Renegotiation in Repeated Oligopoly Interaction

Joseph Farrell†

March 14, 2000

(First Version: October 1992)

Abstract: Standard repeated-interaction theories of oligopoly make collusion seem much easier than a "structural consensus" suggests that it is. I show that more intuitive results can emerge if colluders could renegotiate after a deviation. In repeated Bertrand oligopoly, if agreements are subject to a certain kind of frictionless renegotiation, then full collusion is impossible with more than three firms, however high the discount factor. With more firms, partial collusion is possible, but its deadweight losses are small compared to monopoly effects. I also analyze results for repeated Cournot interaction. Finally, I try to relate the game-theoretic literature, the structural consensus, and the concern over collusion.

Forthcoming in the *Festschrift* in honour of James A. Mirrlees, edited by G. Myles and P. Hammond: Oxford University Press, 2000.

† Professor of Economics and Director of the Competition Policy Center, University of California, Berkeley, CA 94720-3880. email: farrell@econ.berkeley.edu. I thank Luis Cabral, Dennis Carlton, Andrew Dick, Richard Gilbert, Matthew Rabin, and Carl Shapiro for helpful comments, Rene Kamita for research assistance, and the National Science Foundation for early financial support (Grant SES–9111095).

1. Introduction

What determines whether an oligopoly can sustain prices above competitive levels, and if so how much above?[1] Modern oligopoly theory offers two general answers to this question. First, there are static models in which price competition is muted because of product differentiation or because firms compete in dimensions less cutthroat than price; the main examples are Cournot oligopoly theory and models of spatial product differentiation. But these models ignore the often obvious dynamic features of oligopoly markets: in particular, firms' actions are based on predictions about competitors' reactions.

Second, there are dynamic models in which price competition can be muted through the prospect of such reactions. These dynamic models are analyzed using the theory of infinitely repeated games, usually with the solution known as subgame-perfection. The folk theorem states that generically almost any outcome — in particular, many outcomes involving joint monopoly pricing — can arise as the subgame-perfect equilibrium path of a repeated game, provided only that there is sufficiently little discounting. Thus, the model can explain the persistence of pricing substantially above cost on the part of a small or moderate number of firms, even when products are undifferentiated, total capacity greatly exceeds demand, and firms choose prices simultaneously within each period.

But (with plausible ancillary assumptions) these models also claim that even with many firms, monopoly pricing is possible. For instance, in repeated n-firm Bertrand oligopoly with constant marginal costs, shared monopoly is sustainable in subgame-perfect equilibrium provided that the per-period discount factor δ is at least $1 - \frac{1}{n}$. With $\delta = 0.99$, representing a plausible interest rate if the detection and response lag is a month,[2] this means that 100 firms can sustain the monopoly price. And Shapiro (1989, page 365–366)

[1] A theoretical sharpening of this question is why competition among the few does not drive prices all the way down to costs. Tirole (1988, p.209) calls this the "Bertrand Paradox."

[2] This assumes that ordinary discounting is the primary component of the discount factor δ. If there is uncertainty (of the right kind) about whether a firm will remain in the market, or the market will survive, then δ might plausibly be considerably lower, even if firms do not know when it is their last active period. If firms do know when it is their last period, and if they can and will then cut price and take a large share, then cooperation breaks down whenever *any* firm is moribund. This may make cooperation predictably short-lived — and thus hard to sustain even if in fact all firms are flourishing — with relatively modest numbers of firms and levels of "ordinary" discounting.

calculates that (with the same assumptions, plus linear demand) a symmetric Cournot industry can sustain a shared monopoly with up to 400 firms, even if punishment consists of Cournot reversion rather than the harsher minimax punishments allowed by the folk theorem. Still more dramatic numbers emerge if we look at Cournot competition with minimax punishments.[3]

These theories therefore suggest that we should be concerned about coordinated pricing even in markets with fifty or a hundred equal-sized firms.[4] Yet in practice there is what I will call a *structural consensus* that coordinated pricing is normally unlikely to be a major problem even in markets considerably more concentrated than that. Thus, the 1992 *Horizontal Merger Guidelines* of the US Department of Justice and Federal Trade Commission indicate that there will usually be little competitive concern about a merger that leaves a market "unconcentrated," in the sense of having a Herfindahl concentration index (HHI) of 1000 or less (the equivalent of at least ten equal competitors).[5] And there is perhaps some evidence that if changes in industry competitive behavior can be systematically associated with concentration levels, the cutoff point may often be at about four-firm concentration ratios of .5 to .6,[6] which would correspond to six to eight equal-sized

---

[3] In order to sustain monopoly pricing with n firms in the linear-demand ($p = 2 - X$) zero-cost Cournot case, for example, δ must be at least equal to $1 - \frac{4n}{(n+1)^2}$; this follows from the fact that (by the Folk Theorem) the worst subgame-perfect punishment for either player is zero. In the Bertrand case this is straightforward, since charging $p = 0$ forever is obviously a subgame-perfect equilibrium. In the Cournot case the minimax punishment is more complicated, since the innocent firm must produce 2 per period if the guilty is to be held down to zero, and if this were meant to be kept up forever then the innocent firm would want to cheat by producing less in any period. But for large enough values of the discount factor a firm can be punished with an "almost" zero continuation value. Note that I am following fairly standard economics practice by describing a defector from coordinated pricing as the "guilty" firm and nondefectors as "innocent;" this language fits well with the general idea of cooperation as desirable and thus fits rather awkwardly with the oligopoly problem.

[4] The role of "size" is really that, in the models, any firm could readily take the whole market.

[5] 1992 Horizontal Merger Guidelines, April 2, 1992, section 1.51(a). Schmalensee (1987, page 50) describes the 1984 version of this provision as "clearly warranted: mergers that leave markets atomistic almost never increase the likelihood of collusion noticeably." Strictly, these statements concern the *incremental* competitive concern from a merger, and do not address the *level* of coordinated pricing to be expected in such an industry. However, I think that most informed observers would justify the statements by saying that even after such a merger there is little threat of coordinated supercompetitive pricing — not (for instance) by saying, in contrast, that coordinated monopoly pricing is to be expected even absent the merger, so the merger would make little difference.

[6] See *e.g.*, White (1987) and references therein.

firms. Summarizing what he describes as numerous case studies, Potter (1991, page 12) claims that "hostility," by which he seems to mean the tendency toward competitive pricing when there is excess capacity, will often stop "when the industry has consolidated down to three or four key players."

The theory that purports to explain cases of collusion therefore explains too much: the structural consensus contrasts sharply with the game theory models, and implies that the standard subgame-perfection condition is not generally the *binding* constraint on collusion.

Focusing attention on a constraint that is not binding is a serious problem. Not only will it wrongly estimate the extent of collusion, but it also risks suggesting wrong lessons about comparative statics and policy. That is, a structural or legal change that tightens the subgame-perfection constraint will appear pro-competitive if analyzed as if subgame-perfection were the binding constraint. Yet logically it may simultaneously weaken the actually binding constraint, and thus in fact be anti-competitive. Thus it is important to try to find the *binding* constraint.

