

UNIVERSITY OF CALIFORNIA

Los Angeles

THREE ESSAYS ON REPEATED GAMES

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Economics

by

Pedro Dal Bó

2002

The dissertation of Pedro Dal Bó is approved.

Sushil Bikhchandani

Hongbin Cai

William R. Zame

David K. Levine, Committee Chair

University of California, Los Angeles

2002

A mis padres

Contents

1	Social Norms, Cooperation and Inequality	1
1.1	Introduction	1
1.2	Relevant literature	5
1.3	The Matching Game	7
1.4	Perfect Information	9
1.5	Local information	15
1.6	Conclusions	20
2	Tacit Collusion under Interest Rate Fluctuations	23
2.1	Introduction	23
2.2	Related literature	26
2.3	Price competition	30
2.3.1	Optimal tacit collusion with a random discount factor	31
2.3.2	The effects of changes in $f(\pm)$	37
2.3.3	The effects of changes in the number of firms	40
2.3.4	Example with uniform distributions	41
2.4	Quantity Competition	46
2.4.1	Optimal tacit collusion with a random discount factor	47
2.4.2	The effects of changes in $f(\pm)$	54
2.4.3	The effects of changes in the number of firms	56
2.4.4	Example with uniform distributions	57
2.5	Extensions	59
2.5.1	Asymmetric equilibrium prices	59
2.5.2	Deterministic discount factor cycles	60
2.5.3	Correlated discount factor and the volatility effect	62

2.5.4	General normal form games	63
2.6	Conclusions	66
3	Cooperation under the shadow of the future: experimental evidence on infinitely repeated games	68
3.1	Introduction	68
3.2	Previous Experiments	70
3.3	Experiment design	74
3.4	Theoretical predictions	79
3.5	Experimental results	81
3.5.1	Does cooperation increase with the shadow of the future?	82
3.5.2	Infinitely repeated games vs. finitely repeated games	85
3.5.3	Do payoff details matter?	86
3.5.4	Do Economics majors behave differently?	87
3.6	Conclusions	88
4	Appendices	90
4.1	Appendix to chapter 1	90
4.2	Appendix to chapter 2	94
4.3	Appendix to chapter 3: Instructions for PD2-Dice-USD Session (4/9/02)	96
	References	106

List of Figures

1	The profit function	32
2	Ranges of tacit collusion	42
3	The effect of levels	43
4	The effect of volatility	44
5	Volatility matters	45
6	Ranges of tacit collusion for different N	45
7	Tacit collusion and the number of firms	46
8	Ranges of tacit collusion	58
9	Tacit collusion profits with N=2	59
10	Profits under cyclical discount factor	61
11	Volatility and general games	65
12	Computer screen that subjects saw before each interaction.	75
13	Equilibrium average payoffs $V(\pm)$	79
14	Cooperation by treatment (matches after third and all rounds)	85

List of Tables

1	Previous studies (% of cooperation in ...rst round)	71
2	Stage game payoffs in points	76
3	Equilibrium outcomes	80
4	Sessions' descriptive data	83
5	Percentage of cooperation by match and treatment*	84
6	Percentage of cooperation by treatment*	84
7	Percentage of cooperation by round and treatment*	86
8	Distribution of outcomes for $\pm = :5^*$	87
9	Percentage of cooperation by treatment and major*	88
10	Distribution of outcomes for $\pm = :5$ and major*	89

ACKNOWLEDGMENTS

I want to express my deepest gratitude to all the people that were essential for this dissertation and made my life at UCLA such a great experience.

David Levine for his invaluable guidance (separating the important from the unimportant, the feasible from the infeasible), ideas, discussions and support. Chapter 3 of this dissertation would not have been possible without his financial support.

William Zame for his insights, support and discussions. Hongbin Cai for his comments, advice, support and for organizing a “study group” during my third year that was crucial for this dissertation. John Riley for his crucial comments and support and to Sushil Bikhchandani, Harold Demsetz, Philip Leslie, Joseph Ostroy, Jean-Laurent Rosenthal for important comments and suggestions. Hugo Hopenhayn for crucial ideas and suggestions and William Brock, Germán Coloma, Walter Cont, David Rahman, Federico Weinschelbaun, José Wynne and seminar participants at UCLA and several universities around the world for very useful comments and discussions.

Lora Clarke and Christine Frank for making administrative issues a pleasure. The people of CASSEL (David Avery, Monica Bathija, Joven Matias, Patricia Wong, Walter Yuan) and Eric Szu-Wen Chou for all their help running the experiments for Chapter 3.

Sebastián Claro, Walter Cont, Eduardo Fajnzylber and Igal Magendzo for sharing the burden of homeworks and exams in the first two years with so much good humor.

Finally, I thank my family for their encouragement and unconditional support. Ernesto Dal Bó for including me in his research agenda, comments and discussions, my parents for their love from ten thousand kilometers away and Anna Aizer for her comments, support and understanding.

I acknowledge financial support from the following institutions: Universidad de Buenos Aires, the Argentine Ministry Education and UCLA.

VITA

- May 17, 1971 Born, Comodoro Rivadavia, Argentina
- 1992 Statistics Assistant
FIEL, Argentina
- 1994 Licenciado, Economics
Universidad de Buenos Aires
Buenos Aires, Argentina
- 1991-1996 Teaching Assistant
Universidad de Buenos Aires
Buenos Aires, Argentina
- 1996 M.A. Economics
Instituto Torcuato Di Tella
Buenos Aires, Argentina
- 1993-1997 Researcher
Universidad Argentina de la Empresa
Buenos Aires, Argentina
- 1997 Assistant Professor
Universidad de Buenos Aires
Buenos Aires, Argentina

ABSTRACT OF THE DISSERTATION

Three Essays on Repeated Games

by

Pedro Dal Bó

Doctor of Philosophy in Economics

University of California, Los Angeles, 2002

Professor David K. Levine, Chair

This thesis consists of three chapters. The first chapter analyzes the outcomes that can be supported in a society through reward and punishment schemes that operate through community enforcement (social norms). I consider a society of infinitely long-lived and very patient agents that are randomly matched in pairs every period to play a given game. I find that any mutually beneficial outcome can be supported by a self-enforcing social norm under both perfect information and a simple local information system. These Folk Theorem results explain not only how social norms can provide incentives to forestall opportunistic behavior and support cooperation in a community but also how they can support outcomes characterized by inequality.

The second chapter studies tacit collusion under interest rate fluctuations. In contrast to the existing literature on repeated games that assumes a fixed discount factor, I study an environment in which it is more realistic to assume a fluctuating discount factor. In a repeated oligopoly, as the interest rate changes, so too does the degree to which firms discount the future. I characterize the optimal tacit collusion

equilibrium when the discount factor changes over time, under both price and quantity competition, and I show that collusive prices and profits depend not only on the level of the discount factor but also on its volatility. These results have important implications not only for the study of cooperation in repeated games but also for empirical studies of collusive pricing and the role that collusive pricing may play in economic cycles.

The third chapter presents experimental evidence on infinitely repeated games. While there is an extensive literature on the theory of infinitely repeated games, empirical evidence on how “the shadow of the future” affects behavior is scarce and inconclusive. I simulate infinitely repeated prisoner’s dilemma games in the lab by having a random continuation rule. The experimental design represents an improvement over the existing literature by including sessions with finite repeated games as controls and a large number of players per session (which allows for learning without contagious effects). I find strong evidence that the shadow of the future reduces opportunistic behavior closely following the theoretical predictions.

1 Social Norms, Cooperation and Inequality

1.1 Introduction

In a stable society, codes of conduct or social norms guide interaction among people. A social norm functions as an implicit rule of behavior that guides the actions of society members under different circumstances. In this sense, social norms organize interaction among the members of a society, thereby reducing the level of uncertainty.

An important feature of social norms is their ability to curtail opportunistic behavior through the establishment of reward and punishment schemes that operate through social reaction. In this way social norms can support high levels of cooperation among the members of a society. Greif [25] presents an example of this found among the Maghribi traders in the Mediterranean during the 11th century. With very little support from formal institutions to enforce contracts governing overseas trading, the Maghribi traders followed a simple social norm: no Maghribi trader would trade with another Maghribi trader who had cheated a Maghribi trader before. In this way the Maghribi punished deviations and reduced the incentives to cheat, hence, supporting efficient levels of commerce that could not have been achieved if they had relied upon personal retaliation alone.

But social norms not only can lead to efficient outcomes, they can also lead to inefficient outcomes or outcomes with inequality. Akerlof [4] presents several examples of social norms that lead to inefficient and unequal outcomes. Among those he considers the traditional Indian caste system, in which a mutually beneficial transaction -like marriage- between two members of different castes may be forestalled by the expected reaction of third parties.

This paper analyzes the outcomes (efficient or inefficient, equal or unequal) that can be supported by social norms. I consider a simple society consisting of a fixed number of infinitely long-lived members. Every period these members are randomly matched in pairs to play a given game. The only restriction on these matches is that the probability of each match is fixed and independent of time and past behavior. I assume that the members are selfish in the sense that they do not care about others or the social norm per se. Hence, no members follow the social norm because they derive any direct utility from doing so. In this paper, social norms are followed because no member can ever profit from deviating from it on his or her own. In other words, I study social norms that are a subgame perfect equilibrium or a sequential equilibrium, depending on the information requirements.

I consider two societies under different information requirements: one characterized by perfect information and another with a more restrictive information process. For the former, I consider the case in which it is possible for every member to observe each other's behavior. In this situation of perfect information, I prove that if the agents are sufficiently patient, any feasible and individually rational outcome can be supported by a social norm that is a subgame perfect equilibrium. Therefore, a society can achieve in equilibrium any feasible and individually rational outcome through community enforcement, that is, with a social norm that specifies rewards and punishments.

To understand the relevance of this result it is important to note the difference between community enforcement and personal enforcement. Under personal enforcement, a cheater will only face retaliation by the victim. On the contrary, under community enforcement all the members of the society react to a deviation. Hence, community enforcement can offer a punishment that is more swift and severe than personal enforcement can. This implies that any outcome that can be supported by

personal enforcement can also be supported by community enforcement. But most interestingly, there are outcomes that can not be supported by personal enforcement that can be supported by community enforcement. In fact, there are highly unequal outcomes that can only be supported by community enforcement, as the next example shows.

Consider a society with four members that are matched every period to play the following prisoner's dilemma game:

	c	d
c	2, 2	-2, 4
d	4, -2	0, 0

Assume that each pair has the same probability of being matched and consider the following social norm: Player 1, who is called "king", always plays d, players 2, 3 and 4, who are called "serfs", play c if no one has deviated and play d if someone has ever deviated. Under this social norm, the king receives a payoff of 4 every period while the serfs receive an average payoff of $\frac{2}{3}$ in equilibrium. It is easy to verify that this social norm is a subgame perfect equilibrium for discount factors greater than $\frac{3}{4}$.¹ It is interesting to note that with personal enforcement the king would only be able to obtain payoffs lower than 3.² Therefore, community enforcement allows the

¹The King has clearly no incentives to deviate. Seeing that the serfs do not have incentives to deviate requires some calculations. If no one has deviated, a serf facing the King would get an expected payoff of $\delta(2 + \delta \cdot \frac{2}{3}) + (1 - \delta) \cdot 0$ for following the social norm, where δ is the discount factor, and he would get 0 for deviating. Hence, if no one has deviated, the serfs will not deviate when playing against the King if $\delta > \frac{3}{4}$. Similar calculations show that for those discount factors a serf would not deviate when playing against another serf. Therefore, serfs have no incentives to deviate if no one has deviated before. If someone has deviated, they do not have incentives to deviate since the prescribed actions correspond to the stage game Nash equilibrium. Therefore, if $\delta > \frac{3}{4}$, no player has an incentive to deviate and, then, the social norm is a subgame perfect equilibrium.

²Under personal enforcement, a payoff higher than 3 for the King would result in a negative expected payoff for the serf in the matches with the King. Under personal enforcement the serf

king to obtain a higher payoff than what he could obtain under personal enforcement. The reason for this is that under community enforcement each serf knows that if he deviates when playing with the king, no other serf will cooperate with him. Hence, serfs accept the negative payoff they receive every time they are matched with the king because this enables them to reap the benefits of full cooperation among the serfs. With personal enforcement the king can not use this threat and therefore he can not achieve the high level of payoffs he receives under community enforcement.

In the social norm used in the previous example all the players are punished for the deviation of one of the members. While societies may use this kind of punishment schemes to reduce opportunistic behavior, schemes that only punish deviators seem to be more realistic. Using social norms that only punish deviators, I show that any feasible and individually rational outcome can still be supported in a subgame perfect equilibrium under some restrictions on the stage game.

But the requirement of perfect information is unrealistic when the size of the society is large. Therefore, I also study environments with less demanding information requirements. As an alternative to perfect information I consider the existence of a local information processing system, following the seminal papers of Kandori [34] and Okuno-Fujiwara and Postlewaite [41]. In this case, in addition to knowledge gained from their own experience, players have access to information from a system that assigns status levels to players depending on their past behavior.

These status levels enable the social norm to establish punishments and rewards. For example, in the case of the king and the serfs presented before there can be two status levels: "good" and "bad" serfs. If failing to cooperate with the king results in becoming a bad serf and nobody cooperates with a bad serf, a good serf may be willing to cooperate with the king, even when the latter never cooperates.

would not accept that, since he can secure for himself a minimum of zero by playing d.

With local information systems added to social norms, I prove that, if players are sufficiently patient, any feasible and individually rational outcome can be supported by a social norm that is a sequential equilibrium. Therefore, even under very limited information, a society can achieve in equilibrium any feasible and individually rational outcome with a social norm that specifies rewards and punishments based on the information provided by a local information system. I also show a similar result, under some restrictions on the stage game, for social norms that only punish deviators.

These Folk Theorem results explain not only how social norms can provide incentives to forestall opportunistic behavior and support cooperation in a community but also how they can support outcomes characterized by inequality.

The following section presents the relevant literature, comparing previous findings with my own. Section 3 presents the model. Section 4 presents the perfect information Folk Theorem results and Section 5 presents the local information system Folk Theorem results. Section 6 concludes.

1.2 Relevant literature

Game theorists have long recognized that repeated playing and the possibility of future retaliation modifies current behavior, for example see Luce and Raïña [39]. In the case in which the same set of players play the same game repeatedly, other studies have found the conditions under which any feasible and individually rational outcome can be supported in equilibrium³. However, in many interesting cases the same players do not meet repeatedly but rather switch partners over time. For example see the cases of the already mentioned medieval trade coalitions studied in Greif [25]. While

³See Aumann and Shapley [6] for the case without discounting, Fudenberg and Maskin [23] for the case of discounting with perfect information and Fudenberg, Levine and Maskin [22] for the case of discounting with imperfect public information.

the changing of partners might seem to make cooperation impossible by reducing the possibility of personal retaliation, that is not necessarily the case. As Kandori [34] and Okuno-Fujiwara and Postlewaite [41] show, social norms can create incentives for players to punish deviators even if the deviation occurred against another player, since failing to punish can be itself punishable.

Both papers consider a society divided in two groups and every period the players in one group are randomly matched with the players in the other group. In addition, the authors restrict players in the same group to have the same equilibrium payoff⁴. Kandori [34] shows that with perfect information any feasible and individually rational payoff pair can be supported as a subgame perfect equilibrium. Therefore, any outcome that can be reached in the long-term relationship of two agents can also be reached by long-term relationship of two groups. In the case of two groups it is possible to construct credible group retaliations that mimic the ones needed for the Folk Theorem with only two players. This paper studies the set of equilibrium payoffs when we abandon those restrictions and we only require the matching procedure to be independent of history and time. In Section 3 I show that under perfect information, community enforcement can support not only those outcomes that can be supported by personal enforcement, but it can support other outcomes as well (as the king example in the introduction). In fact, I show that under perfect information any feasible and individually rational outcome can be supported by a social norm if the players are patient enough.

Since the requirement of perfect information is unrealistic when the size of the society is large, Kandori [34] and Okuno-Fujiwara and Postlewaite [41] study the

⁴This characterization of the game has the appealing graphical property that we can represent an equilibrium outcome of the game in \mathbb{R}^2 . Since all the members of a group receive the same payoff in equilibrium and there are only two groups, equilibrium payoffs can be written as the pair (v_1, v_2) , where v_1 and v_2 denote the utility received by each player in group 1 and 2, respectively.

existence of Folk theorem results with a local information processing system. Okuno-Fujiwara and Postlewaite [41] prove a Folk Theorem for a weak (non-perfect) equilibrium concept (Norm equilibrium) and Kandori [34] proves a Folk Theorem for sequential equilibrium under certain assumptions of the stage game. While Kandori [34] shows that infrequent transactions and limited information can still be overcome to achieve a Folk Theorem, he does so in the restrictive environment of two groups in which all the members of a group receive the same payoffs in equilibrium. Without these restrictions, Section 4 explains not only how social norms can support cooperation in a community but also how they can support outcomes characterized by inequality⁵.

The idea that social norms may support inequality is not new in the literature. Akerlof [4] presents several examples of social norms that support unequal (and inefficient) payoffs in (non-perfect) equilibrium. Axelrod [7] also presents examples in which social norms support inequalities. While those examples show that social norms can support inequality in equilibrium, the punishments specified in them are not credible because the equilibria studied are not perfect. In contrast, the equilibria presented in this paper are perfect. In addition my results are general to any stage game and not limited to particular examples.

1.3 The Matching Game

The society consists of N players, where N is an even number. In each stage, each of the players is matched with another player to play the stage game g_j . I assume that the matching of players is independent of past actions or time: the probability that player i is matched with player j is θ_{ij} , $0 \leq \theta_{ij} \leq 1$, for every period and history.

⁵In addition, the proof of Theorem 4 shows how theorem 2 in Kandori [34] could be proved without restrictions on the stage game.

This definition allows for partitions of players as in Kandori [34] and Okuno-Fujiwara and Postlewaite [41].

The stage game g_i is a symmetric game played by two players, with actions $a \in A$, for both players and payoffs $g : A^2 \rightarrow \mathbb{R}^2$, with the property that $g_{\text{row}}(a^0; a) = g_{\text{col}}(a; a^0)$, from the symmetry of the game. Given that the stage game is symmetric and that the row or column positions are not important I simplify notation writing $g(a; a^0) = g_{\text{row}}(a; a^0) = g_{\text{col}}(a^0; a)$. Therefore $g(a; a^0)$ denotes the payoff for the player that is playing a when the other is playing a^0 .

To minimax the other player the prescribed strategy is $m = \arg \min_{a \in A} \max_{a^0 \in A} g(a; a^0)$. I normalize the payoffs to have the minimax payoffs, not $g(m; m)$, equal to zero.⁶ If both players play m , each obtains a payoff of $g_m = g(m; m)$. Since m may not be the best response to m , it is the case that $g_m \leq 0$. The maximum payoff that can be obtained in the stage game is $\bar{g} = \max_{a; a^0 \in A} g(a; a^0)$ and the minimum payoff is $\underline{g} = \min_{a; a^0 \in A} g(a; a^0)$. I assume that players can condition their actions on public randomization devices, that is, they can play correlated strategies. Define $\mathcal{C}(A^2)$ as the set of possible correlated strategies in the stage game. Abusing notations I denote an element of that set for player i and j as $(a_{ij}; a_{ji}) \in \mathcal{C}(A^2)$.

Now I proceed to define the set of feasible payoffs of the random matching game. I define first the "play" of the stage game: the play of the stage game describes what profile of actions would be played by each possible matching of players in each period, that is play $\sigma : \text{pair} \in \{0; 1; 2; \dots; g\} \rightarrow \mathcal{C}(A^2)$.⁷ The play indicates what should be played by each possible pair in each period, for example a_{ij}^t denotes what i should play when matched with j in period t . Therefore, the expected stage payoff for i in period t is

⁶I assume for convenience that m is not a mixed action.

⁷If there are N players the number of possible pairs is $\binom{N}{2} = \frac{N!}{2!(N-2)!}$.

$\prod_{j \in i} g(a_{ij}^t; a_{ji}^t)$. If a_{ij}^t is the "play" for every period and δ is the discount factor, the average expected payoff of player i is $v_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \prod_{j \in i} g(a_{ij}^t; a_{ji}^t)$. Each play defines an expected payoff for every player $v = (v_1; v_2; \dots; v_N)$. Therefore the set of feasible payoffs are defined by the payoffs that result from every possible play, denote this set as $V \subseteq \mathbb{R}^N$.⁸

1.4 Perfect Information

In this section I consider societies in which it is possible for every member to observe each other's behavior. In this situation of perfect information, I show that if the agents are sufficiently patient, any feasible and individually rational outcome ($v \in V : v \geq 0$) can be supported by a social norm that is a subgame perfect equilibrium.

Theorem 1 (Folk theorem with perfect information) With perfect information any feasible and individually rational payoff ($v \in V : v \geq 0$) can be supported by a subgame perfect equilibrium for δ large enough.

Proof. Consider the following social norm to support v : if no one has deviated in the last T periods follow the "play" that yields v , if someone has deviated in the last T periods, play m .

First I check that no player has incentives to deviate if no one deviated in the last T periods. For player i the expected utility of conforming with the equilibrium is at least $(1 - \delta)g_i + \delta v_i$ while he would get at most $(1 - \delta)g_m + \delta v_i^p$ by deviating, where $v_i^p = (1 - \delta^T)g_m + \delta^T v_i$. Choose T so that, no matter the value of δ , δ^T is always

⁸Note that the set of feasible payoffs V is different from the set that arises in a two-player repeated game or in a repeated random matching game with two groups with all the members of each group receiving the same payoff as in Kandori [34] and Okuno-Fujiwara and Postlewaite [41]. In fact, the set V has a different dimension in general.

equal to some fixed number $d \in (0; 1)$. Then, given that $v_i > 0 \leq g_m$, it is true that $v_i > v_i^p$ independently of \pm . Therefore for \pm large enough it is true that $(1 - \pm)g + \pm v_i > (1 - \pm)g + \pm v_i^p$.

Now I consider the incentives to deviate if someone has deviated in the last T periods. In this case the incentives to deviate are the greatest when the future reward for facing the present punishment is as far in the future as possible, that is when all the players have to face T periods of punishment. If player i plays m , as the strategy prescribes, he receives a payoff of $v_i^p = (1 - \pm^T)g_m + \pm^T v_i$. If he deviates he receives at most $\pm v_i^p$. Choosing $\pm^T = d$ large enough for v_i^p to be positive (it is here that v_i strictly greater than zero is required), the player has no incentive to deviate during the punishment stage for $\pm < 1$.

Note that there is no contradiction in the requirements made on \pm and T in the two parts of the proof, as it is only required that $\pm^T = d$ and \pm are large enough. ■

The social norm used in the proof is a subgame perfect equilibrium given that regardless of the history of the game no player has incentives to deviate. In particular note that during the punishment stage no player has incentives to deviate since doing so only restarts the punishment stage. Therefore all the members of the society have incentives to enforce the social norm when someone deviates and, then, the punishments are credible.

While it may seem that Theorem 1 is a consequence of Fudenberg and Maskin [23] result for games with N players, this is not the case. While their results apply to games in which N players play the same stage game in all periods, in this case players may change partners every period. In addition, Theorem 1 can not be derived from their results by defining "meta-actions," that is functions from possible matches to actions for each player, given that Fudenberg and Maskin [23] assumes observable actions. However, by defining those "meta-actions" it can be shown that Theorem 1

is a consequence of Fudenberg, Levine and Maskin [22] folk theorem under imperfect public information. Nevertheless, the proof presented here has two advantages: one, its simplicity, and two, the fact that the implicit lower bound on the discount factor to support an individually rational and feasible outcome does not depend on the size of the population as may be the case in Fudenberg, Levine and Maskin [22].

Note that during the punishment stage in the proof of Theorem 1 all the players receive the same low payoff (g_m) regardless of who has deviated. In this social norm, then, all the players are punished for the deviation of one of the members. While societies may use this kind of punishment schemes to reduce opportunistic behavior, punishment schemes that only punish deviators seem to be more appealing.

Definition 2 A social norm displays personal punishment⁹ if the prescribed actions for two players who have not deviated in the past do not depend on the past actions of the rest of the players.

In this way, under personal punishment, only the deviators can be punished. Of course punishing may impose a cost to the player facing the deviator even under personal punishment

But personal punishment introduces new problems to the design of social norms that support a folk theorem. In fact, the social norm used in the proof of Theorem 1 may not be an equilibrium. Given the possible inequality of payoffs with personal punishment some players may have an expected payoff lower than g_m during the punishment stage, and therefore may have incentives not to punish, breaking in that way the credibility of the social norm¹⁰.

⁹Personal punishment should not be confused with personal enforcement. Personal enforcement means that cheaters are only punished by the player that was cheated, while personal punishment means that only cheaters are punished.

¹⁰For example consider the case of the kingdom presented in the introduction but with personal punishment instead: only the players that have deviated are punished. Imagine now that two of the

Fortunately under some conditions it is possible to support any feasible and individually rational payoff with social norms with personal punishment.

Assumption 1: $\exists r \in A$ such that $g(m; r) > g_m$, $g(r; m) = \underline{g}$.

