

Invitation Games and the Price of Stability

Uriel Feige^{*}
Department of Computer Science and Applied
Mathematics
Weizmann Institute
Rehovot, Israel
uriel.feige@weizmann.ac.il

Moshe Tennenholtz
Microsoft Research, Herzliya
and Technion, Haifa
moshet@microsoft.com

ABSTRACT

Given an arbitrary 2-player game G that we refer to as the basic game, we propose a notion of a multiplayer *invitation game* that proceeds for a fixed number of rounds, where in each round some player (whose identity is determined by a scheduler) gets to invite a player of his choice to play a match of the basic game. The question that we study is how does the price of stability of the invitation game compare to that of the basic game. For a wide range of schedulers we prove a dichotomy result, showing that there are only two types of basic games, those that we call *invitation resistant* in which the price of stability of the invitation version is equal to that of the basic game, and those that we call *asymptotically efficient* in which the price of stability tends to 0 as the number of rounds grows.¹ In particular, when the basic game is the prisoners dilemma the game is asymptotically efficient if and only if the payoff when both players defect is nonzero.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiaгент systems*

Keywords

Prisoners dilemma; multi-player games; sub-game perfect equilibrium

1. INVITATION GAMES

Theoretical games can be viewed as models of real life interaction among selfish agents. Some aspects of the real life situation are included in the abstract model, whereas others, judged not to have significant impact on the applicability of

^{*}Work done in Microsoft Research, Herzliya.

¹In our notation price of stability of 0 is best possible, and 1 is worse possible.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ITCS'14, January 12–14, 2014, Princeton, New Jersey, USA.
Copyright 2014 ACM 978-1-4503-2698-8/14/01 ...\$15.00.
<http://dx.doi.org/10.1145/2554797.2554808>.

the abstract mathematical model to the real life motivation, are abstracted away. In some cases there appears to be a contradiction between the mathematical conclusions offered by game theory and real life experience, and this apparent contradiction can in general be attributed to some aspect of the real life situation that is not captured by the game theoretical model. The purpose of the current work is to consider one aspect that is typically abstracted away in game theoretic models. This aspect is the fact that often it is not the case that the players involved just happen to be playing the game, but rather the actual composition of players in the game is a result of strategic decisions on behalf of the players themselves. This aspect comes up naturally in diverse situations. Here are some examples.

1. Games of leisure, such as Bridge. A player may choose which friends to invite over to a game of Bridge depending on how much he enjoys playing with them.
2. Economic interactions that are modeled as games, such as bargaining games. A house owner may prefer inviting one plumber over another, depending not only on the quality of work of the plumber, but also on how stressful is the experience (game) of reaching agreement on a price.
3. Social interactions, e.g., in the sense of *games people play* [5]. People may choose their friends and spouses, and this affects who the other players will be in their subsequent social interactions that are modeled as games.

This aspect of having a choice with whom to play is captured explicitly in some economic models. For example, models of economic competition between firms involve the option for clients to choose which firm to do business with. However, it is our belief that this aspect has ramifications extending also to games in which this aspect is not traditionally modeled. To illustrate the effects of this aspect, we introduce an abstract model that we call *invitation games*, which is a variation on the previously studied model of repeated games. The main phenomenon that we consider is that of the effect of the invitation framework on the *price of stability* of games.

The price of stability of a given game measures the gap in welfare between the optimal social outcome of a game (the one that maximizes the sum of payoffs of the players) and the strategic outcome of the game (the one obtained in the most favorable Nash equilibrium). We prove that under quite mild conditions (for which we provide an exact charac-

terization), this gap tends to 0 in invitation games.² Consequently, one may conclude that in the invitation framework socially optimal outcomes are natural results of selfish behavior.

