LUMPS IN ANTITRUST LAW

SEAN P. SULLIVAN[*]

The framework of aggregation and division that Lee Fennell develops in *Slices and Lumps* is both elegant and encompassing. Through the simple device of questioning how ideas and individuals are grouped together, or split apart, Fennell is able to explain and challenge concepts from diverse areas of law. While few subjects would not benefit from the clarity of this approach, the framework developed in *Slices and Lumps* strikes me as especially tailored to the study of antitrust law.

The importance of aggregation and division in modern antitrust policy cannot be overstated. Illegal acts of collusion are defined by the agreement of separate competitors to join together in acting as though they were a single firm in a collusive scheme. Tying arrangements are only potentially anticompetitive when the tying and tied products could be sliced apart and purchased separately. Anticompetitive concerns with mergers arise from the economic aggregation of separate competitors into a single entity. The inability of a parent company to collude with its wholly owned subsidiary reflects a unity of economic interest that is masked by the formal slicing of institutions at the boundaries of their incorporation.

The identification of relevant markets in antitrust analysis is an exercise solely defined by aggregation and division. As I have recently argued in another context, relevant markets in antitrust have little to do with lay concepts of industry, markets, or lines of trade. Rather, the modern and economically defensible concept of an antitrust relevant market is a cross section of trade in which a hypothesized competitive injury could potentially occur. This functional definition requires care to avoid both overinclusive and underinclusive market concepts. By aggregating up too much trade, an overinclusive market may mask potential competitive harm by including too many competitors in the mix and thus overstating the existence of

---

constraining competitive influences. On the other hand, by slicing trade too thinly, an underinclusive market may either overstate potential harm by omitting important sources of competition, or understate it, by obscuring broader patterns of interdependence and potential anticompetitive coordination among competitors.

Though any of these examples (and many more) could be productively developed within Fennell's framework, I use the remainder of this short essay to consider two questions of a somewhat more fundamental nature. First, how far does the lumpiness of trading partners go in dictating the limits of antitrust policy? Second, what does antitrust miss under the now common practice of lumping together price effects, consumer welfare, total welfare, and other concerns? Discussion of these questions is clarified and sharpened by reliance on Fennell's framework.

I.   LUMPS AS ANTITRUST DELIMITERS

Imagine all voluntary exchange as taking place somewhere on a two-dimensional grid defined by the number of potential trading partners on either side of the transaction. Though these dimensions are in principle continuous, we can for expositional simplicity limit our discussion to the four discrete cases identified in Table 1, below.

Table 1. Buyer/Seller Numerosity and Related Economic Theories

	Large # Buyers	Small # Buyers
Large # Sellers	Classical Price Theory	Monoposony, Oligopsony
Small # Sellers	Monopoly, Oligopoly	Bargaining Theory

One way to understand Fennell's flexible concept of lumpiness is in terms of market concentration. A small number of potential sellers constitutes a lumpy supply side of the market. Consumers have only a few lumpy options when seeking to complete a transaction. By contrast, a large number of potential sellers constitutes a smooth supply side. The same idea

applies on the buyer side of the market. As the following discussion ex-
plains, this type of lumpiness delimits much of the scope of antitrust law,
and also highlights a disquieting conceptual weakness at its core.

Start with the top-left quadrant of Table 1. This situation reflects a trans-
action that could take place between any pair of a large number of potential
buyers and sellers. All else equal, antitrust policy has little to contribute
here. Simply put, trade in this setting usually works well enough on its
own.

The reasoning comes directly from classical economic theory. In classic
price theory models, a single commodity is traded between a large number
of potential buyers and sellers. Assuming away certain externalities and
transaction costs, economic theory predicts that market-clearing prices
will maximize allocative efficiency in this setting. This is a strong claim,
but empirical studies seem to bear it out.