Under this assumption there is an action r that allows a deviator to “ask for forgiveness” by taking the lowest possible payoff in the game \underline{g} , and giving the punisher a payoff higher than g_m while playing m . As such, it is possible to create punishment schemes in which the punisher earns a higher payoff than the deviator and in which the deviator by refusing to take the punishment can at most obtain a payoff of zero. This allows me to construct a social norm that ensures that players have incentives to follow it even when some of the other players have deviated¹¹. This assumption is satisfied, for example, by the prisoner’s dilemma game, in which m stands for d and r stands for c .

Theorem 3 (Folk theorem with perfect information and personal punishment) Under perfect information, Assumption 1 and personal punishment, any feasible and individually rational payoff ($v \in V : v \geq 0$) can be supported by a subgame perfect equilibrium for ϵ large enough.

Proof. Consider the following social norm that yields v in equilibrium: if player i meets player j and neither has been the last player to deviate in the last T periods, they play $(a_{ij}; a_{ji})$ (which yields v in equilibrium); the last player to have deviated in serfs have deviated and have to be punished for ζ periods. In this case the third serf will earn zero each time he meets with one of the other serfs and $\frac{2}{3}$ each time that he meets with the king. Then, he would get an expected payoff of $\frac{2}{3}$ each period during the punishment stage. By deviating he can get a zero payoff during T periods and, then, he may be willing to deviate just to avoid the negative payoff he earns if he does not deviate.

¹¹Without Assumption 1, Theorem 2 is still true under alternative conditions like uniform random matching games and N large.

the last T periods (simultaneous deviations are ignored) plays r and his match plays m .

First, I check that no player has incentives to deviate if no one has deviated in the last T periods. For player i the expected utility of conforming with the equilibrium is at least $(1 - \epsilon)g + \epsilon v_i$ while he would get at most $(1 - \epsilon)\bar{g} + \epsilon v_i^p$ by deviating, where $v_i^p = (1 - \epsilon^T)g(r; m) + \epsilon^T v_i$. Choose T to have $\epsilon^T = d \geq 2(0; 1)$. Then, given that $v_i > 0 \leq g(r; m)$, it is true that $v_i > v_i^p$ independently of ϵ . Therefore for ϵ large enough it is true that $(1 - \epsilon)g + \epsilon v_i > (1 - \epsilon)\bar{g} + \epsilon v_i^p$.

Second, I consider the case in which player i has been the last player to deviate in the last T periods. In this case the incentives to deviate for i are the greatest when the future reward for taking the present punishment is as far in the future as possible, that is when he has to face T periods of punishment. If player i plays r , as the strategy prescribes, he receives a payoff of $v_i^p = (1 - \epsilon^T)g(r; m) + \epsilon^T v_i$. If he deviates he receives ϵv_i^p at most. Choosing $\epsilon^T = d$ large enough for v_i^p to be positive, the player has no incentive to deviate during the punishment stage since $\epsilon < 1$.

Third, I consider the case in which player j has been the last player to deviate in the last T periods. In this case the incentives for i to deviate depend on the payoff for i when meeting j when the latter has not deviated, say g_{ij} . If $g_{ij} > g(m; r)$ it can be easily shown that the incentives for i to deviate are higher when j has to be punished for T periods. In this case, the expected utility of conforming with the equilibrium for i is at least $(1 - \epsilon)g + \epsilon v_i^0$, where $v_i^0 = (1 - \epsilon^{T+1})g_{ij} + \epsilon^{T+1}g(m; r) + (1 - \epsilon^T)g + \epsilon^T v_i$, while he would get at most $(1 - \epsilon)\bar{g} + \epsilon v_i^p$ by deviating, where $v_i^p = (1 - \epsilon^T)g(r; m) + \epsilon^T v_i$. Choose T to have $\epsilon^T = d \geq 2(0; 1)$. Then, given that $v_i > 0 \leq g(r; m) = g$ and $g(m; r) \leq g(r; m)$, it is easy to see that $v_i^0 > v_i^p$ independently of ϵ . Therefore for ϵ large enough it is true that $(1 - \epsilon)g + \epsilon v_i^0 > (1 - \epsilon)\bar{g} + \epsilon v_i^p$. If on the contrary, $g_{ij} \leq g(m; r)$, the incentives for i to deviate are higher when j has not deviated. But

we have already proven in the ...rst part of the proof that in that case i does not want to deviate.

Note that there is no contradiction in the requirements made on ϵ and T in the different parts of the proof, as it is only required that $\epsilon^T = d$ and ϵ are large enough.

■

As I mentioned before, with personal punishment a player, say player 1, may earn a very low payoff during the punishment stage (it may be the case that the player being punished, say player 2, is the only player that gives player 1 a positive payoff on the equilibrium path). Assumption 1 allows me to construct a social norm that makes sure that by deviating during the punishment stage, player 1 can only reduce his payoff even more and, then, has no incentives to deviate and the punishment is credible.

The social norms used in the proofs of the former theorems, in which players punish deviators since failing to do so is itself punished, resemble, in some ways, the enforcement of castes in India. When describing marriage customs in India, Akerlof [4] says: "The caste rules dictate not only the code of behavior, but also the punishment for infractions: violators will be outcasted; furthermore, those who fail to treat outcastes as dictated by caste code will themselves be outcasted."

While perfect information may be plausible in a small community it will certainly be implausible in a large one: it would be difficult for each player to know what every other player has done in the past if the number of players is very large. Therefore, Theorem 1 and 3 would not apply to the study of social norms and their impact on cooperation and inequality in large communities. In the next section I study the outcomes that can be supported by social norms under lower information requirements.

1.5 Local information

Even though in particular cases it is possible to forestall opportunistic behavior without the players having more information than their own experience, as the “contagious equilibrium” in Kandori [34], in more general cases extra information is necessary to provide the needed structure of punishments. In this section I assume that, in addition to their own experience, players have access to a local information processing system that gives players some information regarding their opponent’s past behavior. Following Okuno-Fujiwara and Postlewaite [41], the local information processing system has the following structure: 1) in period t agent i has a “status” or “tag” $z_i(t) \in Z_i$, where Z_i is a finite set (without loss of generality I can assume that $Z_i \subseteq \mathbb{N}_+$); 2) if players i and j are matched in period t and play $a_i(t)$ and $a_j(t)$, the update of status follows a transition mapping $(z_i(t+1); z_j(t+1)) = \zeta_{ij}(a_i(t); a_j(t); z_i(t); z_j(t))$; 3) if at time t player i is matched with player j , the former only knows his own history and $(z_i(t); z_j(t))$.

Based on the local information processing system, a social norm prescribes the behavior for each player as a function of his past history, his status and the status of the matched player. I show that any feasible and individually rational outcome can be supported by a social norm in equilibrium.

Theorem 4 (Folk theorem with local information): With local information any feasible and individually rational payoff $(v \in V : v \succeq 0)$ can be supported by a sequential equilibrium for δ large enough.

Proof. Consider the following social norm that yields v in equilibrium: if player i meets player j , and both are “nice” ($z_i = z_j = 0$), they play $(a_{ij}; a_{ji})$ (which yields v in equilibrium); if any of the two is “guilty” ($z_i \neq 0$ or $z_j \neq 0$) they both play m . The local information system, which assigns the “nice” and “guilty” labels, works as

follows: if a nice player conforms he keeps the $z = 0$ tag; if a nice player meets a guilty player with $z = \zeta$, he gets the $\zeta - 1$ tag next period; if a player deviates, he and his match get a tag $z = T$; if a guilty player conforms, he has the tag reduced one unit. Hence, $z_i = 0$ denotes that i has not deviated, has not seen a deviation in the last T periods and has not met someone (that met someone that met someone....) that deviated in the last T periods. Instead, $z_i > 0$ denotes that i has deviated in the last T periods or is aware that someone has deviated in the last T periods. Summarizing, the social norm is:

$$s_i(t) = \begin{cases} a_{ij} & \text{if } z_i(t) = z_j(t) = 0 \\ m & \text{if } z_i(t) \text{ or } z_j(t) \neq 0 \end{cases} \quad \text{and}$$

$$z_i(t+1) = \begin{cases} T & \text{if } a_i(t) \neq s_i(t) \text{ or } a_j(t) \neq s_j(t) \\ \zeta - 1 & \text{if } \max\{z_i(t); z_j(t)\} = \zeta \text{ and } a_{i,j}(t) = s_{i,j}(t) \\ 0 & \text{if } z_{i,j}(t) = 0 \text{ and } a_{i,j}(t) = s_{i,j}(t) \end{cases}$$

To prove that this social norm is a sequential equilibrium I show that no player has incentives to deviate in any possible information set if he believes the rest of the players will follow the social norm. I assume that outside the equilibrium path players believe that there are no more "guilty" players than those they have evidence there are. That is, if player i has been matched with player j with $z_j = \zeta > 0$, and no other guilty player, he believes that there are no more players with "guilty" tags than those that played with j in the last $T + 1 - \zeta$ periods (that is, all the players that have been matched with j when j obtained a "guilty" tag or after). It is easy to see that this beliefs are consistent (as defined in Kreps and Wilson [37]): if the probabilities of trembles are converging to zero the probability that other deviations happened (besides the ones observed by the player) also converges to zero.

Next I check that in every information node no player has incentives to deviate. First consider the case in which player i has the $z_i = 0$ tag (he has not deviated and

he has not seen a deviation or a “guilty” tag in the last T periods. Then he believes that all players have “nice” tags. In this situation the player i expects to earn at least $(1 - \epsilon)g + \epsilon v_i$ by conforming, and at most $(1 - \epsilon)\bar{g} + \epsilon v_i^p$ by deviating, where $v_i^p = (1 - \epsilon^T)g_m + \epsilon^T v_i$. Choose T to have $\epsilon^T = d/2 \in (0; 1)$. Then, given that $v_i > g_m$, it is true that $v_i > v_i^p$ independently of ϵ . Therefore for ϵ large enough it is true that $(1 - \epsilon)g + \epsilon v_i > (1 - \epsilon)\bar{g} + \epsilon v_i^p$.

Second consider the case in which player i has the $z_i = \zeta$ tag. This could be because either i has deviated in the past or because he has been matched with someone with a guilty tag. Then, player i believes that in ζ periods all players will be nice and he will earn v_i every period. The incentives for i to deviate and try to avoid the punishment are larger the farther away the end of the punishment phase is, that is when $\zeta = T$. In this case i obtains $(1 - \epsilon^T)g_m + \epsilon^T v_i = v_i^p$ by conforming, and ϵv_i^p by deviating. Given $v_i > 0$, I can choose $\epsilon^T = d$ large enough for v_i^p to be positive and, then, he has no incentive to deviate during the punishment stage since $\epsilon < 1$.

Third consider the case of a player i with $z_i = 0$ who is matched with a player j with $z_j = \zeta$. The analysis of this case coincides with the case above and i has no incentives to deviate.

Note that there is no contradiction in the requirements made on ϵ and T in the different parts of the proof, as it is only required that $\epsilon^T = d$ and ϵ are large enough.

■

Note that any player that deviates, sees a deviation or knows that a deviation occurred will be punished as the deviator until the end of the punishment stage. In this way, when a player knows that there has been a deviation his incentives to enforce the social norm do not depend on who has deviated. Whoever has deviated, once out-of-the-path beliefs are specified, it is easy to check that every player will enforce the punishment. The lack of personal punishment in this social norm allows me to

prove the folk theorem without restrictions on the stage game. A similar social norm could be used to prove Theorem 2 of Kandori [34] without the restriction imposed in that paper on the stage game.

But the social norm in the proof of Theorem 4 may be criticized, precisely, because not only the deviator is punished. The next result shows that with local information and Assumption 1 any feasible and strictly individually rational outcome can be supported in a sequential equilibrium with personal punishment.

Theorem 5 (Folk theorem with local information and personal punishment) Under local information, Assumption 1 and personal punishment any feasible and individually rational payoff $(v \geq v^* : v \geq 0)$ can be supported by a sequential equilibrium for δ large enough.

Proof. In Appendix. ■

The local information systems needed in the previous two proofs in this section are “simple”, in the sense that the number of tags needed is finite and does not increase with time or the number of deviations. Since the punishment stage consists of T periods an information mechanism with at least $T + 1$ tags per player is needed: one for each period of punishment and one for when the player is not in the punishment stage.

While the number of tags needed in Theorem 4 and 5 is finite, it can be very large. A way to drastically reduce the number of needed tags is to allow for a random transition rule of tags. In that case we can have two types of tags per player: guilty and nice, and, every period, all the guilty players that have conformed with the punishment are forgiven with probability $p \in (0, 1)$ and they become nice. In this way, p can be used to establish the severity of punishment, as T was doing before, with the need of only two tags per player. Random forgiveness eliminates the need of

counting the number of periods of punishment. As the next proposition shows, under Assumption 1 and personal punishment, any payoff vector that is strictly individually rational and feasible can be supported by a sequential equilibrium with only two tags (nice and guilty) if a random transition rule is allowed.

Proposition 6 (Folk theorem with local information, personal punishment and random transition rule) Under local information processing, Assumption 1, personal punishment and random transition rule, any feasible and individually rational payoff $(v \geq v^* : v \geq 0)$ can be supported by a sequential equilibrium for ϵ large enough.

Proof. In Appendix. ■

The equilibria described in this paper present some characteristics that are worth mentioning. First, in the equilibria in this section, the best response of any player in any situation depends only on his own and his match's tags. Any other information that players may have is irrelevant for making decisions: the best response is to follow the social norm, which tells players what to do under every combination of tags. In this way, the tags are sufficient statistics for the players decision making since they summarize all the relevant information.¹²

Second, the long run behavior of the community is not affected by any finite sequence of deviations. Contrary to some proofs of the Folk Theorem for N players¹³, if there have been deviations the prescribed actions revert to the original ones after T periods of punishment in the equilibria of this paper. Hence, the actions on the equilibrium path are globally stable: regardless of how many deviations have been up today, in the future the play of the game will return to the equilibrium play. This property is of special importance when studying societies with a large number of members. If a single deviation may take the community out of the equilibrium path

¹²This property of equilibria is called "straightforward" in Kandori [34].

¹³See Theorem 2 in Fudenberg and Maskin [23] or Theorem 1 in Abreu, Dutta and Smith [2].

for ever, it would be difficult to observe the equilibrium behavior in a large community in which each member has a small probability of making mistakes.

Third, the equilibria described in this paper are robust to small perturbations of the payoffs matrix of the stage game (of course this perturbations can not violate Assumption 1 in the cases in which this assumption is needed). Given that in the proofs of this paper all the inequalities are strict, if a social norm is an equilibrium under a given payoff matrix, it will also be an equilibrium with a payoff matrix that is arbitrarily close to the original one. Therefore, the equilibria presented here do not depend on a precise characterization of the players payoffs.

1.6 Conclusions

This paper analyzes the outcomes that can be supported by social norms in a society of infinitely long-lived and very patient agents that are randomly matched in pairs every period to play a given game. Unlike previous work that considered a society divided in two groups and all the members of each group receiving the same payoff, I only restrict this matching procedure to be independent of history and time. I find that any feasible and individually rational outcome can be supported by a self-enforcing social norm under both perfect information and a simple local information system. I also find that the same result holds, under some restrictions on the stage game, if the social norms can only punish deviators.

To show the richness of the equilibria analyzed in this paper I present here several outcomes that can be supported in equilibrium by social norms in a simple community. I consider a community of ten members that are matched uniformly to play the following prisoner's dilemma:

c	d
c	2, 2 -1, 4
d	4, -1 0, 0

I present first a society in which social norms support an equal and efficient outcome.

Optimal egalitarian society: In equilibrium all the players play c and receive a payoff of 2. This outcome is feasible and individually rational and, hence, can be supported by a self-enforcing social norm under either perfect information or local information. Therefore, a social norm, with its promise of punishment to deviators (and the consequent inequality under personal punishment) can support an egalitarian outcome that Pareto dominates the inefficient egalitarian equilibrium of the one shot game.

But the results in this paper explain not only how social norms can provide incentives to curtail opportunistic behavior and support cooperation in a community, but also how they can support outcomes characterized by inequality as the next two examples illustrate.

Kingdom: As in the example in the Introduction, consider a society in which in equilibrium one player, the “king”, always plays d and the rest of the players, the “serfs”, play c. In equilibrium the king receives a payoff of 4 and each of the serfs receive $\frac{5}{3}$. This outcome is feasible and, since both payoffs are positive, it is also individually rational and, then, can be supported by a self-enforcing social norm under either perfect information or local information. The king gets the maximum payoff of the game since each serf prefers to be exploited by the king instead of rebelling and suffering the future punishment.

Caste (or Class) society: Consider a society divided in three castes: one player belongs to the high caste and in equilibrium he always plays d; three players belong

to the middle caste and they play c when matched with a member of the same or higher caste and d otherwise; and the remaining six players belong to the lower caste and they always play c in equilibrium. Then, the high caste member receives 4, the middle caste members receive 3 and the low caste members receive $\frac{2}{3}$. This outcome is feasible and individually rational and, then, can be supported by a self-enforcing social norm under either perfect information or local information.

These examples show that social norms can support unequal outcomes even when all the members of the community are basically equal. While in these examples the division of members among the different groups is arbitrary, in reality it may correspond to differences in race, religion or gender. In this way, the results in this paper show how self-enforcing social norms may perpetuate discrimination among members of society even when all of them are intrinsically equal. These results show that discrimination and inequality can exist even when there are no differences in human capital or productivity and no taste for discrimination.

2 Tacit Collusion under Interest Rate Fluctuations

2.1 Introduction

It is well known that oligopolies can use the threat of future price wars to sustain prices above competitive levels if firms care enough about the future (Friedman [19]). The extent to which firms care about the future depends primarily on the interest rate if the firms' objective is to maximize the present value of profits. The firms' discount factor may also depend on other (secondary) forces such as the probability that the product may become obsolete. Given that the interest rate and other variables that affect the discount factor are constantly changing, it is important to study tacit collusion under discount factor fluctuations.

I characterize collusive prices and profits when the discount factor changes over time, under both price and quantity competition, and I show that collusive prices and profits increase with both present and future levels of the discount factor, but decrease with its volatility. These results have important implications not only for the study of collusion but also for repeated game theory in general.

Repeated game theory has until now largely considered the discount factor as a fixed preference parameter¹⁴. Oligopoly games are one example among many of an environment in which it is natural to assume that the discount factor changes over time. Another example would be exogenous changes in the probability that a partnership might end. Thus, the volatility of the discount factor may be an important determinant of cooperation for many kinds of repeated games, not just oligopoly.

With respect to the study of collusion, previous literature has looked at the effect of demand fluctuations on prices, but not discount factor fluctuations. In a seminal

¹⁴The exception is Baye and Jansen [9] that provides folk theorem results for repeated games with stochastic discount factors.

paper, Rotemberg and Saloner [46] show that collusive prices may be countercyclical. In this paper, I not only introduce the role of volatility to the repeated game theory literature but I also show that under discount factor fluctuations the results are less ambiguous and more robust than under demand fluctuations. This paper also presents several new comparative static results that can be used in empirical studies of collusive pricing. In addition, this paper underscores the role of interest rates and imperfect competition in aggregate fluctuations. Any change in policy, technology or preferences that affects the real interest rate (either in level or volatility) may have an impact on aggregate production through changes in collusive behavior.

The environments I study and the specific results I find are as follows. I consider first the case in which the discount factor, identical for all firms, is randomly and independently drawn every period. I characterize the maximum symmetric tacit collusion prices and profits that can be supported in an environment in which firms are identical and they compete repeatedly on either price or quantity. The three main results derived from this characterization, with the third one the most interesting, are as follows.

First, the higher the discount factor in a given period, the higher the collusive prices and profits that can be supported in equilibrium in that period. The intuition behind this is straightforward: the higher the discount factor, the stronger the threat of future price wars and the higher prices and profits can be without firms deviating.

Second, the greater the probability of high discount factors, the higher the collusive prices and profits that can be supported in equilibrium. Again the intuition is straightforward. From the first result we know that the higher the realization of the discount factor, the higher collusive prices and profits will be. Hence, a shift in the distribution function to higher discount factors would result in an increase in the expected value of collusive profits and an increase in the threat of future punishment,

allowing higher equilibrium prices and profits.

Third and more interestingly, I show that the higher the volatility of the discount factor, the lower the collusive prices and profits that can be supported in equilibrium. The reason for this is twofold. First, given that the combination of the incentive compatibility and feasibility constraint results in a concave collusive profit function (as a function of the discount factor), an increase in volatility leads to a decrease in expected profits. Second, this decrease in expected profits reduces the size of future punishment and hence results in a decrease in equilibrium profits and prices. This volatility effect is not secondary to the first two level effects. I show that it plays an important role in determining collusive prices and profits.

It is important to note that allowing for the more realistic case of positively correlated discount factors will not affect the main results per se, given that both a high discount factor today and in the future make it easy to support collusion.

Two other results of this paper are worth noting. First, I show that under quantity competition the optimal symmetric punishment has a simple stick-and-carrot characterization (the punishment takes only one period and is as big as possible in equilibrium), extending the results of Abreu [1] from the fixed discount factor case.

Second, I show that under price competition an increase in the number of firms reduces collusive prices and profits. The reason is that the greater the number of firms the greater the share of the market that can be captured by a deviation, and, hence, the lower equilibrium profits and prices must be to avoid deviations. In the case of quantity competition, more work is needed to assess the validity of this result, since not only do the incentives to deviate change with the number of firms, but so may the threat of future punishment.¹⁵

¹⁵To my knowledge, the effect of the number of firms on tacit collusive prices under quantity competition remains to be solved also for the case of fixed discount factors.

The rest of the paper is organized as follows. In Section 2, I relate this paper to the previous literature. In Sections 3 and 4, I study optimal tacit collusion under price and quantity competition, respectively. In Section 5, I analyze some extensions to the basic model. In Section 6, I conclude.

2.2 Related literature

The related literature falls into six categories: 1) studies of the effects of demand fluctuations on optimal tacit collusion, 2) customer markets and oligopolistic pricing, 3) empirical studies of collusive pricing, 4) studies of the role of oligopolies in macroeconomic fluctuations, 5) studies of optimal punishment schemes under quantity competition, and 6) repeated games with mixed discount factors.

Demand fluctuations and optimal tacit collusion: The well known paper by Rotemberg and Saloner [46] offers interesting results with respect to tacit collusion that also follow from changes in the relative importance of present and future profits. In their paper, however, those changes are driven by changes in demand, not the discount factor. This difference in the source of the changes in the relative importance of future and present profits is not trivial and leads to significantly different results.

First, in this paper an increase in the discount factor always has a nonnegative effect on the equilibrium price, while in Rotemberg and Saloner [46] an increase in demand may result in either an increase or a decrease in price. In their model, the threat of a future price war, which depends on the expectation of future equilibrium profits, results in an upper bound to equilibrium collusive profits. Hence, at this upper bound on profits, increases in demand do not result in increases in profits but a decrease in prices. If instead the demand is so low that the upper bound to profits is not binding, a small increase in demand will result in an increase in prices. In addition, contrary to discount factor fluctuations, the effect of demand fluctuations on prices

may not be robust to assuming quantity competition instead of price competition, as Rotemberg and Saloner [46] note, or to the existence of capacity constraints, as Staiger and Wolak [53] note.

Second, while in this paper an increase in the volatility of the discount factor always results in a decrease in profits and prices, in Rotemberg and Saloner's model an increase in the volatility of demand is again ambiguous -it may result in an increase in profits and prices.¹⁶ Therefore, in contrast to fluctuating demand, changes in the level or volatility of the discount factor have unambiguous effects.

The third difference between the two models lies in the effect that present and future shocks have on collusive prices. In Rotemberg and Saloner's model, a high demand today makes it difficult to support collusion since it offers greater incentives to deviate, while a high demand in future periods makes it easy to collude today given that a future price war becomes a bigger threat. In contrast, in this model both high discount factors today and in the future make it easy to support collusion given that both increase the threat of future punishment.

The different effects that present and future levels of demand have on collusive pricing in Rotemberg and Saloner [46] led to several studies of whether their results were robust to correlation on demand shocks. Kandori [33] finds conditions under which demand correlation does not affect the result of countercyclical collusive pricing. Haltiwanger and Harrington [28] study tacit collusion under deterministic cyclic fluctuations of demand and find that higher collusive prices can be supported when demand is increasing than when it is decreasing. Bagwell and Staiger [8] study tacit collusion when demand shifts stochastically between high and low growth rates and

¹⁶Rotemberg and Saloner [46] do not provide this comparative static result but straightforward examples can be obtained from their model. In their model the profit function may be convex in the fluctuating parameter so that an increase in volatility increases the expected profits and moves up the incentive compatibility constraint.

...nd that collusive prices are higher for high rates of demand growth if demand growth rates are positively correlated through time.

Under discount factor fluctuations, the issue of positive correlation is less important than under demand fluctuations, given that both high discount factors today and in the future increase today's collusive prices. However, I show that the discount factor volatility may be important in understanding how more general discount factor fluctuations affect the basic results.

