

Supporting Cooperation via Agreement Equilibrium*

Mehmet S. Ismail[†]

Ronald Peeters[‡]

October 17, 2019

Abstract

We introduce ‘agreement equilibrium’ as a novel solution concept that can explain the abundance of cooperative behavior that is often observed in laboratory experiments in various contexts. The main idea of the agreement equilibrium is to identify behaviors that individuals can (tacitly) agree on while being ambiguous about their opponents’ intentions to respect or to betray this (tacit) agreement. We investigate properties of the agreement equilibrium and illustrate the agreement equilibrium in a series of famous applications.

Keywords: noncooperative games; tacit cooperation; agreement equilibrium; social dilemmas.

JEL classification codes: C72.

*This paper, which is a novel contribution of the authors, partly builds on “Maximin equilibrium” by the first author. We are grateful to Jean-Jacques Herings for his valuable comments and suggestions. We also thank Matthew Embrey, János Flesch, Willemien Kets, Bezalel Peleg, and Arkadi Predtetchinski, as well as the seminar participants at Maastricht, Stony Brook (Game Theory Conference), São Paulo (Workshop on Game Theory), Rotterdam (FUR Conference), Paris School of Economics, and CEREC Workshop for their valuable comments on an early version.

[†]Department of Political Economy, King’s College London, UK. E-mail: mehmet.s.ismail@gmail.com. Corresponding author.

[‡]Department of Economics, University of Otago, New Zealand. E-mail: ronald.peeters@otago.ac.nz

1 Introduction

Experimental research has provided a wealth of evidence of individual behavior often being more cooperative than is predicted by Nash equilibrium. In particular, but not exclusively, in the context of oligopoly games (see, e.g., Dolbear, 1968; Holt, 1995) and finitely repeated social dilemma games (see, e.g., Ledyard, 1995) experimental subjects have shown to be able to tacitly coordinate on more cooperative outcomes, without having to rely on promises and threats or this requiring explicit communication. Given that excessive cooperation levels arise even when played among complete strangers, cooperative outcomes must have a very strong and natural appeal.

The goal of this paper is to introduce, to investigate and to illustrate a novel equilibrium concept, the *agreement equilibrium*, as an alternative to Nash equilibrium that predicts more cooperative behavior, while still being selective. One interpretation of Nash equilibrium has been that it is a self-enforcing agreement—a strategy profile with the property that if agreed upon there is no unilateral profitable deviation from it. We also consider a strategy profile as an agreement, and when it is not a Nash equilibrium some players’ best-response deviations would strictly improve their payoffs. Anticipating such deviations, we assume that players would be cautious when evaluating agreements. Roughly, a profile of strategies is an agreement equilibrium in case no player prefers agreeing on something else, given his or her uncertainty about the true intentions of the opponents to respect or to betray the agreement. Hence, agreement equilibria are strategy profiles where all players can mutually cautiously agree on conforming to.

The agreement equilibrium is not the first concept that incorporates cautious behavior in games. Perhaps the most well-known such concept is that of maximin strategy in which players maximize the minimum payoff irrespective of the opponent’s choice (von Neumann, 1928). Under Hurwicz’s (1951) criterion, players choose strategies that yield the highest payoff that is derived from a convex combination of the player’s minimum and maximum payoffs; hence, it leads to less pessimistic solutions than the maximin strategy. While both maximin strategy and Hurwicz’s criterion are non-equilibrium concepts, equilibrium concepts that incorporate cautiousness into games have been proposed as early as Selten (1975). The agreement equilibrium concept is related to the equilibrium concepts that model cautious

behavior under ambiguity, which have been a growing literature. Notably, Marinacci (2000) introduces an equilibrium concept with a parameter that captures the ambiguity aversion of players, which may rule out equilibria that involve “riskier” strategies. Other models of ambiguity in games include, but are not limited to, Dow and Werlang (1994), Klibanoff (1996), Lo (1996), Renou and Schlag (2010), and Riedel and Sass (2014). The essential difference between our approach and these approaches is that we do not aim to model (the degree of) ambiguity and instead focus on the worst-case scenario of agreements with a potential to produce non-Nash cooperative behavior in some non-cooperative games.

After having formally defined the agreement equilibrium concept, we illustrate the agreement equilibrium in a series of applications, including the traveler’s dilemma, differentiated Bertrand duopoly, differentiated Cournot duopoly, multi-player spatial competition à la Hotelling/Downs, and finitely repeated (n -player) social dilemma games. Finally, we conclude with suggesting two possible directions to extend this new equilibrium concept.

2 Agreement Equilibrium and its Properties

Let N be a finite set of players, each player $i \in N$ having a finite set of actions S_i , and let their payoffs resulting from their joint actions be given by the function $u_i : S \rightarrow \mathbb{R}$ with $S = \times_{i \in N} S_i$. We consider the mixed extension of this game and denote by $\Sigma_i = \Delta(S_i)$ the set of mixed actions of player $i \in N$ and extend their payoff functions u_i to the domain $\Sigma = \times_{i \in N} \Sigma_i$ multi-linearly. We denote this game by $\Gamma = (N, \Sigma, u)$. Although we use the mixed extension framework for the sake of completeness, we focus on pure strategies in our applications.

