THE OLIGOPOLY PROBLEM, TRIGGER STRATEGIES, AND "COORDINATED EFFECTS"

BY JOSEPH FARRELL[1]

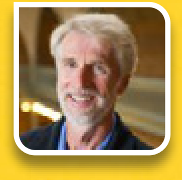

1 Professor of the Graduate School in the Department of Economics, University of California, Berkeley.

THE OLIGOPOLY PROBLEM, TRIGGER STRATEGIES, AND "COORDINATED EFFECTS"

By Joseph Farrell

In earlier work, Jonathan Baker and I explored what we called "non-purposive coordination" and the 2010 Horizontal Merger Guidelines call "parallel accommodating conduct." Such conduct arises when oligopolists respond to one another's competitive initiatives in ways that undermine competitive incentives but that are not driven by a goal of converging on a mutually understood outcome or penalizing departures from one. Here I explore and modernize an approach that was hinted at sixty years ago, elucidating the roles of concentration and diversion ratios without fully predicting conduct. While the conclusions may point in the same general direction, that analysis differs starkly from the analogous findings in the standard repeated-game framework for oligopoly responses. That framework's findings concern sticking to and enforcing, not reaching, a common understanding; and the presence of such an understanding is an assumption, not a finding, in standard game theory. I clarify the role of trigger strategies in repeated-game "folk theorems," and observe that purposive trigger strategies do require careful coordination (from oligopolists' point of view) but non-purposive strategies need not. Finally I observe that the use of language in this context, particularly "coordination" and its cognates, risks being unhelpful in the antitrust analysis of non-purposive responses.

A price-setting oligopolist's competitive incentive to cut its price is undermined if it expects its rivals to respond by cutting theirs; similarly, its incentive to raise its price is buttressed if it believes rivals will respond by raising theirs. The resulting "oligopoly problem" weakens competition among the oligopolists and raises prices; analogous mechanisms can similarly harm customers in non-price dimensions.

More than 60 years ago Donald Turner described a robust and intuitive source of such oligopoly responses:[2]

> The rational oligopolist [takes into account] the reactions of his competitors to any price change that he makes. He must take them into account because his competitors will inevitably react. They will inevitably react, for example, to a price cut on his part because otherwise the price cut will make a substantial inroad on their sales…

Turner continues by claiming a close link between this and concentration:

> …if, for example, there are only three producers of equal size and a price cut by one doubles his sales, the sales of each of his two competitors will be cut in half. The rational seller in an industry with a very large number of competitors does not calculate their reactions to a price cut by him, because they are not likely to be sufficiently affected by the price cut to react; if, for example, there are one hundred producers of equal size, a doubling of sales by one, evenly drawn from his competitors, would cut their sales by only one ninety-ninth. …

Turner's discussion of the role of concentration is not entirely satisfactory: it is analytically risky to "round down" small quantities to zero. As differential calculus teaches us, one often needs to keep track of the relative smallness of quantities all of which are small. But I believe Turner's instincts served him well in this case.

To see this, take Turner's example and suppose that initially each of the 100 equal producers sells 990 units.[3] Firm 1 cuts its price by an amount that, absent responses by the other firms, would double its output, to 2 x 990 = 1980 units. Assuming (with Turner) symmetry and an aggregate diversion ratio of 100 percent (all of firm 1's additional sales are at the expense of the other firms in the market), each of the other 99 firms would lose 10 units of its sales and sell 990 – 10 = 980 units, absent responses. But each of them in fact responds. How does (say) firm 2 respond to a price cut by firm 1 that, absent any response, would cut firm 2's sales by (as Turner notes) one ninety ninth, or from 990 units down to 980?

We do not have a unique rigorous answer to that question. But that shouldn't mean defaulting to an implied answer that firm 2 won't respond at all (as in the one-shot models of oligopoly). Nor is it particularly helpful to observe that in some dynamic frameworks it seems that almost anything can happen (as, more loosely, in repeated-game models). This is related to the issue of "conjectural variation," which has run into similar problems and seems to have prompted similar less-than-helpful reactions. We need some guidance about what is relatively apt to happen, and in our current state of knowledge this appears to mean dialing down any expectations of perfection and being pragmatically satisfied with sensible and intelligent.