Customer markets and oligopolistic pricing: There are other environments in which changes in the discount factor may affect oligopoly prices. In models of customer markets, as in Phelps and Winter [43] and Gottfries [27], and models of competition when consumers have switching costs, as in Klemperer [35] and Chevalier and Scharfstein [14], firms face a trade-off between charging high prices to extract the surplus from current customers and charging low prices to attract new customers (whose surplus can be extracted later). In these models an increase in the discount factor increases the incentives to invest in new customers and results in lower prices, as Rotemberg and Woodford's [47] and Klemperer [35] note. In contrast, in the model of tacit collusion presented here, an increase in the discount factor results in higher prices. The higher the discount factor, the stronger the threat of future price wars and the higher the prices that can be supported in equilibrium.

Empirical literature on collusive pricing: Based on the frameworks established by Rotemberg and Saloner [46] or Porter [44] and Green and Porter [24], there is an extensive literature that concentrates on changes in demand as sources of changes in collusive pricing. Those papers do not include the interest rate in their studies, see for example Porter [45], Domowitz, Hubbard and Petersen [16], Slade [51], Ellison [17] and Borenstein and Shepard [10].¹⁷ An exception can be found in Rotemberg

¹⁷For a review of empirical studies up to the 1980s see Bresnahan [11].

and Woodford's [47] study of markups and the economic cycle. Working with aggregate log-linearized data around the steady state of an intertemporal macroeconomics model, they use rates of return to "instrument" for the firm's expectations of future profits and find that high interest rates result in low markups. In this paper I present additional comparative static results arising from interest rate movements that may be used in empirical studies of collusive pricing.

Collusive pricing and macroeconomic fluctuations: Previous literature has related tacit collusive pricing with macroeconomic fluctuations. For example, Rotemberg and Saloner [46] present a simple two-sector general equilibrium model in which one sector is oligopolistic and the other one is perfectly competitive. They show that exogenous shifts in demand towards the oligopolistic sector induce a decrease in collusive prices (since it increases the short run incentives to deviate) and may result in an increase in aggregate production. Rotemberg and Woodford [48] present a real business cycle model with tacitly colluding oligopolistic producers. In their model, an increase in government expenditure raises the short run incentives to deviate and results in a decrease in collusive prices. This, in turn, increases real wages, employment and output. In addition, the authors note that the increase in government expenditure may result in an increase in interest rates (since consumers must postpone consumption), which reinforces the first effect by lowering the threat of future punishments.

In this paper I present another way in which tacit collusion may result in aggregate fluctuations. Any change in policy, technology or preferences may have an impact on aggregate production through changes in collusive behavior, not only by affecting the real interest rate level, but also by affecting its volatility.

Optimal punishment schemes under quantity competition: Abreu [1] provides a simple stick-and-carrot characterization of optimal symmetric punishments for a fixed discount factor under quantity competition: "...the most efficient way to provide low

payoffs, in terms of incentives to cheat, is to combine a grim present with a credibly rosy future.”¹⁸ In this paper I show that the stick-and-carrot characterization extends to the case of discount factor fluctuations, with both the size of the stick and the size of the carrot depending on the realization of the discount factor.

The level effect and repeated games with fixed discount factors: It is well known that, for repeated games with fixed discount factors, the higher the discount factor, the bigger the set of equilibrium outcomes will be (see for example, Abreu, et al. [3]). In this paper I show that under discount factor fluctuations it is not only the level of the discount factor that matters, but also its volatility.

2.3 Price competition

Consider a market with N identical firms with a constant marginal cost of c and facing a demand function $D(p)$ ($D'(p) < 0$). Firms compete repeatedly on price and the demand is divided equally among the firms charging the lowest price in each period. Firms only care about profits and are risk neutral and, hence, their objective is to maximize the discounted stream of profits. The distinctive feature of this model is that the discount factor δ_t , which discounts earnings from $t + 1$ to t , is a continuous, independent and identically distributed random variable, between a and b , with p.d.f. $f(\delta_t)$ and c.d.f. $F(\delta_t)$.

The timing of the game in a given period t is as follows: the firms observe the realization of the discount factor, δ_t , then they choose the price for that period and finally they observe the market clearing price, quantities and payoffs. All the characteristics of the environment are common knowledge.

Given that firms cannot commit to charge a given price or sign contracts amongst themselves or with third parties regarding prices, any equilibrium of the model must

¹⁸Abreu [1], pg. 206.

be a subgame perfect equilibrium of the infinitely repeated oligopoly game. I restrict my attention to equilibria in which all the firms charge the same price p . In this symmetric case, I can write the profits of each firm as $\pi_i(p) = \frac{(p - c)D(p)}{N}$ and total industry profits as $\Pi(p) = (p - c)D(p)$. I assume that there exists a price p^m that maximizes the total industry profits, that is, p^m is the monopoly (or perfect collusion) price. Denote $\pi^m = \pi_i(p^m)$ as the monopoly profit per firm.

2.3.1 Optimal tacit collusion with a random discount factor

It is well known that in repeated oligopoly games, prices above the marginal cost can be supported in equilibrium if any price undercutting triggers future price wars. In the case of price competition, the best price war, in terms of punishment, is the reversion forever to the Bertrand equilibrium after any deviation. This punishment gives a discounted payoff of zero. Any other punishments that would result in a lower payoff are not enforceable given that any firm can make sure to earn zero profits by charging a price equal to the marginal cost in every period.

Given this punishment, I look for symmetric optimal tacit collusion strategies - strategies without price differences among firms and that in equilibrium support the maximum present value of profits. Since the environment in which the firms interact does not change over time, with the exception of the discount factor, the optimal tacit collusion solution will consist of the highest equilibrium price that the firms can charge in a period given the discount factor in that period. Therefore the solution will consist of a function $p^*(\delta) : [a; b] \rightarrow [c; p^m]$ which gives the highest equilibrium price that can be supported for each discount factor. This in turn defines a function $\pi^*(\delta) : [a; b] \rightarrow [0; \pi^m]$, which denotes the optimal tacit collusion equilibrium profits as a function of the period discount factor.

Fortunately, in the search for the optimal tacit collusion behavior it is enough

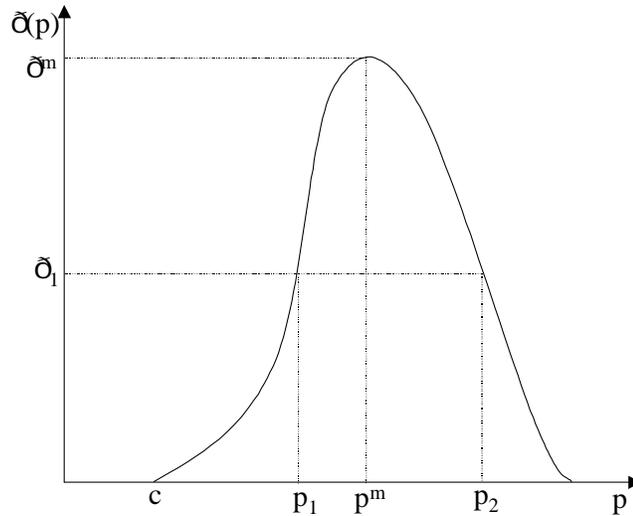


Figure 1: The profit function

to work with $\check{\alpha}^{\pm}$. As we see in Figure 1, a given level of profits, for example $\check{\alpha}_1$, can result from different prices, such as p_1 and p_2 . Given that I am interested in the optimal levels of profits that can be supported under tacit collusion and the fact that $\check{\alpha}_1$ may be supported more easily by p_1 than by p_2 ,¹⁹ I only consider the increasing part of the profit function. In this way, for every profit lower than $\check{\alpha}^m$ corresponds one and only one price lower than p^m . Therefore, I can define the function $\check{A}(\check{\alpha}) = \check{\alpha}^{-1}(\check{\alpha}) : [0; \check{\alpha}^m] \rightarrow [c; p^m]$, and once I solve for $\check{\alpha}^{\pm}$, I can recover p^{\pm} as $p^{\pm} = \check{A}(\check{\alpha}^{\pm})$. Note that $\check{A}(\check{\alpha})$ is increasing on $\check{\alpha}$.²⁰

Given the simplicity of the optimal punishment (reversion to Bertrand) and the fact that we are able to uniquely relate profits to prices, I concentrate on the charac-

¹⁹As it will be clear soon, $\check{\alpha}_1$ can be supported more easily by p_1 than by p_2 since the optimal deviation from p_1 yields $N\check{\alpha}(p_1)$ which is lower than $N\check{\alpha}(p^m)$ which can be obtained deviating from p_2 .

²⁰For simplicity, I will assume for the rest of the section that $\check{A}(\check{\alpha})$ is differentiable and, hence, $\check{A}'(\check{\alpha}) > 0$:

terization of the equilibrium optimal tacit collusion profits $\pi^*(\alpha)$ without relying on the strategies that result in that equilibrium path. I study next the restriction on collusive profits for then characterizing the optimal tacit collusion solution.

Using the recursiveness of the problem, the present value at t of an infinite stream of profits can be written as:

$$V(\alpha_t) = \pi(\alpha_t) + \alpha_t \int_a^b V(\alpha_{t+1}) f(\alpha_{t+1}) d\alpha_{t+1} \quad (1)$$

where $\pi(\alpha_t)$ denotes the profits that the firms receive at time t if the discount factor is α_t . Integrating over equation 1 and rearranging we have $\int_a^b V(\alpha_t) f(\alpha_t) d\alpha_t = \frac{1}{1 - \alpha} \int_a^b \pi(\alpha_t) f(\alpha_t) d\alpha_t$, where $\bar{\alpha}$ is the expected value of α_t . Plugging this into (1), the present value of profits can be written as:

$$V(\alpha_t) = \pi(\alpha_t) + \frac{\alpha_t}{1 - \bar{\alpha}} \int_a^b \pi(\alpha_{t+1}) f(\alpha_{t+1}) d\alpha_{t+1} \quad (2)$$

Since these firms cannot commit to a given price, in equilibrium they must be unwilling to charge a price different from the equilibrium price. How much can a firm gain from deviating? If all the firms are charging the same price above marginal cost, a single company can decrease its price by a penny and capture the whole market. Therefore, if the equilibrium profit is $\pi(\alpha_t)$, a single company can gain $(N - 1)\pi(\alpha_t)$ by deviating (if we forget about pennies). For firms to be unwilling to deviate, punishment must follow a deviation. How much can a firm lose from being punished? As described before, the best punishment is to revert forever to the Bertrand equilibrium (the Nash equilibrium of the one stage game). Under this threat if one firm deviates it will earn the total industry profit the period of deviation but then it will earn zero profits forever. Then, for no firm to have an incentive to

deviate, the following must hold:

$$\pi(\pm_t) \geq \frac{\pi(\pm_{t+1})}{(N-1)} + \int_a^{\pm_t} \pi(\pm_{t+1}) f(\pm_{t+1}) d\pm_{t+1} \quad (3)$$

In addition, the profits per firm cannot be greater than under monopoly pricing:

$$\pi(\pm_t) \leq \pi^m \quad (4)$$

Therefore it is clear that under the optimal symmetric tacit collusion equilibrium firms will choose profits as large as possible without violating the incentive compatibility constraint (3) and the feasibility constraint (4).²¹ Then, dropping the subindexes for simplicity, the optimal tacit collusion profit levels $\pi^*(\pm)$ is a function from $[a; b]$ to $[0; \pi^m]$ subject to the following equation:

$$\pi^*(\pm) = \min \left\{ \frac{\pi^*(\pm)}{(N-1)} + \int_a^{\pm} \pi^*(\pm^0) f(\pm^0) d\pm^0; \pi^m \right\} \quad (5)$$

Note that this equation does not provide the optimal tacit collusion profit since $\pi^*(\pm)$ appears in both sides of it. Equation (5) is just a necessary condition for optimal tacit collusion. In fact, choosing profits equal to zero for every discount factor solves this equation. From the possible many solution to equation (5), the one that provides the highest profit for each discount factor is the optimal tacit collusion solution: $\pi^*(\pm)$. The following proposition fully characterizes the function $\pi^*(\pm)$.

Proposition 7 The function $\pi^*(\pm)$ depends on $f(\pm)$ and N in the following way:

- 1) if $\frac{a}{N-1} \leq \pm$, $\pi^*(\pm) = \pi^m$;

²¹It could be argued that that is not necessary since having profits lower than possible in a finite subset does not affect the expected value. But if we want the solution to be independent of the discount factor of the first period, profits must be as high as possible for every possible value of the discount factor.

- 2) if $\frac{N_i-1}{N} \bar{\pm} < 1 < \frac{a}{N_i-1}$, $\frac{1}{4}^{\pm}(\pm) = \frac{1}{4}^m$ for $\pm \geq \underline{b}$ and $\frac{1}{4}^{\pm}(\pm) = \frac{\pm}{\underline{b}} \frac{1}{4}^m$ for $\pm < \underline{b}$, for a number $\underline{b} \in [a; b]$ that solves the following equation: $\underline{b} = (N_i - 1) \int_a^{\underline{b}} F(\pm) d\pm + \int_{\underline{b}}^b F(\pm) d\pm$;
- 3) if $\bar{\pm} < \frac{N_i-1}{N}$, $\frac{1}{4}^{\pm}(\pm) = 0$.

Proof. Case 1): $\bar{\pm} \geq 1 < \frac{a}{N_i-1}$ implies that $\frac{1}{4}^m = \frac{a}{(N_i-1)(1-\bar{\pm})} \frac{1}{4}^m = \frac{\pm}{(N_i-1)(1-\bar{\pm})} \frac{1}{4}^m$ and perfect collusion, $\frac{1}{4}^{\pm}(\pm) = \frac{1}{4}^m$, can be supported for every discount factor.

Case 2): Consider the case in which the two terms inside the brackets in equation (5) are binding for different ranges of \pm . Given that the first term is increasing in \pm , it would be binding for $\pm < \underline{b}$, the second term would be binding for $\pm > \underline{b}$, and both terms equal and binding for $\pm = \underline{b}$, where $\underline{b} \in [a; b]$. In this case, integrating over equation (5) and denoting the expected profit as A :

$$A = \int_a^{\underline{b}} \frac{z^{\underline{b}}}{(N_i - 1) \int_a^{\pm} \frac{1}{1 - \bar{\pm}}} Af(\pm) d\pm + \int_{\underline{b}}^b F(\pm) \frac{z^{\underline{b}}}{\underline{b}} \frac{1}{4}^m$$

In addition, given that for $\pm = \underline{b}$ both terms of equation (5) are equal, the expected profit can be also written as:

$$A = \frac{\frac{1}{4}^m (N_i - 1) \int_a^{\underline{b}} \frac{z^{\underline{b}}}{1 - \bar{\pm}}}{\underline{b}}$$

Combining these two equations and by the fact that (integrating by parts) $\int_a^{\underline{b}} \pm f(\pm) d\pm = sF(s) \Big|_a^{\underline{b}} - \int_a^{\underline{b}} F(\pm) d\pm$, the number \underline{b} solves the following equation:

$$\underline{b} = (N_i - 1) \int_a^{\underline{b}} \frac{z^{\underline{b}}}{1 - \bar{\pm}} + \int_a^{\underline{b}} F(\pm) d\pm \tag{6}$$

It remains to be shown that, under the conditions of case 2), the number \underline{b} that solves equation (6) exists and is unique. Write $H(r) = (N_i - 1) \int_a^r \frac{z^{\underline{b}}}{1 - \bar{\pm}} + \int_a^r F(\pm) d\pm - r$. Then $H(\underline{b}) = 0$. If $\frac{N_i-1}{N} < \bar{\pm} < 1 < \frac{a}{N_i-1}$, it can be easily seen that $H(a) = (N_i - 1) \int_a^a \frac{z^{\underline{b}}}{1 - \bar{\pm}} + \int_a^a F(\pm) d\pm - a > 0$ and $H(b) = (N_i - 1) \int_a^b \frac{z^{\underline{b}}}{1 - \bar{\pm}} + \int_a^b F(\pm) d\pm - b < 0$. In addition, $H(r)$ is continuous and strictly

decreasing ($\frac{\partial H(r)}{\partial r} = F(r) - 1 < 0$ for $a < r < b$). Then, there exists a unique number \underline{r} , between a and b , that makes $H(\underline{r}) = 0$. If $\bar{r} = \frac{N_i - 1}{N}$, $H(b) = 0$ and $\underline{r} = b$ is the unique solution since $H(\cdot)$ is strictly decreasing.

Case 3): From the analysis of the previous two cases follows that when $\bar{r} < \min\left\{\frac{N_i - 1}{N}; 1 - \frac{a}{N_i - 1}\right\}$ neither a solution with perfect collusion for all or some discount factors is feasible, nor a solution with imperfect collusion is feasible. Then, the only possible solution to equation (5) is $\frac{1}{4}^a(\pm) = 0$. Since $\frac{N_i - 1}{N}$ can be greater than $1 - \frac{a}{N_i - 1}$ only if $a > \frac{N_i - 1}{N}$, in which case \bar{r} can never be lower than $\frac{N_i - 1}{N}$, it follows that $\frac{1}{4}^a(\pm) = 0$ if $\bar{r} < \frac{N_i - 1}{N}$. ■

Proposition 1 shows that, depending on the distribution of the discount factor and the number of firms, there are three mutually exclusive cases that result in three different types of optimal tacit collusion. In case 1), $\bar{r} > 1 - \frac{a}{N_i - 1}$, any possible realization of the discount factor is high enough for each firm to value the future monopoly profits more than the one stage profits of deviation, and, hence, perfect collusion is an equilibrium for any discount factor. On the contrary, in case 3), $\bar{r} < \frac{N_i - 1}{N}$, all the realizations of the discount factor are too low to be able to support any level of collusion. In between these two cases, case 2), perfect collusion can be supported for a range of high realizations of the discount factor while only lower levels of profits can be supported for a range of low realizations. The reason for this is that while for low discount factors it is not possible to support full collusion, it may still be possible to satisfy the incentive compatibility constraint by reducing the present incentives to deviate. For this, the present profits should be lowered so that no firm has an incentive to deviate. In this case, an increase in the discount factor results in an increase in the optimal tacit collusion profits and, hence, in prices. Given that in the other two cases changes in the discount factor have no effect on profits, the next theorem follows.

Theorem 8 $\frac{dW^*(\pm)}{d\pm} \geq 0$ and $\frac{dp^*(\pm)}{d\pm} \geq 0$.²²

Note that the characterization of optimal tacit collusion under discount factor fluctuations includes the case of a fixed discount factor. For the fixed discount factor case, $a = b$, Proposition 1 coincides with the text book solution: perfect collusion if $\pm \geq \frac{N_i - 1}{N}$ and no collusion otherwise.

2.3.2 The effects of changes in $f(\pm)$

The characterization of the optimal tacit collusion equilibrium leads to interesting comparative statics results with respect to changes in the distribution function of the discount factor: 1) the higher the probability of high discount factors, the higher the equilibrium prices and profits, and 2) the higher the volatility of the discount factor, the lower the equilibrium prices and profits.

As an intermediate step to these results, I study first how changes in the distribution function modify the range of perfect collusion under case 2) of Proposition 1. For a cumulative distributions functions F define $\bar{\pm}_F$ as the expected discount factor and \underline{b}_F as the solution limit to perfect collusion if case 2) applies.

Lemma 9 Consider two cumulative distributions functions, F and G , such that $\frac{N_i - 1}{N} < \bar{\pm}_{F;G} < 1$, $\frac{a}{N_i - 1}$ and F second-order stochastically dominates²³ G , then $\underline{b}_F \geq \underline{b}_G$.

Proof. From the definition of \underline{b}_F : $H_F(\underline{b}_F) = (N_i - 1) \int_{\underline{b}_F}^{\bar{\pm}_F} F(\pm) d\pm + \int_{\underline{b}_F}^{\bar{\pm}_F} F(\pm) d\pm = 0$. By second-order stochastic dominance $\int_a^b F(\pm) d\pm \geq \int_a^b G(\pm) d\pm$ and $\bar{\pm}_F \geq \bar{\pm}_G$.

²²I omit straightforward proofs.

²³For two cumulative distributions functions $F(\pm)$ and $G(\pm)$, F second-order stochastically dominates G if for any r , $a < r < b$, $\int_a^r F(\pm) d\pm \geq \int_a^r G(\pm) d\pm$, and the inequality is strict in some range. In that case, it can be proven that $\bar{\pm}_F \geq \bar{\pm}_G$ and $\int_a^b u(\pm) f(\pm) d\pm \geq \int_a^b u(\pm) g(\pm) d\pm$, for any increasing concave twice-piecewise-differentiable function $u(\pm)$. See Hirshleifer and Riley [29].

Therefore, $H_G(\underline{b}_F) = (N-1) \int_{\underline{b}_F}^{\bar{c}} G(\pm) d\pm \geq 0$ and, given that $H_G(\cdot)$ is strictly decreasing and the conditions on \bar{c}_G , there exists $\underline{b}_G \geq \underline{b}_F$ such that $H_G(\underline{b}_G) = 0$. ■

Denote $\pi_F^a(\pm)$, $E\pi_F^a$ and $p_F^a(\pm)$ as the optimal tacit collusion profit, its expected value and optimal collusion prices under F , respectively.

Theorem 10 Consider two cumulative distribution functions, F and G , such that F second-order stochastically dominates G , then $\pi_F^a(\pm) \geq \pi_G^a(\pm)$ and $p_F^a(\pm) \geq p_G^a(\pm)$ for every \pm . In addition, $E\pi_F^a \geq E\pi_G^a$.

Proof. By second-order stochastic dominance $\bar{c}_F \geq \bar{c}_G$. So, from Proposition 1, we can see that if the solution under F belongs to case 1), the solution under G can belong to any of the three cases. If the solution under F belongs to case 2), the solution under G can belong to cases 2) or 3). And if the solution under F belongs to case 3), the solution under G must belong to the same case. For most of this combinations it is straight forward to see that $\pi_F^a(\pm) \geq \pi_G^a(\pm)$ for every \pm . The situation in which both the solution under F as under G belong to case 2) needs more analysis. Since F second-order stochastically dominates G , by Lemma 3, $\underline{b}_F \geq \underline{b}_G$. Then, $\pi_F^a(\pm) = \frac{\pm}{\underline{b}_F} \pi^m \geq \pi_G^a(\pm) = \frac{\pm}{\underline{b}_G} \pi^m$ if the incentive compatibility constraint is binding in both cases, $\pi_F^a(\pm) = \pi^m \geq \pi_G^a(\pm) = \frac{\pm}{\underline{b}_G} \pi^m$, if the incentive compatibility constraint binds for G but not for F , and $\pi_F^a(\pm) = \pi_G^a(\pm) = \pi^m$ if it is not binding for any of the two. Therefore, $\pi_F^a(\pm) \geq \pi_G^a(\pm)$ for every \pm .

The result with respect to prices follows directly from the positive relationship between profits and prices.

Note that $\pi_F^a(\pm)$ is increasing and concave, hence, by second-order stochastic dominance and $\pi_F^a(\pm) \geq \pi_G^a(\pm)$ for every \pm we have that $E\pi_F^a = \int_a^{\bar{c}} \pi_F^a(\pm) f(\pm) d\pm \geq$

$$\int_a^R \frac{1}{4}_F^{\pi}(\pm)g(\pm)d\pm \succeq \int_a^R \frac{1}{4}_G^{\pi}(\pm)g(\pm)d\pm = E \frac{1}{4}_G^{\pi}. \blacksquare$$

The intuition of this result becomes clear if we consider two particular cases of second order stochastic dominance: when F first-order stochastically dominates²⁴ G and when G is a mean preserving spread of F.

From Theorem 2 we know that given a distribution of the discount factor, say G, equilibrium prices and profits are increasing in the realization of the discount factor. Then, a shift in the distribution function to higher values (which yields a cumulative distribution function F that first-order stochastically dominates G), would result in an increase in expected profits. This, in turn, increases the threat of future punishments and increases equilibrium prices and profits.

Corollary 11 If F first-order stochastically dominates G, then $\frac{1}{4}_F^{\pi}(\pm) \succeq \frac{1}{4}_G^{\pi}(\pm)$ and $p_F^{\pi}(\pm) \succeq p_G^{\pi}(\pm)$ for every \pm . In addition, $E \frac{1}{4}_F^{\pi} \succeq E \frac{1}{4}_G^{\pi}$.

From Proposition 1 we know that given a distribution factor, say F, the optimal tacit collusion profit function is concave in the discount factor. Therefore, a mean preserving spread (which yields G), would result in a reduction in expected profits. This, in turn, reduces the threat of future punishment and results in lower equilibrium prices and profits.

Corollary 12 If G is a mean preserving spread of F, then $\frac{1}{4}_F^{\pi}(\pm) \succeq \frac{1}{4}_G^{\pi}(\pm)$ and $p_F^{\pi}(\pm) \succeq p_G^{\pi}(\pm)$ for every \pm . In addition, $E \frac{1}{4}_F^{\pi} \succeq E \frac{1}{4}_G^{\pi}$.

Therefore, the volatility of the discount factor is inversely related to the firms' profits. This result might seem somewhat counterintuitive given that the firms are

²⁴For two cumulative distributions functions $F(\pm)$ and $G(\pm)$, F first-order stochastic dominates G if for all r , $a \leq r \leq b$, $F(r) \leq G(r)$, and the inequality is strict in some range. In that case, it can be proven that F second-order stochastic dominates G and $\int_a^R u(\pm)f(\pm)d\pm \succeq \int_a^R u(\pm)g(\pm)d\pm$, for any increasing piecewise differential function $u(\pm)$. See Hirshleifer and Riley [29].

risk neutral, but the intuition is in fact simple. The combination of the incentive compatibility constraint with the feasibility constraint yields a profit function which is concave in the discount factor even when firms are risk neutral. Hence, an increase in volatility reduces expected profits reducing the threat of future punishment and lowering equilibrium prices and profits.