What, then, is the binding constraint? Why is coordinated high pricing much harder — if the structural consensus rather than the textbook theory is correct — than the repeated-interaction models predict? In this paper I explore one possible candidate: the punishments for deviation specified in the folk theorem (or under Nash reversion), although subgame-perfect, are not really credible if the firms can renegotiate. I explore the extent to which oligopolists can sustain above-cost pricing in equilibria that are not only subgame-perfect but "renegotiation-proof." Using an apparently reasonable definition of renegotiation-proofness (when renegotiation is perfect), I show that collusion is severely limited by the number of firms, *even in the limit* as discounting becomes very unimportant ($\delta \to 1$). These results seem appealingly consistent with the structural consensus.

In most of this paper, following economists' sometimes casual usage, I use the term "collusion" and its cognates to denote oligopoly pricing (far) above cost. Such pricing may be achieved by explicit collusion, or might occur through other means ("conscious parallelism").[7] To the extent that they differ, the paper's analysis, which assumes a

[7] I am grateful to Richard Gilbert for encouraging me to pursue this distinction. The issue of how these

communication-rich environment, would be primarily about explicit collusion and thus perhaps only loosely linked to the structural consensus, which is arguably much more about parallelism. For this reason, as I discuss in the Conclusion, I am not convinced how much weight to put on the results' similarities to the structural consensus. I therefore offer the paper mainly in the hope of provoking more exploration.

2. Renegotiation as a Constraint on Collusion

The textbook theory of repeated games implicitly assumes that players can coordinate on any of the enormous range of subgame-perfect equilibria in the repeated game. This approach may seem reasonable if the players have ample opportunities for communication. But this very assumption also subverts the standard subgame-perfect analysis. If, indeed, players can efficiently coordinate on an equilibrium, what would really happen after (out of equilibrium) one player cheated? If the pre-specified punishment would hurt the innocent as well as the guilty, then we might well expect *renegotiation*, perhaps in the simple form of an agreement to ignore the transgression "this time". Such undiscriminating punishments may therefore lack credibility, even if they are subgame-perfect. This suggests that credible punishments must be renegotiation-proof, at least in the sense that not all players are hurt by them. Levenstein (1997), for instance, argues that the bromine cartel before World War I was unable to use Abreu-Pearce-Stacchetti (1986) punishments because they were not renegotiation-proof.

This line of thought led Bernheim and Ray (1989) and Farrell and Maskin (1989) to develop a theory of "weak renegotiation-proofness" in repeated games.[8] As the name suggests, weak renegotiation-proofness (WRP) is best conceived as a *necessary* condition for credibility of an equilibrium in a repeated game when (re)negotiation is completely

differ has been recognized in principle for a long time: see *e.g.*, Asch and Seneca (1976) — but not adequately studied.

[8] I use the Farrell-Maskin terminology: Bernheim and Ray called the same concept "partial Pareto perfection". Both of these papers also developed further, stronger renegotiation-proofness conditions, but those are unsatisfactory in various ways. Pearce (1987) and Rabin (1991) have also developed different theories of renegotiation in infinitely repeated games. Benoit and Krishna (1993) develop a theory for finitely repeated games, but (like the subgame-perfection theory) it predicts no collusion at all in finitely repeated versions of the simple oligopoly games we consider.

smooth. Attempts to develop sufficient conditions have been less successful, but this paper focuses on necessary conditions and shows that collusion among scores or hundreds of firms, which can be consistent with subgame-perfection as discussed above, is not even weakly renegotiation-proof. Specifically, I show that with more than three symmetrically placed firms in Bertrand oligopoly, full collusion (or anything close to it) is inconsistent with what might seem a reasonable form of renegotiation-proofness, even if the discount factor δ is arbitrarily close to 1. I also show that in Cournot oligopoly, full collusion is unsustainable (in the same sense) among more than nine firms with linear demand.

I shall repeatedly use the result (and the method of proof) of a characterization theorem for WRP equilibrium in two-player games, derived by Farrell and Maskin (1989). Their Theorem 1 states that in a two-player game a necessary condition for a pair of average payoffs,[9] $v \equiv (v_1, v_2)$, to be weakly renegotiation-proof (WRP) for large enough values of the discount factor δ is that there exist (possibly mixed) punishment-phase action-pairs $a^1 \equiv (a_1^1, a_2^1)$ and $a^2 \equiv (a_1^2, a_2^2)$, to punish players 1 and 2 respectively. These must be such that the punishment for each player is sufficiently severe, yet the "innocent" player is not tempted to renegotiate. Formally, for instance, the punishment action-pair a^2 must be such that, in the stage-game while player 2 is being punished, player 2 cannot (even by cheating) get a payoff of more than v_2 when player 1 plays a_1^2, and must be such that player 1's payoff $g_1(a_1^2, a_2^2)$ from (a_1^2, a_2^2) is at least v_1. (In a useful notation, we require that $c_2(a_1^2) \equiv \max_a g_2(a_1^2, a) \leq v_2$, and that $g_1(a_1^2, a_2^2) \geq v_1$.) Of course a similar condition must hold for the action-pair (a^1) that is used to punish player 1. Intuitively, the action-pair a^2 is used for a suitably chosen number of periods after player 2 deviates from prescribed behavior; then normal play resumes. Farrell and Maskin also showed that this condition (with strict inequalities replacing weak) is sufficient for v to be WRP for large enough δ.

It is important to note that although the innocent player's action during the punishment phase might hold the guilty player perhaps only slightly below his normal-phase payoff (that is, he could perhaps get almost his normal-phase payoff in a punishment period if he

[9] In a repeated game with discount factor δ, a player who gets payoff x_t in period t has a discounted payoff of $D \equiv \sum_0^\infty \delta^t x_t$ and an "average" payoff of $(1 - \delta)D$; the term makes sense because, for instance, if $x_t = x$ for all t then the average payoff is also x.

cheated on his punishment), he may get much less than that according to the prescribed actions in a punishment period. In the simple oligopoly context, it is easy to specify actions that hold down the guilty player's *intended* punishment-phase payoff, by telling him to stay out of the market; this also helps to make the innocent player unwilling to renegotiate. As we will see, it is a lot harder to hold down the guilty player's payoff from cheating during the punishment phase (without also hurting the innocent player).

3. Repeated Bertrand Oligopoly

Consider n identical price-setting firms selling undifferentiated products, each with constant marginal costs c and no capacity constraints. The demand side is represented by a demand curve $X = D(p)$, where X is total industry quantity and p is the price that consumers face (*i.e.*, the lowest price set by any firm). Let $\pi(p) \equiv (p - c)D(p)$ be industry profits when consumers must pay price p. We assume for simplicity that the function $\pi(\cdot)$ is strictly increasing in the range (c, p^m), where of course p^m is the monopoly price, *i.e.*, the price that maximizes $\pi(p)$. Write $\pi^m \equiv \pi(p^m)$.

