An Experimental Test of Tradeoffs between Discount Rates and Number of Firms in Supporting Collusion

Robert M. Feinberg
American University

Thomas A. Husted
American University

Matthew Waskiewicz
Cardiff University

One prediction of oligopoly theory is that there should be a tradeoff between discount rates (rates of time preference) and the number of competitors in a market, in supporting the possibility of collusive equilibria. Here we conduct a series of laboratory experiments with markets of 2, 3, and 4 firms, and discount rates explicitly accounted for, and examine whether the tradeoffs predicted in theory occur in the behavior of our subjects. We find that an increased number of firms in a market is associated with larger market output (and lower prices), reflecting the generalized Cournot result throughout. We fail to observe an impact of higher discount rates in further limiting collusive behavior.

INTRODUCTION

Feinberg and Husted (1993) were among the first to investigate the role of time preferences in determining collusive behavior in markets. As predicted, laboratory subjects (playing a role as business firms) faced with more heavily discounted future rewards were more inclined to compete in the present period, less likely to attempt to collude (though the effects were not large). However, one prediction of oligopoly theory is that there should be a tradeoff between discount rates (rates of time preference) and the number of competitors in a market, in terms of supporting the possibility of collusive equilibria. Essentially, as there are more firms in a market, tacit collusion gets harder to sustain and this can only occur with lower discount rates (more patience, greater willingness to wait for future rewards). To our knowledge, no research examining this tradeoff has been done.

Feinberg and Husted (1993) varied discount rates but kept the number of market participants fixed at two. On the other hand, the few works to examine the issue of varying numbers of firms (Fouraker and Siegel, 1963; Dufwenberg and Gneezy, 2000; Huck, et al., 2004), while supporting the view that increasing numbers beyond two reduces the likelihood of tacit collusion, have not considered the impact of discounting. In this paper we conduct a series of laboratory experiments with markets of 2, 3, and 4 firms, and discount rates explicitly accounted for in some 2-firm markets, and examine whether the tradeoffs predicted in theory occur in the behavior of our subjects. We also examine the role of differing experimental conditions in determining outcomes.

PREVIOUS LITERATURE

Starting with Stigler (1964), economists have noted that the number of firms should matter in determining cartel success (whether tacit or overt), with less likelihood of longevity as the number of firms increases; the ease of detecting "cheaters" declines as numbers increase, leading to less cartel cohesion. Friedman (1971) formalized this notion in a game-theoretic setting involving "trigger strategies" whereby the collusive solution is chosen until one defection occurs and the static Cournot solution is played thereafter (sometimes known as the "grim strategy"). Collusion is supported for sufficiently low discount rates, r (or sufficiently high discount factors, δ = 1/1+r).[1]

One of the earliest empirical investigations of characteristics associated with price-fixing rings was Hay and Kelley (1974), examining US criminal (horizontal) price-fixing cases which were won or settled by the U.S. Department of Justice for the years 1963-72. They find that almost 80% of the conspiracies involved 10 or fewer firms. Marquez (1994) estimates the determinants of the duration of 52 international cartels covering a very wide period from the late 1880s through the early 1980s, using maximum likelihood techniques; the only statistically significant factor is the Herfindahl index, which (as expected) increased duration. In a more sophisticated study of 81 recent (active after 1990, ending by 2007) international cartels, Levenstein and Suslow (2011) examine – among other factors – number of participants and market interest rates (as a proxy for discount rates), failing to find significant impacts on cartel stability of either of these.[2]

One limitation of such studies is their reliance on reported (and prosecuted) cartels, while the theoretical predictions relate to tacit (non-cooperative game) collusion. There have been a number of experimental studies of the effects of numbers of firms, but fewer on the effects of discount rates on the ability of players to tacitly collude. In terms of numbers of firms, early work by Fouraker and Siegel (1963) suggested increasing numbers made collusion less feasible. More recently, Dufwenberg and Gneezy (2000) found this as well, and Huck et al. (2004) suggested that tacit collusion was unlikely with as few as four firms in experimental oligopoly markets. Very recently Roux and Thoni (2015) confirm the Huck et al. findings in the usual setup where "punishment" after defection affects all firms; however they do find the interesting result that where punishment can be targeted only to the particular "cheater" collusion can be supported with larger numbers of players.