2.3.3 The effects of changes in the number of firms

With N firms in the market a single firm may steal a fraction $\frac{N-1}{N}$ of the market by undercutting the price. Since this fraction is increasing in the number of firms, the higher the number of firms the higher is the present profit from deviation for a given profit, and the more difficult it will be to support collusion. In fact, it can be easily seen from Proposition 1 that for any distribution of the discount factor, there is large enough number of firms above which it is not possible to support any collusion.²⁵ Define $\pi_N^*(\delta)$ and $p_N^*(\delta)$ as the optimal tacit collusion profits and prices for N firms.

Theorem 13 If $N > \frac{1}{1-\delta}$, then $\pi_N^*(\delta) = 0$ and $p_N^*(\delta) = c$.

In addition, it can be easily shown that increases in the number of firms reduce prices and profits (at both industry and firm levels). The next theorem follows from restatement Proposition 1 in terms of industry profits π_N^* and noting that the range of perfect collusion in case 2) shrinks with increases in the number of firms.

²⁵It is interesting to note that this result does not depend on fixing the size of the market while changing the number of firms. If both the size of the market and the number of firms increase in the same proportion (that would consist on multiplying the demand function $D(p)$ and the number of firms N by a positive integer), the same result holds. Since an increase in the number of firms and size of the market leaves monopoly profits per firm unchanged but increases the incentives to deviate, the scope of collusion diminishes up to a point in which it disappears.

Theorem 14 Consider two different number of firms N and M , $N < M$, then $p_N^a(\pm) \geq p_M^a(\pm)$, $\frac{1}{4}^a_N(\pm) \geq \frac{1}{4}^a_M(\pm)$ and $p_N^a(\pm) \geq p_M^a(\pm)$ for every \pm .

2.3.4 Example with uniform distributions

The particular case in which the discount factor is distributed uniformly between a and b , $0 < a < b < 1$, provides clear examples of the previous results.

In the uniform case, taking into consideration that $\bar{\pm} = \frac{a+b}{2}$, I can restate Proposition 1 in the following way:

Proposition 15 If $\pm \gg U(a; b)$, the function $\frac{1}{4}^a(\pm)$ depends on a , b and N in the following way:

- 1) if $b \geq 2 \left(a \frac{N+1}{N-1} \right)$, $\frac{1}{4}^a(\pm) = \frac{1}{4}^m$;
- 2) if $\frac{2(N-1)}{N} \left(a \frac{N+1}{N-1} \right) < b < 2 \left(a \frac{N+1}{N-1} \right)$, $\frac{1}{4}^a(\pm) = \frac{1}{4}^m$ for $\pm \geq \underline{b}$ and $\frac{1}{4}^a(\pm) = \frac{\pm}{\underline{b}} \frac{1}{4}^m$ for $\pm < \underline{b}$, with $\underline{b} = b \frac{N}{N(b^2 - a^2) - 2(b-a)(N-1)}$;
- 3) if $b < \frac{2(N-1)}{N} \left(a \frac{N+1}{N-1} \right)$, $\frac{1}{4}^a(\pm) = 0$.

Therefore, in the case of uniform distribution of the discount factor, the level of profits for each discount factor depends on the magnitudes of a , b and N . If $b \geq 2 \left(a \frac{N+1}{N-1} \right)$ perfect collusion can be supported for any realization of the discount factor. If $\frac{2(N-1)}{N} \left(a \frac{N+1}{N-1} \right) < b < 2 \left(a \frac{N+1}{N-1} \right)$, perfect collusion can be supported only for high discount factors and only lower levels of profits can be supported for lower discount factors. Finally, if $b < \frac{2(N-1)}{N} \left(a \frac{N+1}{N-1} \right)$ a no collusion can be supported.

Figure 2 shows the different ranges of a and b for the three cases of tacit collusion when $N = 2$.

Since $b > a$, the relevant portion of the figure is above the 45 degree line. That part of the graph shows the ranges of a and b that result in different kinds of tacit

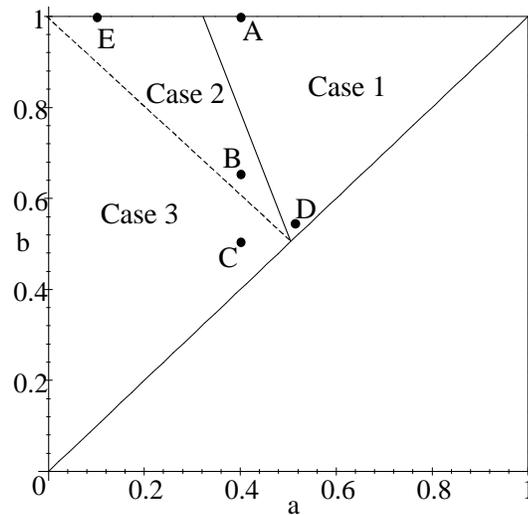


Figure 2: Ranges of tacit collusion

collusion. For example, to the northeast of the solid black line are the combinations of a and b that results in perfect collusion (case 1) when there are two firms in the market. Between the solid and dashed black lines we see the combinations that result in perfect collusion for high discount factors and imperfect collusion for low discount factors (case 2), and below the dashed black line are the combinations that cannot support any collusion (case 3).

Consider the distributions of δ represented in Figure 2 by the points A, B and C (the discount factor is distributed $U(0;4; 1)$, $U(0;4; 0;65)$ and $U(0;4; 0;5)$, respectively). Each of the points falls in a different region and hence will result in a different tacit collusion solution. The distribution denoted by point A results in perfect collusion, the distribution denoted by point B results in perfect collusion for high discount factors and imperfect collusion of lower discount factors and the distribution denoted by C results in no collusion at all. There are two additional things to note from this example. First, profits are (weakly) increasing in the realization of the discount

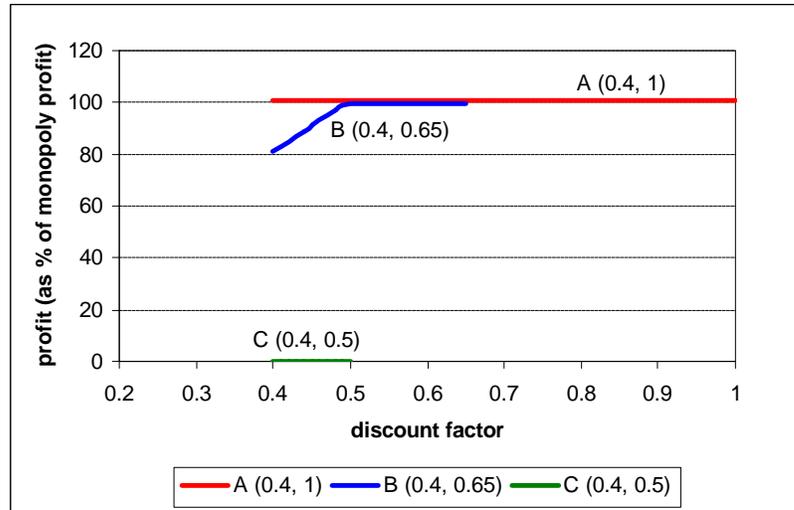


Figure 3: The effect of levels

factor, as Theorem 2 proves. While for A and C the tacit collusion profits do not depend on the realization of the discount factor, for B increases in realization of the discount factor may result in an increase of profits and prices. Second, the “more to the right” the distribution function is, the higher profits and prices are. Figure 3 shows that profits under A are larger than under B or C, as Corollary 5 proves.

Consider now the distribution function denoted by point D in Figure 2, $U(0.5; 0.55)$. This distribution function has the same expected value but a lower volatility than the distribution function denoted by point B. We can see from Figure 2 that if there are only two firms in the market, perfect collusion can be supported at point D, while perfect collusion can only be supported for a range of high discount factors for point B. Figure 4 shows the tacit collusion profit functions for these two cases as a percentage of monopoly profits per firm. Consistent with Corollary 6, Figure 4 shows that a mean preserving spread in the distribution of the discount factor reduces the expected value of profits.

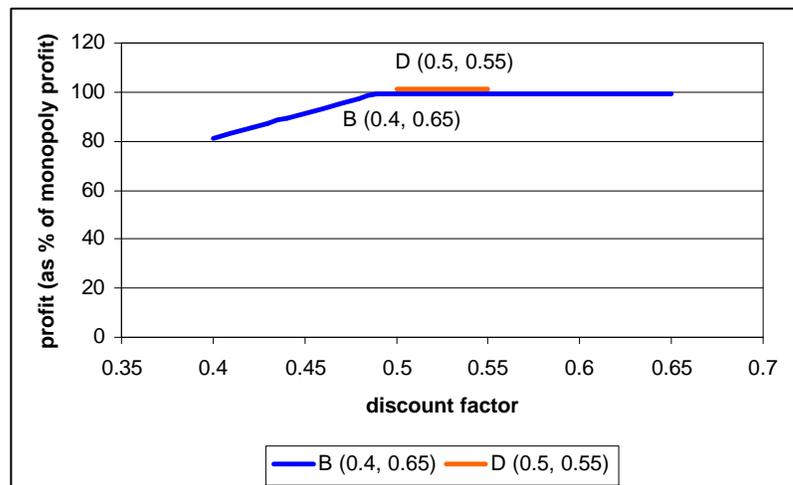


Figure 4: The effect of volatility

To make clear that the volatility effect is not a second order effect consider the distributions denoted by point E in Figure 2, $U(0;1;1)$. This distribution has a higher expected discount factor than the distribution denoted by point D but it also has a higher volatility. Figure 5 shows that the distribution function with the highest expected discount factor and volatility results in lower collusive profits.

Figure 6 shows the limits to the three cases of tacit collusion for $N = 2, 4, 8$ and 16. We see that the greater the number of firms, the smaller the set of distribution functions for which some collusion is possible.

Consider now the distribution of δ represented in Figure 6 by point E: the discount factor is distributed $U(0;0.52;1)$. From Figure 6, we see that perfect collusion can be supported if $N = 2$, while perfect collusion can only be supported for a range of high discount factors if $N = 4$, and cannot be supported at all for $N = 8$. Figure 7 shows the tacit collusion industry profits (as a percentage of industry monopoly profit) for these three cases and, consistent with Theorem 8, shows that the profits decrease

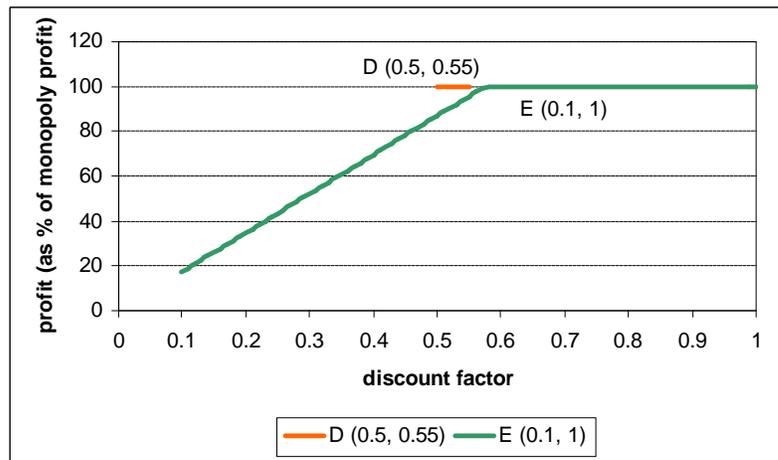


Figure 5: Volatility matters

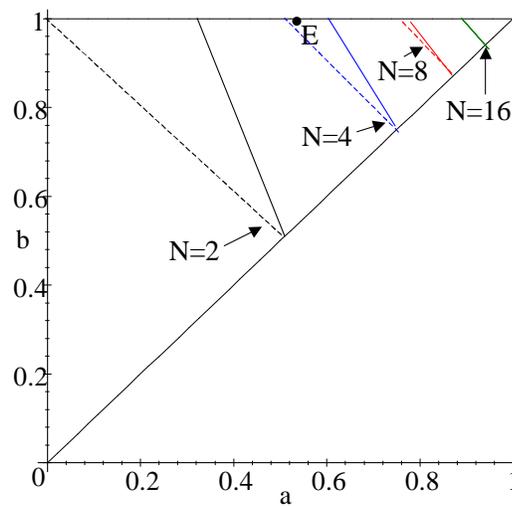


Figure 6: Ranges of tacit collusion for different N

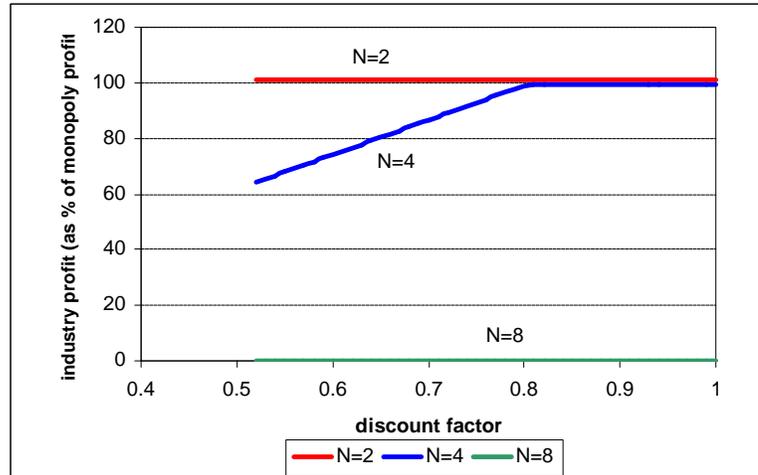


Figure 7: Tacit collusion and the number of firms

with the number of firms.

2.4 Quantity Competition

In this section I show that, under certain assumptions, the three main results that hold under price competition also hold under quantity competition. Namely, first, the higher the discount factor in a period, the higher the collusive prices and profits in that period, second, the higher the probability of high discount factors, the higher the collusive prices and profits, and third, the higher the volatility of the discount factor, the lower the collusive prices and profits that can be supported in equilibrium.

However, to prove this I have to characterize the optimal punishment scheme, which was not necessary under price competition. This is interesting because I show that, while punishment schemes can be extremely complex under quantity competition, the optimal punishment has a simple stick-and-carrot characterization.

I consider the same model of section 2 with one main difference: firms compete on quantities. In addition, and only for the sake of generality, I also assume that firms

have a continuous and differentiable cost function $c(q)$ instead of the linear cost of section 2.

As in section 2, I restrict my attention to symmetric equilibria: all the firms produce at a given period the same quantity q . In this symmetric case, I can write the profits of each firm as $\pi_i(q) = P(Nq)q - c(q)$ and total industry profits as $\Pi(q) = N\pi_i(q)$. I assume that there exists a quantity q^m that maximizes the total industry profits, that is the perfect collusion quantity (q^m would be the N th part of a monopolist optimal production if there are no fixed cost per factory and increasing returns to scale). Denote $\pi_i^m = \pi_i(q^m)$ as the perfect collusion profit per firm.

2.4.1 Optimal tacit collusion with a random discount factor

In the case of quantity competition the Cournot reversion is not necessarily the best available punishment since it may be possible to generate subgame perfect threats that lower the profits below the Cournot level. Therefore, to characterize the optimal tacit collusion solution it is also necessary to define the optimal punishment scheme. In this section I characterize the optimal equilibrium punishment and collusion under certain assumptions. The first assumption is that there exists a symmetric Cournot equilibrium.

Assumption 1: There exists a quantity q^c that is the unique symmetric Cournot equilibrium.

In this equilibrium each firm earns a profit of π_i^c and it can be proven that $\pi_i^m > \pi_i^c$, and $q^m < q^c$.

The second assumption concerns the profits from deviation. In the case of quantity competition, if $N - 1$ firms are each producing a quantity q , the remaining firm can obtain at most a profit of $\pi_i^d(q) = \max_{s \geq 0} [P(s + (N - 1)q)s - c(s)]$ by producing some other quantity. The second simplifying assumption establishes that both $\pi_i^d(q)$ an

$\frac{1}{4}(q)$ are decreasing with the former having a bigger slope than the latter, in absolute terms, for quantities below q^c while the opposite occurs for quantities above q^c .

Assumption 2: For $q \in [q^m; q^c]$, $\frac{d\frac{1}{4}^d}{dq} < \frac{d\frac{1}{4}}{dq} < 0$ and for $q \in (q^c; +\infty)$, $\frac{d\frac{1}{4}}{dq} < \frac{d\frac{1}{4}^d}{dq} < 0$.

These assumptions are valid, for example, in a market with a linear demand function and constant marginal cost. In addition, in the linear case there is a unique quantity that maximizes industry profits (q^m) and a unique and symmetric Cournot equilibrium (q^c). Hence there is no contradiction between the assumptions made in this section.²⁶

As in section 2, the optimal symmetric tacit collusion equilibrium can be characterized by the maximum level of profits per firm that can be supported for each discount factor, which I denote $\frac{1}{4}^\pi(\pm) : [a; b] \rightarrow [q^c; q^m]$, abusing notation from section 2. Since assumption 2 ensures that there is a one to one relationship between profits and quantities produced in the relevant range, once $\frac{1}{4}^\pi(\pm)$ is obtained, the optimal tacit collusion quantities $q^\pi(\pm) : [a; b] \rightarrow [q^m; q^c]$ are also obtained. From the demand function we can obtain the optimal tacit collusion profits $\pi^\pi(\pm) = P(Nq^\pi(\pm))$.²⁷

As in section 2, I use the recursiveness of the problem to write the present value of profits:

$$V(\pm) = \frac{1}{4}(\pm) + \frac{\delta}{1 - \delta} \int_a^b \frac{1}{4}(\pm^0) f(\pm^0) d\pm^0 \quad (7)$$

²⁶In addition, these assumptions, as the assumption presented in the next subsection, could be obtained from assumptions regarding the demand and cost functions. Since those assumptions would be only sufficient ones and would not provide a better intuition I prefer to present conditions regarding $\frac{1}{4}(q)$ and $\frac{1}{4}^d(q)$ that yield the desired results.

²⁷Note that this functions denote equilibrium outcomes and not strategies. The supporting strategies are not explicitly defined due to their lack of peculiarities.

In addition, the feasibility condition can be written as:

$$\frac{1}{4}(\pm) \geq \frac{1}{4}^m \quad (8)$$

The incentive compatibility constraint differs from that in the previous section since neither the short run incentives to deviate nor the future punishments are the same. Under price competition, a firm can capture the whole market by a small price deviation, obtaining $(N - 1)\frac{1}{4}(\pm)$ in profits from deviation. Under quantity competition, the maximum profit from deviation is $\frac{1}{4}^d(q(\frac{1}{4}(\pm))) - \frac{1}{4}(\pm)$, where $q(\frac{1}{4})$ is the quantity that every firm has to produce to get a per firm profit of $\frac{1}{4}$. In addition, the possible punishment from deviation may not be the same as in price competition. In price competition reverting to a situation of zero profits is a credible threat, since that is the Bertrand equilibrium. Instead, under quantity competition a punishment of zero profits forever may not be credible. What is credible depends on the biggest credible threat. This threat would consist of punishing the deviator with the lowest equilibrium discounted payoff, denoted by $\underline{V}(\pm)$, while rewarding compliance with the equilibrium with the highest equilibrium discounted payoff, denoted by $\overline{V}(\pm)$, if tomorrow's discount factor is \pm . Assume for now that the extreme discounted equilibrium payoff functions $\overline{V}(\pm)$ and $\underline{V}(\pm)$ exist, as it is proven later, and define their expected values as $E\overline{V}$ and $E\underline{V}$, respectively. Therefore, for $\frac{1}{4}(\pm)$ to be incentive compatible, it must be the case that no player has incentives to deviate if conforming is rewarded with the highest possible expected continuation payoff $E\overline{V}$ and deviating is punished with the lowest possible expected continuation payoff $E\underline{V}$:

$$\frac{1}{4}^d(q(\frac{1}{4}(\pm))) - \frac{1}{4}(\pm) \geq \pm E\overline{V} - E\underline{V} \quad (9)$$

For simplicity, write the left hand side of equation (9) as $\mathcal{C}(\frac{1}{4}(\pm))$ and denote $E\overline{V} - E\underline{V}$ on the right hand side of the equation as B . As such, for a given B , the

incentive compatibility constraint can be written as

$$\omega(\frac{1}{4}(\pm)) \leq \pm B \leq \frac{1}{4}(\pm) \quad (10)$$

Note that $\omega(\frac{1}{4}^c) = 0$ and that $\omega(\frac{1}{4})$ increases as $\frac{1}{4}$ separates from $\frac{1}{4}^c$. Then, for a given amount of threat $\pm B$, there is a highest and lowest amount of profit that can be supported. Next I characterize the incentive compatible upper bound to profits, and its interaction with the feasibility constraint, and then characterize the incentive compatible lower bound to profits.

Lemma 16 Under Assumptions 1 and 2, for a given B , the incentive compatible upper bound to profits is not binding for any \pm if $aB > \omega(\frac{1}{4}^m)$. If instead $aB \leq \omega(\frac{1}{4}^m)$, there exists a number $\underline{b}(B) \in [a; b]$ such that the upper bound can be written as $\frac{1}{4}(\pm) = \omega_+^{-1}(\pm B)$ for $\pm \leq \underline{b}(B)$, where $\omega_+^{-1}(\pm B)$ is the inverse of $\omega(\frac{1}{4})$ if we restrict its domain to $[\frac{1}{4}^c; \frac{1}{4}^m]$, and it is not binding for $\pm > \underline{b}(B)$. In addition, the incentive compatible upper bound is increasing in \pm for $\pm \leq \underline{b}(B)$, and $\underline{b}(B)$ and $\omega_+^{-1}(\pm B)$ are continuous.

Proof. In Appendix. ■

Therefore, for low discount factors the maximum level of profits that can be supported is bounded by the incentive compatible upper bound, while for high values it is bounded by the feasibility constraint. Combining both we have the IC⁺-F constraint:

$$\frac{1}{4}(\pm) \begin{cases} \leq \omega_+^{-1}(\pm B) & \text{if } aB \leq \omega(\frac{1}{4}^m) \text{ and } \pm \leq \underline{b}(B) \\ \leq \frac{1}{4}^m & \text{otherwise} \end{cases} \quad (11)$$

In the optimal symmetric tacit collusion equilibrium, firms will choose profits as large as possible given the incentive compatible upper bound and the feasibility constraint. In addition, given that conforming with the equilibrium strategy must be rewarded with the highest equilibrium payoff, the highest equilibrium discounted

payoff $\bar{V}(\pm)$ has a simple relationship with the optimal tacit collusion solution. If $\frac{1}{4}^a(\pm)$ is the optimal tacit collusion profit function, then $\bar{V}(\pm) = \frac{1}{4}^a(\pm) + \frac{\pm}{1 \pm a} \int_a^1 \frac{1}{4}^a(\pm^0) f(\pm^0) d\pm^0$ and its expected value is $E\bar{V} = \frac{E\frac{1}{4}^a}{1 \pm a}$. Therefore, given the lowest expected equilibrium payoff $E\underline{V}$, the optimal tacit collusion solution is subject to the following equation:

$$\frac{1}{4}^a(\pm) = \begin{cases} \frac{1}{4}^m & \text{otherwise} \\ \frac{1}{4}^a(\pm) & \text{if } a \pm \frac{E\frac{1}{4}^a}{1 \pm a} \geq E\underline{V} \end{cases} \quad \text{for } \frac{1}{4}^a(\pm) \geq \frac{1}{4}^m \text{ and } \pm \geq \frac{1}{4}^a(\pm) \quad (12)$$

The following lemma characterizes the incentive compatible lower bound to profits.

Lemma 17 Under Assumptions 1 and 2, for a given B , the incentive compatible lower bound to profits can be written as $\frac{1}{4}(\pm) = \frac{1}{4}^i(\pm; B)$, where $\frac{1}{4}^i(\pm; B)$ is the inverse of $\frac{1}{4}$ if we restrict its domain to $(\frac{1}{4}^i; 1; \frac{1}{4}^c]$. In addition, the incentive compatible lower bound is decreasing in \pm , and $\frac{1}{4}^i(\pm; B)$ is continuous.

Proof. In Appendix. ■

Having characterized the incentive compatible lower bound to profits, I must still characterize the lower discounted continuation payoff $\underline{V}(\pm)$. I show that the optimal punishment scheme, which yields $\underline{V}(\pm)$, has a simple stick-and-carrot characterization (the punishment takes only one period and is as big as possible in equilibrium), extending the results of Abreu [1] from the fixed discount factor case.

Lemma 18 Given $E\frac{1}{4}^a$ and $E\underline{V}$, the lowest equilibrium payoff function is $\underline{V}(\pm) = \frac{1}{4}^i(\pm; E\underline{V}) + \frac{\pm}{1 \pm a} E\frac{1}{4}^a$.