A nice illustration is a "pit market" experiment described by Charles
Holt, in which a large numbers of buyers and sellers are cast into a trading
pit to try to organically find and negotiate trades. The process is noisy and
chaotic as the student subjects in the experiment rush around the trading
pit trying to find willing trading partners. Yet even with minimal infor-
mation about market conditions, and with little time to find and negotiate
trades, subjects in this experiment often achieve results surprisingly close
to the efficiency predictions of perfect competition.

The point is that, for all their artificial properties, textbook models of
perfect competition appear to be quite robust, at least when there are many
potential trading partners on both sides of a transaction. Problems only
emerge—and antitrust policy only becomes important—when lumps form
on either the buyer or seller side of an exchange.

Start with a lumpy seller side: the bottom-left quadrant of Table 1. In
monopoly and oligopoly models, trade takes place between a large number
of potential buyers and a small number of potential sellers. A mature liter-
ature in industrial organization economics predicts that trade in this setting
will not generally exhibit the allocative efficiency of perfect competition.
Often, but not always, lumpy sellers will seek to maximize their profits by
driving prices above the allocatively efficient level. Compared to a perfect
competition benchmark, sellers do better, buyers do worse, and overall so-
cial welfare declines.

Roughly the same holds true when the lumps form on the buyer side of a transaction. In the top-right quadrant of Table 1, a large number of potential sellers attempt to trade with a small number of potential buyers. The efficiency implications of monopsony and oligopsony models are the mirror image of monopoly and oligopoly. Here, the lumpy buyers will often, though not always, seek to pad their wallets by driving prices down below the allocatively efficient level. Compared to a perfect competition benchmark, buyers do better, sellers do worse, and overall social welfare again declines.

Traditional antitrust policy is largely preoccupied with these two forms of trade (the anti-diagonal of Table 1). Among other things, antitrust law can often be interpreted as protecting the smoother side of a transaction against the lumpier side. Thus, the efforts of a small number of oligopolists to collude may be illegal under Section 1 of the Sherman Act. Actions that would further increase trading lumpiness by driving a competitor from the market may be illegal under Section 2 of the Sherman Act. Mergers that enhance or exploit certain forms of lumpiness may be illegal under Section 7 of the Clayton Act.

More could be said about antitrust policy in the three quadrants discussed so far, but I would rather shift focus to the fourth category of trade. The often overlooked lower-right quadrant of Table 1 involves exchange between small numbers of potential buyers and sellers. In the limit, this category of trade converges to negotiated exchange between a single buyer and a single seller: a condition sometimes termed "bilateral monopoly" in the law and economics literature. The question raised by discussion so far is what role, if any, antitrust policy should play in this quadrant.

The answer is complicated by the uncertain economics of this form of trade. In the extreme case of individual negotiation and bilateral monopoly, simple bargaining models can rationalize any division of the benefit of trade that does not make either the buyer or seller worse off than if the transaction had not occurred at all. Put another way, there is no clear prediction about the division of buyer and seller surplus in the lower-right quadrant. Allocatively efficient exchange can occur at any of an infinite number of potential "prices" in this setting. But does allocatively efficient exchange occur at all?

Again, the answer is uncertain. Most basic bargaining models predict that mutually beneficial exchanges will occur wherever feasible. But this turns out to be empirically doubtful. Impasse and failed exchange are often

observed in experimental research on bilateral bargaining—at least when it takes place between individuals. Indeed, the holdout problem, the subject of much concern in law and economics, is at root an empirical prediction that beneficial exchange may not take place at all if it must be negotiated between a small number of potential trading partners.

Backing away from the extreme of individual negotiation, the situation becomes less opaque in some forms of auctioned exchange, though many of the uncertainties of trade between few potential trading partners remain. An adequate treatment of action economics is impossible in the space of this essay. For present purposes, it suffices to note that auction prices can depend on factors such as differences in relative valuations, auction rules, and the number of bidders. Outcome efficiency is similarly complicated, potentially depending on information availability, auction rules, and bidder strategies, among other things.