Suppose that, prior to playing this game in all independence, players can make an agreement (explicitly or implicitly) on how they are going to play this game; i.e., to agree on a strategy profile $\alpha \in \Sigma$. Then, when actually playing the game, all players have the choice to play in accordance to their agreement, or to deviate from it. In the process of making the agreement, players are aware of their agreement being non-binding in that other players, possibly multiple of them in an unorchestrated manner, are able to betray the agreement. We are interested in the strategy profiles that all players are able to safely agree on, in the sense that it guarantees a ‘reasonable’ payoff under the possibility of any of the other players betraying the agreement.

In the *agreement game* of the original game Γ , the payoffs to the players at each possible agreement $\alpha \in \Sigma$ are the players' worst-case payoffs of the original game under the possibility of any combination of opponent players to deviate to their unilateral optimal deviations from the agreement α .

Definition 1 (agreement game). *The agreement game of the game Γ , denoted $g(\Gamma)$, is the game Γ with for each player $i \in N$ the payoff function u_i replaced by the function g_i defined as*

$$g_i(\alpha) = \min_{\sigma_{-i} \in \tilde{\Sigma}_{-i}(\alpha)} u_i(\alpha_i, \sigma_{-i}),$$

where $\tilde{\Sigma}_{-i}(\alpha) = \times_{j \in N \setminus \{i\}} \tilde{\Sigma}_j(\alpha)$ with for each player $j \in N \setminus \{i\}$ the set $\tilde{\Sigma}_j(\alpha)$ defined as $\tilde{\Sigma}_j(\alpha) = \{\alpha_j\} \cup \{\sigma_j \in \Sigma_j \mid u_j(\sigma_j, \alpha_{-j}) \geq u_j(\sigma'_j, \alpha_{-j}) \text{ for all } \sigma'_j \in \Sigma_j \text{ and } u_j(\sigma_j, \alpha_{-j}) > u_j(\alpha)\}$.

Given an agreement $\alpha \in \Sigma$, for each player $j \in N$, the set $\tilde{\Sigma}_j(\alpha)$ comprises the possibility to stick to the agreement (i.e., play strategy α_j) or to best respond given that the others stick to the agreement, provided that the best response leads to a strict improvement. Next, for each player $i \in N$, the set $\tilde{\Sigma}_{-i}(\alpha)$ is the set of strategy profiles by all other players $j \in N \setminus \{i\}$, where each opponent player $j \in N \setminus \{i\}$ either sticks to the agreement, or best responds to the agreement α holding the other players' strategies, α_{-j} fixed. The payoff $g_i(\alpha)$ is player i 's most pessimistic payoff expectation when evaluating the agreement $\alpha \in \Sigma$ under the constraint that the other players play according to $\tilde{\Sigma}_{-i}(\alpha)$. Note that for each $\alpha \in \Sigma$ and player $i \in N$, this value $g_i(\alpha)$ is uniquely defined. In particular, relevant for games with three or more players, the set $\tilde{\Sigma}_{-i}(\alpha)$ does not allow for optimal coalitional deviations by a non-singleton subset of the opponent players. That being said, there are n -person games in which correlation in beliefs can be reasonable—e.g., player i may believe that other players display similar or correlated behavior even if they actually act independently. Under such circumstances, the modeler may modify the set $\tilde{\Sigma}_{-i}(\alpha)$ to allow such correlations. What may happen when taking also such deviations into account will be shortly addressed in Subsection 3.4.

To illustrate the derivation of the agreement game $g(\Gamma)$, consider the left game in Figure 1 as the original game Γ . This game Γ has a unique Nash equilibrium with both players playing strategy C . The situation where both play strategy A does not constitute a Nash equilibrium, but would be an outcome players could agree on playing. A player who sticks to this agreed-on strategy would only be “disappointed” in case the other player would move to a strategy

that involves an irrational deviation, C , from the agreement (A, A) . The only deviation to consider is the opponent best responding with B ; so, $\tilde{\Sigma}_{-i}(A, A) = \{A, B\}$. Regardless of the opponent sticking to the agreement or optimally deviating from it, players safely expect to receive 100 by playing the agreed-on strategy A ; so, $g_i(A, A) = 100$.

	A	B	C
A	100 , 100	100 , 105	0 , 0
B	105 , 100	95 , 95	0 , 210
C	0 , 0	210 , 0	5 , 5

	A	B	C
A	100 , 100	100 , 0	0 , 0
B	0 , 100	0 , 0	0 , 5
C	0 , 0	5 , 0	5 , 5

Figure 1: Illustrative example with the original game to the left and its corresponding agreement game to the right. Both players playing strategy A constitutes an agreement equilibrium. The other agreement equilibrium is (C, C) , which is the unique Nash equilibrium of the original game.

In contrast, an agreement for both to play strategy B and earn a payoff of 95 would, under consideration of rational deviations, only guarantee a payoff of 0 to both players, since both players are aware that the other player may betray the agreement by playing C ; so, $g_i(B, B) = 0$. The corresponding agreement game on the right of Figure 1 displays the payoffs for these two possible agreements, and all other seven possible agreements.