Proof. Consider any punishment scheme consisting of a profit of $\frac{1}{4}(\pm)$ in the first period and an expected continuation payoff of $E\underline{V}$. Define the present value of the game in that case as $\underline{V}(\pm) = \frac{1}{4}(\pm) + \pm E\underline{V}$. For this punishment scheme to be credible it must be the case that $\underline{V}(\pm) \geq \frac{1}{4}^d(q(\frac{1}{4}(\pm))) + \pm E\underline{V}$. Choose now the first

payoff of a two phase punishment $\frac{1}{4}^0(\pm)$ so that $\frac{1}{4}^0(\pm) + \frac{\pm}{1_i \pm} E \frac{1}{4}^a = \underline{V}(\pm)$. Given that $\frac{E \frac{1}{4}^a}{1_i \pm} \geq E \underline{V}$, $\frac{1}{4}^0(\pm) \leq \underline{V}(\pm)$ and by $\frac{1}{4}^d(q(\cdot))$ being increasing, $\underline{V}(\pm) \geq \frac{1}{4}^d(q(\frac{1}{4}^0(\pm))) + \pm E \underline{V}$ and the two phase punishment is credible. Therefore any equilibrium punishment can be matched with a two phase punishment that yields the best continuation payoff in the second phase. Then, choosing the lowest equilibrium present payoff, I obtain the lowest equilibrium discounted payoff for a given discount factor, and $\underline{V}(\pm) = \frac{1}{1_i \pm} \left(\frac{E \frac{1}{4}^a}{1_i \pm} + E \underline{V} \right) + \frac{\pm}{1_i \pm} E \frac{1}{4}^a$. ■

Therefore, given the optimal tacit collusion solution, the lowest possible continuation payoffs are subject to the following equation:

$$\underline{V}(\pm) = \frac{1}{1_i \pm} \left(\frac{E \frac{1}{4}^a}{1_i \pm} + E \underline{V} \right) + \frac{\pm}{1_i \pm} E \frac{1}{4}^a \quad (13)$$

The solution to the problem of finding the optimal tacit collusion profits and the optimal punishment that support that collusion consists of finding the functions $\frac{1}{4}^a(\pm)$ and $\underline{V}(\pm)$ that solve equations (12) and (13) simultaneously and choosing the solution with the highest expected profit $E \frac{1}{4}^a$. The next proposition shows that this problem has a unique solution.

Proposition 19 Under Assumptions 1 and 2, $\frac{1}{4}^a(\pm)$ and $\underline{V}(\pm)$ exist. In addition $\frac{1}{4}^a(\pm)$ is unique.

Proof. Taking the expected value over (12), for any possible solution $\frac{1}{4}(\pm)$ it has to hold that:

$$E \frac{1}{4} = \frac{1}{1_i \pm} \left(\frac{E \frac{1}{4}^a}{1_i \pm} + E \underline{V} \right) + \frac{\pm}{1_i \pm} E \frac{1}{4}^a \quad (14)$$

In the same way, taking the expected value over (13), for any possible solution

$\underline{V}(\pm)$ it has to hold that:

$$E\underline{V} = \frac{z^b}{1 - \delta} \mu \pm \frac{E\underline{V}^a}{1 - \delta} + \frac{f(\pm)d \pm}{1 - \delta} E\underline{V}^a \quad (15)$$

Note that there is a one to one relationship between the profit functions that satisfy equation (12) and the expected values that satisfy equation (14). That is, if $\underline{V}^a(\pm)$ satisfies equation (12), then $E\underline{V}^a$ must satisfy equation (14), and if the value $E\underline{V}^a$ satisfies equation (14), $\underline{V}^a(\pm)$ satisfies equation (12) with $E\underline{V}^a$ in the right hand side. The same is true for equations (13) and (15). Therefore, we can find $\underline{V}^a(\pm)$ and $\underline{V}(\pm)$ by choosing the solution to equations (14) and (15) with the highest $E\underline{V}^a$. Note that $E\underline{V}^a = \underline{V}^c$ and $E\underline{V} = \frac{\underline{V}^c}{1 - \delta}$ solve the pair of equations and, hence, there is at least one solution.

$$\text{Let } H(r; s) = \frac{z^b}{1 - \delta} \mu \pm \frac{r}{1 - \delta} \int f(\pm) d\pm + \frac{F}{1 - \delta} \frac{r}{1 - \delta} \int \underline{V}^m \pm r$$

Since $\underline{V}^a(\pm)$, $\int f(\pm) d\pm$, and $F(\cdot)$ are continuous, $H(r; s)$ is also a continuous function. Then, the set of numbers that make $H(r; s) = 0$ is closed, given that the inverse images of closed sets are closed for continuous functions. In addition it must be bounded since $r \in [\underline{V}^c; \underline{V}^m]$ and $s \in [0; \frac{\underline{V}^c}{1 - \delta}]$. Therefore, the set of solutions is non-empty, closed and bounded. Then, among the solutions there exists one with the highest r that gives $(E\underline{V}^a; E\underline{V})$. Plugging this into equations (12) and (13) we obtain $\underline{V}^a(\pm)$ and $\underline{V}(\pm)$. Uniqueness is clear from the fact that there is a one to one relationship between $E\underline{V}^a$ and $\underline{V}^a(\pm)$. ■

Optimal tacit collusion must fall in one of the following three cases, depending on which restriction is binding. First, it may be that only the feasibility constraint binds for every discount factor. In this case, the value of the future monopoly profits outweighs the profits from deviation, and perfect collusion is an equilibrium for any discount factor. Second, it may be possible that the incentive compatibility constraint

binds for low discount factors while the feasibility constraint binds for high discount factors. Third, it may be possible that the incentive compatible upper bound to profits binds for every value of the discount factor. While in the first case changes in the discount factor do not affect profit and prices, in the last two cases, an increase in the discount factor results in an increase in collusive profit and prices. The reason for this is that a higher discount factor results in a higher threat of punishment, so that higher profits can be achieved without firms having incentives to deviate, and the next Theorem follows.

Theorem 20 Under Assumption 1 and 2, $\frac{d\pi^c(\delta)}{d\delta} \geq 0$ and $\frac{dp^c(\delta)}{d\delta} \geq 0$.

As in section 2, the equilibrium profits and prices are increasing in the discount factor.

2.4.2 The effects of changes in $f(\delta)$

In section 2, the comparative static results with respect to the distribution function of the discount factor depend on the optimal tacit collusion profit function being increasing and concave. If that is the case, shifts to the left of the distribution function or increments in volatility reduce the expectation of future profits and result in lower equilibrium profits and prices. Because under quantity competition the level of punishment is not independent of the discount factor, it is not enough to look at the shape of the optimal tacit collusion profits to obtain a comparative static result with respect the distribution function. What is important is the shape of the threat of future punishments: $\bar{V}(\delta)$; $\underline{V}(\delta)$.

The stick-and-carrot property of the optimal punishment implies that streams of payoffs leading to the highest and lowest discounted equilibrium payoff differ only in the first period. As a result, the threat of future punishment is simply the maximum

difference in payoffs that can be supported in equilibrium in one period. Since I have already proved that the upper bound to profits is increasing and the lower bound to profits is decreasing, it only remains to be shown that the upper bound is concave while the lower bound is convex. The following assumption is a sufficient condition for that.

Assumption 3: For $q \geq [q^m; + 1)$, $\frac{d^2 \bar{v}}{dq^2} \leq 0 \leq \frac{d^2 \underline{v}}{dq^2}$.

As Assumptions 1 and 2, this assumption is valid in a market with a linear demand function and constant marginal cost. Hence, there is no contradiction among the assumptions made in this section.

Lemma 21 Under Assumptions 1-3, $\bar{V}(\pm)$; $\underline{V}(\pm)$ is increasing and concave.

Proof. In Appendix. ■

From the future threat being increasing and concave in the next period discount factor, the desired comparative static result with respect to the distribution function of the discount factor follows.

Theorem 22 Consider two cumulative distribution functions, F and G , such that F second-order stochastically dominates G , then $\bar{v}_F^a(\pm) \geq \bar{v}_G^a(\pm)$ and $p_F^a(\pm) \geq p_G^a(\pm)$ for every \pm . In addition, $E_F \bar{v}_F^a \geq E_G \bar{v}_G^a$.

Proof. Let $\bar{V}_j(\pm)$ and $\underline{V}_j(\pm)$, $j = F; G$, be the highest and lowest equilibrium discounted payoff under j . First, I show that $E_F \bar{V}_F(\pm) \geq E_G \bar{V}_G(\pm)$; $E_F \underline{V}_F(\pm) \geq E_G \underline{V}_G(\pm)$. Suppose not, then $E_G \bar{V}_G(\pm) \geq E_F \bar{V}_F(\pm)$; $E_G \underline{V}_G(\pm) \geq E_F \underline{V}_F(\pm)$. Since F second-order stochastically dominates G and $\bar{V}_G(\pm)$; $\underline{V}_G(\pm)$ is increasing and concave by Lemma 16, $E_F \bar{V}_G(\pm) \geq E_G \bar{V}_G(\pm)$; $E_F \underline{V}_G(\pm) \geq E_G \underline{V}_G(\pm)$. But then, the strategies that yield $\bar{v}_G^a(\pm)$ and $\underline{v}_G^a(\pm)$ under G do not violate the

incentive and feasibility constraints under F , by ∂_{β}^i being increasing and ∂_{β}^{i-1} being decreasing on $\beta \in \left(\frac{E_i}{1+\beta}, E_i\right)$. Therefore $\beta_F^i(\beta)$ is not an optimal tacit collusion solution under F , which is a contradiction.

Second, given that $E_F \frac{\partial \beta_F^i(\beta)}{\partial \beta} > E_G \frac{\partial \beta_G^i(\beta)}{\partial \beta}$ and ∂_{β}^i is increasing, $\beta_F^i(\beta) > \beta_G^i(\beta)$ for every β . The last two results follow from the positive relationship between profits and prices and the relationship between F and G , respectively. ■

The intuition behind this results is simple. Given that the threat of future punishment is increasing and concave in the discount factor, both increases in the probability of low discount factors and increases in its volatility reduce the expected value of the punishment and result in a reduction of collusive profits and prices.

2.4.3 The effects of changes in the number of firms

In section 2 I showed that under price competition, increases in the number of firms increase the incentives to deviate, decreasing equilibrium profits. This result may not be valid when firms compete on quantities since not only do the incentives to deviate change with the number of firms, but so may the threat of future punishment. In fact, the higher the number of firms the easier it is to support low profits -a consequence of which is that industry Cournot profits fall with the number of firms- and the higher the threat of punishment for deviation. Therefore, while under price competition it is enough to study the effect of the number of firms on the incentives to deviate, this is not sufficient under quantity competition.

While more work is needed to characterize general conditions under which increases in the number of firms decrease equilibrium profits and prices, the next subsection presents an example of such a situation.²⁸

²⁸To my knowledge this issue also remains to be solved for the case of a fixed discount factor.

2.4.4 Example with uniform distributions

I study next the case in which the discount factor is distributed uniformly between a and b , $0 < a < b < 1$, and the inverse demand function -net of a constant marginal cost- is $P = 12 - Q$, and I provide clear examples of the previous results.

Three different types of optimal tacit collusion exist. If a and b are high, relative to the number of firms, perfect collusion can be supported for any realization of the discount factor. If a and b are low, relative to the number of firms, perfect collusion cannot be supported for any realization of the discount factor, but in contrast to what happens under price competition, some collusion can still be supported. If a and b fall in a middle ground, perfect collusion can be supported only for high discount factors and only lower levels of profits can be supported for low discount factors.

Figure 8 shows the different ranges of a and b for the three cases of tacit collusion, for $N = 2$ and $N = 16$.

Since $b > a$, the relevant portion of the figure is above the 45 degree line. That part of the graph shows the ranges of a and b that result in different kinds of tacit collusion. For example, to the northeast of the solid black line are the combinations of a and b that result in perfect collusion when there are two firms in the market. Between the solid and dashed black lines are the combinations that result in perfect collusion for high discount factors and imperfect collusion for low discount factors, and below the dashed black line are the combinations that cannot support perfect collusion. For example consider the distributions depicted by points A, B and C.

The closest related paper is Brock and Scheinkman [12] which studies the effect of the number of firms on tacit collusion for a fixed discount factor, price competition and an exogenous capacity per firm. They find that changes in the number of firms have a non-monotone effect on optimal collusive prices. Note that the capacity is exogenous and the link to quantity competition from Kreps and Scheinkman [38] does not apply.

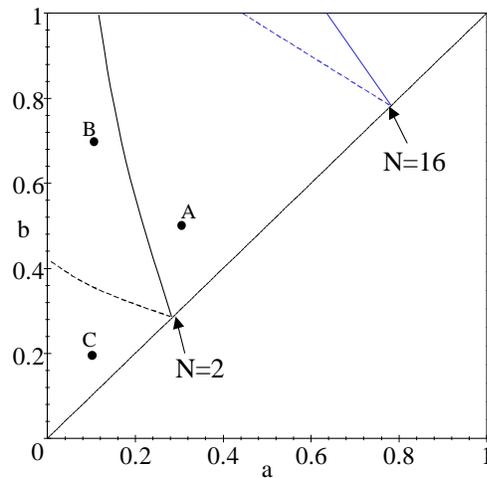


Figure 8: Ranges of tacit collusion

While A results in perfect collusion, B can only support perfect collusion for high discount factors and lower profits for low discount factors. Finally, C cannot support perfect collusion for any discount factor but can still support some collusion. See Figure 9.

From Figure 9 it is clear that profits are increasing in the discount factor. The comparison between the optimal tacit collusion profits for points A and C is an example of the result that the higher the probability of high discount factor, the higher collusive profits and prices. The comparison between the collusive profits for points A and B is an example of the result that the higher the volatility of the discount factor, the lower profits and prices.

One can see the limits to the three types of tacit collusion for $N = 16$ in Figure 8. For the distribution function depicted by point A, perfect collusion can be supported if $N = 2$, but perfect collusion cannot be supported at all -but lower levels of collusion can- for $N = 16$, therefore, profits must be lower in the latter case.

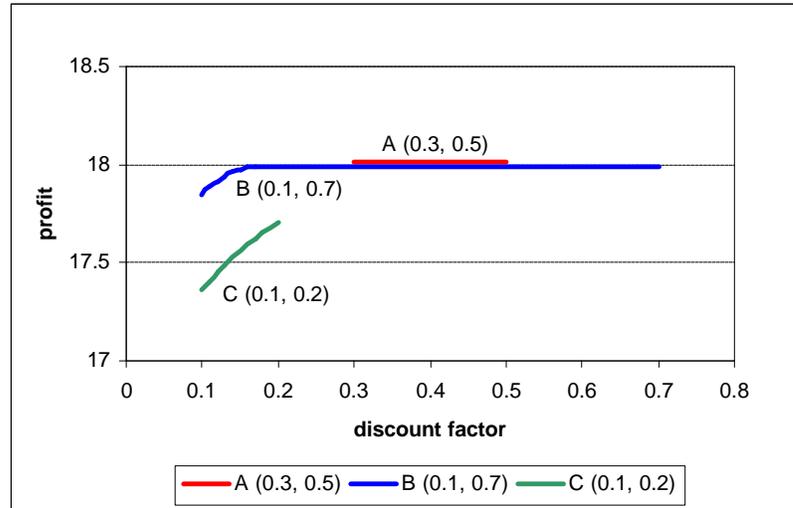


Figure 9: Tacit collusion profits with $N=2$

2.5 Extensions

In this section I analyze the restrictiveness of the assumption of symmetric equilibria and study some extensions. As an extension, I modify the assumption of independently distributed discount factor in two ways. I consider first deterministic discount factor cycles and show that increasing discount factors make easier to support collusion. Second, I consider the case in which the distribution of tomorrow's discount factor depends on today's value and show that an increase in the discount factor may result in a decrease in equilibrium prices and profits (since the increase in the discount factor may lead to an increase in its future volatility). Finally, I study the validity of the three main results of this paper for general repeated games.

2.5.1 Asymmetric equilibrium prices

In this paper I only consider symmetric equilibrium collusive prices and quantities. This assumption may not be that restrictive given that joint overall profits to firms

are generally higher when all the firms charge the same price or produce the same quantity in equilibrium. The existence of asymmetries in firms' equilibrium collusive behavior can only reduce prices and total industry profits, since it is the incentive compatibility constraint of the less favored firm that binds.

In addition, in the case of price competition, this asymmetry effect is strengthened by an intrinsic discontinuity of the Bertrand model. With price competition, if firms offer different prices there will be a group of firms that will not provide goods to the market and will get zero profits. These firms will have large incentives to deviate. Thus, under price competition, the impact of even small price asymmetries on the incentive compatibility constraints can be significant.

Therefore, there is a compelling reason to restrict ourselves to symmetric equilibrium behavior in this paper: it is the equilibria that maximizes the industry's total profit. Introducing asymmetries would reduce the industry's profits by increasing the incentives to deviate for those less favored firms that get a small share of the market.²⁹

2.5.2 Deterministic discount factor cycles

In this section I consider deterministic discount factor cycles and show that higher collusive prices and profits can be supported when the discount factor is increasing. The reason is simple: the higher the future discount factors, the higher future collusive profits and the larger the threat of punishment. Hence, the higher the future discount factors, the higher present collusive prices and profits, as the next example shows.

Example 23 An increasing discount factor facilitates collusion: For the discount

²⁹Nevertheless, under quantity competition, it may be useful to allow for asymmetric behavior on the equilibrium path. The optimal punishment schemes characterized in this paper may only be optimal under the restriction of symmetry on the equilibrium path. It is possible that asymmetries during the punishment stage generate bigger punishments and higher symmetric collusion, as it is the case under a fixed discount factor (Abreu [1]).

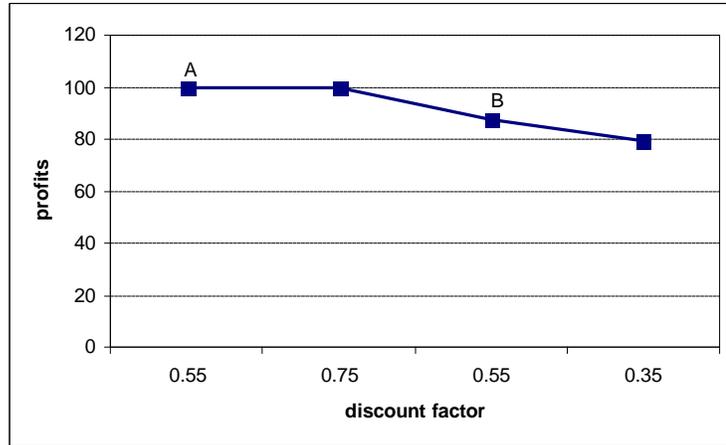


Figure 10: Profits under cyclical discount factor

factor cycle $\delta: 0.55; 0.75; 0.55; 0.35$, price competition and two firms in the market, the optimal tacit collusion solution is represented in Figure 10 as a percentage of monopoly profits. We can see that for $\delta = 0.55$ the optimal tacit collusion is higher when the discount factor is increasing (point A) than when it is decreasing (point B). Therefore, it is easier to support collusion for a given discount factor when the discount factor is increasing.

Under cyclical discount factor fluctuations, both high discount factors today and in the future make it easy to support collusion given that both increase the threat of future price wars. In contrast, Haltiwanger and Harrington [28] found that under cyclical demand fluctuations, a high demand today makes it difficult to support collusion since it offers high incentives to deviate, while high demand in future periods makes it easy to collude today because it increases the threat of future price wars.

2.5.3 Correlated discount factor and the volatility effect

Given that both a high discount factor today and in the future make it easy to support collusion, allowing for the more realistic case of positively correlated discount factors will not affect the main results. But these results may be modified if changes in today's discount factor affect its future volatility. In this section I present an extension to the basic model to illustrate that an increase in the discount factor does not necessarily lead to higher collusive prices and profits if the increase in the discount factor also raises the volatility of future discount factors. When the value of the present discount factor affects the distribution of the future discount factor, the solution to the optimization problem cannot be found easily. Nevertheless, under price competition and a discrete distribution of the discount factor, the problem can be solved as a linear programming problem (see appendix)³⁰.

Example 24 Consider the case in which the discount factor can take only three values ($\frac{1}{4}$, $\frac{1}{2}$ and $\frac{3}{4}$), there are two firms and the monopoly profit per firm is 18. The distribution function of the discount factor depends on the past discount factor in the following way:

$p(\pi_t \pi_{t-1})$	π_t			
		$\frac{1}{4}$	$\frac{1}{2}$	$\frac{3}{4}$
π_{t-1}	$\frac{1}{4}$	3/5	1/5	1/5
	$\frac{1}{2}$	0	1	0
	$\frac{3}{4}$	12/25	0	13/25

³⁰Note that this example is not equivalent to an extension to three states of Bagwell and Staiger [8] model of correlated demand shocks. In that model changes in present demand growth only affected the future through changes in the expectation of future growth rates. In this example, changes in the discount factor affect both expectations and the present valuation of future profits making the analysis more complicate.

Solving the linear programming problem we found that the optimal symmetric tacit collusion equilibrium yields profits equal to 4:8, 18 and 15:8 for the discount factor being $\frac{1}{4}$, $\frac{1}{2}$ and $\frac{3}{4}$, respectively.

This example shows that an increase in the discount factor, while increasing the expectation of the future discount factor, may still result in a reduction of profits and prices. The reason is that not only does the expectation of future discount factors matter, but so does its volatility. In this case, given that future discount factors have higher volatility when $\pm = \frac{3}{4}$ than when $\pm = \frac{1}{2}$, equilibrium profits are lower under the former than under the latter.

2.5.4 General normal form games

In this section I study whether the main results of this paper can be extended to general infinitely repeated games with discount factor fluctuations. I consider an infinitely repeated simultaneous move game in which the discount factor is independently and identically distributed. As in the rest of this paper, players observe the realization of the discount factor before choosing an action.

In this more general environment the following results regarding discount factor levels can be shown: first, the higher the realization of the discount factor the larger the set of equilibrium outcomes, and second, the higher the probability of high discount factors the larger the set of equilibrium outcomes.³¹ In contrast, it is not true that an increase in the volatility of the discount factor always results in a decrease

³¹Given that the discount factor is i.i.d., before discounting it, the threat of future punishment is independent of the present realization of the discount factor. Hence, the higher the discount factor the more important that threat is and the bigger the set of equilibrium outcomes. Given this first result, shifts of the distribution function to the right result in increments in the threat of punishment and an increase in the set of outcomes for every realization of the discount factor.

in the set of equilibrium outcomes. The next example shows that an increase in the volatility of the discount factor may increase the set of equilibrium outcomes for some discount factors.

Example 25 An increase in volatility of the discount factor may increase the set of equilibrium outcomes:

Consider the following stage game:

		Column		
		a	b	c
Row	A	5, 5	0, 0	-2, 10
	B	0, 0	4, 4	-2, 5
	C	10, -2	5, -2	0, 0

In this stage game there is a unique Nash equilibrium (C,c), which is Pareto dominated by either (A,a) or (B,b). The infinite repetition of the stage game opens the possibility that these outcomes can be supported in equilibrium. Note that (A,a) yields a higher payoff than (B,b) but offers higher incentives to deviate. With discount factor fluctuations, we should expect that in the optimal symmetric equilibrium, (A,a) is played for realization of the discount factors that are high enough, (B,b) for lower ones and, finally, (C,c) when the realization of discount factor is too low to be able to support any cooperation by the threat of future punishment. In fact, if the discount factor is distributed $U(0,1)$, the outcomes of the optimal symmetric equilibrium are the following³²: (A,a) if $\delta \geq 0.65$, (B,b) if $0.13 \leq \delta < 0.65$, (C,c) if $\delta < 0.13$, as shown in Figure 11. In this case the expected utility equals 3.82.

Consider now a modification of the distribution of the discount factor. From the original $U(0,1)$ distribution take the mass of the segment $[0.45; 0.55]$ and add it to

³²The problem consists of finding the minimum discount factors for which (A,a) and (B,b) can be played in a symmetric equilibrium.

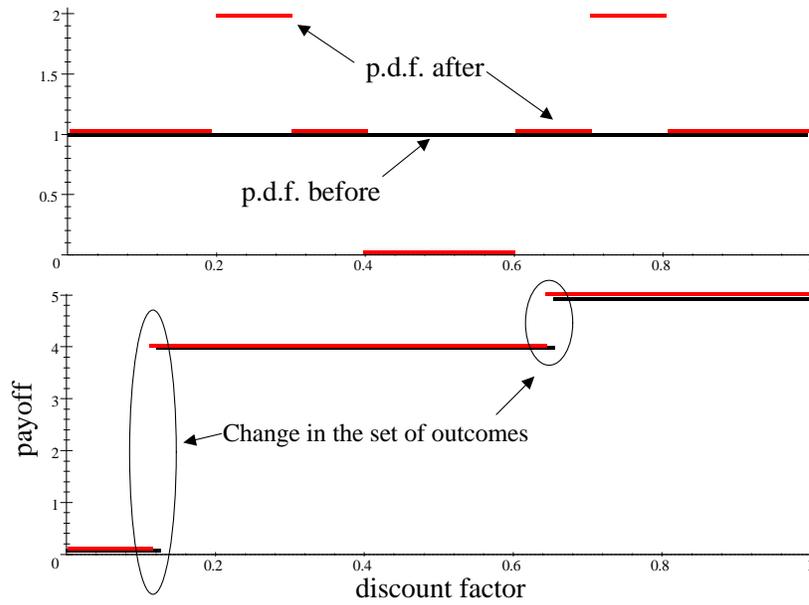


Figure 11: Volatility and general games

the area between $[0:25; 0:3]$ and $[0:7; 0:75]$. This modification adds volatility to the discount factor but at the same time adds weight to the discount factors that yield a high payoff. Hence, the change in the distribution function increases the equilibrium expected utility and relaxes the incentive compatibility constraint, increasing the set of discount factors for which (A,a) can be supported in equilibrium. In fact, under the modified distribution function, the outcomes of the equilibrium that maximizes the players expected utility is the following: (A,a) if $\delta \geq 0:63$, (B,b) if $0:126 \leq \delta < 0:63$, (C,c) if $\delta < 0:126$. In this case the expected utility equals 3:97.