While Turner does not say this, I read him as envisioning, along those lines, that firm 2's response would be apt to be proportionate: its responsive price cut would regain (or protect) some of the 10 units that it would otherwise lose, or maybe all of them, but probably not more. I think this is quite sensible and intuitive as a description of natural or non-collusive oligopoly behavior.[4] Let us label the sales volume that firm 2 regains with its response, divided by the 10 units that it would lose if it did not respond, as its regain ratio, or R. So firm 2 regains 10R units and on net loses 10(1 – R) units of sales. It is not guaranteed that R is between 0 and 1, but that seems an intuitively sensible benchmark and is consistent with Stackelberg responses (the rational prediction if players do not look further ahead than that response) in many simple demand systems.

In regaining 10R units, firm 2 takes (or takes back) from each of its rivals, and hence in particular from firm 1, one ninety-ninth of that 10R, or 10R/99 of a unit. In aggregate the effect on firm 1's sales of such responses by firms 2 to 100 is thus 99 x (10R/99) or 10R units. Summing up, firm 1's price cut, which absent responses would have gained it 990 new sales, actually gains it "only" 990 – 10R new sales, which even if R=1 is only slightly different (980 versus 990) from the initial estimate ignoring responses. Indeed, even if say R=2 one might say the

2 Donald Turner, "The Definition of Agreement Under the Sherman Act," *Harvard Law Review* (1962).

3 This makes some of the arithmetic simpler and of course has no separate substantive force.

4 This thinking views firm 2 as responding to the decline in its sales, or more precisely to the downward or (ten unit) leftward shift in its residual demand curve, without regard to the source of the shift; and Turner is clearly also envisioning that firm 1's initiating price cut remains in force long enough for the other firms' responses to apply in the market. In modern words, he is considering a Markov response, and assuming, like Charlie Brown, that the football is not yanked away moments before he kicks it. See the discussion in Eric Maskin & Jean Tirole "A Theory of Dynamic Oligopoly, III: Cournot Competition," *European Economic Review* (1987).

same thing. One can see Turner's point when he argues that firm 1 need hardly take such responses into account in evaluating such a price cut, although a more modern discussion would instead point out that taking those responses into account would not much change firm 1's incentive for the initiating price cut.

By contrast, consider the highly concentrated or three-firm version of Turner's example. Again, consider a price cut that would, absent responses, double firm 1's sales. Continuing (with Turner) to assume a 100 percent aggregate diversion ratio, that would take half of 990, or 495, units from each of firms 2 and 3 if they did not respond. If instead each responds with a price cut that would in isolation (re)gain it a fraction R of that 495 units, then half of that regained volume is taken from firm 1, so the responses cost firm 1 a fraction of the initially gained 990 units. Although the unknown parameter R appears in both cases (and it is not guaranteed that the same value of R would apply in both), here values of R well below 1 are consistent with a substantial dilution effect, in contrast to the many-firm case where even with R=1 (or somewhat higher) the responses still don't much matter to firm 1.

Some readers will recognize that the impact of firm 1's initial price cut on each of its rivals depends on the diversion ratio from firm 1 to that rival, and the impact in turn on firm 1 of that rival's response is proportional to the strength of that response and to the diversion ratio from that rival to firm 1.[5] Thus with *n* symmetric competitors the dilution of firm 1's sales gain is a sum of (n – 1) terms of the form , or in total . Thus as Turner claimed, firm 1's incentives become more competitive as n increases (concentration falls).

This sketch of a model is obviously incomplete in various ways. For instance, the analysis terminates after one round of responses to firm 1's initiating price change. That might seem to suggest (though it doesn't require) that those responses be privately optimal, as in Stackelberg (but with multiple followers) equilibrium: Jonathan Baker and I have argued that Stackelberg seems a reasonable starting point for analyzing "natural" oligopoly dynamics,[6] but there are plenty of unresolved questions, such as whether the followers foresee one another's responses.