Few studies have explicitly incorporated payoffs discounted over time. Feinberg and Husted (1993) present participants in a quantity-setting duopoly setting with payoffs declining at a slow or fast rate (plus no rate of decay) and find that collusion is most-often supported with no discounting and least-often with a high rate of decay. Recently, Frechette and Yuksel (2016) have incorporated payoff discounting in a duopoly setup, again finding reduced cooperation resulting. More common, however, is the use of random termination methods (i.e., toss of a die at the end of each round to determine continuation) in order to simulate infinitely repeated games.[3] While in theory a continuation probability (of less than 100%) with no explicit payoff discounting should be equivalent to an infinitely repeated game with discounting, in practice experimental subjects may view the two differently; Frechetter and Yuksel (2016) examine this and find that explicit payoff discounting seems to reduce cooperation more than does an equivalent random termination rate.

THEORETICAL FOUNDATIONS AND EXPERIMENTAL APPROACH

As noted above, there has been no examination of the tradeoffs between numbers of firms and rates of discounting (either through variation in explicit payoff reductions over time or random termination probabilities at different levels). This is what we address here. To illustrate – in a simple infinitely repeated Bertrand context – the predicted tradeoff, suppose **n** firms share collusive profits Π_M if they each choose the collusive price P_M but take the full collusive profits if they set a price below P_M. However, such "cheating" would then be assumed to lead all firms to defect forever (the "grim strategy") implying zero profits in all future periods.

Assuming all profits are received at the end of each period, with a discount rate (or "rate of time preference") = **r**, the present discounted value of profits from choosing the collusive price = Π_M/nr. On the other hand, the present value of cheating (full collusive profits after the first period, but nothing thereafter) = $\Pi_M/(1+r)$. Tacitly collusive pricing can then be sustained as an equilibrium if $\Pi_M/nr > \Pi_M/(1+r)$. But this will occur *iff* $1 + r > nr$, which implies $r(n-1) < 1$. Clearly as the number of firms increases, the required discount rate allowing the inequality to hold must fall (intuitively, only very patient firms will be able to sustain a tacitly collusive equilibrium with a large number of them). The relationship between **n** and **r** required to sustain collusion is more complicated in the Cournot approach we model below, but the general pattern remains.

The general structure of the series of experiments is as follows: a quantity-setting (Cournot, homogenous good) oligopoly format, with subjects presented with a profit table based on their and their competitors' choices. Initially we had subjects make repeated plays in a simple game setup with three options, corresponding to the perfectly competitive, Nash/Cournot, and tacitly collusive choices, in setups facing 1, 2, and 3 other players; we realized, however that this approach could be biasing subjects to the Nash solution, which was always the middle of the three options. For this reason, in a subset of the results we present here, subjects faced the full continuum of output choices ranging from the tacitly collusive choice in the 4-person game (q = 11) to the perfectly competitive/Bertrand choice in the duopoly setup (q=42).

Experimental subjects in six sessions were recruited from principles of macroeconomics and microeconomics courses at American University in 2015 and 2016 (no student participated in more than one session). The first two sessions were conducted entirely with "pencil and paper" while the latter utilized in some or all experiments the VeconLab experimental software available through the University of Virginia's Department of Economics (Holt 2015). Upon arrival, subjects were each handed \$5 in cash to demonstrate the potential for real payoffs, and a packet that included instructions, profit matrices, and for paper-based sessions, a sheet on which to record decisions.[4] The profit matrices are tables that detail the subject's profit in a market given their choice of output and the output choice(s) of the other subject(s) in the same market. During each experimental session, subjects made output decisions in two-firm, three-firm, and/or four-firm quantity-setting homogenous good markets for which each firm's behavior was determined by set demand and cost functions. Subjects seemed well-motivated and earnings averaged approximately \$40 for sessions of about 1.5 hours.