This example shows that a mean preserving spread of the discount factor may result in an increase on the expected payoff of the players and an expansion of the set of possible outcomes for some discount factors.

For repeated games in general it is no longer true that increases in the volatility

of the discount factor reduce the set of equilibrium outcomes. The reason is that the maximum level of utility, as a function of the discount factor, is not necessarily a concave function (nor is the minimum level of utility necessarily convex). As result, increases in volatility may increase the expected equilibrium utility, increasing the threat of future punishment. This, in turn, increases the set of outcomes that can be supported for each discount factor.

2.6 Conclusions

To my knowledge, this paper represents the first effort to examine the effect of discount factor fluctuations in repeated games. In a repeated oligopoly, I characterized the optimal symmetric collusion and found that collusive prices and profits increase with both present and future discount factor levels and decrease with discount factor volatility. These results stress the importance that discount factor levels have on repeated games and introduce a new element to the literature: the volatility of the discount factor.

This work has several important implications for future study. While most of the existing empirical literature on collusive pricing has largely ignored the role of the interest rate, this paper suggests that both the level and the volatility of the interest rate are important determinants of collusive pricing. Thus, to be complete, future empirical work should consider these forces.

This paper also has implications for the study of aggregate fluctuations. I show that any change in policy, preferences or technology may have an impact on the aggregate level of activity through changes in collusive behavior, not only by affecting the real interest rate, but also by affecting its volatility.

Finally, it would be interesting to study extensions of this work to general repeated games. While I show here that volatility reduces the scope for cooperation in repeated

oligopolies, I also show that this is not necessarily true for general repeated games. Determining conditions under which higher volatility reduces the set of equilibrium outcomes for general repeated games remains for future work.

3 Cooperation under the shadow of the future: experimental evidence on infinitely repeated games

3.1 Introduction

Game theorists have long recognized that repeated playing and the possibility of future interaction may modify current behavior. The possibility of future interaction may enable players to establish punishment and reward structures to prevent or limit opportunistic behavior. In a complete information game it is necessary that players do not know when the game ends: each period players must believe that there may be future interaction with positive probability. The higher this probability, the bigger “the shadow of the future” is and the easier it is to reduce opportunistic behavior in theory.

While there has been a large number of theoretical papers on this subject, the empirical and experimental evidence is scarce and in most cases inconclusive or characterized by methodological problems. Given this scarcity of conclusive empirical results, we run a series of experiments to study whether the possibility of future interaction modifies players’ behavior, allowing them to prevent opportunistic actions.

Infinitely repeated prisoner’s dilemma games are simulated in the experiment by having a random continuation rule. The experimental design represents an improvement over the existing literature by including sessions with finite repeated games as controls and a large number of players per session (which allows for learning without contagious effects).

We found strong evidence that the higher the probability of continuation, the higher the levels of cooperation. While in one shot prisoner’s dilemma games there is only 9% of cooperation, for a probability of continuation of $3/4$ there is 38% of cooperation.

The effect of the shadow of future on the levels of cooperation is greater than previous studies have shown and points out that self-enforcing reward and punishment schemes that eliminate opportunistic behavior are important not only in theory.

In addition, the results from the infinitely repeated games are compared with the results from finitely repeated games to test whether cooperation depends on “the shadow of the future,” as theory predicts, or merely on the length of the games. The lengths of the finitely repeated games were chosen to coincide with the expected lengths of the infinitely repeated ones. In the finitely repeated games the levels of cooperation are significantly lower than in the infinitely repeated ones. For example, in repeated games with a finite horizon of 4 rounds the probability of continuation is of 21% against the 38% of cooperation in an infinitely repeated game with a probability of continuation of $3/4$.

Finally, to study how close the behavior of the subject is to the theoretical predictions, we use the fact that different prisoner’s dilemma payoff matrices result in different sets of equilibrium outcomes. We used two different payoff matrices in the experiment with the peculiarity that, for a probability of continuation of $1/2$, one of them can have both players cooperating in equilibrium while the other can not. Consistently with this, we found that the percentage of outcomes in which both subject cooperate is almost 19% when it is an equilibrium, while it is less than 3% when it is not. Then, these experimental results show that behavior closely, but not perfectly, follows the theoretical predictions that are dependent on the payoff details of the stage game, providing further support to the theory of repeated games.

The experimental evidence presented here show that the shadow of the future matters, it significantly reduces opportunistic behavior and it does it in a way that closely follows the theoretical predictions.

The next section summarizes previous experimental research on the topic and

discusses some of its shortcomings. Section 3 describes the experimental design and section 4 describes the theoretical predictions. Section 5 presents the results of the experiment and the last section concludes.

3.2 Previous Experiments

While there exist extensive experimental literature on infinitely repeated games, the experimental literature on ininitely repeated games is limited and presents some methodological shortcomings.

Previous experiments on ininitely repeated games are of two types: 1) experiments with a random continuation rule known to the subjects and 2) experiments with a finite number of repetitions known to the experimenter but unknown to the subjects. In the first type, subjects knew that there was always a positive probability of continuation. In the second type, games were not infinitely repeated since there was a final round, but this round was unknown for the players. Therefore, in each round the subjects may have assigned a positive probability of continuation.

Experiments that fall into the first category are those of Roth and Murnighan [50] and Murnighan and Roth [40]. These two papers present results of experiments with ininitely repeated prisoner's dilemma for different continuation probabilities. Roth and Murnighan [50] find that the higher the probability of continuation, the higher the number of players that cooperated in the first period of the game, see Table 1. Murnighan and Roth [40] present results for experiments with twelve different variations of prisoner's dilemma. Considering the results of the twelve variations together, one can see that, in contrast to the results of Roth and Murnighan [50], higher probabilities of continuation did not result in more cooperation in the first round, see Table 1.

In addition to presenting contradictory evidence (and offering little hope that

Table 1: Previous studies (% of cooperation in ...rst round)

	Probability of continuation		
	0.105	0.5	0.895
Roth and Murnighan [50] ^a	19	29.75	36.36
Murnighan and Roth [40] ^b	17.83	37.48	29.07

a) Over 121 subjects. b) Over 252 subjects

opportunistic behavior can be limited by increases on the shadow of the future), these two papers display several methodological problems. In both experiments, subjects played against the experimenter instead of playing against each other. While in Roth and Murnighan [50] subjects were told that they were playing against the experimenter they were not told that the experimenter was following the tit-for-tat strategy. In Murnighan and Roth [40] subjects were told that they were playing against each other while in fact they were playing against the experimenter who was following either the tit-for-tat or grim strategy. In addition, in both experiments subjects were not paid proportionately to the “points” they earned during the experiments. In Roth and Murnighan [50] “the best player” in the experiment, as they called it, received a \$10 prize, while in Murnighan and Roth [40] the player with the highest average score in each session received \$40 and the second one \$20.

Another experiment that employed a random continuation rule is Feinberg and Husted [18]. They combine a mixed continuation probability with different discount factors (they shrink the payoffs in every round) to study the effect of repetition in the levels of cooperation in a prisoner’s dilemma game disguised as a duopoly game. They found that the levels of cooperation increase as the discount factor increases. Nevertheless, that increase is small and far from the increase needed to fully exploit the possible benefits from cooperation even when the experiment and its instructions were

purposely designed to facilitate cooperation. In addition, these results are weakened because the payments made to the subjects were quite low and the basic payoffs were not the same in all treatments³³.

Another experiment in which subjects faced a random continuation rule was conducted by Palfrey and Rosenthal [42]. This paper presents an experiment designed to test whether the possibility of future interaction leads to greater cooperation in public good games with incomplete information. They compare the rate of contribution for a public good when players meet only once with each other and when they meet repeatedly with a probability of continuation of 0.9³⁴. They found that repetition leads to more cooperation than one shot games but this increase is small (the percentage of contribution goes from 29% to 40%). They concluded that "This contrast between our one-shot and repeated play results is not encouraging news for those who might wish to interpret as gospel the oft-spoken suggestion that repeated play with discount rates close to one leads to more cooperative behavior. True enough it does-but not by much."³⁵ As the authors suggest later, the power of repeated play may be more evident in a simpler environment.

A problem with the experiments with a random continuation rule is that it is not clear that any increase of cooperation as the probability of continuation increases is due to an increase in the importance of the future as theory predicts or if it is merely due to an increase of the repeated game expected horizon. There is experimental evidence that subjects cooperate more in infinitely repeated prisoner's dilemma games than in one-shot ones (see Andreoni and Miller [5] and Cooper et al. [15]). A reason

³³Another experiment that used a random continuation rule to study repeated oligopoly games is Holt [30]. Since this experiment was designed to test for consistent-conjectures, the results do not provide information regarding cooperation.

³⁴There were at least 20 round, after which the probability of continuation was 0.9.

³⁵Palfrey and Rosenthal [42], pag. 548.

for this effect is that in infinitely repeated games subjects may have incentives to build reputations if there is incomplete information (see Kreps et al. [33]). Therefore, if we observe an increase in cooperation as the probability of continuation increases, it could be due to an increase on reputation effects as the expected horizon of the game increases and not to the infinitely repeated feature of the game.

There exist some old repeated oligopoly experiments, like the ones presented in Fouraker and Siegel [20], that fall into the second category of games in which the number of repetitions was known to the experimenter but not to the subjects. In each round the subjects may assign a positive probability of continuation and we may consider these experiments as experiments on infinitely repeated games, at least in the minds of the subjects. Fouraker and Siegel [20] found some cooperation in Cournot duopoly markets but not in triopoly markets.

A more recent paper in this category is Brown Kruse et al. [13] which presents an experiment on repeated price competition in an oligopoly market with fixed capacity constraints. While they observe prices above competitive levels, those prices are far below the monopoly price. In addition, in the treatments in which collusion is more easily supported (requires a lower belief of continuation) the prices are lower. This contradicts what we would expect from infinitely repeated game theory, which predicts that when collusion is easier to support, higher prices should be observed if some of the subjects coordinate in collusive equilibria.

All experiments with a fixed number of rounds unknown to subjects raise the problem that the experimenter can not control for the players' beliefs with respect to the continuation of the game³⁶.

Previous experimental results do not provide much support for the theory of

³⁶This type of design also adds a source of incomplete information since subjects may not know what others subjects beliefs are.

in...nitely repeated games or much hope that self-enforcig reward and punishment schemes are used to overcome opportunistic behavior. But given the shortfalls of some experiments' design, (i.e. no real interaction among subjects, ...nal earnings that are not proportional to the payoffs during the game, low earnings, ...xed number of rounds unknown to the subjects and lack of control sessions), and the complicate environment of others (i.e. environments of incomplete information), previous experimental evidence is insu¢cient to assess the degree in which the theory of in...nitely repeated games is supported empirically. This paper presents results from an experiment that was designed to overcome the mentioned shortcomings and shows that not only the shadow of the future matters, but that its effect is signi...cant and that it closely, but not perfectly, follows the theoretical predictions.

3.3 Experiment design

As mentioned before, this experiment was designed to overcome the shortcomings of the previous literature and allow for a better testing of the theory of in...nitely repeated games. Then, we use simple stage games: prisoner's dilemma games. The subjects interacted with each other through computer terminals anonymously (see Figure 1) and the pairing of subjects was done such that there was no possibility of interaction or contagious effects among the different repeated games. We controlled for the subjects' discount factor by having a know probability of continuation. The subjects' ...nal payoff were proportional to the points earned during the experiment plus a show up fee. The exchange rate points/\$ made sure that subjects had a signi...cant incentives to try to increase their earnings.

In addition, the experimental design incorporates new elements that allow for a better testing of the theory of repeated games. First, in addition to the random continuation rule sessions, we run sessions with ...xed ...nite horizon games. The length

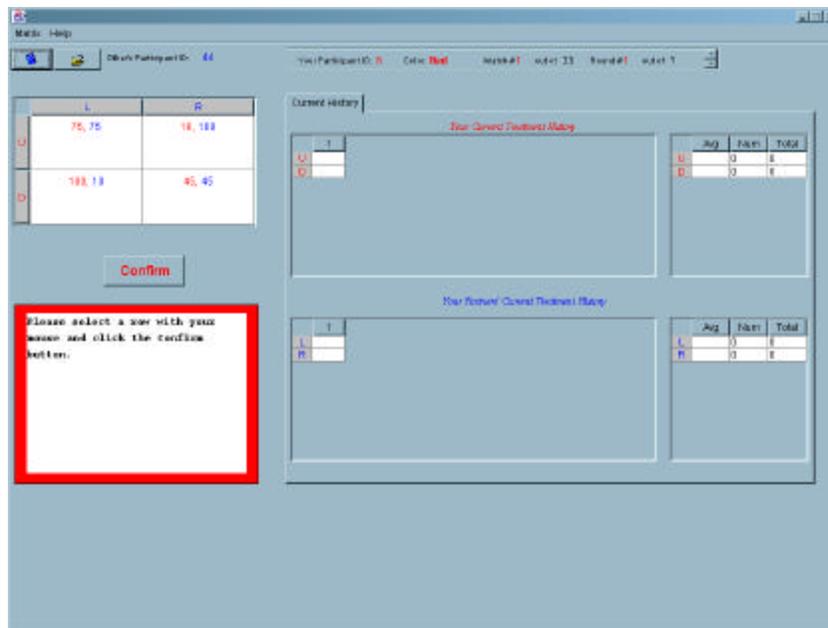


Figure 12: Computer screen that subjects saw before each interaction.

of the ...xed ...nite horizon sessions were chosen to coincide with the expected length of the random continuation rule ones. Therefore, for the ...rst time in the literature, the experimental design allows us to compare the results from the in...nite repeated sessions with the results from ...xed ...nite horizon repeated games to test whether cooperation depends on “the shadow of the future”, as theory predicts, or merely on the length of the games.

Second, we considered two diærent prisoner’s dilemma games that results in different set of equilibrium outcomes for some discount factors. In this way we can study how closely the experimental results follow the theoretical predictions.

I explain next the main characteristics of the experiment in greater detail.

Stage game payoffs: We consider two diærent stage games payoffs, denoted PD1

and PD2³⁷:

Table 2: Stage game payoffs in points

		PD1		PD2			
		Blue player		Blue player			
		C	D				
Red player	C	65 , 65	10 , 100	Red player	C	75 , 75	10 , 100
	D	100 , 10	35 , 35		D	100 , 10	45 , 45

The set of equilibrium outcomes for the infinitely repeated version of these games are described in the next section.

Public randomization device: To allow subjects to coordinate actions and rotate through different outcomes, every ten seconds a random number between 1 and 1000 was displayed on a screen at front of the room.

Subjects' total earnings: All payoffs in the game were in points. At the end of each session, the points earned by each subject were converted into dollars at the exchange rate 200 points=\$1 and paid privately in cash. In addition, subjects were paid a 5 dollar show up fee. In this way, subjects' real earnings in dollars are proportional (up to a constant) to the "points" obtained during the experiment. In addition, these amounts seem significant enough to influence subjects' behavior. In a session with mixed finite horizon games and 30 (60) subject the difference between the maximum and minimum possible earning is above 15 (31) dollars³⁸.

Infinitely repeated games: In one half of the sessions a random continuation rule

³⁷While in the experiment the actions were called U and D for Red subjects and L and R for Blue subjects, we will use here the usual names C and D.

³⁸In the sessions with a random continuation rule this difference depends, of course, on the realization of the random continuation rule.

was used to induce infinitely repeated games. This was done by having one of the subjects -who had been selected as the monitor- roll publicly a four sided die after each round. The randomization generates an infinitely repeated game given that at the moment of choosing an action there is always the possibility of interacting in future rounds with the same partner.

The probability of continuation \pm , of which three different values were considered, is the principal treatment variable in these sessions. One treatment corresponds with the one-shot game case: $\pm = 0$ and the rest corresponds with positive probability of continuations: $\pm = 0.5; 0.75$. This treatment variable allows us to control for the subjects' beliefs on the probability of continuation. We call these sessions "Dice" sessions.

Finitely repeated games: In the other half of the sessions fixed finite horizon games were used. We considered three treatments that allows me to compare results with the infinitely repeated games experimental results: 1) one-shot game: $H = 1$, 2) two rounds repeated game: $H = 2$ and 3) four rounds repeated game: $H = 4$. Note that each of this treatments coincides in the number of rounds with the expected number of rounds in one of the random continuation rule treatments³⁹. The number of rounds was common knowledge among the subjects. We call these sessions "Finite" sessions.

Order of treatments: To control for learning effects from one treatment to another, two sessions were run for each kind of continuation rule and payoff matrix. For example, for PD1 and Dice we run one session with the order ($\pm = 0, \pm = .5, \pm = .75$) and another with the inverse order ($\pm = .75, \pm = .5, \pm = 0$). We call the first kind of session "Normal" and the last kind "USD" (up-side-down).

³⁹In the infinitely repeated game with a continuation probability of \pm , the expected number of stages is equal to $\frac{1}{1-\pm}$. Therefore, the expected number of stages in the random continuation session will be 1, 2 and 4 for \pm equal to 0, 0.5 and 0.75, respectively.

Matching procedure: A rotation matching scheme was used to avoid potential interaction and contagious effects between the different repeated games⁴⁰. In each session subjects were divided in two groups: Red and Blue. In each match every Red subject was paired with a Blue subject. No pair consisted twice of the same subjects. In addition, subjects were not paired with someone that had played with someone that had played with him or her or with someone that had played with someone that had played with someone that had played with him or her, and so on. Thus, the pairing was done in such a way that the decisions one subject made in one match could not affect, in any way, the decision of subjects he or her would meet in the future. All these features were explained and made clear to the subjects.

Given that each subject was only matched once with each subject of the other color, the total number of matches in a session is $\frac{N}{2}$, where N is the number of subject in a session. Given that there are three treatment per session, in each treatment there are $\frac{N}{6}$ matches. The size of the experimental lab CASSEL allowed us to run experiments with up to sixty subjects, what provided up to ten matches per treatment per subject. This provided enough matches for the subjects to familiarize with the game in each treatment leaving enough observations to analyze.

Sessions: Given the two stage games (PD1 and PD2), the different continuation rules (Dice and Finite), the different treatments ($\pm = 0; :5; :75$, and $H = 1; 2; 4$), and the change in the order of the treatments (Normal and USD), this experiment consists of eight sessions with three treatments each. Each treatment, or part, consists of one practice match for which subjects are not paid and $\frac{N}{6}$ real matches, where N is the number of subjects in the session. Each match consists of as many rounds as the continuation rule indicates.

⁴⁰Note that, given Kandori [?, kandori92]'s contagious equilibrium, random matching is not enough to isolate the different games.

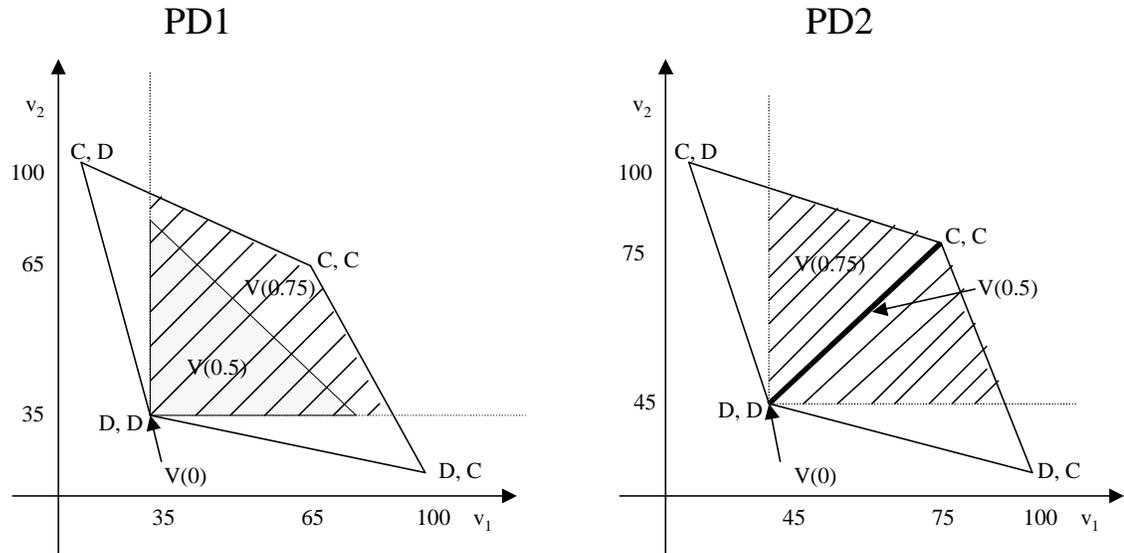


Figure 13: Equilibrium average payoffs $V(\pm)$

3.4 Theoretical predictions

As discussed in the previous section, the random continuation results in an infinitely repeated game in the sense that there is always the possibility of future interaction. If we assume that the payoffs in Table 2 are the actual total payoffs that the subjects obtain from the game and this is common knowledge, that is if we abstract from problems of interdependent utilities, altruism, taste for cooperation and reputation effects, the set of subgame perfect equilibria can be calculated using the results in Stahl [51]. Figure 2 shows the set of equilibrium average payoffs ($V(\pm)$) in each game for each of the discount factors used in the experiment. The outcomes that can be supported in equilibrium -and therefore the outcomes accordingly with theory we should observe- for the different discount factors are presented in Table 3.

New equilibria appear as the discount factor increases, allowing the subjects to reach -in principle- higher levels of cooperation and payoffs. We can think that some

Table 3: Equilibrium outcomes

δ	PD1	PD2
0	DD	DD
0.5	DD, CD, DC	DD, CC
0.75	DD, CD, DC, CC	DD, CD, DC, CC

subjects will make the most of this opportunity to cooperate, regardless of the fact that DD remains an equilibrium for high discount factors. Therefore, we have the following testable hypothesis⁴¹:

Hypothesis 1: The larger δ , the larger the levels of cooperation.

In the finite and infinite horizon sessions the theoretical prediction is that no cooperation is possible. Therefore I have the following testable hypothesis⁴²:

Hypothesis 2: In infinitely repeated games ($\delta = .5$ and $\delta = .75$) result in larger levels of cooperation than finitely repeated games ($H = 2$ and $H = 4$).

From Table 3 we see that the set of equilibrium outcomes is different for PD1 and PD2 for $\delta = .5$. Under that discount factor, CC can be played in equilibrium for PD1

⁴¹It is important to note that for this hypothesis it is not necessary to assume that the subjects' only payoffs from the stage game are the ones in Table 2. With different payoffs the predictions presented in Figure 2 and Table 3 may not be appropriate, but Hypothesis 1 can still be true. Abreu, Pearce and Stacchetti [3] showed that the set of equilibrium payoffs (and consequently the set of outcomes) that can be observed in a infinitely repeated game (even with imperfect monitoring), can not decrease when the discount factor increases. Then, for any stage game in which DD is the only Nash equilibrium, increases in the discount factor result in increases in the levels of cooperation.

⁴²As mentioned before, the levels of cooperation in a infinitely repeated game may be above the levels from one shot games given reputation effects. Unfortunately, there is no theoretical result that allow us to compare the set of equilibrium outcomes between finitely and infinitely repeated games under incomplete information. Therefore, the following proposition is based solely on the theory for repeated games under complete information.

while that is not the case under PD2⁴³. Therefore, we have the following testable hypothesis:

Hypothesis 3: For $\delta = 0.5$, the payoffs PD2 result in more outcomes CC than PD1.

The first two hypothesis are quite general in the sense that they can be summarized by the statement that “the shadow of the future matters.” The last hypothesis, instead, is quite specific in the sense that it is closely based on the specified payoff matrixes. Because of this close dependence on the payoff matrixes, Hypothesis 3 may seem more likely to be rejected and, then, a more strict test of the theory of infinitely repeated games.

3.5 Experimental results

Given the two stage games (PD1 and PD2), the different continuation rules (Dice and Finite), the different treatments, and the change in the order of the treatments (Normal and USD), this experiment consists of eight sessions with three treatments each. The main descriptive statistics of these sessions are in Table 4. The experimental sessions were run between November 2001 and April 2002 with an average length of one hour (without counting the time to pay subjects). Without counting the subjects selected to be monitors, 390 subjects participated in the experiment, an average of 48.75 subjects per session with a maximum of 60 and a minimum of 30. The subjects were UCLA undergraduates recruited through advertisement in university webpages and signs posted on campus. A 22.31% of these undergraduates declared that they were in one of the Economics major programs (Economics, Business Economics, Mathematics/Economics and Economics/International Area Studies).