More generally, what determines the parameter R? One might try to derive a response ratio R in a truly dynamic model. Maskin & Tirole (1987) do this or a closely related exercise for a quantity-setting game in which two firms alternate in setting quantities and each such decision is binding for two periods.[7] But they (even they!) limit themselves to the symmetric duopoly case, and find even there that calculating a rigorously derived response rate requires a good deal of analysis followed by solving a quartic equation. And I have not seen subsequent work that offers us a workhorse model, or empirics, that would give us a good sense of how R varies with market conditions such as transparency, demand elasticity, pass-through rates, and concentration—or indeed whether those are even the key conditions to look at.

Looking for a way to dial down ambitions of perfection and make some progress, one might follow Turner's footsteps and assume something like reasonable proportionality, for instance R between 0 and 1 and relatively stable, although recognizing that none of that is guaranteed. If R stays stable and in a reasonable range, we just saw that one can say something about the effects of concentration and diversion ratios on the extent to which oligopoly incentives depart from the unilateral (no responses) incentives that have become the primary focus of so much antitrust economics.

The bottom line that rivals' responses may raise oligopoly prices more in a more concentrated market is of course a staple of antitrust economics, but the analysis that an average 21st-century antitrust economist or economically-savvy lawyer would offer to support that bottom line would scarcely overlap at all with Turner's logic or a refinement or modernization thereof—a logic that would be radically unfamiliar to such an interlocutor. This fact is strange and somewhat perturbing, and I want to explore here how it comes about.

As most readers will be aware, modern oligopoly price theory largely bifurcates between static models on the one hand, and repeated-game models on the other. An archetypal static model involves simultaneous one-shot choices of price or quantity, and is the usual approach to modeling "unilateral effects." An archetypal repeated-game model consists of an infinite series of periods; in each period rival firms play a copy of a given "stage-game." The stage-game is usually just like its static cousin and consists of simultaneous moves by the rivals that, importantly, remain in effect only during the period.

In a one-shot model or in such a repeated-game model, it is never possible for one firm to "respond" to seeing another's price in the natural sense that both prices will apply at the same time. This is entirely clear when one looks at the model with that question in mind, but (in

5 The "from" and "to" terminology originated with discussions of merger-induced price increases, and gets turned around in analyzing a price cut, but is too well-established for me to mend it here.

6 Joseph Farrell & Jonathan Baker, "Natural Oligopoly Responses, Repeated Games, and Coordinated Effects in Merger Analysis: A Perspective and Research Agenda," *Review of Industrial Organization* (2021). Also generally see Jonathan Baker & Joseph Farrell, "Oligopoly Coordination, Economic Analysis, and the Prophylactic Role of Horizontal Merger Enforcement," *Penn Law Review* (2020).

7 Eric Maskin & Jean Tirole "A Theory of Dynamic Oligopoly, III: Cournot Competition," *European Economic Review* (1987).

my experience) seldom explicitly pointed out. By the time firm 2 sees firm 1's price in the one-shot model the game is over; less obviously, in the repeated-game model, by the time firm 2 sees firm 1's period-t price and can choose its own, potentially in response, it is period (t+1) and firm 1's period-t price has already been superseded.

It may not be clear how long firms in practice are committed to their prices versus how quickly other firms can react, but the prevailing repeated-game literature takes an extreme position on that question: it assumes that firm 2 can never respond to firm 1's price while the latter still applies.[8] In this way intuitive notions of what might well be "inevitabl[e]" responses, such as Turner's, are *ignored by definition* and the consequences of a surprise price cut and of the responses to it are evaluated in different periods. It is a startling thing for a literature on price responses to assume.