The subjects played the same rivals within each experiment during a session, but faced a different combination of rivals for every new experimental condition. Of course, some randomness in end-period is required to simulate a supergame (infinitely repeated game) format. In the earliest sessions, play involved a minimum number of periods, continuing beyond that for additional periods with a 67% probability for each additional period (determined by a throw of a die); in the more recent experimental sessions, we were vague about how long the sessions would go, and simply stopped them after differing numbers of periods (between 7 and 12). As early sessions suggested little ability to deviate from the single-period Nash equilibrium, collusive behavior was encouraged in the last several sessions through instructions that stated the combination of choices for each session that would yield the largest combined payoffs to the subject and the other(s) in the market, and noting a more aggressive approach would likely yield a competitive response by the rival. Given the focus of the paper on situations in which tacit collusion would likely break down, it seemed appropriate to educate participants in strategies which could achieve collusion. We distinguish in our results those experimental outcomes for which consistent instructions and other conditions were maintained.

In our experiments reported here we have a linear demand structure with constant costs behind the payoff tables presented to subjects. The market demand equation is $P = 100 - Q$, $AC = MC = 16$. This implies a collusive optimum at $Q = 42$, a perfectly competitive equilibrium at $Q = 84$, and the Nash/Cournot solution at $Q = 84\, n/(n+1)$. Payments to experimental subjects per period were 1/10 (in cents) of those implied; with some rounding, per period (per-firm) payoffs were 88 cents in the two-firm collusive solution, 59 cents in the three-firm collusive solution; 78 cents in the two-firm Nash/Cournot equilibrium, 44 cents in the three-firm Nash/Cournot solution. In a four-firm setup, each firm's collusive payoff would be 44 cents per period, the Nash/Cournot payoff would be 27 cents. One rationale for a Cournot (vs. Bertrand) framework is the more gradual increase in defection incentives, allowing a greater opportunity to observe tradeoffs between impacts of increased numbers and higher discount rates on participant choices; also, as the Cournot model is widely used in merger analysis, our results may be more applicable to policy.

We would anticipate movements away from tacit collusion as we move from 2 to 3 and then 4-firm markets, controlling for discount rates. And similarly, we would expect to see such effects, controlling for number of firms, as we move from discount rates of zero, to 11% and then 50%.[5] Of particular interest, though, is the effect of varying both numbers and discount rates – frankly; one would not expect to see successful attempts at tacit collusion in four-firm markets regardless of the discount rate, and similarly little success with discounting regardless of numbers. However, intuitively, the likelihood of collusive or Cournot equilibria in 4-firm markets with zero discount rates should be comparable to those in 2-firm markets with some discounting.

In Shapiro (1989, p.365), a formula is presented for the relationship between numbers of firms and the highest discount rate which would support a tacitly collusive equilibrium (with "Cournot reversion" following defection) in the repeated Cournot framework we investigate: clearly, as the number of firms increases the maximum discount rate which can support such an equilibrium falls – for **n**=2, this is 89% (a discount factor, or rate of decay of 0.53), for **n**=3, 75% (or rate of decay of .57), for **n**=4, 67% (or rate of decay of .60). In practice, rates this high are unlikely to support tacit collusion. Nevertheless, one would expect tradeoffs to emerge between numbers and rates in experimental markets.