1.1 The model

Consider an arbitrary game G with row player R and column player C . The payoff matrices for the players will also be denoted by R and C . Hence if R plays action i and C plays action j the payoff to R is $R_{i,j}$ and the payoff to C is $C_{i,j}$. Given G as above and a scheduler S , we define an *invitation* version of G , denoted by $G_{n,S}$, as follows. There are $n > 2$ players denoted by P_1, \dots, P_n . The game $G_{n,S}$ proceeds sequentially in rounds, starting at round 1. In any given round t , the scheduler S may either choose to end the invitation game, or to output an index $s \in \{1, \dots, n\}$, thus selecting a player P_s . (We shall shortly provide more details on how the scheduler makes his choices.) Then in round t , player P_s gets to invite a player of his choice (say, player P_q , where $q \neq s$) to a *match*, where a match is a game of G , with P_s playing the role of R and P_q playing the role of C . Hence the actions available to players in round t are as follows:

1. Player P_s (selected by the scheduler) has two actions to perform: to choose an index $q \neq s$, $1 \leq q \leq n$, thus inviting P_q to a match, and to choose an action i for player R in the game G .
2. Player P_q invited by P_s has one action to perform: to choose an action j for player C in game G .
3. Every other player $P_{q'}$ with $q' \notin \{s, q\}$ has no action to perform.

In a given round, let P_s be the player selected by the scheduler, let P_q be the player invited by P_s , let i be the action of player P_s in game G , and let j be the action of P_q in game G . The payoffs for the players are as follows

1. Player P_s gets payoff R_{ij} .
2. Player P_q gets payoff C_{ij} .
3. Every other player $P_{q'}$ gets payoff 0.

The total payoff of a player in $G_{n,S}$ is the sum of his payoffs over all rounds (when the scheduler ends the game).

In a given round, let P_s be the player selected by the scheduler, and let P_q be the player invited by P_s . Then the information revealed to the players at the end of the round is:

1. Player P_s learns the action played by P_q (and hence can infer the payoffs of all players in this round).
2. Player P_q learns that he was invited by P_s and the action of P_s (and hence can infer the payoffs of all players in this round).
3. In the *full information* version of the game, also every player $P_{q'}$ with $q' \notin \{s, q\}$ learns that the round took place, and also who P_s and P_q are and what they played. There are two *local information* (also referred

to as *partial information*) versions of the game. In the *asynchronous* version $P_{q'}$ with $q' \notin \{s, q\}$ learns nothing (not even that the round took place). In the *synchronous version*, $P_{q'}$ with $q' \notin \{s, q\}$ learns that the round took place, that he himself was neither selected by the scheduler nor invited to a match, but nothing else.

Let us now present more details on the nature of the scheduler. Recall that in a given round t the scheduler needs to decide whether to end the game at that round, and if not, to decide which player P_s to select. In this work we shall restrict attention to *nonadaptive* schedulers, for which both decisions depend only on the round number t , but do not depend on the history of the game up to this round. (It makes sense to consider also adaptive scheduler, but this is beyond the scope of the current paper.) The nonadaptive scheduler may be either deterministic or randomized, where in the latter case its decision may also depend on random coin flips. For a nonadaptive deterministic scheduler, the total number of rounds in an invitation game is fixed in advance. In this paper, the total number of rounds of a randomized schedule will be fixed in advance, though we mention here that one may also consider the so called *unknown horizon* case in which the number of rounds is a random variable that depends on the randomness of the scheduler.

A special case of the nonadaptive deterministic scheduler is the *round robin* scheduler, for which the number of rounds T is divisible by n , and the rounds are partitioned into $B = T/n$ blocks of consecutive rounds. For any given block, in round s (for $1 \leq s \leq n$) within the block, the round robin scheduler selects player P_s (we say that it is P_s 's *turn* to invite). The game $G_{n,S}$ with a round-robin scheduler and T rounds is denoted by $G_{n,T}$. For round robin schedulers with local information, there is no essential distinction between the synchronous version and the asynchronous versions of the game (because once a player gets to play a match, either by reaching his turn or by invitation, he gets synchronized again).

Throughout, given a 2-player game G , we refer to G as the *basic game* and to $G_{n,S}$ (or $G_{n,T}$ in the case of round robin scheduler) as the *invitation game*. Let us clarify that a nonadaptive scheduler is not a player in the invitation game, but rather part of the description of the rules of the game. Hence the scheduler is common knowledge to all players.

1.2 Strategies and solution concepts

Consider a 2-player game G , let A_R be the set of actions available to the row player R and let A_C be the set of actions available to the column player C . A *pure strategy* for player R (C , respectively) is a choice of a single action $i \in A_R$ ($j \in A_C$, respectively). A *mixed strategy* for player R (C , respectively) is a probability distribution over actions in A_R (in A_C , respectively). A *pure Nash equilibrium* (pure equilibrium, for short) is a pair of pure strategies, one for R and one for C , such that no player can increase his own payoff by deviating from his prescribed strategy, given that the other player has not deviated. A *mixed Nash equilibrium* (mixed equilibrium, for short) is a pair of mixed strategies, one for R and one for C , such that no player can increase his own expected payoff by deviating from his prescribed strategy, given that the other player has not deviated. A finite game has at least one equilibrium (possibly mixed).