The economic complexity and uncertainty of trade between small numbers of potential trading partners is substantial. At least for now, economic theory simply has more that it confidently can say about the properties of trade in aggregate (the first three quadrants of Table 1) than it has to say about the properties of specific instances of trade between small numbers of potential trading partners (the fourth quadrant of Table 1). But where does this leave antitrust in the lower-right quadrant?

There are plausible arguments that antitrust has little to offer here. First, in contrast to the bottom-left and top-right quadrants, there may be no asymmetry of lumpiness to create a disadvantaged side of the transaction in need of protection. Indeed, roughly equal lumpiness might protect each side of an exchange against attempted exercises of market power by the other. This thinking has motivated calls for relaxed antitrust enforcement in the literature, at least where market power would be exercised against a single entity, if at all. The treatment of "powerful buyers" as a mitigating factor in competitive effects analysis in the current Merger Guidelines evinces similar thinking. Second, as will be discussed in greater detail in the next section, allocative inefficiency does not necessarily follow from surplus appropriation in this setting, suppressing at least one of the traditional justifications for antitrust intervention.

But there are also plausible arguments that antitrust is indeed needed in the bottom-right quadrant of Table 1. First, and foremost, there is no principled reason to think that traditional antitrust injuries cannot occur in transactions involving only a few potential trading partners. The Merger

Guidelines identify mergers of competing bidders as a possible source of anticompetitive harm. And mergers affecting only a small number of buyers are sometimes opposed by the federal agencies. A recent example is the Federal Trade Commission's move to prevent Staples from acquiring Office Depot on grounds that this merger would weaken the negotiation posture of large business-to-business buyers of office supplies. To the extent that antitrust law already protects the negotiation posture of large corporate entities as they negotiate multi-million-dollar supply contracts, it may already reach far into the bottom-right quadrant of Table 1.

Second, the empirical evidence on bargaining failure may suggest an alternative and independent basis for opposing extreme lumpiness under antitrust law. Past a certain point, a reduction in the number of potential trading partners may lead to allocative inefficiency simply as a result of an increased likelihood of bargaining impasse and failed exchange. This is admittedly not a traditional basis for antitrust enforcement, but it tracks the underlying concern of allocative efficiency and total welfare theories of antitrust policy, and it probably deserves at least some consideration in the context of heavily concentrated markets.

The point of this discussion is not to suggest that antitrust policy must take any specific course of action in situations like the lower-right quadrant of Table 1. Neither economics nor antitrust law are mature enough to recommend or exclude any approach as a bright-line rule. But trade between small numbers of potential trading partners is a common and important form of commerce, and clearer antitrust policy in this area is something to which we might aspire in years to come.

While much more could be said about the four quadrants of Table 1 and how lumps in trading partners influence antitrust policy, I now turn to how the framework of *Slices and Lumps* facilitates discussion of another latent issue in modern antitrust: imprecision in the fundamental goals of antitrust policy.

II.   LUMPING PRICE, OUTPUT, AND EFFICIENCY

Many antitrust textbooks start by comparing equilibrium outcomes in models of perfect competition and single-price monopoly, roughly paralleling some of this essay's earlier discussion. Takeaways from the usual textbook treatment also track aspects of earlier discussion. Relative to the many small sellers that make up the supply curve in a model of perfect competition, a monopolist internalizes the effect of each incremental price

reduction on all units sold. This gives the monopolist a profit motive to raise the price of a good or service above the competitive level.

In the typical comparison, the consequences for consumers and society provide a number of related justifications for antitrust policy. The market price is higher under monopoly than under perfect competition. Some consumers pay the higher monopoly price, and their benefit of trade is reduced in exact proportion to the enhanced profit margin of the monopolist. Other would-be consumers are unwilling or unable to trade at the higher monopoly price. Their failure to obtain the goods or service leads to a lower total quantity of trade, and also implies an allocative inefficiency. The goods or service could have been supplied to these would-be consumers—as proven by their ability to buy it at the lower perfect-competition price—and because the profit maximizing behavior of the monopolist obstructs these trades, it necessarily deprives society of beneficial transactions.