To explain one more example, consider the agreement on playing (B, A) . The first player's worst-case agreement payoff is 0 because her opponent may rationally deviate to C ; i.e. $\tilde{\Sigma}_2(B, A) = \{A, C\}$ and $g_1(B, A) = \min\{u_1(B, A), u_1(B, C)\} = \min\{105, 0\} = 0$. Likewise the second player's worst-case agreement payoff is 100, because the first player is already at her best response in the profile (B, A) such that $\tilde{\Sigma}_1(B, A) = \{B\}$.

To draw the parallel to the level- k model (Stahl, 1993), if we were to define level-0 behavior as complying to the agreement α , then the set $\tilde{\Sigma}_i(\alpha)$ comprises player i 's level-0 strategy and her level-1 strategies, subject to a status quo bias, in the sense that a best response is only used if it induces a strict improvement. To draw analogy to the level- k reasoning, we consider best-response deviations from agreements; though, stricter or weaker deviations can be considered as suggested in Ismail (2014a,b). Next, each player i evaluates the agreement as the payoff the strategy α_i guarantees, $g_i(\alpha)$, when being unsure whether other players are to defect to a best response strategy (i.e., when being uncertain whether they are level-0 or level-1 types). Hence, players are assumed to make a cautious evaluation of the agreement while being ambiguous about opponent players being opportunistic or not (Gilboa and Schmeidler, 1989).

We make use of the conversion to the agreement game and the Nash equilibrium concept in order to define the *agreement equilibrium* of a given game. In short, a strategy profile is an agreement equilibrium if all players prefer agreeing on their part of this strategy profile given that others agree on their part while knowing that each opponent player may breach the agreement by best responding to others respecting the agreement. That is, an agreement equilibrium is a strategy profile where all players can mutually cautiously agree on conforming to.

Definition 2 (agreement equilibrium). *A strategy profile $\alpha \in \Sigma$ is an agreement equilibrium of the game Γ if it is a Nash equilibrium of the agreement game $g(\Gamma)$. That is, $\alpha \in \Sigma$ is an agreement equilibrium if for every player i , $\alpha_i \in \arg \max_{\alpha'_i \in \Sigma_i} g_i(\alpha'_i, \alpha_{-i})$, or expressed in the primitives of the original game,*

$$\alpha_i \in \arg \max_{\alpha'_i \in \Sigma_i} \min_{\sigma_{-i} \in \tilde{\Sigma}_{-i}(\alpha'_i, \alpha_{-i})} u_i(\alpha'_i, \sigma_{-i}).$$

As is clear from the definition, the agreement equilibrium has a ‘maximin’ flavor. Von Neumann’s (1928) maximin strategy is based on maximizing a minimum payoff irrespective of the opponents’ choices, whereas under Hurwicz’ criterion a player chooses a strategy that yields the highest payoff that is derived from a convex combination of the player’s minimum and maximum payoffs. Agreement equilibrium is a strategy profile from which no player can unilaterally deviate and increase his or her worst-case payoff in the agreement game.

To illustrate the derivation of the agreement equilibrium, we go back to the left game in Figure 1. The agreement equilibria of this game are precisely the Nash equilibria of the corresponding agreement game, which is the right game in the figure. We find that, in addition to both players playing strategy C , also the situation in which both players play strategy A constitutes an agreement equilibrium. The agreement equilibrium (A, A) is not a Nash equilibrium and, indeed, each player has an incentive to deviate to strategy B . But the worst-case payoff that results if a player deviates (i.e., non-deviator receiving 100) is as great as possible among all potential agreements. No matter whether the opponent (or, in general, any of the opponents) respects the agreement or breaks it by best responding, respecting the agreement guarantees a good payoff (given the agreement of others).

Proposition 1. *Nash equilibria are agreement equilibria.*

Proof. Let $\sigma \in \Sigma$ be a Nash equilibrium of the original game. Then, for each player $i \in N$, it holds that $\tilde{\Sigma}_i(\sigma) = \{\sigma_i\}$ and hence that $g_i(\sigma) = u_i(\sigma)$. Since σ is a Nash equilibrium we know that for all $i \in N$, and all of his alternative strategies $\sigma'_i \neq \sigma_i$ it holds that $u_i(\sigma'_i, \sigma_{-i}) \leq u_i(\sigma)$. By construction of the agreement game we know that $g_i(\sigma'_i, \sigma_{-i}) \leq u_i(\sigma'_i, \sigma_{-i})$, since the worst payoff is at most the amount the player gets when all respect the agreement. So, for all players $i \in N$ and all his strategies $\sigma'_i \in \Sigma_i$ we have that $g_i(\sigma'_i, \sigma_{-i}) \leq g_i(\sigma)$, implying that σ is an equilibrium of the agreement game and thus, by definition of it, an agreement equilibrium. \square

Nash equilibria being agreement equilibria implies existence and possible multiplicity. Moreover, in the illustrative example we have already seen that there may exist agreement equilibria (both playing strategy A) that are not Nash equilibria. Therefore, agreement equilibrium is a coarsening of Nash equilibrium.¹ However, the next two propositions show that the two concepts are identical in the special case where every player has a dominant strategy and in case of a two-player zero-sum game. One immediate corollary of the latter is that an agreement equilibrium is not guaranteed to exist in pure strategies.

