Price Discrimination and Irrational Consumers

Discussion of Armstrong and Ellison

2005 Econometric Society World Meeting

University College, London

Kenneth Hendricks

University of Texas at Austin

February 4, 2006

1. Overview

I have the formidable task of commenting on two outstanding surveys of two important literatures. Armstrong (2005) examines developments in the theory of price discrimination and Ellison (2005) examines the use of bounded rationality in industrial organization. Most of the recent work on price discrimination is concerned with oligopoly markets. The main issues in Armstrong's survey are: does competition eliminate price discrimination and, if not, what is the impact of a ban on price discrimination? General results on these questions are difficult to obtain. Thus, one of the roles of theory is to classify the kinds of oligopoly markets where price discrimination is likely to occur, the form that it is likely to take, and impact that it is likely to have on profits and welfare. Armstrong describes the progress that has been made on developing this taxonomy. I will highlight what I view are the main insights and make a couple of comments on outstanding issues.

The theme of firms exploiting consumers is also present in Ellison's survey, which focuses primarily on irrational consumers. However, the main issues are methodological. The rational game-theoretic approach has dominated theoretical and empirical work in IO for the past twenty years. Indeed, the literature on bounded rationality in IO is so sparse that Ellison's survey is in part a discussion of the potential of bounded rationality approaches for the field. A leading theorist recently told me that he thought that IO economists are so heavily invested in the rational game-theoretic approach that we are unwilling to consider other approaches. We either ignore behavior that is inconsistent with the paradigm or we build complicated, rational game-theoretic models (typically incorporating private information) to explain it. Ellison's very thoughtful, and possibly prophetic, essay challenges the field to reexamine its traditional approach. I try to meet his challenge in my remarks.

2. Price Discrimination

Armstrong's basic approach is to use a simple, parameterized Hotelling model to present the main insights of the recent work and the intuition behind them. Two firms are located at either end of a unit interval. Marginal costs of production are zero. Consumers have unit demands, and their preferences are indexed by three parameters: v, the consumer's utility from consuming the unit; x, the consumer's location in the unit interval; and t, the transport cost per unit of distance. Here v can also be interpreted as a measure of quality; x is a measure of the consumer's relative preference for the product of the firm that is located at the left-end point, and t represents how much he cares about the difference between the products, what Armstrong calls his level of

"choosiness". The three parameters are independently distributed among consumers according some distribution.

One important result is that competition eliminates price discrimination on product dimensions where firms have no comparative advantage. In the Hotelling model, firms are symmetric in their ability to supply quality. Fixing t, and assuming x is not observable, this symmetry implies that the location of the marginal consumer at any pair of prices does not depend on v, and the equilibrium is uniform pricing (i.e., p = t). This result is a special case of the cost-plus-fixed fee result of Armstrong & Vickers (2001) and Rochet and Stole (2002). It suggests, for example, that as firms in the U.S. cell phone market become more symmetric in their capacity to provide quality service, their calling plans will simplify to two-part tariffs. Indeed, the two-part tariff was the standard tariff schedule in the U.S. during the 1980's when the cell phone markets were more local and arguably more symmetric.

Each firm has a comparative advantage supplying customers who prefer their product. Thus, when firms observe t or x, or a signal about them, they will charge personalized prices in much same way as a monopolist would. In these cases, the question is whether price discrimination increases equilibrium profits relative to uniform pricing. The answer turns out to depend upon what is observable. If firms observe t, the consumer's "choosiness" level, or signals about t, price discrimination softens competition, and leads to higher profits. Moreover, firms will want to share their information about consumers. On the other hand, if firms observe x, the consumer's location, or a signal about x, then price discrimination intensifies competition, profits fall, and firms will prefer to keep their information private.

What is the difference between these two cases? Corts (1998) shows that the key issue is whether the firms' best replies shift across market segments in the same direction (best response symmetry) or in opposite directions (best response asymmetry). When firms discriminate on t, or a signal about t, their best replies shift in the same direction, since the high t segment is a stronger market for both firms than the low t segment. By contrast, when firms discriminate on x, or a signal about x, their best replies shift in opposite directions, since each firm's strong market (i.e., consumers who prefer its product) is the other firm's weak market. Best response symmetry/asymmetry is clearly an important new theoretical tool for studying competition in differentiated markets.