In each of infinitely many periods, each firm i chooses its price p_i (all firms choose simultaneously within a period). Firm i's within-period payoff is $\pi(p_i) \equiv (p_i - c)D(p_i)$ if p_i is uniquely the lowest price, $\pi(p_i)/m$ if it is one of m equal lowest prices, and zero otherwise. Each firm's action becomes common knowledge before the beginning of the next period; thus, firms' actions can depend on what others have done in the past. The firms share a common and constant discount factor $\delta \in (0, 1)$. We shall be concerned with equilibrium (average) payoffs as $\delta \to 1$.

Following Farrell and Maskin (1989), we ask to what extent innocent firms can hold down the best-response payoff of a guilty firm in punishment periods, while simultaneously getting a good payoff themselves. Farrell and Maskin addressed this issue with the following duopoly result for the linear-demand case. We prove it here for a general demand curve, and then extend it to oligopoly with more than two firms.

Lemma 1. *In symmetric Bertrand duopoly with constant costs and no capacity constraints, the following results characterize an innocent firm's ability to minimize its guilty* *rival's best-response expected payoff v while maximizing its own expected payoff u. (a) If the innocent firm is restricted to pure strategies (that is, its price is known to its rival), then it cannot hold its rival's best-response payoff strictly below its own intended payoff: thus $u \leq v$. (b) If the innocent firm can commit to randomizing its price (in such a way that the distribution of its price, but not the price itself, is observable by its rival before the latter chooses its own price), then it can hold its rival's best-response expected payoff down to any $v \in (0, \pi^m)$, while itself earning an expected payoff of $u = v(1 + \log(\pi^m/v)) > v$.*

Proof. First, note that in choosing strategies a^2 to satisfy conditions of the form $g_1(a^2) \geq v_1$ and $c_2(a^2) \leq v_2$, we can assume without loss of generality that firm 2 plays only prices $p > p^m$ in a_2^2. That is, it is meant to set prices high enough that it gets no business while firm 1 sets prices $p \in [c, p^m]$. Of course this implies that $g_2(a^2) = 0$, but it is important to understand that a punishment strategy must limit $c_2(a^2) \equiv \max_a g_2(a_1^2, a)$, not (only) $g_2(a^2)$.

First consider pure strategies on the part of the innocent firm, firm 1. If it charges $a_1^2 = p$, its intended stage-game payoff $g_1(a^2) = u$ cannot exceed $\pi(p)$ (and any $a_2^2 > p^m$ will lead to that value of u). For any $\epsilon > 0$, the guilty firm can undercut p by a small enough amount to get a payoff that exceeds $u - \epsilon$. Thus it cannot be held to any best-response punishment-phase payoff strictly less than u. Except for an uninteresting opennness problem, we have $c_2(a_1^2) \geq g_1(a^2)$ (with equality for $c \leq a_1^2 = p \leq p^m$). This proves part (a).[10]

Now consider mixed price strategies a_1^2 for the innocent firm. To hold its rival's best-response expected payoff $c_2(a_1^2)$ down to v, the innocent firm must ensure, for every p, that if the guilty firm sets price p then it will have no more than a $v/\pi(p)$ chance of getting the whole market (we deal with ties below). Thus, if $F(\cdot)$ is the distribution function of the

[10] The reader may ask: if the innocent firm chooses a pure strategy p that would yield strictly positive $u > 0$, why can't the guilty firm undercut p infinitesimally and thus get almost u but give the innocent firm zero? But we are not comparing payoffs of a Stackelberg equilibrium. Rather, the question is, if (say) firm 1 chooses an intended punishment action-pair a^2 that involves an unmixed a_1^2, then how do $g_1(a^2)$ and $c_2(a_1^2)$ compare? The same clarification addresses a question a number of readers have asked: if v is nearly π^m, doesn't the Lemma claim that each firm can get the full monopoly profit? No: the Lemma says (in that case trivially) that firm 1 can get π^m (if firm 2 stays out of the market) using a pricing strategy under which firm 2's *cheating best response* wouldn't get it more than π^m.

innocent firm's price, we require that for all p such that $\pi(p) > v$,

$$F(p) \geq 1 - \frac{v}{\pi(p)}. \tag{1}$$

Since the punishment-design task is to maximize the innocent firm's payoff (with the guilty firm out of the market) while holding the guilty firm's best-response payoff down to v, and since $\pi(\cdot)$ is increasing in the relevant range, we can assume without loss that condition (1) holds with equality for $p \in (p(v), p^m)$, where $p(v)$ is defined by $\pi(p(v)) \equiv v$, and the remaining mass of v/π^m is at p^m.

By construction, this implies that firm 2 cannot achieve an expected payoff strictly above v by setting any price other than the monopoly price. To show that the mixed strategy a^2 in question truly holds firm 2 down to an expected payoff of v, we need only check that its expected payoff from $p = p^m$ is no more than v. But this expected payoff is equal to v even if all ties are resolved in firm 2's favor. Under this strategy, moreover, when firm 2 sets prices above p^m, firm 1's profits are

$$u(v) = \int_{p(v)}^{p^m} \pi(p) \mathrm{d}F(p) + \frac{v}{\pi^m}\pi^m.$$

Substituting for $\mathrm{d}F$ from (1), integrating by parts, and changing variable from p to π, we find that the innocent firm's expected profits are

$$u(v) = v\left(1 + \log\frac{\pi^m}{v}\right).$$

This proves part (b) of the Lemma. ∎

Lemma 1 implies the following result, which was shown by Farrell and Maskin (1989, pp. 339–441) for the linear-demand case:

Proposition 2. *In the simplest repeated Bertrand duopoly, symmetric collusion is possible in* WRP *equilibrium for large enough discount factors; so is asymmetric collusion provided that it is not too asymmetric.*

Since all the relevant calculations are in profit space, the calculation of how asymmetric the profit shares can be is no different from that in Farrell and Maskin: each firm must get no less than a fraction z of the shared monopoly profits, where $z(2 - \log z) = 1$, or $z \approx 0.32$. As we will see, it is important that this number is less than a third but more than a quarter.

4. Beyond Duopoly: Asymmetric and Quasi-Symmetric Punishments

With two players, it is natural to suppose that renegotiation is blocked if the innocent party would benefit from implementing the agreed-upon punishment. With more than two players, we have to ask who can block renegotiation — equivalently, how many of the innocent would have to join the guilty party in order to push renegotiation through.

Obviously, a wide variety of answers would be possible. One candidate would be that any player can block renegotiation: with this assumption, weak renegotiation-proofness means that no continuation equilibrium of the equilibrium can Pareto-dominate any other. This is perhaps the most obvious generalization of the two-player definition. It turns out to imply that firms "should" use highly asymmetric punishments and that by doing so they can sustain a great deal of collusion.