⁴³Note that cooperation can still be observed in equilibrium for PD1 given that the outcomes CD and DC can be part of an equilibrium.

The subjects performed a total of 22,482 actions with an average of approximately 2810 actions per session and 58 actions per subject. They earned an average of \$18.94 with a maximum of \$25.85 and a minimum of \$12. The total payment to subjects was of \$7387.55 (without considering the payments to the monitors and students that showed up but did not participate in the experiment).

Even when subjects participated in a practice match before the real matches of each treatment, we should expect to see during the ...rst matches of each treatment some learning regarding the treatment characteristics and other's subjects behavior. As you can see in Table 5, there is clear learning regarding the possibilities of cooperation in the $\pm = 0$ treatment of the Dice sessions and all the treatments of the Finite sessions (that is, in all the treatments with ...xed and ...nite repeated games). For example, in the $\pm = 0$ treatment, cooperation goes from above 26% in the ...rst match to less than 11% in the third match.

Given this learning regarding the treatments, in the analysis of the experimental result we focus on the matches after the third.

3.5.1 Does cooperation increase with the shadow of the future?

Our ...rst objective is to study how changes in the probability of future interaction affect the levels of cooperation. The experimental results show that the greater the shadow of the future, the larger the levels of cooperation. Considering the aggregate results for the Dice sessions (matches after third and all rounds) we see that cooperation is just above 9% for the one shot treatment, while it is above 27% and 37% for $\pm = :5$ and $\pm = :75$, respectively -see Table 6. These differences are statistically significant with p values of less than 0.001. Therefore, the experimental results support Hypothesis 1: the larger \pm , the larger the levels of cooperation.

In addition, these results show that the effect of the shadow of the future is of an

Table 4: Sessions' descriptive data

		PD1		PD2	
		Dice	Finite	Dice	Finite
Normal	Date	11/6/01	11/13/01	2/7/02	4/18/02
	Time*	2:30-3:28	4:45-5:31	1:45-2:56	5:15-6:25
	Subjects	42	30	54	48
	Any Econ ⁺	23.81%	23.33%	12.96%	18.75%
	Actions	2268	1050	3132	2688
	Ave Earning	17.09	13.03	19.91	19.36
	Max Earning	19.40	15.23	22.18	21.88
	Min Earning	13.48	12.05	15.98	15.48
	Total \$	717.7	390.85	1075.10	929.20
USD	Date	11/29/01	11/20/01	4/9/02	4/15/02
	Time*	5:10-6:05	5:10-6:05	4:45-5:53	4:45-5:54
	Subjects	42	54	60	60
	Any Econ ⁺	16.67%	12.96%	31.67%	35%
	Actions	1722	3402	4020	4200
	Ave Earning	14.37	17.77	23.09	22.11
	Max Earning	16.23	21.55	25.85	25.10
	Min Earning	12.18	12	19.93	17.15
	Total \$	603.65	959.45	1385.10	1326.50

*Starting Scheduled time and ...nal actual time.

⁺ Percentage of all Economics majors in the session.

Table 5: Percentage of cooperation by match and treatment*

Dice	Match									
	1	2	3	4	5	6	7	8	9	10
$\pm = 0$	26.26	18.18	10.61	11.62	12.63	12.63	5.56	5.26	5.26	5
$\pm = :5$	28.36	27.12	34.58	35.53	21.60	19.08	29.84	35.96	28.16	50
$\pm = :75$	40.44	28.57	27.78	32.92	46.51	33.09	44.05	53.51	42.26	45.83
Finite										
H=1	26.56	18.23	16.67	17.19	11.98	8.02	6.79	10.49	6.14	6.67
H=2	19.79	15.89	14.84	9.64	11.46	10.80	12.04	10.19	6.58	6.67
H=4	31.64	30.34	30.47	25.52	25.13	23.77	16.36	19.75	14.91	20.83

*All rounds.

Table 6: Percentage of cooperation by treatment*

Dice	Finite	
$\pm = 0$	9.17	H=1 10.34
$\pm = :5$	27.41	H=2 10.11
$\pm = :75$	37.64	H=4 21.43

*Matches after third and all rounds.

important magnitude: the percentage of cooperation for $\pm = :75$ is almost three times greater than for the one shot treatment. This magnitude is greater than previously studies have indicated. For example, in the public good experiments with incomplete information of Palfrey and Rosenthal [42] the percentage of contributions increases only from 29% to 40% when the treatment changes from one shot games to a random continuation rule with $\pm = :9$. This is also the case if we compare the results of this experiment with the results from Roth and Murnighan [50] and Murnighan and Roth [40]⁴⁴. While in those experiments the percentage of cooperation in the ...rst round

⁴⁴Given the described characteristics of these experiments we compare the results for the ...rst round of each match.

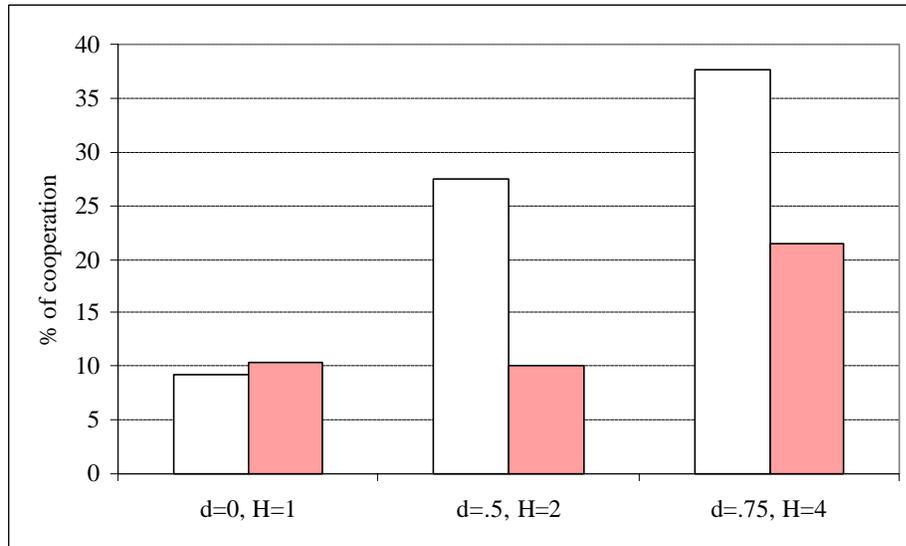


Figure 14: Cooperation by treatment (matches after third and all rounds)

less than doubles when the probability of continuation increases from 0.105 to 0.895, in this experiment the percentage cooperation is four times higher with a probability of continuation of 0.75 than in one shot games. These results support the idea that infinitely repeated interaction can significantly reduce opportunistic behavior.

3.5.2 Infinitely repeated games vs. finitely repeated games

Our second objective is to compare the levels of cooperation in the Dice and Finite sessions. As Table 6 shows, the percentage of cooperation is similar for the one shot treatments in both types of sessions (p value = 0.507), showing that there are no significant differences in the “kind” of people that participated in each session. More importantly, the percentage of cooperation is greater in the infinitely repeated games than in the finitely repeated games (with p values of less than 0.001). Therefore, the experimental results support Hypothesis 2: infinitely repeated games result in larger levels of cooperation than finitely repeated games.

Studying the levels of cooperation by round for each of the treatments (Table 7) results in further support for the theory of repeated games. In the fourth round of the $\pm = :75$ treatment the level of cooperation is significantly greater than in the fourth (and last) round of the $H = 4$ treatment (34.58% against 10.63%, with p value of less than 0.001). The level of cooperation in the final round of the $H = 4$ treatment is similar to the level of cooperation in one shot games. Therefore, the absence of a future affects subjects behavior in the final round of the fixed finite horizon games. They cooperate less when there is no future. This seems to be understood by the subjects at the beginning of the game. The levels of cooperation in the first round are much greater in infinitely repeated games than in finitely repeated games (46.20% against 34.58%, with p value of less than 0.001). Similar reasoning applies to the comparison of the behavior for $\pm = :5$ and $H = 2$.

Table 7: Percentage of cooperation by round and treatment*

	Round											
Dice	1	2	3	4	5	6	7	8	9	10	11	12
$\pm = 0$	9.17											
$\pm = :5$	30.93	26.10	19.87	12.50	12.96							
$\pm = :75$	46.20	40.76	38.76	34.58	33.04	27.27	24.75	26.28	29.17	26.04	32.29	31.25
Finite												
H=1	10.34											
H=2	13.31	6.90										
H=4	34.58	21.55	18.97	10.63								

*Matches after third.

3.5.3 Do payoff details matter?

Our third objective is to compare the outcomes under PD1 and PD2 when $\pm = :5$. Remember that CC is not an equilibrium outcome under PD1 but it is under PD2.

Consistent with that, the percentage of outcomes in which both players cooperate is significantly lower under the payoff matrix PD1 than under PD2 (3.17% against 18.83% with a p value of less than 0.001). Then, the experimental results support Hypothesis 3: For $\delta = .5$, the payoffs PD2 result in more outcomes CC than PD1.

Table 8: Distribution of outcomes for $\delta = .5$ *

	PD1	PD2
CC	3.17	18.83
CD	16.67	11.00
DC	11.90	14.50
DD	68.25	55.67

*Matches after third and all rounds.

The percentage of outcomes in which only one subject cooperates (CD and DC) is greater under PD1 than under PD2 (28.57% against 25.50%) as theory predicts. Nevertheless, this difference is not statistically significant (p value of 0.19) pointing out the difficulty of coordinating on alternating asymmetric outcomes even when there is a public randomization device available.

3.5.4 Do Economics majors behave differently?

It is important to note that the support for all three hypothesis does not depend on the major of the subjects. All three hypothesis are supported by the experimental results for students in any of the Economics majors and students in the rest of the majors. With respect to the first two hypothesis, for both Economics majors and Non-Economics majors cooperation increases as the probability of future interaction increases and cooperation is greater in infinitely repeated games than in finitely repeated games -see Table 9.

Notwithstanding this, there are differences in behavior across majors. Economics

majors cooperate significantly less than Non-Economics majors in games with a fixed infinite horizon (this difference is significant for the Finite sessions with p values of 0.009, 0.042 and less than 0.001 for H=1, 2 and 4, respectively, but it is not significant for $\delta = 0$ with p value of 0.45). Instead, the evidence is contradictory regarding infinitely repeated games. While Economics majors cooperate more than Non-Economics majors for $\delta = .5$ (p value of less than 0.001), that is not the case for $\delta = .75$. In fact, the percentage of cooperation is lower for Economics majors, but this difference is not significant (p value of 0.159).

Table 9: Percentage of cooperation by treatment and major*

	Dice			Finite	
	Non-Econ	Econ		Non-Econ	Econ
$\delta = 0$	9.68	7.41	H=1	12.19	4.44
$\delta = .5$	26.65	29.97	H=2	11.12	6.85
$\delta = .75$	38.93	33.15	H=4	23.81	13.81

*Matches after third and all rounds.

With respect to the third hypothesis, for both types of majors PD1 results in a lower percentage of CC than PD2 when $\delta = .5$ (p values of less than 0.001 for both types of majors). Nevertheless, it is interesting to note that this effect is stronger for Economics majors.

3.6 Conclusions

The experimental evidence presented in this paper provides strong support for the extensive use of the theory of infinitely repeated games by showing that the shadow of the future matters, it significantly reduces opportunistic behavior and it does it in a way that closely follows the theoretical predictions.

Table 10: Distribution of outcomes for $\pm = :5$ and major*

	PD1		PD2	
	Non-Econ	Econ	Non-Econ	Econ
CC	2.55	5.36	14.02	30.81
CD	18.88	8.93	10.51	12.21
DC	12.76	8.93	14.25	15.12
DD	65.82	76.79	61.21	41.86

*Matches after third and all rounds.

The data produced in this experiment deserves further study. It remains for future work to analyze the reward and punishment schemes used by the subjects. It would also be important to study whether, given these schemes, subjects' behavior constitutes an equilibrium, or how close they are to an equilibrium, by measuring subjects' average losses as in Fudenberg and Levine [21].

4 Appendices

4.1 Appendix to chapter 1

Proof of Theorem 5: Consider the following social norm that yields v in equilibrium: if player i meets player j , and both are nice ($z_i = z_j = 0$), they play $(a_{ij}; a_{ji})$; if a nice player i meets a guilty player j ($z_i = 0, z_j \neq 0$) the former plays m and the latter plays r , (that is, the nice player punishes the guilty one, and this "asks" for forgiveness); and if two guilty players ($z_i \neq 0, z_j \neq 0$) meet they both play m . The local information system works as follows: if a player deviates he gets a tag $z = T$, denoting that he has to be punished for T periods; if a guilty player conforms he has his tag reduced one unit and if a nice player conforms he keeps the $z = 0$ tag.

To prove that the former is a sequential equilibrium I show that no player has incentives to deviate after any possible history if the rest of the players follow the social norm. As in the previous proposition, this makes beliefs unimportant, since regardless of what has happened in the past no player has incentives to deviate. First I check that a nice player has no incentives to deviate and second, I check that a guilty player has no incentives to deviate.

When a nice player studies if he should conform or deviate he must not only consider today's profit from the deviation but also the loss in the future T periods of punishment. This loss depends on whether the other players are nice or guilty and the payoff that the player receives upon meeting each nice player. Define $g_{ij} = g(a_{ij}; a_{ji})$ as the payoff that a nice player i receives for playing with a nice player j . If j is nice, player i would face a future loss for deviating today equal to $g_{ij} - \underline{g}$ every time he meets player j in the next T periods (remember that $g(r; m) = \underline{g}$). If j is guilty the future loss by deviation for i , when he meets j , would be $g(m; r) - g_m$. Hence for a nice player i the future loss for deviating today each time he meets player j would

be greater when the latter is nice if $g_{ij} \geq \underline{g} > g(m; r) \geq g_m$. Therefore when $g_{ij} > g(m; r) \geq g_m + \underline{g}$ player i has more incentives to deviate when j is guilty than when j is nice.

If the former inequality holds for every j with $\theta_{ij} > 0$, a nice player i would have the greatest incentive to deviate when the other players are T-guilty. In that case, if he conforms with the prescribed strategy he would get at least $(1 - \epsilon)g(m; r) + \epsilon v_i^0$ where $v_i^0 = \epsilon \sum_{j \in I} \theta_{ij}^{-1} \epsilon g(m; r) + \epsilon \sum_{j \in I} \theta_{ij}^{-1} v_j$. By deviating he would get at most $(1 - \epsilon)\bar{g} + \epsilon v_i^{p1}$, where $v_i^{p1} = \epsilon \sum_{j \in I} \theta_{ij}^{-1} \epsilon g_m + \epsilon \sum_{j \in I} \theta_{ij}^{-1} (1 - \epsilon)\bar{g} + \epsilon \sum_{j \in I} \theta_{ij}^{-1} v_j$. Choose T so that $\epsilon^T = d \geq 2(0; 1)$. Then, given that $g(m; r) > g_m \geq \underline{g}$, it is true that $v_i^0 > v_i^{p1}$ independently of ϵ . Therefore, for ϵ large enough it is true that $(1 - \epsilon)g(m; r) + \epsilon v_i^0 > (1 - \epsilon)\bar{g} + \epsilon v_i^{p1}$.

If $g_{ij} < g(m; r) \geq g_m + \underline{g}$ for every j with $\theta_{ij} > 0$, player i has more incentives to deviate when all the other players are nice. In this case he would get at least $(1 - \epsilon)\underline{g} + \epsilon v_i$ by conforming and at most $(1 - \epsilon)\bar{g} + \epsilon \sum_{j \in I} \theta_{ij}^{-1} \epsilon \underline{g} + \epsilon \sum_{j \in I} \theta_{ij}^{-1} v_j$ by deviating. Since $v_i > \underline{g}$, choosing T so that $\epsilon^T = d \geq 2(0; 1)$, the former is greater than the latter for ϵ large enough.

It could also be the case that player i faces players with whom the inequality is satisfied with probability θ , and with probability $(1 - \theta)$ faces players with whom it is not satisfied. In this case the incentives for i to deviate will be the highest when the players belonging to the first group are T-guilty and the rest are nice. By conforming he would get at least $(1 - \epsilon)\underline{g} + \epsilon v_i^{p0}$, where $v_i^{p0} = \epsilon \sum_{j \in I} \theta_{ij}^{-1} \epsilon \theta g(m; r) + (1 - \theta)\bar{g} + \epsilon \sum_{j \in I} \theta_{ij}^{-1} v_j$ and by deviating he would get at most $(1 - \epsilon)\bar{g} + \epsilon v_i^{p2}$, where $v_i^{p2} = \epsilon \sum_{j \in I} \theta_{ij}^{-1} \epsilon \theta g_m + (1 - \theta)\bar{g} + \epsilon \sum_{j \in I} \theta_{ij}^{-1} (1 - \epsilon)\bar{g} + \epsilon \sum_{j \in I} \theta_{ij}^{-1} v_j$. Choose T so that $\epsilon^T = d \geq 2(0; 1)$. Then, given that $g(m; r) > g_m$, it is true that $v_i^{p0} > v_i^{p2}$ independently of ϵ . Therefore, for ϵ large enough it is true that $(1 - \epsilon)\underline{g} + \epsilon v_i^{p0} > (1 - \epsilon)\bar{g} + \epsilon v_i^{p2}$.

Consider the case of a ζ_i guilty player i . Define θ_t as the probability that i will be matched with a guilty player at time t . The least he can make by conforming with

the prescribed strategy is:

$$(1 - \delta) \underline{g} + \sum_{t=2}^{\infty} \delta^{t-1} i_{\otimes_t} g_m + (1 - \delta_{\otimes_t}) \underline{g} + \sum_{t=\zeta+1}^{\infty} \delta^{t-1} i_{\otimes_t} g(m; r) + (1 - \delta_{\otimes_t}) \underline{g} + \delta^T v_i$$

(remember that $g(r; m) = \underline{g}$) and by deviating he can get at most:

$$(1 - \delta) \sum_{t=2}^{\infty} \delta^{t-1} i_{\otimes_t} g_m + (1 - \delta_{\otimes_t}) \underline{g} + \sum_{t=\zeta+1}^{\infty} \delta^{t-1} i_{\otimes_t} g_m + (1 - \delta_{\otimes_t}) \underline{g} + \delta^T \underline{g} + \delta^{T+1} v_i$$

Therefore the gains from conforming are at least:

$$(1 - \delta) \delta^T v_i + \sum_{t=\zeta+1}^{\infty} \delta^{t-1} i_{\otimes_t} (g(m; r) - \underline{g}_m) + \delta^{T-\zeta} \underline{g} \quad .^{45}$$

By Assumption 1 the second term is non negative for any sequence of \otimes_t . Hence, it is enough to have $\delta^T v_i + \delta^{T-\zeta} \underline{g} > 0$ for the gains from conforming to be positive. This can be done by choosing $\delta^T = d$ large enough.

Note that there is no contradiction in the requirements made on δ and T in the two parts of the proof, it is only required that $\delta^T = d$ and δ are large enough. ¥

Proof of Proposition 6: Consider the following social norm that yields v in equilibrium: if player i meets player j , and both are nice ($z_i = z_j = 0$), they play $(a_{ij}; a_{ji})$; if a nice player i meets a guilty player j ($z_i = 0, z_j = 1$) the former plays m and the latter plays r ; and if two guilty players ($z_i = z_j = 1$) meet they minmax each other. The local information system works as follows: if a player deviates he gets a tag $z = 1$, denoting that he has to be punished; with probability $p \in (0; 1)$ all the guilty players that have conformed last period are forgiven and get a tag $z = 0$ or remain guilty ($z = 1$) with probability $(1 - p)$; and if a nice player conforms he keeps the $z = 0$ tag.

First I check that a nice player has no incentives to deviate and second, I check that a guilty player has no incentives to deviate.

As in the proof of Theorem 5, for a nice player i the future loss by deviating today

⁴⁵This formulas are only valid for the case of $\zeta < T$. For the case in which $\zeta = T$, a similar formula for the gains from conforming results. The only difference is that in the latter case the second term disappears.

each time he meets player j would be greater when the latter is nice if $g_{ij} > \underline{g} > g(m; r) > g_m$. Therefore if $g_{ij} > g(m; r) > g_m + \underline{g}$ player i has more incentives to deviate when j is guilty than when j is nice.

If the former inequality holds for every j with $\theta_{ij} > 0$, a nice player i would have the greatest incentive to deviate when the other players are guilty. Using the recursiveness of the problem it can be found that the expected utility player i receives by conforming is $\frac{1}{1 \pm p} [(1 \pm \theta)g(m; r) + \pm(1 \pm \theta)v_i]$ while by deviating he receives at most $(1 \pm \theta)\bar{g} + \frac{\pm(1 \pm \theta)}{1 \pm p} pg_m + (1 \pm \theta)\underline{g} + \frac{\pm^2(1 \pm \theta)}{1 \pm p} v_i$.

Calculating the difference between them and simplifying, for player i not to have incentives to deviate it must be the case that $\frac{1}{1 \pm p} [g(m; r) \pm \theta pg_m + (1 \pm \theta)\underline{g} + \pm(1 \pm \theta)v_i] \leq \bar{g}$. From Assumption 1 and $v_i > 0$, it must be true that the term in brackets is positive, therefore, the inequality is satisfied for $\pm p$ large enough.

If $g_{ij} < g(m; r) > g_m + \underline{g}$ for every j with $\theta_{ij} > 0$, player i has more incentives to deviate when all the other players are nice than when they are guilty. In this case he would get at least $(1 \pm \theta)\underline{g} + \pm v_i$ by conforming and (using the recursiveness of the problem) at most $(1 \pm \theta)\bar{g} + \pm(1 \pm \theta)\underline{g} + \frac{\pm^2}{1 \pm p} p(1 \pm \theta)\underline{g} + (1 \pm \theta)v_i$ by deviating. Then, for player i not to have incentives to deviate it must be the case that $\frac{\pm}{1 \pm p} v_i > \bar{g} - \underline{g}$. Given that $v_i > \underline{g}$, the inequality is satisfied for $\pm p$ large enough.

It could also be the case that player i faces players with whom the inequality is satisfied with probability θ , and with probability $(1 - \theta)$ faces players with whom it is not satisfied. In this case the incentives for i to deviate will be the highest when the players belonging to the first group are guilty and the rest are nice. In this case it can be shown that player i has no incentives to deviate if $\frac{\pm}{1 \pm p} p\theta(g(m; r) > g_m) > (1 \pm \theta)\underline{g} + \pm(1 \pm \theta)v_i > \bar{g} - \underline{g}$. From Assumption 1 and $v_i > 0$, it must be true that the term in brackets is positive (remember that $\underline{g} > 0$),

therefore, the inequality is satisfied for δ large enough.

If player i is guilty it can be shown that, regardless of the other players' tags, he does not have incentives to deviate if $\delta(1 - p)v_i + (1 - \delta)g_i > 0$. This inequality is satisfied for δ large enough.

Note that there is no contradiction in the requirements made on δ and p in the four different parts of the proof: it is only required that $\delta, p \in (0, 1)$ with δ and p large enough. \square

4.2 Appendix to chapter 2

Proof of Lemma 10: It is straight forward to see that if $aB > c(\frac{1}{4}^m)$ perfect collusion can be supported for any discount factor. If that condition does not hold, define

$$\underline{b}(B) = \begin{cases} \infty & \\ \min_{j \in [a; b]} f_j & \text{if } c(\frac{1}{4}^m) \leq bB \\ b & \text{if } c(\frac{1}{4}^m) > bB \end{cases}$$

If $aB \leq c(\frac{1}{4}^m) \leq bB$ there exists a number $\underline{b}(B)$ that makes $c(\frac{1}{4}^m) = \underline{b}(B)B$ by continuity of a linear function. If $bB < c(\frac{1}{4}^m)$, $\underline{b}(B) = b$. Therefore, $\underline{b}(B)$ exists (and is continuous).

When $\delta > \underline{b}(B)$, the incentive compatibility constraint is not binding since $\frac{1}{4}^m$ could be supported with even a lower discount factor. On the opposite case, when $\delta < \underline{b}(B)$, the incentive compatibility constraint is binding given that under Assumption 2, $\frac{d c(\frac{1}{4}^m)}{d \frac{1}{4}^m} = \frac{d \frac{1}{4}^m}{d q} \frac{d q}{d \frac{1}{4}^m} > 1 > 0$.