The possibility that profits fall when firms practice price discrimination distinguishes oligopoly markets from monopoly markets. Mixed bundling by multi-product firms is another case discussed by Armstrong where discriminating firms are worse off. These cases raise the question of whether firms can commit not to engage in price discrimination. Armstrong surveys a

number of papers that model this issue as a two-stage game in which firms in effect choose the space of pricing strategies for the second stage. But, it is not always clear what actions firms can credibly and visibly take that make it prohibitively costly for them to charge personalized prices or bundled discounts. The commitment issue also arises in models of inter-temporal price discrimination. Despite the apparent absence of explicit commitment mechanisms, firms nevertheless frequently adopt uniform pricing rules when theory would imply price discrimination. For example, CD prices do not vary much across artists or over time. One possible mechanism is reputation but I am not aware of any papers that have explored the circumstances under which firms can build reputations for uniform pricing.

As Armstrong mentions, the development of scanner technologies and online markets has made it easier for firms to obtain information about consumer preferences from their purchase histories. These developments can enhance consumer welfare insofar as firms use this information to make consumers aware of other products. The worry, however, is that firms will use personal information to charge their customers higher prices. Armstrong's review of the literature on this issue shows, however, that price discrimination in oligopoly markets may mitigate this concern. Firms compete more aggressively when they can offer lower prices to their rivals' customers than their own customers. Furthermore, price discrimination across periods means that any rents earned by a firm on its customers in subsequent periods will be reflected in its initial offer. Thus, consumers may be able to earn back some or all of these rents when she becomes a customer of the firm. In fact, competition for new customers may force firms to commit to sharing customer information, forgoing the rents from private information. For example, credit agencies agree to share information on their consumers' credit histories. These results have important implications for antitrust policy on price discrimination.

3. Bounded Rationality

Ellison distinguishes three kinds of bounded rationality in IO models: agents who use simple rules to make choices; agents who find it costly to make decisions; and agents who exhibit behavioral biases. The first two are closely related since agents with high decision or information costs typically find it optimal to adopt simple rules-of-thumb. Both of these approaches have a long tradition in IO. The last approach is new, and originates in the recent work of psychologists and behavioral economists. Papers in this literature retain the maximizing calculus of the rational game theoretic models but replace the utility functions with objective functions that incorporate a behavioral bias. The models are then solved using standard game-theoretic solution concepts.

Most IO economists would probably view the rule-of-thumb approach as complementary to the rational, game-theoretic approach. For example, in Varian's (1980) model of sales, some consumers use a simple rule of randomly selecting one supplier. The rule is rationalized by assuming these consumers have high search costs, but the costs can easily represent other kinds of frictions. The important issue is the impact of these consumers on the equilibrium pricing behavior of firms. Firms charge different prices for the same good, and the price dispersion is substantial with the support of the price distribution including the monopoly price. The theme here is a familiar one: a small amount of friction at the individual level can have big impact on outcomes. Another example is reputation models. A small probability of an incumbent firm being unwilling to accommodate entry has a dramatic effect on entry decisions in a multi-market, sequential entry model; a small probability that a supplier is committed to providing high quality service has a dramatic effect on the behavior of suppliers who find it costly to provide high quality service. In each of these cases, the goal is to explain behavior that is not consistent with equilibrium in a fully rational, game-theoretic model. But, instead of modifying the solution concept (e.g., ε-equilibrium), the approach has been to incorporate the "irrational" behavior directly into the model in the form of player types and then solve the game using standard solution concepts. I view this form of bounded rationality as basically a perturbation of the rational, game-theoretic model.

The behavioral bias approach represents a more fundamental shift in modeling. Della Vigna and Melmendier's (2004, 2005) fascinating work on health club memberships illustrates the main issues. They assume that health club members have quasi-hyperbolic discounting preferences. Members incur a short-run disutility with each exercise visit and obtain a long-term reward in the form of better health. The hyperbolic discounting implies that consumers will prefer a higher visitation rate ex ante than they will actually choose ex post. D&M show that the monopolist's optimal two-part tariff contract for these kinds of consumers involves a per-usage fee that is below marginal cost. The low fee helps consumers overcome their self-control problems and increases their willingness to pay for the membership, which the club extracts with the membership fee. The empirical prediction is that the usage rates of members on the monthly plan, where visits are free, are often so low that they would have been better off ex post using the per visit plan. D&M find that this is indeed the case for the majority of members in their sample. However, the hyperbolic discounting model fails to explain the delays in contract cancellation. To account for this feature of the data, D&M assume that some of the members are overconfident of their degree of self-control and consistently overestimate their usage rate.