Indeed, with this assumption, punishments can be designed as follows: if firm i has cheated, nominate some other firm $j \neq i$ to be the renegotiation-blocker, and (during the punishment phase) concentrate all production in firm j, specifying that firm i and all firms $k \neq j$ produce nothing. This makes weakly renegotiation-proof (in this sense) collusion very easy. Formally, a simple extension of the two-firm analysis implies the following:

Proposition 3. *Consider repeated Bertrand oligopoly and assume that asymmetric punishments are allowed and that any one firm can block renegotiation. Then, the average payoffs* $(\pi_1, \ldots, \pi_n)$ *are* WRP *for large enough* δ *if for every* i *there is a* $j \neq i$ *such that* $\pi_j < \pi_i(1 + \log \dfrac{\pi^m}{\pi_i})$. *In particular, symmetric division of the monopoly profit among* n *symmetric firms is* WRP *(with large enough* δ*) for all* n.

Proof. Lemma 1 tells us that a Bertrand duopolist can hold its rival to a maximum expected payoff of v (where $v < \pi^m$ is given) while itself earning up to $v(1 + \log[\pi^m/v])$ if the rival keeps out of the market (as we can assume it is meant to do during punishment).

Hence, for firm i to be effectively punished without hurting firm j requires that we can find v such that $v(1+\log[\pi^m/v]) \geq \pi_j$ while $v < \pi_i$. This condition ensures that punishment is possible and that firm j would block renegotiation. We have in effect reduced the problem to a duopoly problem by telling all firms other than j — including both the "guilty" firm i and all "third" firms — to stay out of the market during i's punishment phase.

We must also consider the incentives of firms other than i to cheat during i's punishment. We can specify that if any firm cheats during another's punishment then the original punishment is canceled and the most recent cheater is punished. With this (standard) assumption, firm j will not cheat during i's punishment if it would not cheat during the normal phase, since it is getting more payoff without cheating and yet it would be punished just as severely for cheating. Other firms $k \neq i, j$ are at worst being asked to remain out of the market, as if they were being punished themselves, and so if no firm would cheat on its own punishment then none will cheat on another's.

The claim then follows on substituting $\pi_i = \pi^m/n$ for all i. ∎

Clearly, unrestricted asymmetric treatment of the innocent firms is a very powerful tool for blocking unanimous renegotiation. Indeed, Horniacek (1996) has applied similar arguments, exploiting asymmetries of treatment among innocent players after a deviation, to argue that a version of weak renegotiation-proofness, and even a weakened version of strong perfect equilibrium (Rubinstein 1980), is little constraint on cooperation when there are three or more players.

In symmetric oligopoly games, however, one might expect that the innocent firms will all be treated alike in the punishment of a guilty firm. In that case, and in some cases even where not all are treated exactly alike, either all the innocent firms gain from that punishment (so that renegotiation will plausibly be blocked), or all lose (making renegotiation likely). Following that idea, we extend the two-player definition of a weakly renegotiation-proof equilibrium to the many-player case as follows:

Definition. *A subgame-perfect equilibrium is quasi-symmetrically weakly renegotiation-proof (*QSWRP*) if, evaluated at the beginning of the period after player i alone deviates* *from prescribed play, every other player's continuation payoff (weakly) exceeds what it would have been had player i not just deviated. Thus no innocent player would want to forget and forgive.*

In the remainder of the paper I will show that symmetric oligopolies with more than a handful of firms cannot sustain full — or, indeed, much — collusion in QSWRP repeated Bertrand equilibrium, even for discount factors very near 1, and even if we allow for verifiable randomizations in punishment periods.

5. Full Collusion in Repeated Bertrand Competition

We begin by revisiting the textbook repeated Bertrand oligopoly model and asking again whether or not full collusion can be sustained when the discount factor is high. Our answer stands in stark contrast to the subgame-perfect theory.

Proposition 4. *In repeated Bertrand competition, full collusion is impossible in* QSWRP *equilibrium if $n > 3$, even for discount factors very close to 1.*

Proof. Again we use the result of Lemma 1. Suppose that full collusion is a QSWRP equilibrium with n firms. Without loss of generality, there is at least one firm, firm 1, say, whose normal-phase payoff is no greater than π^m/n; let $\pi_1 \leq \pi^m/n$ be its normal-phase average payoff.

Since the claim is that for $n > 3$ collusion is impossible, let us tactically suppose that the innocent firms can correlate their strategies and are not constrained by issues of distribution of profits among themselves. This reduces their problem to choosing a distribution of the lowest price quoted by any of the innocent firms: call this lowest price p_0. They must choose the distribution of p_0 so as to give themselves expected profits of at least $\pi^m - \pi_1$ while holding down an opportunistic Bertrand rival to expected profits of less than π_1. The problem is that making it impossible for the guilty firm to make high expected profits even by "cheating on its punishment" requires that with high probability p_0 must be small: but this conflicts with the need (so as to avoid the temptation to renegotiate) to make profits for themselves.

By assuming that the innocent firms can correlate their strategies, we reduce the problem analytically to that of a single innocent firm who must make a profit of $\pi^m - \pi_1$ while holding a guilty firm down to π_1. Thus, by Proposition 1, a necessary condition is that

$$\pi^m - \pi_1 \leq \pi_1 \left(1 + \log \frac{\pi^m}{\pi_1}\right).$$

Since $\pi_1 \leq \pi^m/n$, this implies that

$$\frac{n-1}{n}\pi^m \leq \frac{\pi^m}{n}(1 + \log n),$$

which reduces to $n - \log n \leq 2$, or (in integer terms) $n \leq 3$.

Thus $n \leq 3$ is a necessary condition for QSWRP full collusion (symmetric or not) in infinitely repeated Bertrand oligopoly, however large the discount factor δ, and even if the innocent firms can correlate their strategies, and their randomizations are observable, in a punishment phase. ∎

The proof essentially notes that each firm must get at least a critical share of the shared monopoly profits. This is just as in the duopoly case, and the same calculation (as in Farrell and Maskin) implies that this critical share is slightly less than a third: four firms or more cannot all get this much. Three can, so we might expect that three can collude according to this criterion.

If the innocent firms cannot correlate their strategies in the punishment phase, then the distribution of p_0 must be the minimum-value distribution of $n - 1$ independently distributed prices, and each innocent firm must have a sufficient chance of winning, at a sufficiently high price, that it would not prefer to renegotiate back to the status quo as if the deviation had not happened.

To see whether symmetric collusion is sustainable in this way, we need only consider the case $n = 3$, since we have shown that such collusion is impossible anyway for $n > 3$ and Proposition 2 showed that it is possible with $n = 2$. Here, we show that correlation of strategies in the punishment period is not essential for collusion with $n = 3$, but that observability of randomization is.