If we compare the discount factor (or rate of decay) implied by discount rates embodied in our experiments ($\boldsymbol{\delta}$) to those which would be required to support tacit collusion for various values of **n** (noted above) – denoting these by ($\boldsymbol{\delta^*}$), we can interpret the ratio of the two as a measure of the likelihood of collusion emerging – keeping in mind that tacit collusion is difficult to induce in prisoner's dilemma experiments (in part because players may not be adopting a strategy quite as "grim" as that assumed in the Shapiro formula referred to above). This ratio is 1.9 for **n**=2/**r**=0, 1.8 for **n**=3/**r**=0, 1.7 for both **n**=2/**r**=0.11 and **n**=4/**r**=0, and 1.3 for **n**=2/**r**=0.5. This suggests, e.g., that we should see comparable rates of collusion in 4-firm markets with no discounting as in duopoly markets with "light" discounting, and more collusion in the triopoly market with no discounting than in the duopoly market with "heavy" discounting.[6]

RESULTS

The primary focus of our empirical analysis is on market results, where all sessions lasted between 7 and 12 periods. We examine market averages over all periods except the first and last. However, in examining the impact of various experimental conditions, we also report results at the level of the individual period-by-period observation.

We have results for a total of 85 markets: 42 two-firm markets with r=0; 15 two-firm markets with r=0.11 (90% decay per period) – what we call "light discounting"; 18 three-firm markets with r=0; and 10 four-firm markets with r=0. We also analyze separately a subset of 47 of these markets with uniform experimental conditions – in particular, consistent continuous strategy choices across all experiments -- involving 22 two-firm markets with r=0, 10 two-firm markets with r=0.11, 10 three-firm markets with r=0, and 5 four-firm markets with r=0. All of these markets had the same market demand and constant costs.[7] We attach as an Appendix a selection of instructions and payoff tables from our experimental sessions.

Our results suggest that subjects are primarily engaging in Cournot-Nash behavior, with larger combined outputs (and of course lower prices) in the 3- and 4-firm markets than in the 2-firm markets. Over the entire sample, the average market output (dropping the first and last period results) for the duopoly markets with no discounting is 54.4 (with the Collusive output at 42 and the Cournot-Nash output at 56); in the uniform experimental conditions sub-sample, the mean output is 54.6. With light discounting, this increases only very slightly to 55.4 in the full sample, to 54.9 in the sub-sample. In the 3-firm markets without discounting, the average combined quantity is 62.3 for the full sample (here the Cournot-Nash output would be 63), and 62.0 for the sub-sample, while in the 4-firm markets without discounting, the average combined quantity is 66.3 for the full sample (here the Cournot-Nash output would be 68) and 66.7 for the sub-sample. What seems clear is that – as is often found in experimental work of this sort – the prisoner's dilemma is a powerful force.

A major issue referred to earlier is that the first few sessions followed most of the prior literature in severely limiting the choice set facing subjects (and using a different set of choices varying with the number of firms in the market; it seems quite likely that this might have biased them towards the Nash-Cournot equilibrium (which was always the middle of 3 options they could pick). In the more recent sessions (analyzed as our "uniform conditions sub-sample") we allow a continuum of choices from 11 to 42, and that same range regardless of the number of firms in the market. We believe this to be an innovation in our work as it includes the full range from collusive to perfectly competitive in all three market structures, with N equal to 2, 3, or 4 firms (21-42, 14-28, and 11-21, respectively). This allows us to see if experimental subjects are able to identify particular equilibria outcomes without being guided there by an easily navigated payoff table.

In these later sessions we had an unknown (to the subjects) stopping point, at between 7 and 12 periods (but did not specify the termination rule) and we conducted the experiments via computer. In contrast, in the earlier experiments we had a specified random termination rule (based on the toss of a die after a fixed number of periods), and had participants fill out paper decision sheets (requiring time between rounds for experimenters to match up market participants, record entries and return decision sheets to subjects). We also explore the implications of changing these experimental approaches in what follows.

To attempt to deal with the various changes in experimental conditions occurring simultaneously, we initially ran some regression specifications (both for all markets and for the "uniform conditions" subsample) to explain average market quantities chosen (over all periods, but dropping the first and last period) as a function of number of firms, and light-discounting,. The first column of Table 1 presents the full sample summary statistics for our dependent variable, the average market output (logged in our estimated models), and the included independent variables, measured as dummy variables and the number of firms (n=3 or n=4 participants, with n=2 as the omitted group) in the market and the presence of discounting. The other variables describe the market conditions such as limited choice restrictions, paper experiments, and random experiment termination. The second column of Table 1 lists the summary statistics for the smaller subsample of observations from the experiments carried out on the computer with continuous output choices.