Proposition 2. *If in a given game every player has a strictly dominant strategy, then the strategy profile where all players play their dominant strategy is the unique agreement equilibrium.*

Proof. Let for each player i , $\tau_i \in \Sigma_i$ denote her strictly dominant strategy. The strategy profile composed of strictly dominant strategies, $\tau = (\tau_i)_{i \in N}$, being an agreement equilibrium follows from Proposition 1. It remains to be shown that it is the unique agreement equilibrium. Suppose, by contraposition, that there exists an agreement equilibrium $\sigma \neq \tau$. We will show that all players i with $\sigma_i \neq \tau_i$ have an incentive to deviate to their strictly dominant strategy in the agreement game. Take such a player i . Since $u_i(\tau_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$ and $g_i(\sigma_i, \sigma_{-i}) \geq g_i(\tau_i, \sigma_{-i})$, there must exist a strategy profile with some of the others playing their strictly dominant strategy ς_{-i} such that $g_i(\tau_i, \sigma_{-i}) = u_i(\tau_i, \varsigma_{-i})$. However, ς_{-i} is also a valid deviation when considering the strategy profile σ , implying $g_i(\sigma_i, \sigma_{-i}) \leq u_i(\sigma_i, \varsigma_{-i})$. Combining all inequalities gives

$$u_i(\tau_i, \varsigma_{-i}) = g_i(\tau_i, \sigma_{-i}) \leq g_i(\sigma_i, \sigma_{-i}) \leq u_i(\sigma_i, \varsigma_{-i}) < u_i(\tau_i, \varsigma_{-i});$$

¹It is easy to construct examples showing that agreement equilibrium is not logically related to correlated equilibrium (Aumann, 1974) and rationalizability (Bernheim, 1984; Pearce, 1984) as two other well-known coarsenings of Nash equilibrium.

a contradiction. □

Proposition 3. *In two-player zero-sum games, a strategy profile is an agreement equilibrium if and only if it is a Nash equilibrium.*

Proof. The ‘if’ part follows from Proposition 1. Left to show is the ‘only if’ part. Let a two-player zero-sum game Γ be given and let (v_1, v_2) be the corresponding values (i.e., the payoffs that the players’ agreement equilibrium strategies guarantees them). Moreover, let an agreement equilibrium $\sigma \in \Sigma$ be given. Suppose, by contraposition, that it is not a Nash equilibrium. First, suppose that the payoffs related to σ do not equal the values. Then, for one of the players, $u_i(\sigma) < v_i$, such that also $g_i(\sigma) < v_i$. Since, player i can guarantee a payoff of v_i in both the original game and the corresponding agreement game, the profile σ cannot be a Nash equilibrium of the agreement game; a contradiction. Now, suppose that the payoffs related to σ do equal the values. Since σ is not a Nash equilibrium, at least one of the players i has a strategy σ'_i such that $u_i(\sigma'_i, \sigma_{-i}) > v_i$. But then, for the other player, $j \neq i$, $g_j(\sigma) < v_j$. This contradicts σ being an agreement equilibrium (i.e., a Nash equilibrium of the agreement game), since player j has a strategy that guarantees a payoff of v_j in the agreement game. □

3 Applications

In this section we present the agreement equilibrium in the context of several interesting and important applications. First, we illustrate that the agreement equilibrium explains experimental data pretty well for the traveler’s dilemma. We continue with a differentiated goods Bertrand duopoly model and show that there exists a symmetric semi-collusive agreement equilibrium. Next, for the equivalent Cournot model we show that the Nash equilibrium is the only symmetric agreement equilibrium, and that there exists a continuum of asymmetric agreement equilibria that includes the Stackelberg equilibria. These latter two applications, for which our results reiterate Melkonyan et al.’s (2018) calculations based on a method by Ismail (2014a,b), illustrate the use of the agreement equilibrium concept in games with continuous actions. Next, we move to the location model of Hotelling/Downs. While played among three players this game is known not to produce a pure strategy Nash equilibrium, it does produce an (efficient) agreement equilibrium in pure strategies. Finally, we show that

agreement equilibrium (like Nash equilibrium) cannot explain cooperation in social dilemmas when played one-shot, but (unlike Nash equilibrium) it can support cooperation when played finitely repeated.

3.1 Traveler’s dilemma

Figure 2 presents the traveler’s dilemma as introduced by Basu (1994). Two players report a number between 2 and 100 and receive a payoff depending on both reported numbers. If both report the same number they both get the dollar amount equal to their reports. If they report different numbers, the one reporting the lower number receives an additional $z > 1$ dollar on top of the dollar amount of her report, while the other receives a z dollar discount on the dollar amount of the other’s report. In the game displayed in the figure, z equals 2.

	100	99	98	...	4	3	2
100	100, 100	97, 101	96, 100	...	2, 6	1, 5	0, 4
99	101, 97	99, 99	96, 100	...	2, 6	1, 5	0, 4
98	100, 96	100, 96	1, 5	0, 4
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
4	6, 2	6, 2	1, 5	0, 4
3	5, 1	5, 1	5, 1	...	5, 1	3, 3	0, 4
2	4, 0	4, 0	4, 0	...	4, 0	4, 0	2, 2

Figure 2: Traveler’s dilemma.