One issue is the portability of the model. It is interesting to contrast D&M findings for health clubs to those obtained by Miravete (2003) for calling plans. He finds that most consumers on monthly metered and flat rate plans had ex post usage rates that justified their choice of plan. Furthermore, when this is not the case, consumers often responded to the potential savings by switching plans quickly. One possible explanation for the difference in results is that calling plans is a choice environment where consumers incur short-run benefits and delayed costs.

A second issue is that models of self-control problems are not necessarily inconsistent with the rational, game theoretic approach. The traditional choice model in economics takes preferences as the primitive and utility functions as representations. Gul and Pesendorfer (2001) formulate a model of self-control preferences and provide a representation theorem. (See also Fudenberg and Levine (2005) for a game-theoretic approach to self-control problems). The real issue here may be one of functional form rather than rationality. The parameterized hyperbolic discounting function is obviously a highly tractable way of modeling self-control problems.

The third issue is more troubling. The naïve consumers in the DM model are clearly bounded rational agents. They do not know their true preferences, are not aware that they do not know, and fail to learn even though actual play reveals their true discount rate. Other behavioral bias models (e.g., Gabaix and Laibson (2005)) also rely upon a substantial fraction of consumers who are naïve, unaware, and learn slowly, if at all, to explain behavior. This feature of the models is a fundamental challenge to the traditional rules. In the past, researchers would have been required to model the perception bias by assuming that consumers are uncertain about their preferences, are aware of their uncertainty, and update their beliefs in response to actual play. Today, I think the rules are less clear. It is difficult to know what to make of statements like "this fact is consistent with a large share of naïve consumers." What constitutes an explanation of behavior? What does it mean for a model to be coherent if some players take actions based on expectations that are not consistent with the outcomes generated by those actions? Finally, when objective functions do not represent a preference ordering or incorporate misperceptions, how does one conduct welfare analysis?

In his paper, Ellison presents an interesting matrix which lists standard IO models as rows and different biases as columns. This matrix will undoubtedly be taught in many graduate IO courses over the next few years, and filling in the boxes is likely to become a major focus of researchers in IO. I think that this research strategy is on firm and fertile ground when it imports non-standard preferences into monopoly and game-theoretic IO models. The DM study of the effects of hyperbolic discounting on pricing in health club markets and Heidhues and Koszegi (2004) study on pricing in oligopoly markets when consumers experience loss aversion are

excellent examples of what can be achieved. But, I am skeptical of models that rely on consumers with persistent perception biases to explain behavior.

Finally, I agree with Ellison that the recent shift from focusing on irrational firms to focusing on irrational consumers is an advance. As Ellison's example of a trip to supermarket illustrates, consumers have much larger choice sets, make many more purchasing decisions, and in far smaller quantities, than most firms. As a result, decision and information costs are likely to matter more for consumers than for firms. The billions of dollars spent on image advertising targeted mostly at consumers also suggest that consumers are more easily influenced than firms. The challenge is to move beyond description and show how incorporating these frictions into a model of consumer choice can generate new insights into phenomena such as image advertising.

References

Armstrong, M. (2005): "Recent Developments in the Economics of Price Discrimination."

Armstrong, M and J. Vickers (2001): "Competitive Price Discrimination, " Rand Journal of Economics, 32(4), 579-605.

Corts, K. (1998): "Third-Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment," Rand Journal of Economics, 29(2), 306-323.

Della Vigna, S. and U. Malmendier (2004): "Contract Design and Self-Control: Theory and Evidence," Quarterly Journal of Economics, 119 (2), 353-402.

Della Vigna, S. and U. Malmendier (2005): "Paying Not to Go to the Gym," American Economic Review, forthcoming.

Ellison, G. (2005): "Bounded Rationality in Industrial Organization."

Fudenberg, D. and D. Levine (2005): "A Dual Self Model of Impulse Control," mimo.

Gabaix, X. and David Laibson (2005): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets," Quarterly Journal of Economics, forthcoming.

Gul, F. and W. Pesendorfer (2001): "Temptation and Self-Control," Econometrica, 69, 1403-1436.

Heidhues, P. and B. Koszegi (2004): "The Impact of Consumer Loss Aversion on Pricing," mimo.

Miravete, E. (2003): "Choosing the Wrong Calling Plan? Ignorance and Learning," American Economic Review, 93(1), 297-310.

Rochet, J-C., and L. Stole (2002): "Nonlinear Pricing with Random Participation," Review of Economic Studies, 69(1), 277-311.