TABLE 1
SUMMARY STATISTICS, MEAN AND STANDARD DEVIATION

Variable	Full (85 obs.)	Subsample (47 obs.)
Average Market Q	57.63 (6.84)	57.52 (6.97)
Two Firms (n=2)	0.671 (0.473)	0.681 (0.471)
Three Firms (n=3)	0.212 (0.411)	0.213 (0.414)
Four Firms (n=4)	0.118 (0.324)	0.106 (0.312)
Light Discounting	0.176 (0.383)	0.213 (0.414)
Paper Experiment	0.224 (0.419)	0
Limited Choice	0.447 (0.500)	0
Random Termination	0.118 (0.346)	0

The ordinary least squares (OLS) regression coefficients for the full sample of markets are presented in Table 2. We find strong statistical support for our prediction on the relationship between the number of firms in the market and the average market outcome. Specifically, markets with three firms have nearly a 14 percent greater average output than the two firm markets and markets with four firms have a 20 percent higher average output level. Discounting has a small positive, but not statistically significant, effect on average chosen output levels.[8]

TABLE 2
OLS ON FULL SAMPLE
DEPENDENT VARIABLE = LOG AVERAGE MARKET OUTPUT

Variable	Coefficient (t-statistic)
Three Firms (n=3)	0.136*** (5.19)
Four Firms (n=4)	0.200*** (6.12)
Light Discounting	0.018 (0.63)
R^2	0.405
Number of Observations	85

Statistically significant at *** = 1%, ** = 5%, * = 10%

TABLE 3
OLS ON SUBSAMPLE WITH UNIFORM EXPERIMANTAL CONDITIONS
DEPENDENT VARIABLE = LOG AVERAGE MARKET OUTPUT

Variable	Coefficient (t-statistic)
Three Firms (n=3)	0.129*** (3.40)
Four Firms (n=4)	0.203*** (4.14)
Light Discounting	0.005 (0.12)
R^2	0.377
Number of Observations	47

Statistically significant at *** = 1%, ** = 5%, * = 10%

The OLS coefficients for the subsample of firm markets where subjects used the computerized version of the game and their output choices were unlimited are reported in Table 3.[9] The coefficients are consistent with the results reported for the full sample. Markets with three firms have about a 13 percent greater average output than the two firm markets and markets with four firms have about a 20 percent higher average output level. Discounting again has no statistically significant impact.[10]

In results not reported here, we limit the sample only to 2-firm markets and investigate the impact of discounting – again, we find the sign in the expected direction, but no statistically significant effect. Another issue is whether the effect of increasing numbers reflects a breakdown of tacit collusion or simply the expected Cournot impact of increased numbers; when we replace the average market quantity with its deviation from the predicted Nash equilibrium value, we find no impact of number of firms suggesting little ability to tacitly collude even with two firms.

We also examine individual period-by-period market outcomes, controlling for differences (noted above) in experimental conditions, as well as participant fixed effects.[11] We have a total of 1,623 observations across 35 participants. Table 4 displays the results of our analysis (with market output in logs – results in levels were similar). As with the average market results, we continue to find that as the number of firms increases, the market output increases significantly – by 14.8% in going from 2 to 3 firms, and by 22.2% in going from 2 to 4 firms; these compare well to the 12.5% and 21.4% predicted increases in the Cournot equilibrium outputs given our market parameters. As in the market average results, we find a small but not statistically significant competitive impact of discounting.