Proposition 5. *Two or three firms can symmetrically collude in* QSWRP *equilibrium in repeated Bertrand oligopoly for high enough discount factors even if the innocent firms cannot correlate their punishment randomizations.*

Proof. Consider a guilty firm that must be held to a maximum (cheating-on-its-punishment) payoff of v, while the other two (who cannot now correlate their randomized prices) try to do as well as they can subject to inflicting this punishment. Suppose that each randomizes with distribution function $F(\cdot)$ on $[p(v), p^m]$ — there is no need to charge prices less than $p(v)$ and it is foolish to charge prices greater than p^m. Define $G(p) \equiv 1 - F(p)$. Then the firm being punished can get, in the short run, $\max_p \pi(p)G(p)^2$. The innocent firms want to choose $G(\cdot)$ to keep this down to v and, subject to that, to maximize their expected profits. Since $\pi(\cdot)$ is monotonic, the innocent firms want to charge high prices as much as possible. Therefore $G(\cdot)$ will satisfy

$$G(p) = \sqrt{\frac{v}{\pi(p)}}, \quad p(v) \le p < p^m,$$

and the remaining $\sqrt{v/\pi^m}$ weight is an atom at price p^m.

Now consider the expected profit of each of the innocent firms. An innocent firm that bids $p < p^m$ wins the market with probability $G(p)$; if it bids $p = p^m$ there is a positive probability of a tie, which I assume leads to equally divided profits. Thus an innocent firm's expected profit (within a punishment-phase period) is

$$\int_{p(v)}^{p^m} \pi(p)G(p)\mathrm{d}F(p) + \sqrt{v/\pi^m}(1 - \tfrac{1}{2}\sqrt{v/\pi^m})\pi^m,$$

We can re-write this in terms of the variable π as

$$\int_{v}^{\pi^m} \pi\sqrt{v/\pi}\ \tfrac{1}{2}\sqrt{v/\pi^3}\,\mathrm{d}\pi + \sqrt{v/\pi^m}(1 - \tfrac{1}{2}\sqrt{v/\pi^m})\pi^m,$$

which is

$$\tfrac{1}{2}v\log\frac{\pi^m}{v} + \sqrt{v\pi^m} - \tfrac{1}{2}v. \tag{2}$$

When $v = \frac{1}{3}\pi^m$, (2) is approximately $.59\pi^m > \frac{1}{3}\pi^m$, so each innocent firm can get more than $\pi^m/3$ while holding a guilty firm strictly below $\pi^m/3$ during a punishment period.

Consequently, symmetric Bertrand collusion with three firms is QSWRP for large enough δ even if the innocent firms cannot correlate their punishment strategies. ∎

Note that this is *not* a mixed-strategy equilibrium in the sense that each of the innocent firms expects an equal payoff from each of the prices over which it randomizes. On the contrary: an innocent firm that sets price $p < p^m$ gets an expected this-period profit of $\pi(p)G(p) = \sqrt{v\pi(p)}$, which is strictly increasing in p (for $p < p^m$). Thus, innocent firms are tempted (in the short-run sense) to deviate from their specified punishment strategy during punishment.

If randomizations are unobservable, collusion is not sustainable at all. For, in order to punish effectively, the innocent firm(s) must randomize with support going all the way down to $p(v)$, or else the guilty firm can get more than v by undercutting. But if the innocent firms were playing a mixed-strategy equilibrium in the sense that their short-run payoffs were equal across realizations of the randomization, this would mean that they (jointly) were not expecting more than v. That is, the innocent firms' *joint* payoff during a punishment phase cannot exceed the guilty firm's best-response payoff during that punishment phase. In turn, this latter must be strictly less than the guilty firm's (per-period) normal-phase payoff, unless the guilty firm is playing a best response during the normal phase. Applying these observations to the firm with the lowest normal-phase payoff, we see that no collusion is sustainable. ∎

6. Partial Collusion in Repeated Bertrand Oligopoly

If full collusion is impossible, can the firms collude on a price between c and p^m? One of the unsatisfactory predictions of the subgame-perfect theory is that it answers no: either full collusion is subgame-perfect, or no collusion is.[11] Here, by contrast, firms who are too numerous to collude fully may be able to collude partially, *i.e.*, on a price greater than c

[11] The reason is that for firms to collude in symmetric subgame-perfect equilibrium on a price p that yields industry profits of $\pi(p)$ requires precisely that the present value of one-nth of a perpetual stream of $\pi(p)$ be at least equal to $\pi(p)$. Clearly, $\pi(p)$ and p drop out of this comparison and all that matters is n versus δ.

but less than p^m.[12] However, in numerical terms, I find that the efficiency consequences (deadweight losses) due to this partial collusion are small, even if δ is close to 1 and even if n is as small as five.[13]

The function $v(1+\log[\pi^m/v])$, which gives the profits the innocent firms together can achieve while effectively punishing the guilty firm by holding its payoff below v, has an infinitely steep slope for small v. Thus, for any n,

$$(n-1)v < v(1+\log\frac{\pi^m}{v})$$

for small enough v. Since this is the condition for a payoff of v per firm to be sustainable in QSWRP equilibrium for large enough discount factors δ, it follows that, for every n, *some* collusion is possible in QSWRP equilibrium for large enough δ. But even for $n \approx 5$, the condition implies a small value of nv/π^m (little collusion), as we now show.

Suppose the firms can collude on a price p that yields industry profits $\pi(p)$. By the same argument as before, for this to be sustainable in QSWRP equilibrium requires that

$$(n-1)\frac{\pi(p)}{n} \leq \frac{\pi(p)}{n}(1+\log\frac{\pi^m}{\pi(p)/n}),$$

which reduces to $\pi(p)/\pi^m \leq ne^{2-n}$.

For example, if $n=4$ (the smallest number of firms that cannot fully collude), they can collude only on prices p such that $\pi(p)/\pi^m \leq 4e^{-2} \approx .54$.[14] A price that yields roughly half of the monopoly profit implies only a smaller fraction of the monopoly markup and a smaller fraction still of the deadweight loss. For instance, with constant-elasticity demand, Table 1 shows that the maximum sustainable partial collusion leads to markups well below

[12] For parameter values for which the renegotiation-proof theory predicts partial collusion, the subgame-perfect theory predicts full collusion, not none. This follows from the fact that every WRP equilibrium is a subgame-perfect equilibrium.

[13] Even if $n=4$ the losses are not large. A previous version of the paper haltingly explored a different efficiency issue: the extent to which QSWRP equilibrium must give substantial market shares even to high-cost firms.