In this textbook treatment of the evils of monopoly, price correlates with many variables. Price and output are directly linked. There is a single price that determines the quantity traded (and vice versa). For the same reason, price correlates with allocative efficiency. Any price above (or below) the competitive price implies fewer trades than would have occurred under competitive pricing, which in turn implies an inefficient allocation of resources. If the market price deviates in any way from the equilibrium price in perfect competition, then some feasible and beneficial exchanges are bound to remain unrealized.

Finally, many texts and antitrust scholars draw an additional correlation between consumer welfare and total welfare in this setting. In the textbook comparison of single-price monopoly with perfect competition, the reduction in consumer welfare is brought about by a higher-than-competitive monopoly price. This monopoly price prevents some beneficial trades from occurring—an allocative inefficiency. And this allocative inefficiency (combined with a welfare neutral redistribution of the gain of trade among those who continue to trade) implies a net reduction in society's overall gains from trade, reducing total welfare. So consumer welfare and total welfare are linked as well.

This chain of inferences, linking reductions in consumer welfare to reductions in total welfare, may help to explain Robert Bork's infamous conflation of the two concepts in the *Antitrust Paradox*. Bork's conflation, and some underlying uncertainty about the differences, may help explain

the willingness of many judges to lump together the concepts of price elevation, output reduction, consumer welfare reduction, and total welfare reduction. Finally, the authoritative language of some judicial opinions has undoubtedly allowed (if not compelled) some practitioners and scholars to adopt this same imprecision in treating these different ideas as though they were basically interchangeable in practice.

But there is a problem with lumping these ideas together too freely. The chain of inferences that links price to output and consumer welfare to total welfare does not generalize to many settings more complicated than the simplistic comparison in which it is often presented. For one thing, the baseline of perfect competition is rarely, if ever, the appropriate benchmark for measuring competitive injury. Almost always, the baseline for assessing a challenged act is itself an imperfectly competitive status quo. Moreover, the removal of a single assumption throws the entire chain of inference into doubt.

The critical assumption is that a firm with market power sets a single per-unit price for the good or service it provides. In perfect competition, this assumption aligns with intuition. When fiercely competitive price-taking agents sit on either side of a commodity exchange, it is hard to imagine anything other than a single per-unit price emerging as the market equilibrium. But monopolists are not price takers. And unless some external force prevents it, the monopolist's power to set prices will generally include the power to set *different* prices for different transactions. This raises an obvious question: what does the comparison to perfect competition, or any appropriate baseline, look like if the monopolist does not charge a single per-unit price?

The answer depends on what pricing model the monopolist adopts—and there are *many* possibilities. To take one extreme, suppose the monopolist charges each customer a customized price. In a model of perfect price discrimination, the seller is assumed to be able to accurately predict the willingness-to-pay of every consumer and to be able to prevent arbitrage by technical or legal constraints (foreclosing resale among consumers). A monopolist in such a position could charge each consumer the most that this person would possibly pay for the quantity of the good or service being traded. The effect is complete appropriation of consumer welfare by the monopolist: the value of trade to all consumers is driven down to almost nothing. But the effect is also complete preservation of the allocative efficiency of perfect competition. The monopolist has a strong profit motive

to make sure that every last efficient trade occurs in this setting, so the total quantity traded is the same for perfect price discrimination as it is for perfect competition, and total welfare is the same as well. In short, the effect of perfect price discrimination is to completely unlink changes in consumer welfare from changes in total welfare.