Since $c(\frac{1}{4}^m)$ is increasing for $\delta > \underline{b}(B)$, its inverse exists in the relevant range and the incentive compatibility constraint can be written as a function of δB , $\frac{1}{4}^m(\delta) = c_{\delta}^{-1}(\delta B)$, for $\delta > \underline{b}(B)$. Since $\frac{d c(\frac{1}{4}^m)}{d \frac{1}{4}^m} > 0$, this constraint is increasing in the discount factor. Finally, given that $c(\frac{1}{4}^m)$ is continuous $c_{\delta}^{-1}(\delta B)$ is also continuous. \square

Proof of Lemma 11: Since $\bar{c}(\frac{1}{4})$ is decreasing in $(\frac{1}{4} \in]\frac{1}{4}^c]$, its inverse exists in that range and the incentive compatibility constraint can be written as a function of $\pm B$, $\frac{1}{4}(\pm)$, $\bar{c}_i^{-1}(\pm B)$, which is decreasing in the discount factor. Finally, given that $\bar{c}(\frac{1}{4})$ is continuous $\bar{c}_i^{-1}(\pm B)$ is also continuous. ¥

Proof of Lemma 15: From the stick and carrot property of the optimal punishment, Lemma 13, the continuation payoffs of both the highest and lowest equilibrium discounted payoff coincide and $\bar{V}(\pm) \in \underline{V}(\pm) = \frac{1}{4}^p(\pm) \in \bar{c}_i^{-1} \pm \frac{E \frac{1}{4}^p}{1 \pm} \in E \underline{V}$. From the characterization of $\bar{V}(\pm)$ and equation (12) we know that the shape of $\frac{1}{4}^p(\pm)$ depends on the shape of the IC⁺-F constraint. In the IC range the concavity of the constraint is determined by the sign of $\frac{d^2 \bar{c}_i^{-1}}{d\pm^2}$. By Assumptions 2 and 3, $\frac{d^2 \bar{c}}{d\frac{1}{4}^2} = \frac{d^2 \frac{1}{4}^d}{d\frac{1}{4}^2} \frac{d\frac{1}{4}}{d\frac{1}{4}} \frac{d^2 \bar{c}}{d\frac{1}{4}^2} + \frac{d\frac{1}{4}^d}{d\frac{1}{4}} \frac{d^2 \bar{c}}{d\frac{1}{4}^2} > 0$, and by Lemma 11, $\frac{d \bar{c}_i^{-1}}{d\pm} > 0$, then $\frac{d^2 \bar{c}_i^{-1}}{d\pm^2} = \frac{\frac{1}{4} B^2}{(\frac{d \bar{c}}{d\frac{1}{4}})^2} \frac{d^2 \bar{c}}{d\frac{1}{4}^2} \frac{d \bar{c}_i^{-1}}{d\pm} > 0$ and \bar{c}_i^{-1} is concave. The F range of the constraint is also concave, since it is a constant. Hence, given that the IC⁺-F constraint is increasing and continuous, the IC⁺-F constraint is increasing and concave and so is $\frac{1}{4}^p(\pm)$. From equation 13 we know that $\underline{V}(\pm)$ is convex if $\bar{c}_i^{-1}(\cdot)$ is also convex. By Assumptions 2 and 3, $\frac{d^2 \bar{c}}{d\frac{1}{4}^2} = \frac{d^2 \frac{1}{4}^d}{d\frac{1}{4}^2} \frac{d\frac{1}{4}}{d\frac{1}{4}} \frac{d^2 \bar{c}}{d\frac{1}{4}^2} + \frac{d\frac{1}{4}^d}{d\frac{1}{4}} \frac{d^2 \bar{c}}{d\frac{1}{4}^2} > 0$, and by Lemma 12, $\frac{d \bar{c}_i^{-1}}{d\pm} < 0$, then $\frac{d^2 \bar{c}_i^{-1}}{d\pm^2} = \frac{\frac{1}{4} B^2}{(\frac{d \bar{c}}{d\frac{1}{4}})^2} \frac{d^2 \bar{c}}{d\frac{1}{4}^2} \frac{d \bar{c}_i^{-1}}{d\pm} < 0$ and \bar{c}_i^{-1} is decreasing and convex. Therefore, $\frac{1}{4}^p(\pm) \in \bar{c}_i^{-1} \pm \frac{E \frac{1}{4}^p}{1 \pm} \in E \underline{V}$ is increasing and concave on \pm . ¥

Optimal collusion and linear programming: Consider the case in which the discount factor takes in every period one of L values: $\pm_1; \pm_2; \dots; \pm_L$. Denote as T the transition matrix, where t_{is} denotes the probability that the future discount factor is \pm_s given that today's is \pm_i . Let V be the column vector of discounted continuation payoffs and $\frac{1}{4}$ the column vector of profits given the discount factor. Define \mathbf{t} as the matrix for which $\mathbf{t}_{is} = \pm_i t_{is}$. Then, $V = \frac{1}{4} + \mathbf{t}V$, the incentive compatibility constraint is $(N \in]1) \frac{1}{4} - \mathbf{t}V$ and the feasibility constraint is $\frac{1}{4} \in \frac{1}{4}^m 1_L$, where 1_L is a column vector of ones. Then, the optimal tacit collusion profits result from the following problem:

$\max \mathbb{1}_i$ subject to: $(N_i - 1) \mathbb{1}_i \leq \mathbb{1}_i^3 - \mathbb{1}_i^2$, $\mathbb{1}_i \geq 0$ and $\mathbb{1}_i \leq \frac{1}{4} \mathbb{1}_L$, where $\mathbb{1}_i$ is any non-negative row vector.

4.3 Appendix to chapter 3: Instructions for PD2-Dice-USD Session (4/9/02)

Welcome

You are about to participate in a session on decision-making, and you will be paid for your participation in cash, privately at the end of the session. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

Please turn off pagers and cellular phones now. Please close any program you may have open on the computer.

The entire session will take place through computer terminals, and all interaction between you will take place through the computers. It is important that you not talk or in any way try to communicate with other participants during the session.

We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the session and will be shown how to use the computers. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear.

General Instructions

In this session one participant will act as a monitor. The monitor will be paid a fixed amount for the session. The monitor will assist in running the session and checking that the session is run correctly. We will select the monitor now.

Open your envelope, and read the record sheet inside. If your sheet says "monitor" you are the monitor. Will the monitor please come to the master computer. If your

sheet does not say "monitor" you will use this sheet later to record your participant number that will be assigned by the computer and your ...nal score. Keep your sheet in a safe place, you will need it at the end of the session to receive your payment.

At this time, please pull out the dividers that separate you from your neighbors. During the course of this session, please refrain from communicating with your neighbors.

Please double click on the Dice Icon.

In the dialog box, please enter your full name and select server #128.97.190.171, as shown on the screen at the front of the room, and click OK. This will log you on to the session. In the upper side of your screen you can see you ID number for this session and your color - please look at the example on the screen in the front of the room. Please write your participation ID number in the record sheet that came in the envelope.

[Wait for people to ...ll in the record sheet]

Any question?

The session you are participating in is broken down into 3 separate parts. At the end of the last part, you will be paid the total amount you have accumulated during the course of the 3 parts in addition to the show-up fee. Everybody will be paid in private after showing the record sheet. You are under no obligation to tell others how much you earned.

During the session all the earnings are denominated in points. Your dollar earnings at the end of the session are determined by the points/\$ exchange rate posted on the board in the front and back of the room. This exchange rate is equal to 200points/\$. Therefore, 200 points are equivalent to \$1.

The participants are divided in two groups: Red and Blue.

Red and Blue participants will be matched together to interact in the following

way. As you see on the screen at the front of the room, the Red participant can choose between U or D and the Blue participant can choose between L and R.

If the Red participant chooses U and the Blue participant chooses L, both earn 75 points.

If the Red participant chooses U and the Blue participant chooses R, the Red participant earns 10 and the Blue participant earns 100 points.

If the Red participant chooses D and the Blue participant chooses L, the Red participant earns 100 and the Blue participant earns 10 points.

If the Red participant chooses D and the Blue participant chooses R, both earn 45 points.

The points of the Red participants are indicated on the screen in red, and the Blue participant points are indicated in blue.

In addition, the screen will show on the right hand side the result of previous rounds of the current match.

Every ten seconds, we will generate a random number between 1 and 1000 and project this number on the screens in the front of the room. You can use this number to select one of the actions, if you want, like the flip of a coin. For example, if you are a Red participant, you can decide to choose U any time the random number is above, say, 200.

Part 1

We will begin the ...rst part now. This ...rst part will consist of 10 matches. In each match every Red participant is paired with a Blue participant. You will not be paired twice with the same participant during the session or with a participant that was paired with someone that was paired with you or with someone that was paired with someone that was paired with someone that was paired with you, and so on. Thus, the pairing is done in such a way that the decisions you make in one match

cannot affect the decisions of the participants you will be paired with in later matches or later parts of the session.

In this part, after each round the monitor will roll a four sided dice. If the numbers 1, 2 or 3 appear, the participants will interact again without changing pairs. If a 4 appears, the match ends and participants are re-matched to interact with other participants. Therefore, in this part, each pair will interact until a 4 appears. After that, a new match will start with different pairs. Therefore you will interact until a 4 appears, with 10 different participants.

But first, we are going to teach you about this part of the session and how to use the computer by going through one practice match. During the practice part do not hit any keys until you are told to do so. You are not paid for the practice match; it is just for you to familiarize yourself with the session and the computer program.

[Begin Treatment 1 - Practice Part 1]

[Press OK to continue and press Start Treatment]

[Important: Tick Abort Current Match!!!!]

Your screen shows the possible actions you can choose, the actions the participant you are matched with can choose, and the points. You may choose your action by pressing the desired action at the side of the matrix now. If you are a Red participant you can press the actions in red, U or D, and if you are a Blue participant you can press the actions in Blue, L or R. Make your choices now. Once everyone in the room has made their selections and pressed confirm, your results from this round will appear on the screen.

[Wait until everybody has made a decision]

Monitor, would you please roll the dice?

[1) If a 1, 2 or 3 appeared press NO] A ___ appeared therefore this match continues. Now you are in the second (third, fourth, ...fth,) round of the same match. You are still

interacting with the same participant. Your screen shows all the same information as before. In addition you can see on your right the result of the previous rounds. [Important: Tick Abort Current Match!!!!] You may choose your action by pressing the desired action at the side of the matrix now. Make your choices now. Once everyone in the room has made their selections and pressed confirm, your results from this round will appear on the screen. Monitor, would you please roll the dice? [If 1, 2 or 3 appeared go to 1). If 4 appeared go to 2)]

[2) If a 4 appeared press YES] A 4 appeared therefore this match ended. On the screen you see a dialog box with the points you earned during the practice match. Press OK to end the practice match.

We have finished with the practice match. Any questions?

We start now with the first part of the session. You will now participate in 10 matches, each match paired with a different participant. In each match you will interact with the same person until a 4 appears. Remember: your decisions in one match cannot affect the decisions of the people you will interact with in future matches. This is not a practice; you will be paid!

[Begin Treatment 2 - Part 1 Paid.]

[Press Start Treatment, press Yes]

[Important: Tick Abort Current Match!!!!]

Make your choices now. Remember to press confirm.

[Wait until everybody has made a decision]

Monitor, would you please roll the dice?

[1) If 1, 2 or 3 appears press NO] A ___ appeared. This match continues. You are still interacting with the same participant. [Important: Tick Abort Current Match!!!!] Make your choices now. Remember to press confirm. Monitor, would you please roll the dice? [If 1, 2 or 3 appeared go to 1). If 4 appeared go to 2)]

[2] If 4 appears press YES] A 4 appeared. This match ends. On the screen you will see a dialog box with the points you earned during this match. Press OK to be matched with the next participant.

.

This is the end of Part 1. On your screen you will see a dialog box indicating your point and dollar points for this part. Press OK to move to the next part.

Part 2

We will begin the second part now. This part will consist of 10 matches. In each match every Red participant is paired with a Blue participant. No pair will consist of the same participants as in Part 1. As before, you will not be paired twice with the same participant during the session or with a participant that was paired with someone that was paired with you or with someone that was paired with someone that was paired with someone that was paired with you, and so on. Thus, the pairing is done in such a way that the decisions you make in one match cannot affect the decisions of the participants you will be paired with in later matches or later parts of the session.

In this part, after each round the monitor will roll a four sided dice. If the numbers 1 or 2 appear, the participants will interact again without changing pairs. If 3 or 4 appear, the match ends and participants are re-matched to interact with other participants. Therefore, in this part, each pair will interact until a 3 or 4 appear. After that, a new match will start with different pairs. Therefore you will interact until a 3 or 4 appear, with 10 different participants.

But first, we are going to teach you about this part of the session and how to use the computer by going through one practice match. During the practice part do not hit any keys until you are told to do so. You are not paid for the practice match; it is just for you to familiarize yourself with the session and the computer program.

[Begin Treatment 3 - Practice Part 2]

[Press Start Treatment, press Yes]

[Important: Tick Abort Current Match!!!!]

As before, your screen shows the possible actions you can choose, the actions the participant you are matched with can choose, and the points. You may choose your action by pressing the desired action at the side of the matrix now. Make your choices now. Once everyone in the room has made their selections and pressed confirm, your results from this round will appear on the screen.

[Wait until everybody has made a decision]

Monitor, would you please roll the dice?

[1) If a 1 or 2 appeared press NO] A ____ appeared therefore this match continues. Now you are in the second (third, fourth, ...fth,) round of the same match. You are still interacting with the same participant. Your screen shows all the same information as before. In addition you can see on your right the result of the previous rounds.

[Important: Tick Abort Current Match!!!!] You may choose your action by pressing the desired action at the side of the matrix now. Make your choices now. Once everyone in the room has made their selections and pressed confirm, your results from this round will appear on the screen. Monitor, would you please roll the dice?

[If 1 or 2 appeared go to 1). If 3 or 4 appeared go to 2)]

[2) If a 3 or 4 appeared press YES] A ____ appeared therefore this match ended. On the screen you see a dialog box with the points you earned during the practice match.

Press OK to end the practice match.

We have finished with the practice match. Any questions?

We start now with the second part of the session. You will now participate in 10 matches, each match paired with a different participant. In each match you will

interact with the same participant until a 3 or 4 appear. Remember: your decisions in one match cannot affect the decisions of the people you will interact with in future matches. This is not a practice; you will be paid!

[Begin Treatment 4 - Part 2 Paid.]

[Press Start Treatment, press Yes]

[Important: Tick Abort Current Match!!!!]

Make your choices now. Remember to press confirm.

[Wait until everybody has made a decision]

Monitor, would you please roll the dice?

[1) If 1 or 2 appear press NO] A ___ appeared. This match continues. You are still interacting with the same participant. [Important: Tick Abort Current Match!!!!]
Make your choices now. Remember to press confirm. Monitor, would you please roll the dice? [If 1 or 2 appeared go to 1). If 3 or 4 appeared go to 2)]

[2) If 3 or 4 appear press YES] A ___ appeared. This match ends. On the screen you will see a dialog box with the points you earned during this match. Press OK to be matched with the next participant.

.

This is the end of Part 2. On your screen you will see a dialog box indicating your point and dollar points for this part and your cumulative total points for the first two parts. Press OK to move to the next part.

Part 3

We will begin the third part now. This part will consist of 10 matches. In each match every Red participant is paired with a Blue participant. No pair will consist of the same participants as in Part 1 or 2. As before, you will not be paired twice with the same participant during the session or with a participant that was paired with someone that was paired with you or with someone that was paired with someone

that was paired with someone that was paired with you, and so on. Thus, the pairing is done in such a way that the decisions you make in one match cannot affect the decisions of the participants you will be paired with in later matches.

In this part, each pair will interact once. After that, a new match will start with different pairs. Therefore, you will interact once with 10 different participants.

But first, we are going to teach you about this part of the session and how to use the computer by going through one practice match. During the practice do not hit any keys until you are told to do so. You are not paid for the practice match; it is just for you to familiarize yourself with the session and the computer program.

[Begin Treatment 5 - Practice Part 3]

[Press Start Treatment, press Yes]

[Important: Tick Abort Current Match]

As before, your screen shows the possible actions you can choose, the actions the participant you are matched with can choose, and the points. You may choose your action by pressing the desired action at the side of the matrix now. Make your choices now. Once everyone in the room has made their selections and pressed confirm, your results from this round will appear on the screen.

[Wait until everybody has made a decision]

You have interacted once so this match ends. On the screen you will see a dialog box with the points you earned during the practice match. Press OK to end the practice match.

We have finished with the practice match. Any questions?

We start now with the third part of the session. You will now participate in 10 matches, each match paired with a different participant. In each match you will interact with the same participant once. Remember: your decisions in one match cannot affect the decisions of the people you will interact with in future matches.

This is not a practice; you will be paid!

[Begin Treatment 6 - Part 3 Paid.]

[Press Start Treatment, press Yes]

[Stage screen 1]

Make your choices now. Remember to press con...rm.

Press OK to be matched with the next participant.

[Press Start Match, Yes]

[Stage screen 2]

Make your choices now. Remember to press con...rm.

Press OK to be matched with the next participant.

[Press Start Match, Yes]

.

[Stage screen ____]

Make your choices now. Remember to press con...rm.

This is the end of Part 3. On your screen you will see a dialog box indicating your point and dollar points for this part and your cumulative total points for the three parts. Press OK to end this part.

Farewell

The session has ended. On your screen you will see a dialog box indicating your total earnings for the session. Please make sure you record the dollar points in your record sheet. Press OK to end the session. Take this sheet to the counter for payment. This sheet will be matched to our computer print out of results for payment. Your payments will be rounded up to the nearest quarter. Thank you for your participation.

References

- [1] Abreu, D. (1986). "Extremal Equilibria of Oligopolistic Supergames," *Journal of Economic Theory*, 39, 191-225.
- [2] Abreu, D., Dutta, P.K. and Smith, L. (1994). "The Folk Theorem for Repeated Games: A NEU Condition," *Econometrica*, 62, 939-948.
- [3] Abreu, D., Pearce, D. and Stachetti, E. (1990). "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica*, 58, 1041-1063.
- [4] Akerlof, G. (1976) "The Economics of Caste and of the Rat Race and other Woeful Tales," *Quart. J. Econ.*, 90, 599-617.
- [5] Andreoni, J. and Miller, J.H. (1993). "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence," *Economic Journal*, 103(418).
- [6] Aumann, R. and Shapley, L. (1976). "Long Term Competition: A Game Theoretic Analysis," mimeo, Hebrew University.
- [7] Axelrod, R. "The Evolution of Cooperation," Basic Books, 1984.
- [8] Bagwell, K. and Staiger, R.W. (1997). "Collusion over the business cycle," *RAND Journal of Economics*, 28, 82-106.
- [9] Baye, M.R. and Dennis W. Jansen, D.W. (1996). "Repeated Games with Stochastic Discounting," *Economica*, 63, 531-541.
- [10] Borenstein, S. and Shepard, A. (1996). "Dynamic pricing in retail gasoline markets," *RAND Journal of Economics*, 27, 429-451.

- [11] Bresnahan, T.F. (1989). "Empirical Studies of Industries with Market Power." In R. Schmalensee and R.D. Willig, eds., *Handbook of Industrial Organization*, Vol. II. North Holland, Amsterdam.
- [12] Brock, W.A. and Scheinkman, J.A. (1985). "Price Setting Supergames with Capacity Constraints," *Review of Economic Studies*, 52, 371-382.
- [13] Brown Kruse, J., Rassenti, S., Reynolds, S.S. and Smith, V.L. (1994). "Bertrand-Edgeworth Competition in Experimental Markets," *Econometrica*, 62(2).
- [14] Chevalier, J.A. and Scharfstein, D.S. (1996). "Capital-Market Imperfections and Countercyclical Markups: Theory and Evidence," *American Economic Review*, 86, 703-725.
- [15] Cooper, R., DeJong, D.V., Forsythe, R. and Ross, T.W. (1996). "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games," *Games and Economic Behavior*, 12(2).
- [16] Domowitz, I., Hubbard, G.R. and Petersen, B.C. (1986). "Business cycles and the relationship between concentration and price-cost margins," *RAND Journal of Economics*, 17, 1-17.
- [17] Ellison, G. (1994). "Theories of Cartel Stability and the Joint Executive Committee," *RAND Journal of Economics*, 25, 37-57.
- [18] Feinberg, R.M. and Husted, T.A. (1993). "An Experimental Test of Discount-Rate Effects on Collusive Behaviour in Duopoly Markets," *The Journal of Industrial Economics*, 41(2).
- [19] Friedman, J.W. (1971). "A Non-Cooperative Equilibrium for Supergames," *Review of Economic Studies*, 38, 1-12.

- [20] Fouraker, L.E. and Siegel, S. (1963). *Bargaining behavior*. New York, McGraw-Hill.
- [21] Fudenberg, D. and Levine, D.K. (1997). "Measuring Players' Losses in Experimental Games," *Quarterly Journal of Economics*, 112(2), 507-36.
- [22] Fudenberg, D., Levine, D.K. and Maskin, E. (1994). "The Folk Theorem with Imperfect Information," *Econometrica*, 62, 997-1039.
- [23] Fudenberg, D. and Maskin, E. (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54, 533-554.
- [24] Green, E.J. and Porter, R.H. (1984). "Noncooperative Collusion Under Imperfect Price Competition," *Econometrica*, 52, 87-100.
- [25] A. Greif, A. (1993). "Contract Enforceability and Economic Institutions in Early Trade- the Maghribi traders coalition," *Amer. Econ. Rev.*, 83, 525-548.
- [26] Greif, A., Milgrom, P. and Weingast, B. (1994). "Coordination, Commitment, and Enforcement - the case of the merchant guild," *J. Polit. Economy*, 102, 745-776.
- [27] Gottfries, N. (1991). "Customer Markets, Credit Market Imperfections and Real Price Rigidity," *Economica*, 58, 317-323.
- [28] Haltiwanger, J. and Harrington, J.E. (1991). "The impact of cyclical demand movements on collusive behavior," *RAND Journal of Economics*, 22, 89-106.
- [29] Hirshleifer, J. and Riley, J.G. (1992). *The Analytics of Uncertainty and Information*. Cambridge University Press.
- [30] Holt, C.A. (1985). "An Experimental Test of the Consistent-Conjectures Hypothesis," *The American Economic Review*, 75(3), 314-325.

- [31] Holt, C.A. (1995) "Industrial Organization: A Survey of Laboratory Research" in Kagel and Roth [32].
- [32] Kagel, J.H. and Roth, A.E. editors (1995). *The Handbook of Experimental Economics*. Princeton University Press.
- [33] Kandori, M. (1991). "Correlated Demand Shocks and Price Wars During Booms," *Review of Economic Studies*, 58, 171-180.
- [34] M. Kandori, *Social Norms and Community Enforcement*, *Rev. Econ. Stud.*, 59 (1992), 63-80.
- [35] Klemperer, P. (1995). "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Review of Economic Studies*, 62, 515-539.
- [36] Kreps, D.M., Milgrom, P., Roberts, J. and Wilson, R. (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*; 27(2).
- [37] Kreps, D.M. and Wilson, R. (1982) "Sequential Equilibrium," *Econometrica*, 50 (1982), 863-894.
- [38] Kreps, D. and Scheinkman, J. (1983). "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics*, 14, 326-337.
- [39] Luce, R.D. and Raiffa, H. (1957). "Games and Decisions," Wiley, New York.
- [40] Murnighan, J.K. and Roth A.E. (1983). "Expecting Continued Play in Prisoner's Dilemma Games," *Journal of Conflict Resolution*, 27(2).
- [41] Okuno-Fujiwara, M. and Postlewaite, A. (1995). "Social Norms and Random Matching Games," *Games Econ. Behav.*, 9, 79-109.

- [42] Palfrey, T.R. and Rosenthal, H. (1994). "Repeated Play, Cooperation and Coordination: An Experimental Study," *Review of Economic Studies*, 61(3).
- [43] Phelps, E. and Winter, S. (1970). "Optimal Price Policy under Atomistic Competition," in E. Phelps, ed., *Microeconomic Foundations of Employment and Inflation Theory*. W.W. Norton, New York.
- [44] Porter, R.H. (1983). "Optimal Cartel Trigger Price Strategies," *Journal of Economic Theory*, 29, 313-338.
- [45] Porter, R.H. (1983). "A Study of Cartel Stability: the Joint Executive Committee, 1880-1886," *Bell Journal of Economics*, 14, 301-314.
- [46] Rotemberg, J.J. and Saloner, G. (1986). "A Supergame-Theoretic Model of Price Wars During Booms," *American Economic Review*, 76, 390-407.
- [47] Rotemberg, J.J. and Woodford, M. (1991). "Markups and the Business Cycle," in *NBER Macroeconomics Annual, Vol 6*, edited by O.J. Blanchard and S. Fischer, MIT Press, Cambridge, Mass.
- [48] Rotemberg, J.J. and Woodford, M. (1992). "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity," *Journal of Political Economy*, 100, 1153-1207.
- [49] Roth, A.E. (1995). "Introduction to Experimental Economics" in Kagel and Roth [32].
- [50] Roth A.E. and Murnighan, J.K. (1978). "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma," *Journal of Mathematical Psychology*, 17.
- [51] Slade, M.E. (1987). "Inter-firm Rivalry in a Repeated Game: an empirical test of tacit collusion," *The Journal of Industrial Economics*, 35, 499-516.

- [52] Stahl, D.O.II. (1991). "The Graph of Prisoner's Dilemma Supergame Payoffs as a Function of the Discount Factor," *Games and Economic Behavior*, 3(3).
- [53] Staiger, R.W. and Wolak, F.A. (1992). "Collusive pricing with capacity constraints in the presence of demand uncertainty," *RAND Journal of Economics*, 23, 203-220.
- Abreu, D., Pearce, D. and Stacchetti, E. (1990). "Towards a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5).