[14] Numerical calculations also indicate that for $\delta = .99$, the four firms can capture only a fraction .51, rather than .54, of the potential monopoly profits.

monopoly markups and to a deadweight loss of no more than an eighth of the monopoly deadweight loss for four firms facing a wide range of demand elasticities.[15]

Table 1 about here

For $n = 5$, the condition is that $\pi(p)/\pi^m \leq 5e^{-3} \approx .25$. Again, results for constant-elasticity demand are shown in Table 1, suggesting a deadweight loss of no more than two percent of the monopoly deadweight loss for elasticities between 1.1 and five.[16] For $n = 6$, $\pi(p)/\pi^m \leq 6e^{-4} \approx .11$, which corresponds to still smaller markups and deadweight losses.[17] For $n \geq 7$, $\pi(p)/\pi^m \leq ne^{2-n} \leq 7e^{-5} \approx .05$.

These prices are quite low by the standards of conventional oligopoly models. For $n \geq 5$, calculations suggest that the imperfectly collusive prices in this model are significantly less than those in a one-shot Cournot model (see Table 2). Moreover, since deadweight loss is of second order in the markup, these fairly small markups — even for $n = 4$ or $n = 5$ — have small welfare consequences, and for larger values of n the maximum deadweight loss declines dramatically.

Table 2 about here

---

[15] Because we relate outcomes to the monopoly outcome, the demand elasticity must exceed 1. A similar calculation assumes linear demand $X = 2 - p$ and zero marginal costs; then the monopoly price is 1. A collusive profit of .54 corresponds to a collusive price of approximately .33, a third of the monopoly markup, and hence a ninth of the monopoly deadweight loss (since loss is proportional to the square of the markup).

[16] In the linear case the maximum collusive markup with five firms is approximately $p \approx .15$, with deadweight loss of $\frac{1}{2}(.15)^2$ or about two percent of the monopoly deadweight loss.

[17] For our linear-demand example this corresponds to a markup of less than six percent of the monopoly markup, and a deadweight loss of less than 0.4 percent of the potential monopoly deadweight loss. For our constant-elasticity cases shown in Table 1, it produces a deadweight loss of about 0.3 percent of the monopoly deadweight loss.

7. Repeated Cournot Competition

I now consider the repeated Cournot case. As above, I first derive conditions on n for shared monopoly to be QSWRP (for discount factors very close to 1); I then ask how much firms can collude if those conditions fail. In this section I work with linear demand $p = 2 - X$: neither the general demand case nor the constant-elasticity case has proven tractable. To keep the notation simple I assume zero marginal costs; this is a mere normalization once one assumes constant and equal marginal costs.

Full Collusion

Proposition 6. *In repeated Cournot oligopoly, full collusion is not* QSWRP *with more than nine firms, even for δ very close to 1.*

Proof. If the innocent firms each produce y during the punishment phase, then (since we can assume without loss of generality that the firm being punished produces nothing) each innocent firm gets a per-period payoff during the punishment phase of $\pi^I = y(2-(n-1)y)$. Meanwhile, by cheating during a punishment period, the guilty firm could achieve a single-period payoff of $\chi \equiv \max_z z(2-(n-1)y-z) = \frac{1}{4}(2-(n-1)y)^2$. Thus Theorem 1 of Farrell and Maskin (1989) requires that it must be possible to choose y so that π^I is at least as large as the normal-phase (collusive) per-period payoff, which is $1/n$, while $\chi < 1/n$. This amounts to requiring that y satisfy $y(2-(n-1)y) \geq \frac{1}{n}$ and $\frac{1}{4}(2-(n-1)y)^2 < \frac{1}{n}$. These two conditions on y can hold simultaneously if and only if $n < 9$. To see this, write $a \equiv 2-(n-1)y$, so that the conditions are $a(2-a) \geq \frac{n-1}{n}$ and $a^2 < \frac{4}{n}$. These can both be made to hold if and only if the first can be made to hold strictly with the second one holding with equality, so we can substitute $a = \frac{2}{\sqrt{n}}$ in the first condition to yield $(\sqrt{n}-2)^2 < 1$, or $n < 9$. ∎

As a technical observation, note that the case $n = 9$ is right on the boundary: the conditions can be made to hold weakly but not strictly. This means that the necessary condition for WRP equilibrium holds, but the sufficient condition does not: we must remain agnostic about the possibility. As we will see next (and as is not surprising given the calculations above), nine firms can all but fully collude in QSWRP equilibrium in repeated Cournot play with our cost and demand assumptions, if δ is close to 1.

Partial Collusion

Now we put a bound on the amount of profit (and welfare loss) that can stem from QSWRP partial collusion in Cournot oligopolies with $n \geq 9$ firms, in which (as we have seen) complete symmetrical collusion is not a QSWRP equilibrium.

Proposition 7. *If $n > 9$, the firms can symmetrically divide in* QSWRP *equilibrium a total profit of less than* $\frac{16n}{(n+3)^2}$, *corresponding to a price of less than*

$$p = 1 - \frac{\sqrt{(n-9)(n-1)}}{n+3}.$$

Proof. To sustain a per-firm profit of π, there must be punishment actions satisfying the now-familiar conditions. Let Y be the aggregate output of the $n-1$ innocent firms in a punishment period; then we require that $Y(2-Y) \geq (n-1)\pi$ and $\frac{1}{4}(2-Y)^2 \leq \pi$. Writing $q \equiv \pi^{1/2}$, we require $Y(2-Y) \geq (n-1)q^2$ and $2 - Y \leq 2q$; hence, $q \leq \frac{4}{n+3}$, and the Proposition follows when we set $\pi = p(2-p)/n$. ∎

For comparison, under one-shot Cournot equilibrium, the price is $\frac{2}{n+1}$ and total profit is $\frac{4n}{(n+1)^2}$; and of course price and total profit is 1 under full collusion. For large n, therefore, roughly four times more profit is sustainable in QSWRP equilibrium than without collusion, but vanishingly less than under full collusion, even as $\delta \to 1$.

8. Efficient Punishments

The results above show that frictionless renegotiation would make self-enforcing collusion impossible for even moderate numbers of firms. It would also be interesting to derive results in the other direction, suggesting when collusion *is* possible. Unfortunately, since weak renegotiation-proofness is not a convincing sufficient condition for credibility, the analysis above cannot help very much with that. But if collusion can be sustained by means of punishments that themselves divide monopoly profit among the innocent firms, then full collusion will be a "strong perfect equilibrium" (Rubinstein 1980).

This can never happen in repeated Bertrand competition, even with $n = 2$, as Farrell and Maskin (1989) noted. The problem is that an effective punishment must involve randomization by the innocent firm(s), which is inconsistent with their collecting the full monopoly profit during a punishment phase.