As another example, a monopolist may engage in imperfect price discrimination, setting different prices in different geographic locations or selling products in different quantity or quality brackets. These strategies increase the monopolist's profits, but their implications for consumer and total welfare are ambiguous. Often, this type of imperfect price discrimination will make one group of consumers better off than if the monopolist had set a single price while simultaneously making another group of consumers worse off. The effect on consumer welfare then turns on the philosophically difficult question of how different groups of consumers should be sliced or aggregated in computing changes in consumer welfare. Total welfare presents fewer conceptual challenges but is no less ambiguous, in that it may be higher or lower than under single price monopoly depending on how much each of the different groups of consumers gains and loses under the imperfect price discrimination scheme. Comparisons to perfect competition are more or less extreme than under the comparison to single-price monopoly, depending on the same considerations.

Other pricing schemes present their own complications. There are other types of price discrimination that a monopolist may pursue. There are also the various pricing strategies now studied in the growing economic literature on nonlinear pricing: things like quantity discounts, peak-load pricing, two-part tariffs, various forms of bundling, among others. And as noted earlier in this essay, individual negotiation and auctioned exchange present further complications—at the most extreme, decoupling prices from efficiency implications over substantial bands of exchange.

The point of this discussion is not to suggest that any particular pricing scheme is good or bad in the abstract. It is also not to suggest that antitrust law cannot handle cases involving nonlinear pricing. The point is simply to highlight the dubious theoretical basis for lumping changes in price, consumer welfare, allocative efficiency, and total welfare together as though they were equivalent in every case. They are not. And the problems inherent in treating them as equivalent are substantial.

For one thing, the uncritical lumping of these concepts stunts the growth of antitrust policy. There are reasoned arguments for why antitrust should

focus on consumer welfare just as there are arguments for why it should focus on total welfare. There are reasons to direct attention to price effects just as there are reasons to let price take a second seat to efficiency. These options, and the tradeoffs between them, are needlessly obscured by the common practice of lumping together all these concerns.

Another problem arising from the conflation of concepts like consumer welfare and total welfare is that it confuses articulation of the principles upon which cases are being decided. In the recent Supreme Court decision of *Ohio v. American Express*, for example, did the majority really mean what it said in defining market power as "the ability to raise price profitably *by restricting output*"? If so, then it would seem to be motivated by concern for allocative efficiency and total welfare—not consumer welfare as it elsewhere suggests. If not, then it confuses readers with a mistaken lumping of allocative efficiency and consumer welfare. In the other direction, when the Eleventh Circuit passed down the important rule that to defend a merger on efficiency grounds any cost savings needed to "benefit competition and, hence, consumers," did the court really mean to articulate a pure consumer welfare standard for merger review? Or was it simply and mistakenly assuming that anything which reduced consumer welfare always reduced total welfare as well?

Finally, the imprecise lumping together of different objectives frustrates basic antitrust analysis. As Steven Salop has observed, the antitrust analysis of harm to competitors differs substantially between the consumer welfare and total welfare standards. Inefficient allocations of output between the members of a cartel may similarly invite harsher treatment under a total welfare norm of antitrust than under the consumer welfare standard. More generally, imprecision about what norms define anticompetitive harm and procompetitive injury complicates the day-to-day weighing of competitive effects, as required in many merger cases, for example. The comparison of competitive costs and benefits is a difficult exercise under the best of circumstances. And when combined with confusion about what counts as a harm or benefit—the predictable result of routine conflation of different antitrust policy objectives—the difficulty and unpredictability of the exercise can simply expand beyond control.

All these problems arise from the assumption that changes in price, consumer welfare, allocative efficiency, and total welfare always (or almost always) move together in simplistic ways. This is not a sound assumption, but even if it was, the uncritical lumping together of conceptually different

policy objectives would remain both unwarranted and counterproductive to modern antitrust practice. Antitrust law would be well-served by a more conscious slicing apart of its different objectives in scholarship, opinions, and advocacy. Helpfully, this exercise is not only invited by the framework Fennell develops in *Slices and Lumps* but facilitated by it as well.