The repeated-game literature instead focuses on a different kind of price response: a response to bygone history and not only to current conditions. In game-theory language, those responses are not Markov (they depend on payoff-irrelevant history, and in many simple expositions they depend only on payoff-irrelevant history).[9] Certainly they are not "inevitabl[e]" in anything like Turner's sense. Turner-esque responses to a still-in-force price make sense for firm 2 at some level irrespective of what firm 1 might know or believe, but a punitive response intended to deter price-cutting makes sense only if firm 1 foresees it or learns to do so.[10] Such a response is "bootstrapped:" it would be completely rational for each firm to ignore the bygone history if other firms did so, and in an intuitive sense that ("bygones are bygones") would be a natural starting point. Getting instead to a point where everyone expects specific punitive responses because they expect that everyone else expects just those responses takes some work. And indeed thwarting that "work," largely by denying the firms the use of open explicit communication, is the agenda of enforcement against price-fixing.

The repeated-game framework and literature captures some important aspects of oligopoly responses and conduct, and I am not disparaging that contribution.[11] But it is also plainly missing something important, and its dominant market share in antitrust economics education risks entrenching a narrow focus on the things that it does stress. So it is worth being conscious of points that it misses.

First, a general point about game theory. Almost all game theory, including the standard repeated-game models, evaluates possible predictions of the game's outcomes to see whether they would fail because of an incentive to deviate: if not, they are called equilibria. A bit more precisely, the theory evaluates strategy profiles—full descriptions of how the game might be played (including following histories that are not supposed to arise)—to see whether at least one player has an incentive to depart from its specified strategy, assuming that others stick to theirs. Non-equilibrium outcomes are treated as not plausible predictions. Thus the existence of a "common understanding" of how everyone is expected to play is baked, or assumed, into any equilibrium and thereby into (almost) any game-theoretic prediction. This common understanding, then, is not a finding but a methodological assumption. If one thinks, as many do, of non-competitive oligopoly dynamics as involving reaching a common understanding and then ensuring incentives not to "cheat," the repeated-game literature focuses on the latter and has rather little to say on the former.

This may be part of why repeated-game theory often diagnoses the existence of a shared-monopoly equilibrium even in very unconcentrated markets.[12] Yes, the theory also points out that fully competitive equilibria also exist, but it would be helpful if it had a lower rate of false positives. One possible approach is to make stronger demands (e.g. renegotiation-proofness, or bargaining perspectives on the allocation of gains among oligopolists) on the equilibrium. But I now suspect that a wiser interpretation is that the internal incentive stability of an equilibrium (that is, being "an equilibrium") is not usually the gating factor: that is more often the "work" of getting there. Game theory unfortunately has not had very much to say about that work and how much more difficult it becomes if, as under effective anti-collusion enforcement, firms must avoid explicit direct communication.

Second, a point more specific to the repeated-game literature. Most discussions of a repeated oligopoly game focus on trigger strategies. Broadly, a firm uses a trigger strategy if it charges a high price so long as every firm in the industry has always done so, but switches to a

8 A similar timing and commitment issue arises with "limit pricing:" if an incumbent can set a high price but cut (or more generally completely re-optimize) its price as quickly as an entrant can challenged it, the threat of entry is (in simple models) toothless. See Paul Milgrom & John Roberts, "Limit Pricing and Entry under Incomplete Information," *Econometrica* (1982). This contrasts with the idea of "contestability," which assumed that an incumbent could not cut its price on entry. Aaron Edlin has suggested policy intervention to bring monopoly markets closer to the latter theory. See Aaron Edlin, "Stopping Above-Cost Predatory Pricing," *Yale Law Journal* (2001).

9 The intuitive idea is simple; Maskin & Tirole give it a rigorous formulation and discuss the relationship with other intuitive properties. Eric Maskin & Jean Tirole, "Markov Perfect Equilibrium I: Observable Actions," *Journal of Economic Theory* (2001).

10 This relates to Thomas Schelling's distinction between "warnings" and "threats." In technical game-theoretic terms, the ex post motivation for a punitive response here is that everyone else is expecting it.

11 Indeed, I have taken part in that literature: see for instance Joseph Farrell & Eric Maskin, "Renegotiation in Repeated Games," *Games and Economic Behavior* (1989).