TABLE 4
PERIOD-BY-PERIOD ANALYSIS OF DETERMINANTS OF MARKET OUTPUT
OLS ON FULL SAMPLE
DEPENDENT VARIABLE = LOG MARKET OUTPUT

Variable	Coefficient (t-statistic)
Three Firms	0.148*** (12.33)
Four Firms	0.222*** (13.88)
Continuous Output	0.056*** (3.29)
Paper Selection	0.062*** (3.44)
Light Discounting, r = 0.11	0.021 (1.31)
Period 1	-0.057*** (4.07)
Period 2	-0.047*** (2.94)
Period 3	-0.025* (1.67)
Period 4	0.017 (1.06)
Period 5	0.008 (0.53)
Last Period	-0.019 (1.19)
Constant	4.088*** (136.27)
Observations	1,623
R-squared	0.271
Participant Fixed Effects	YES

Statistically significant at *** = 1%, ** = 5%, * = 10%

Looking at the within-market timing of output choices, we see some evidence of an initial effort towards tacit collusion (possibly as a short-term response to our instructions, which were somewhat biased in this direction), breaking down period-by-period after that. We also see no evidence of last-period "end-game" effects.

Not surprisingly, we do find that experimental conditions matter (though our results presented earlier suggest these are may not be affecting the impact of our variables of interest – i.e., numbers of firms and discounting). Switching to games that utilize continuous output choice over a wide range seems to lead to significantly more competitive market outcomes than experienced when a sharply more limited output choice was given to subjects. And games using a pencil-and-paper method of recording output choices rather than the computerized Veconlab program also elicit a more competitive market outcome.

CONCLUSION

While our results are somewhat disappointing in terms of identifying impacts of discounting, we do find that, facing an unchanging market demand and costs, experimental subjects generally move towards the predicted Cournot-Nash equilibrium in markets of 2, 3, and 4 firms. While an increased number of firms in a market is associated with larger market output (and lower prices), this does not reflect tacit collusion (with N=2) breaking down as N increases, but rather the generalized Cournot result throughout. We find it especially interesting that our relatively inexperienced subjects are able to locate the Nash equilibrium results when they face a wide range of continuous strategic choices; despite instructions somewhat biased towards tacit collusion, any efforts in this direction seem to break down within several periods.

Positive discount rates do not seem to have the predicted competitive impact, though in part this may reflect the difficulty of simulating this in the laboratory. Finally, perhaps not surprisingly, differences in experimental conditions do seem to matter in determining participant response (though not necessarily affecting tests of our variables of interest). As there is little theoretical guidance for which of these to choose in conducting experiments, what seems best is simply transparency in reporting experimental findings.

ACKNOWLEDGMENTS

We thank Christian Fischer and Artie Zillante for helpful comments on earlier drafts, which were presented at Washington & Lee University, the 2015 Southern Economic Association meetings, the 2016 Eastern Economic Association meetings, and the 2016 International Industrial Organization Conference. Funds were provided by the College of Arts and Sciences at American University.

ENDNOTES

1. Abreu, Pearce, and Stacchetti (1990) show that a monotonic relationship can exist between duopoly outcomes and the discount rate, even within the set of r that can sustain collusive equilibria.
2. What do seem to matter are internal mechanisms of the cartel (such as market allocation, punishment and compensation procedures). More recently (Levenstein and Suslow, 2016), they do find both market interest rates and firm-level measures of financial fragility – both seen as (imperfect) proxies for discount rates – to have an effect on cartel formation and breakup in a sample of prosecuted US cartels over the past 50 years.
3. Dal Bo (2005) employs various continuation probabilities, finding that a higher continuation probability (analogous to a lower discount rate) leads to a greater likelihood of cooperation. Normann and Wallace (2012) discuss the pros and cons of various experimental termination strategies – finite, unknown (to the participants), and random termination (with known continuation probabilities); examining the implications in a series of prisoner's dilemma experiments, they find that average cooperation rates are relatively unaffected.
4. Instructions and related materials are available on request from the authors.
5. In what follows we do not report on experimental sessions with discount rates of 50%. We found that the payoffs declined too rapidly for students to react in a consistent manner.
6. While either a random termination rule or an unknown stopping point implies subjects will likely discount future rewards to some extent, it remains true that our cases with explicit payoff discounting should lead subjects to adopt a higher rate of discounting than without this.
7. Our attempt to bring in "heavy discounting," with an r=.50 in some 2-firm markets led to odd outcomes – and not consistent between sessions or participants. The problem is that subjects can view the entire pattern of round-by-round payoffs and in this case, by the 6th period the Nash payoffs (starting out at almost 80 cents per period) are down to 10 cents. While we had thought

