attendance on the left-hand side. After an elasticity has been calculated from the regression coefficient, a puzzle often arises. The price searcher should never price in the inelastic portion of a demand curve, yet the elasticity implied in the regression is less than one. Consideration of price discrimination might explain this puzzle in two ways.¹² First, as a matter of pure theory, the result that the monopolist will not choose a price at which demand is inelastic comes from the uniform-price monopoly model. With price discrimination, all such bets are off. Second, as an empirical matter, omitting a price-discrimination variable from the right-hand side of the equation biases the coefficient of whatever price variable is used, calling into question the relevance of the estimated coefficient.

References

- Borenstein, S. 1985. Price Discrimination in Free Entry Markets. *Rand Journal of Economics* 16 (autumn): 380–97.
- Chiang, Raymond, and Chester S. Spatt. 1982. Imperfect Price Discrimination and Welfare. *Review of Economic Studies* 49 (April): 153–81.
- Clerides, Sofranis K. 2004. Price Discrimination with Differentiated Products: Definition and Identification. *Economic Inquiry* 42 (July): 402–12.
- Dana, James D., Jr. 1998. Advance-Purchase Discounts and Price Discrimination in Competitive Markets. *Journal of Political Economy* 106 (April): 395–422.
- Fort, Rodney. 2004. Inelastic Sports Pricing. *Managerial and Decision Economics* 25 (March): 87–94.

12. Fort (2004) explains the puzzle by considering additional revenue streams, such as parking charges and concession sales. Such revenue-generating actions are tantamount to price discrimination using market segmentation via quality differentiation and tie-in sales.

- Katz, Michael L. 1984. Price Discrimination and Monopolistic Competition. *Econometrica* 56 (November): 1453–71.
- Levine, Michael. 2002. Price Discrimination Without Market Power. *Yale Journal on Regulation* 19 (winter): 1–36.
- Lott, John R., Jr., and Russell D. Roberts. 1991. A Guide to the Pitfalls of Identifying Price Discrimination. *Economic Inquiry* 29 (January): 14–23.
- Norman, George. 1981. Spatial Competition and Spatial Price Discrimination. *Review of Economic Studies* 48 (January): 97–111.
- Oi, Walter. 1971. A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly. *Quarterly Journal of Economics* 85 (February): 77–96.
- Oren, S., S. A. Smith, and R. B. Wilson. 1983. Competitive, Nonlinear Tariffs. *Journal of Economic Theory* 29: 49–71.
- Pigou, A. C. 1920. *The Economics of Welfare*. London: Macmillan.
- Robinson, Joan. 1933. *Economics of Imperfect Competition*. London: MacMillan.
- Shmanske, Stephen. 1991a. Price Discrimination and Monopolistic Competition. *Studies in Economic Analysis* 14 (spring): 25–48.
- . 1991b. Two-Part Pricing in Differentiated Duopoly. *Journal of Institutional and Theoretical Economics* 147 (September): 441–58.
- . 1998. Price Discrimination at the Links. *Contemporary Economic Policy* 16 (July): 368–78.
- Spulber, Daniel F. 1979. Non-cooperative Equilibrium with Price Discriminating Firms. *Economics Letters* 4: 221–27.
- Stigler, George. 1968. *The Organization of Industry*. Homewood, Ill.: Richard D. Irwin.
- Wright, J. F. 1965. Some Reflections on the Place of Discrimination in the Theory of Monopolistic Competition. *Oxford Economic Papers* 17 (July): 175–87.

Acknowledgments: Previous versions of this article were presented at the meetings of the Association of Private Enterprise Education, April 2003, in Las Vegas and at the Workshop in Economic Research, California State University, East Bay.