12 Carl Shapiro, "Theories of Oligopoly Behavior," *Handbook of Industrial Organization* (1989).

much lower price if any firm has ever deviated from the high price. In this sense a trigger strategy is discontinuous or at least has a very sharp increase in R or a conjectural variation or similar response parameter at or near the anticipated price. An equilibrium featuring trigger strategies is the oligopoly analogue of thermonuclear "mutual assured destruction." Like MAD, it is in principle powerful discipline, but it is also easy to see how it can go horribly wrong. Why is it so pervasive in discussions of oligopoly dynamics?

One reason, I conjecture, is the role that trigger strategies play in the proofs of what are called the repeated-game folk theorems. While there are several folk theorems, the most familiar discussion in antitrust economics shows (simplifying slightly) that any allocation of e.g. shared-monopoly profits that gives each oligopolist some of the excess profits above the one-shot equilibrium level can be sustained as the payoffs from a trigger-strategy equilibrium of the repeated game. Moreover, although it is more complicated to state the result precisely, any allocation that can be sustained as the payoffs from *any* "subgame-perfect" equilibrium of the repeated game can also be sustained using trigger strategies. Thus students of repeated-game equilibrium learn to consider what can and cannot be achieved in theory using trigger strategies. But those strategies' central role in those *proofs* does not imply any recommendation or any particular prominence in use. Much of what can be sustained using trigger strategies may equally be sustained using price-matching or similarly non-trigger strategies. Equally, the somewhat-less-than-competitive outcome of Turner's model could also be sustained using trigger strategies—but that of course does not mean that an oligopoly behaving in Turner style is in any sense using trigger strategies behind the scenes.

This matters, I think. In particular, consider the role played by what game theorists call coordination. Literally the word means arranging things together. A "coordination game" is a game, such as technology adoption with network effects, in which the key thing is to do "the same thing" as other players.[13] Classic game-theoretic examples include pure coordination games (in which nothing else matters):

1,1	0,0
0,0	1,1

Another classic coordination game is the "battle of the sexes" (in which coordination incentives coexist with conflict over what to coordinate on):

2,1	0,0
0,0	1,2

In oligopoly, coordination in this sense — picking the same, or precisely-matched, strategies—is *sometimes but not always* important for the outcome to feature above-competitive pricing. In particular, if one focuses on trigger strategies, it does indeed become crucial to pick consistent triggers — otherwise triggering will indeed ensue. For instance, if oligopolist 1 picks a trigger of $100 and oligopolist 2 picks a trigger of $101, a price war ensues after period 1, because oligopolist 2 will interpret oligopolist 1's $100 price as "cheating."[14]

Thus the term "coordination" risks being confusing. In its game-theoretic meaning it applies sometimes but not always when oligopolists reach what antitrust might call successful coordination. And Jonathan Baker and I adopted the phrase "non-purposive coordination" to emphasize that it is the incentive-sapping effect of responses (as in Turner) that provides the economic essence of the harm, whether they are illegally "coordinated" or are "natural and non-purposive" or neither. It is late in the day to rethink the most helpful usage for terms such as coordination, agreement, collusion, and so on. But we can at least remain uncomfortably aware of the risks of confusion, and remind ourselves to think very carefully around these issues.

The theory of repeated games, and in particular the folk theorems, are a powerful piece of analysis and play a prominent role in understanding the boundaries of *possible* oligopoly and similar cooperation. But they need to be complemented by a better understanding of which of the very many possible outcomes are relatively likely. We do not yet have a rigorous workhorse model of the latter, so it is important to use what insights we have and work pragmatically to develop those, and not ignore them merely because they are incomplete and imperfect. Understanding the dynamics of firm responses under relatively simple credible, if not "inevitable," behavior seems likely to be among the better ways forward.

13 It is usually clear from context what "the same thing" means. One can also think of it as a game with multiple equilibria in which each player would like the other player(s) to know what he is doing.

14 That assumes that each oligopolist's price is equal to the threshold that it applies to another oligopolist's price; one could alternatively separate those two but the basic point would remain.