this would lead to "cheating" immediately – and a rapid move to the Nash equilibrium -- it just seemed to confuse students and some may have figured they should just try to jointly maximize profits before they disappeared completely. In future work, we will attempt to deal with this issue more fully.

8. Results in levels were quite similar.
9. All other experimental conditions were the same in all of these (in particular, instructions and termination rule).
10. In results not reported here, results for the sample of observations with varied experimental conditions (i.e., those not in the "uniform conditions" subsample) were virtually identical to those in Tables 2 and 3.
11. While our focus was somewhat different, Feinberg and Husted (1999) also estimated time-series of participant choices in a prisoner's dilemma setting.

REFERENCES

Abreu, D., Pearce, D. & Stacchetti, E. (1990). Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica*, 58, 1041-1063.

Dal Bo, P. (2005). Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games. *American Economic Review*, 95, (5), 1591-1604.

Dufwenberg, M. & Gneezy, U. (2000). Price Competition and Market concentration: An Experimental study. *International Journal of Industrial Organization*, 18, (1), 7-22.

Feinberg, R.M. & Husted, T.A. (1993). An Experimental Test of Discount Rate Effects on Collusive Behavior in Duopoly Markets. *Journal of Industrial Economics*, 41, 153-160.

Feinberg, R.M.& Husted, T.A. (1999). Estimating Reaction Functions in Experimental Duopoly Markets. *International Journal of the Economics of Business*, 6, (1), 57-63.

Fouraker, L. & Siegel, S. (1963). *Bargaining Behavior*, New York: McGraw-Hill.

Frechette, G.R. & Yuksel, S. (2016). Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination. *Experimental Economics*, doi:10.1007/s10683-016-9494-z.

Friedman, J.W. (1971). A Noncooperative Equilibrium for Supergames. *Review of Economic Studies*, 38, 1-12.

Hay, G.A.& Kelley, D. (1974). An Empirical Survey of Price Fixing Conspiracies. *Journal of Law & Economics*, 17, (1), 13-38.

Huck, S., Normann, H.-T. & Oechssler, J. (2004). Two are Few and Four are Many: Number Effects in Experimental Oligopolies. *Journal of Economic Behavior and Organization*, 53, (4), 435-446.

Levenstein, M.C. & Suslow, V.Y. (2011). Breaking Up is Hard to Do: Determinants of Cartel Duration. *Journal of Law & Economics*, 54, (2), 455-492.

Levenstein, M.C. & Suslow, V.Y. (2016). Price-Fixing Hits Home: An Empirical Study of U.S. Price-Fixing Conspiracies. *Review of Industrial Organization*, 48, (4), 361–379.

Marquez, J. (1994). Life Expectancy of International Cartels: An Empirical Analysis. *Review of Industrial Organization*, 9, (3), 331-341.

Normann, H.-T. & Wallace, B. (2012). The Impact of the Termination Rule on Cooperation in a Prisoner's Dilemma Experiment. *International Journal of Game Theory*, 41, 707-718.

Roux, C. & Thöni, C. (2015). Collusion Among Many Firms: The Disciplinary Power of Targeted Punishment. *Journal of Economic Behavior & Organization*, 116, 83–93.

Shapiro, C. (1989). Theories of Oligopoly Behavior. In *Handbook of Industrial Organization*,
Schmalensee & Willig, eds., Amsterdam: Elsevier Science Publishers, 329-414.

Stigler, G.J. (1964). A Theory of Oligopoly. *Journal of Political Economy*, 72, (1), 44-61.