Turning to repeated Cournot oligopoly (with our linear demand and zero cost assumptions), when can full collusion be sustained with efficient punishments? In a punishment phase, the $n-1$ innocent firms divide the monopoly output (*i.e.*, 1) among them. By cheating in this phase, the guilty firm can achieve a per-period payoff of $\max_y y(2-1-y) = \frac{1}{4}$. Hence, each firm must get *strictly* more than this per period in the normal phase; if a firm got $\frac{1}{4}$ or less in the normal phase then it should cheat for a period and could not be effectively punished, given that only Pareto-efficient punishments are being considered. To give each firm strictly more than $\frac{1}{4}$ in the normal phase requires $n \leq 3$. At the same time, it is clear that with $n \leq 3$ this punishment (suitably repeated) works to sustain collusion for δ close enough to 1. Thus we have:

Proposition 8. *There is a strong perfect equilibrium (in which both the equilibrium path and punishments involve the monopoly price) in infinitely repeated Cournot competition with linear demand and with sufficiently little discounting, if and only if $n \leq 3$. No such equilibrium is possible in repeated Bertrand competition.*

9. Rebuttal and/or Conclusion

In this somewhat inconclusive concluding section, I try to distinguish more carefully between explicit collusion and conscious parallelism — a distinction that most game theory ignores — and in the process raise some doubts about the interpretation of the results above.

A priori, we surely should analyze renegotiation as well as negotiation of collusive agreements. If one likes its details (in particular the QSWRP assumption), the model above therefore seems appealing on the input side. On the output side, the model says that a few firms can sustain significantly super-competitive prices, but only a few. It also says that oligopolies can sustain somewhat super-competitive prices when they cannot sustain a monopoly price, but only to a (rapidly) decreasing extent as the number of firms grows. These predictions resonate appealingly with the structural consensus. Dare one thus hope that this is broadly the "right" theory of self-enforcing collusion in oligopoly?

Unfortunately, I am skeptical. To discuss why, let us become more careful in our language: *collusion* is explicit agreement or communication to raise prices; *coordination* is raising prices with or without such communication. The theory above predicts that (in unconcentrated industries) it is hard to coordinate on high prices when there is (and will always be) explicit collusion, whereas the structural consensus is that (in such industries) it is hard to coordinate on high prices if explicit collusion is *not* possible. Since collusion surely affects firms' ability to coordinate, these are not directly comparable predictions.[18] The discussion below is not very deep, but it took me some time to clarify to this point, so I hope it will be helpful to others.

It is helpful to distinguish three negotiation environments for oligopoly pricing. The first is the purely theoretical benchmark of *perfect negotiation.* This is plainly assumed above, and also by at least the more advanced forms of the textbook "subgame-perfect" theory: both equilibrium pricing and punishments can be negotiated in a sophisticated way, and everyone understands what is going on, sometimes even down to coordination of random pricing.[19] As this paper assumes, in this environment one might also expect perfect renegotiation.

In the second environment, there is *imperfect negotiation.* Imperfections arise for many

[18] Of course, in cross-section, explicit collusion (where it's illegal) will tend to occur where other factors make non-collusive coordination at least somewhat difficult (otherwise why risk prison terms?). Thus, in principle, collusion might even be negatively correlated with success at raising price. One should not confuse this correlation with the effect of collusion. I thank Dennis Carlton for this observation. See also Werden and Baumann (1986).

[19] In the literature, the oldest papers assumed reversion to one-shot Nash behavior, while more recent work (including this paper) tends to construct optimal punishments, which can be quite complex. Researchers of course recognize informally that the latter impose heavier demands on negotiation, and thus more thoroughly assume perfect collusion, but there does not seem to be analysis following up on this point.

reasons, often including asymmetries among firms, imperfect information about one another's costs and about demand, and perhaps also the need to keep negotiations (collusion) secret. Nevertheless, explicit negotiation does take place in this environment: firms explicitly agree on prices and quite possibly on what will happen if deviations are detected. Any renegotiation will presumably also be imperfect. This environment represents the case in which firms collude.

In the third environment, *no explicit negotiation* takes place. Firms recognize their interdependence, and may engage in some relatively costly and clumsy signaling of intentions, but signaling bandwidth is quite limited.

In this framework, the model above, and the textbook models, clearly concern the perfect-negotiation environment. The structural consensus, however, concerns the observed mix of the imperfect-negotiation and the no-negotiation environments. Thus, to make much of the assonance between my results and the structural consensus would require (something like) that the three negotiation environments do not differ very much in terms of the prevalence of successful price coordination. That is not something one would want to assume. Although there is some academic debate on the point, I think it is natural to believe that the laws against collusion matter, which means that the imperfect-negotiation and no-negotiation environments differ substantially. Let me call this natural hypothesis the *concern over collusion*. It does not contradict the structural consensus if the observed regime is one in which prohibitions on collusion are often obeyed. Then it just says, very plausibly, that coordinated high pricing can become much easier with explicit negotiation, whether this is when explicit collusion is legal (see Dick 1996 on Webb-Pomerene cartels and Suslow 1991 on international cartels) or when prohibitions are violated (see Fraas and Greer 1977 on price-fixing prosecutions).[20]

[20] Very briefly, some kinds of evidence: (a) there continue to be convictions and confessions in *per se* price-fixing cases, indicating that executives are sometimes willing to risk prison in order to collude; (b) where collusion is legal, cartels of substantial size seem to spring up: see for instance Dick (1996) and Suslow (1991); (c) the structural consensus strongly suggests that it would be rare for cartels of comparable size to arise absent collusion; (d) both lawmakers and business people seem to believe that the laws matter.

Is Coordination Monotonic in Coordination Costs?

Absent renegotiation, one can (and most people do) take a simple monotone view of the role of the three negotiation environments. Coordinated pricing would be rather easy in the theoretical perfect-negotiation environment (the textbook models), is harder with imperfect negotiation (how much harder depending on how imperfect), and is harder still (especially if more than a handful of firms are involved) with no explicit negotiation. This view seems consistent with (a) the calculations discussed in the Introduction about how many firms can collude if negotiation is perfect, (b) limited evidence on the extent of successful collusion with imperfect negotiation, and (c) the structural consensus, if the legal prohibitions on explicit collusion are mostly obeyed (at least outside the most concentrated industries). The great gulf between the textbook calculations and the structural consensus, which I presented in the Introduction as a problem or puzzle, is then viewed just as a big difference between the perfect-negotiation regime and the observed regime. Antitrust enthusiasts might attribute more of the difference to the difference between imperfect negotiation and the observed regime; antitrust cynics might attribute more of it to the difference between perfect and imperfect negotiation.

Thinking about renegotiation complicates this picture considerably. As McCutcheon (1997) has stressed, when obstacles to negotiation also obstruct renegotiation (more precisely when they are expected to do so), it is not so obvious that more obstacles will always lead to more competitive pricing. In a moderately unconcentrated industry, the model above says that successful collusion will be rare in the perfect-negotiation environment, while the concern over collusion says that it will be more common in the imperfect-negotiation environment than in the no-negotiation environment.

This view implies that when renegotiation is possible, the extent of supercompetitive pricing is *not everywhere monotonic* in the ease of negotiation. Because of this non-monotonicity, it is hard to regard the model above as "the right explanation" for the structural consensus, even if they give similar predictions.