While regardless of the value of z , the unique Nash equilibrium is for both players to report the number 2, which is also the only strategy that survives the process of iterated elimination of strictly dominated strategies, data gathered in laboratory experiments have revealed that individuals do not behave in accordance to this Nash equilibrium prediction and that their reports are not invariant to the value of the parameter z . Goeree and Holt (2001) found that when the value of z is high, 80% of the individuals choose the Nash equilibrium strategy, but when the value of z is small, about the same percent of the individuals choose the highest number. This finding confirms the findings of Capra et al. (1999) where over time play converged towards Nash equilibrium for high values of z and to the other extreme for low values of z . Finally, in a web-based experiment without payments and low value of z , Rubinstein (2007) found 55% of 2,985 individuals choosing the highest amount and only 13% choosing the Nash equilibrium strategy.

Figure 3 presents the agreement game corresponding to the traveler’s dilemma game in

Figure 2. This resulting game possesses two (pure strategy) Nash equilibria, which are the agreement equilibria of the original traveler’s dilemma game. In one agreement equilibrium both payers report the lowest number; in the other both players report the highest number. Although the figures only illustrate this for the situation $z = 2$, in fact, these are the agreement equilibria for all values of $z > 1$. We may conclude that agreement equilibrium can explain experimental findings. The key property of the agreement equilibrium driving the result is that if players were to agree on both reporting the highest number and earning 100 dollar, being betrayed leads to a payoff of 97, which is a reasonably high worst-case payoff and worth settling on such an agreement—a reasoning process that individuals may have been following in the reported experiments.

	100	99	98	...	4	3	2
100	97, 97	97, 96	96, 95	...	2, 1	1, 0	0, 2
99	96, 97	96, 96	96, 95	...	2, 1	1, 0	0, 2
98	95, 96	95, 96	1, 0	0, 2
...
4	1, 2	1, 2	1, 0	0, 2
3	0, 1	0, 1	0, 1	...	0, 1	0, 0	0, 2
2	2, 0	2, 0	2, 0	...	2, 0	2, 0	2, 2

Figure 3: Agreement game of the traveler’s dilemma.

3.2 Bertrand competition

Consider the symmetric Bertrand game in which two firms produce with zero cost and firm i ’s demand is given by $q_i = a - bp_i + dp_j$ with prices p_i and p_j as strategic variables. The best response of firm j to the price of firm i is given by $p_j^*(p_i) = \frac{a+dp_i}{2b}$, and the two best response functions intersect in the Bertrand-Nash equilibrium with both firms setting their price at $p^* = \frac{a}{2b-d}$. We are interested if there exist additional symmetric agreement equilibria; in particular, agreement equilibria where firms display collusive behavior.

Firm i ’s payoff at the price pair (p_i, p_j) in the agreement game is given by $g_i(p_i, p_j) = \min\{\pi_i(p_i, p_j), \pi_i(p_i, p_j^*(p_i))\}$. Since firm i ’s profit is increasing in firm j ’s price, which of the two profit levels is larger depends on the location of the point (p_i, p_j) relative to the best response curve of firm j . If $p_j > p_j^*(p_i)$, then $g_i(p_i, p_j) = \pi_i(p_i, p_j^*(p_i))$, and $g_i(p_i, p_j) = \pi_i(p_i, p_j)$ otherwise. Figure 4 illustrates how a firm’s isoprofit curves, and consequently its best-response curve, changes when moving from the original game to the agreement game.

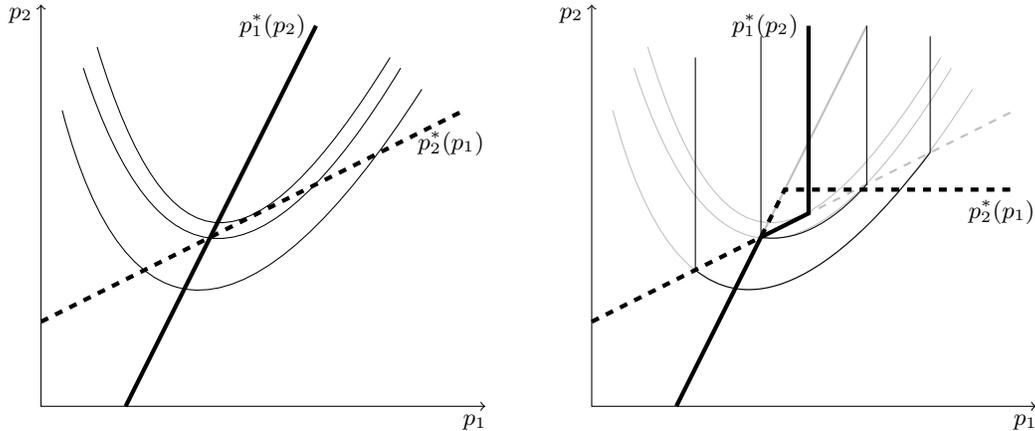


Figure 4: Isoprofit curves of firm 1 and both firms' best response curves in the Bertrand game (left) and in the corresponding agreement game (right).

We see that the two best response curves intersect, in addition to the Nash equilibrium, in the point where both firms set prices as Stackelberg leaders; that is, each firm i sets the price $p^a = \frac{2ab+ad}{4b^2-2d^2}$ that maximizes $\pi_i(p_i, p_j^*(p_i))$. In this additional agreement firms set supra-competitive prices. Since the price is below the joint profit maximizing price $p^m = \frac{a}{2b-2d}$, we can conclude that the agreement equilibrium just derived is semi-collusive.

3.3 Cournot competition

Consider the symmetric Cournot duopoly with firms producing at zero cost and firm i 's inverse demand is given by $p_i = a - bq_i - dq_j$. Similar to Figure 4 for the Bertrand game, Figure 5 shows how the isoprofit curves and best-response curves change when moving from the original game to the agreement game.

Unlike in the Bertrand game we do not find a symmetric agreement equilibrium in addition to the Nash equilibrium. However, now we find a continuum of asymmetric agreement equilibria that connects the two Stackelberg equilibria when tracking the best response curves via the Nash equilibrium. A typical feature of these asymmetric agreement equilibria is that, relative to the Nash equilibrium, they yield a larger profit to the firm with the larger output and a lower profit to the firm with the lower output, and a lower joint profit.

Interestingly, symmetric agreement equilibrium in comparison to Nash equilibrium being more collusive in Bertrand and not in Cournot is consistent with experimental findings, as is shown in the meta-studies by Engel (2007) and Suetens and Potters (2007).

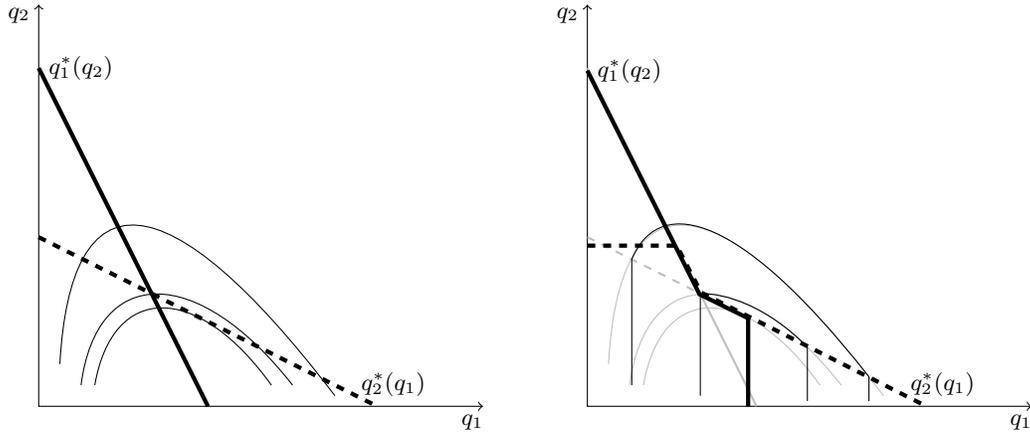


Figure 5: Isoprofit curves of firm 1 and both firms' best response curves in the Cournot game (left) and in the corresponding agreement game (right).

3.4 Hotelling/Downsian model of spatial/political competition

Consider Rubinstein's (2016) variation of the Hotelling/Downs game in which there are three ice cream vendors selling ice cones to sunbathers on the beach and seven feasible spots where vendors can park their carts. A volume of 1 potential ice eaters is evenly distributed between each pair of adjacent locations. The total volume of 6 potential ice eaters all will purchase their ice cones at the nearest ice cream cart. For example, in case two vendors are located at location C and one at location F , as illustrated in Figure 6, the two vendors at location C will each sell half of a 3.5 volume of ice cones and the vendor at location F sells to the remaining 2.5 volume of customers.

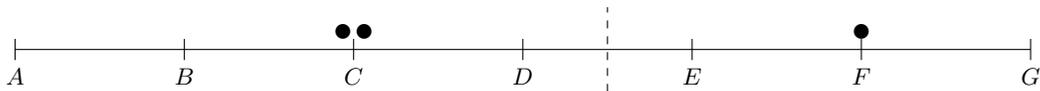


Figure 6: Hotelling/Downs game. The seven possible locations. In this example two vendors are at location C and one at location F ; the customers to the left of the dashed line will go to location C and those to the right of the dashed line to location F .

We argue that there is no pure strategy Nash equilibrium. First, locations A and G are never optimal given any pair of location choices of the other two vendors. Second, given that none of these two locations are used by any of the vendors, the choice of location B or F can only be rationalized by the two other vendors being at locations C and E . It is easily seen that one of these other vendors would be willing to move its cart to location D . Out of the remaining three locations, location D can only be rationalized by the other two vendors both

being situated either at location C or at location E . Again, it is easily seen that both these vendors prefer to move to the location just beyond location D . Finally, in case only locations C and E are used, all vendors want to move to location D if they are all at one location, and in case they are not all on one location, the vendor that does not share its location has an incentive to move to location D .

Interestingly, this game produces an agreement equilibrium in pure strategies with the three vendors being located at locations B , D and F and all vendors selling an equal amount of ice cones. The vendor located in the middle is at its best response and has no incentive to deviate from the agreement. The other two vendors have an incentive to betray by shifting their ice cream carts one location more to the middle. The worst-case situation for the middle vendor is obtained if both outer vendors move simultaneously to their best response locations. Still, given that the other vendors agree to locate at B and F , there is no other location that this middle vendor can better agree on. Furthermore, the other two vendors are guaranteed to sell a volume of 2 ice cones in the considered agreement. Observe that the vendor at location B , does not want to change to location C in the agreement game, since it only guarantees selling a volume of 1.5 ice cones, since in the worst case the other two vendors both effectuate their best responses of moving one location nearer. Similarly, the vendor at F has no incentive to unilaterally deviate to E in the agreement game.