This non-monotonicity is also superficially consonant with McCutcheon's concern that anti-collusion policy might be counterproductive. However, here, the implied paradoxical effect (in which banning smoke-filled meetings leads to higher prices) occurs between the perfect-negotiation environment and the imperfect-negotiation environment. This is not where anti-collusion policy operates: recall that perfect negotiation is only a theoretical benchmark. Legal prohibitions on explicit collusion, to the extent that they are respected, shift us from the regime of imperfect negotiation to the regime of no explicit negotiation; to the extent that they are covertly violated, they leave us in the second regime but presumably make negotiation harder, sketchier, and more imperfect.

All of this makes it seem important to distinguish, as does the law, between explicit collusion and coordination without collusion. Game-theoretic models typically gloss over this difference: by staying firmly within the perfect-negotiation environment, they do not address the difference in outcomes between the imperfect-negotiation and no-negotiation environments. Game theory generally proceeds by asking whether or not *there is* an equilibrium with certain (as for instance collusive) properties. For the most part, it does not address whether such an equilibrium can readily be *attained*.[21]

The 1984 *Guidelines* made a very explicit distinction on these lines. In section 3.11(a), they stated of markets with HHI below 1000 that "implicit coordination among firms is likely to be difficult and ...the prohibitions of section 1 of the Sherman Act [forbidding collusion] are usually an adequate response to any explicit collusion that might occur [...]." There was no claim that explicit collusion would be prohibitively difficult (if attempted) simply because of the moderately large number of firms in such a market.[22] Thus they incorporated the *concern over collusion* in their discussion of the *structural consensus*.

Having now put my model in its place (I think), what is the bottom line? Clearly, neither negotiation nor renegotiation is perfect in reality. The model does show that taking

---

[21] Parts of the literature on "cheap talk" (see *e.g.*, Farrell and Rabin 1996, especially page 114) address this point, but so far it has probably conferred intellectual respectability on the question more than it has provided answers. Obviously analysis of talk is particularly relevant to collusion. Other refinement concepts might, generously, be seen as addressing the question in the form of "which equilibrium will happen?"

[22] It is not part of the argument here, but I would note my own skepticism at the 1984 suggestion that merger policy should not be used simply because explicit collusion is already illegal. Presumably nobody in the non-Antitrust divisions of the Justice Department was claiming that locks were unnecessary because burglary was illegal.

account of the prospects for renegotiation can affect *ex ante* incentives very greatly. The dramatic difference between the renegotiation-proof and the subgame-perfect predictions (in the perfect-negotiation environment where they live) may suggest that neither one is a reliable policy guide in its precise predictions, although studying both will surely help one to assess real-world threats of collusion. Comparison of the model with what we believe is true in practice (which presumably reflects the possibility of imperfect renegotiation or its no-negotiation analogue)[23] further suggests that the interaction with the negotiation environment is quite subtle.

Finally, (within its world, but in contrast to the textbook models) the model says that it is difficult to sustain collusion using only threats of price retaliation. This difficulty might suggest a strong private incentive to develop richer cartel institutions that would hamper deviations, ease the punisher's pain (or exacerbate the punished firm's), or retard renegotiation.[24]

[23] Dick (1996) studies attempts at "reorganization" of (legal) cartels after a breakdown of discipline. This might be a promising source of information on renegotiation, with the usual caveat that the observed cases represent *actual* defections, while the theory concerns how the prospect of defection and renegotiation affects how cartels might be designed so as to eliminate defections.

[24] I thank Dennis Carlton for this point.

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1986) "Optimal Cartel Equilibria with Imperfect Monitoring," *Journal of Economic Theory*, **39**, 251–269.

Asch, Peter, and Seneca, J. J. (1976) "Is Collusion Profitable?," *Review of Economics and Statistics*, **58**, 1–12.

Benoit, Jean-Pierre, and Vijay Krishna (1993) "Renegotiation in Finitely Repeated Games," *Econometrica*, **61**, March, 303–324.

Bernheim, B. Douglas, and Debraj Ray (1989) "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, **1**, 295–326.

Dick, Andrew (1996) "When Are Cartels Stable Contracts?," *Journal of Law and Economics*, **39**, 241–283.

Farrell, Joseph and Eric Maskin (1989) "Renegotiation in Repeated Games," *Games and Economic Behavior*, **1**, 327–360.

Farrell, Joseph and Matthew Rabin (1996) "Cheap Talk," *Journal of Economic Perspectives*, **10**, 103–118.

Fraas, Arthur and Douglas Greer (1977) "Market Structure and Price Collusion: An Empirical Analysis," *Journal of Industrial Economics*, **26**, 21–44.

Horniacek, Milan (1996) "The approximation of a strong perfect equilibrium in a discounted supergame," *Journal of Mathematical Economics*, **25**, 85–107.

Levenstein, Margaret (1997) "Price Wars and the Stability of Collusion: A Study of the Pre–World War I Bromine Industry," *Journal of Industrial Economics*, **45**, 117–137.

McCutcheon, Barbara (1997) "Do Meetings in Smoke-Filled Rooms Facilitate Collusion?," *Journal of Political Economy*, **105**, 330–350.

Pearce, David (1987) "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation," mimeo, Yale.

Potter, Donald V. (1991) "Success Under Fire: Policies to Prosper in Hostile Times," *California Management Review*, **33**, 1–15.

Rabin, Matthew (1991) "Reneging and Renegotiation," University of California, Berkeley, working paper.

Rabin, Matthew (1994) "A Model of Pre-Game Communication," *Journal of Economic Theory*, **63**, 370–391.

Rubinstein, Ariel (1980) "Strong Perfect Equilibrium in Supergames," *International Journal of Game Theory*, **9**, 1–12.

Schmalensee, Richard (1987) "Horizontal Merger Policy: Problems and Changes," *Journal of Economic Perspectives*, **1**, 41–54.

Shapiro, Carl (1989) "Theories of Oligopoly Behavior," in *Handbook of Industrial Organization*, ed. Richard Schmalensee and Robert Willig. Amsterdam: North Holland.

Suslow, Valerie (1991) "Cartel Contract Duration: Empirical Evidence from International Cartels," Working paper, University of Michigan Business School.

Tirole, Jean (1988) *The Theory of Industrial Organization.* Cambridge: MIT Press.

United States Department of Justice and Federal Trade Commission (1984, 1992) "Horizontal Merger Guidelines," Available at ftc.gov and usdoj.gov/atr websites.

Werden, Gregory J., and Michael G. Baumann (1986) "A Simple Model of Imperfect Competition in which Four are Few but Three are Not," *Journal of Industrial Economics*, **34**, 331–335.

White, Lawrence J. (1987) "Antitrust and Merger Policy: A Review and Critique," *Journal of Economic Perspectives*, **1**, 13–22.