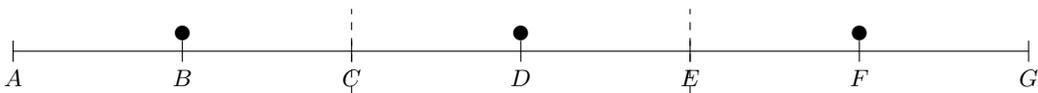


Figure 7: Illustration of the agreement equilibrium in the Hotelling/Downs game.

Figure 7 illustrates the agreement equilibrium just derived. Interestingly, in the agreement equilibrium presented, all vendors make equal numbers of sales. Furthermore, it is efficient in walking distance: none of the sunbathers has to pass by an unused location when walking to the nearest vendor.

Notably, the agreement equilibrium described in Figure 7 is robust to “correlated” individual profitable deviations by others. Suppose that the vendor at location B considers the possibility that both other vendors jointly deviate to location C , as a result of which this vendor’s sales drop from 2 to 1.5. Analogously, also the sales of the vendor at location F would drop from 2 to 1.5 if the vendors at locations B and D would jointly deviate to location

E . Likewise, an orchestrated move of the vendors at locations B and F to locations C and E would decrease the sales of vendor D from 2 to 1.5. Since the volume of 1.5 sales is also obtained in the Nash equilibrium of the agreement game for each of the vendors, none of the vendors can improve when considering correlated deviations.

As is argued in Rubinstein (2016), this location game has a unique symmetric Nash equilibrium in mixed strategies, where vendors locate at the middle location, D , with probability 20% and at each of the adjacent locations, C and E , with probability 40%. In his experiment he finds that of the 8,329 subjects, 43% chose the middle position. While the fraction of subjects choosing the middle position is double the amount expected if subjects were to randomize in accordance to the Nash equilibrium strategy, this fraction is just about a third more than expected if subjects would play the identified agreement equilibrium.

3.5 Social dilemmas

3.5.1 One-shot n -person public goods game

Consider the public goods game with n citizens simultaneously deciding how much of $\bar{x} > 0$ resources to contribute to the public good, knowing that all citizens will experience a benefit of μ from each unit contributed. We assume $\mu < 1 < n\mu$ such that the social benefit of each unit contributed exceeds the social cost ($n\mu > 1$) while the private costs exceed the private benefits ($1 < \mu$), as a result of which it is a dominant strategy for the citizens not to contribute.

Given a profile of contributions (x_i, x_{-i}) , the payoff to citizen i in this game is given by $u_i(x_i, x_{-i}) = -x_i + \mu \sum_{j=1}^n x_j$. Since the best response to any agreement (x_i, x_{-i}) is for all other citizens to contribute nothing, the payoff to citizen i in the agreement game is given by $g_i(x_i, x_{-i}) = \min\{-x_i + \mu \sum_{j=1}^n x_j, -(1 - \mu)x_i\} = -(1 - \mu)x_i \leq 0$, which is for all x_{-i} maximized by setting $x_i = 0$. We can conclude, but this also follows directly from Proposition 2, that there is a unique agreement equilibrium in which none of the citizens contributes anything.

3.5.2 Repeated prisoner's dilemma

Consider the prisoner's dilemma game in Figure 8 with $b > a > d > c$ and let the two players play this game repeatedly over the course of $k \geq 2$ periods. While there is a unique (subgame

perfect) Nash equilibrium in which both player defect in all periods, cooperative behavior can be supported as a agreement equilibrium using grim-trigger strategies.

	<i>C</i>	<i>D</i>
<i>C</i>	<i>a, a</i>	<i>c, b</i>
<i>D</i>	<i>b, c</i>	<i>d, d</i>

Figure 8: Prisoner's dilemma.

Let (α, α) be the strategy profile where both players play grim-trigger, giving each player a payoff of $u(\alpha, \alpha) = k \cdot a$. The best response of the opponent to this agreement is to cooperate for $k - 1$ periods and to defect in the last period. The payoff of each player in the agreement game is therefore $g(\alpha, \alpha) = (k - 1) \cdot a + 1 \cdot c$. Now, consider the possibility of a player deviating to some $\alpha' \neq \alpha$ in the agreement game. First, let $2 \leq \ell < k$ be the first time this player defects when playing in accordance to α' . Then $u(\alpha', \alpha) = (\ell - 1) \cdot a + 1 \cdot b + (k - \ell) \cdot d$. Since the best response of the other player to α' is to defect one period earlier, the respective payoff in the agreement game $g(\alpha', \alpha)$ is at most $(\ell - 2) \cdot a + 1 \cdot c + (k - \ell - 1) \cdot d$. Second, let the deviating player already defect in the first period when playing according to α' . Then, $u(\alpha', \alpha) = 1 \cdot b + (k - 1) \cdot d$ and $g(\alpha', \alpha)$ is at most $k \cdot d$. The profile (α, α) is an agreement equilibrium if $(k - 1) \cdot a + 1 \cdot c$ is not less than $(\ell - 2) \cdot a + 1 \cdot c + (k - \ell - 1) \cdot d$ for all $\ell \geq 2$ and $k \cdot d$. The former comparison is always satisfied by $a > c$; the latter requires the parameters to satisfy $(k - 1)a + c \geq kd$. Hence, cooperation can be sustained in a finitely repeated prisoner's dilemma provided that efficiency gains of cooperation are sufficiently large.

In general, like the Nash equilibrium concept, the agreement equilibrium concept is not able to explain cooperative behavior in social dilemmas, when played only once. However, as we have seen, while cooperation is also known not to be possible in (subgame perfect) Nash equilibrium when a social dilemma game is played finitely repeated, this is possible in agreement equilibrium. Also here agreement equilibrium is compatible with empirical findings using data retrieved in finitely repeated prisoner's dilemma experiments, including Andreoni and Miller (1993).

4 Conclusion

To shortly summarize, in this paper we propose the agreement equilibrium as an alternative solution concept for noncooperative games. Profiles of strategies are agreement equilibria

if players can safely agree on playing in accordance to these profiles in the sense that the potential harm of betrayal is minimized. Agreement equilibrium is a coarsening of the Nash equilibrium and may allow players to coordinate on more cooperative outcomes, as is illustrated in the traveler's dilemma, differentiated Bertrand and finitely repeated prisoner's dilemma.

We view the extension of the agreement equilibrium, as defined here, interesting in two mutually non-exclusive directions. The first concerns the consideration of better response deviations by some or all of the opponents, which will generally result to a further decrease of the payoffs in the agreement game, reflecting players being even more cautious when tacitly coordinating on an agreement. The second direction relates to the interpretation of the agreement game payoffs as players being ambiguous about the opponents being of level-0 or level-1, where level-0 types hold on to the agreement. Extending this ambiguity to higher cognitive levels would, like the consideration of better response deviations, lead to players deem more possible deviations by the opponents likely and more cautious agreement behavior. It is not obvious how these extensions would affect predictions and whether the effect may vary across game types.

References

- Andreoni, J. and J. H. Miller (1993). Rational cooperation in the finitely repeated prisoner's dilemma: Experimental evidence. *The Economic Journal* 103(418), 570–585.
- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1(1), 67–96.
- Basu, K. (1994). The traveler's dilemma: Paradoxes of rationality in game theory. *The American Economic Review* 84(2), 391–395.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica* 52(4), 1007–1028.
- Capra, C. M., J. K. Goeree, R. Gomez, and C. A. Holt (1999). Anomalous behavior in a Traveler's Dilemma? *American Economic Review* 89(3), 678–690.
- Dolbear, F. T., L. B. Lave, G. Bowman, A. Lieberman, E. Prescott, F. Rueter, and R. Sherman (1968). Collusion in oligopoly: An experiment on the effect of numbers and information. *The Quarterly Journal of Economics* 82(2), 240–259.
- Dow, J. and S. R. d. C. Werlang (1994). Nash equilibrium under Knightian uncertainty: Breaking down backward induction. *Journal of Economic Theory* 64(2), 305–324.
- Engel, C. (2007). How much collusion? A meta-analysis of oligopoly experiments. *Journal of Competition Law and Economics* 3(4), 491–549.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Goeree, J. K. and C. A. Holt (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91(5), 1402–1422.
- Hurwicz, L. (1951). The generalized Bayes minimax principle: A criterion for decision making under uncertainty. *Cowles Commission Discussion Paper: Statistics No. 335*, 1–7.
- Ismail, M. (2014a). Maximin equilibrium. Unpublished manuscript.
- Ismail, M. (2014b). Maximin equilibrium. GSBE Research Memoranda; No. 037. Maastricht: GSBE.

- Klibanoff, P. (1996). Uncertainty, decision, and normal-form games. Unpublished manuscript.
- Ledyard, J. O. (1994). Public goods: A survey of experimental research. In J. Kagel and A. Roth (Eds.), *Handbook of Experimental Economics*. Princeton University Press, Princeton.
- Lo, K. C. (1996). Equilibrium in beliefs under uncertainty. *Journal of Economic Theory* 71(2), 443–484.
- Marinacci, M. (2000). Ambiguous games. *Games and Economic Behavior* 31(2), 191–219.
- Melkonyan, T., H. Zeitoun, and N. Chater (2018). Collusion in Bertrand vs. Cournot competition: A virtual bargaining approach. *Management Science* 64(12), 5599–5609.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52(4), 1029–1050.
- Renou, L. and K. H. Schlag (2010). Minimax regret and strategic uncertainty. *Journal of Economic Theory* 145(1), 264–286.
- Riedel, F. and L. Sass (2014). Ellsberg games. *Theory and Decision* 76(4), 469–509.
- Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times. *The Economic Journal* 117(523), 1243–1259.
- Rubinstein, A. (2016). A typology of players: Between instinctive and contemplative. *The Quarterly Journal of Economics* 131(2), 859–890.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4(1), 25–55.
- Stahl, D. O. (1993). Evolution of smart_n players. *Games and Economic Behavior* 5(4), 604–617.
- Suetens, S. and J. Potters (2007). Bertrand colludes more than Cournot. *Experimental Economics* 10(1), 71–77.
- von Neumann, J. (1928). Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* 100, 295–320.