²In our notation price of stability of 0 is best possible, and 1 is worse possible.

Given a basic game G , we consider its invitation version $G_{n,S}$. The notions of pure and mixed strategies extend naturally to $G_{n,S}$, but let us elaborate on this. A pure strategy for player P_i in the game $G_{n,S}$ is a function $H \rightarrow ([n] \setminus \{i\}, A_R, A_C)$ mapping a history H of the game (the part observable by P_i in case of local information) to an action in the given round. The action in a round t has three components. The first component, a choice of an index from $[n]$ other than i , is the player that P_i would invite in round t given history H if the scheduler S dictates that it is P_i 's turn to invite. This aspect of the action can be left empty if it is not P_i 's turn to invite. The second component is P_i 's choice of action in the basic game G , if P_i pays the role of R . Also this aspect of the action can be left empty if it is not P_i 's turn to invite. The third component is P_i 's choice of action in the basic game G , if P_i plays the role of C (that is, if P_i is invited and sees history H). This aspect of the action can be left empty if it is P_i 's turn to invite. Mixed strategies are probability distributions over pure strategies.

For invitation games, we assume that every player P wishes to maximize the expected sum of payoffs that P receives from all rounds. As in the case of the basic 2-player game G , an *equilibrium* for $G_{n,S}$ is a profile of strategies, one for each player, such that no player can gain by unilaterally deviating from his strategy. However, due to the sequential aspect of $G_{n,S}$, the notion of equilibrium deserves further discussion. We envision a situation in which prior to the commencement of $G_{n,S}$ a profile of n (mixed) strategies is declared, with the i th strategy associated with player P_i . The profile of strategies is an equilibrium in the sense that if all players but one follow their declared strategies, the one deviating player does not gain by his deviation (compared to his expected payoff had he followed his declared strategy). Moreover, the declared profile needs to satisfy the additional property of *subgame perfect equilibrium* (see [15]). Namely, at any point of the game, regardless of the history of the game up to that point, the strategies when restricted to the remaining subgame form an equilibrium for that subgame. This requires strategies to be defined for every possible history H , including histories that lie outside the equilibrium path.³

Convention. The term *equilibrium* (whether pure or mixed) for an invitation game will always refer to a subgame perfect equilibrium.

1.3 Our results

Given a finite game G , let $W^+(G)$ denote the maximum welfare (sum of payoffs for all players) of an outcome of G , and let $W^-(G)$ denote the minimum welfare of an outcome of G . For example, if G is a basic two player game, then $W^+(G) = \max_{ij} [R_{ij} + C_{ij}]$ and $W^-(G) = \min_{ij} [R_{ij} + C_{ij}]$.

Definition 1. The price of anarchy (*PoA*) of a finite game G is 0 if $W^+(G) = W^-(G)$, and otherwise

$$PoA(G) = \max_N \frac{W^+(G) - W_N(G)}{W^+(G) - W^-(G)}$$

where N ranges over all (mixed) equilibria of G (recall also that if G is an invitation game, the equilibria need to be

³One may think of deviations of players from their declared strategies as being accidental rather than deliberate. The declared strategy needs to include information on how players intend to act also in situations in which there were previous accidental deviations.

subgame perfect), and W_N denotes the expected welfare in equilibrium N . The price of stability (*PoS*) of a game G is 0 if $W^+(G) = W^-(G)$ and otherwise

$$PoS(G) = \min_N \frac{W^+(G) - W_N(G)}{W^+(G) - W^-(G)}.$$

Definition 1 is a natural variation on previous definitions of PoA and PoS [9, 14, 1], though we alert the reader that we took the liberty of rearranging some aspects of previous definitions. According to Definition 1, a price of anarchy (or stability) of 0 is best possible (no price needs to be paid), whereas a price of 1 is worst possible (all welfare is lost). In contrast, earlier work defines PoA and PoS (when payoffs are nonnegative) as $\max_N \frac{W^+(G)}{W_N(G)}$ and $\min_N \frac{W^+(G)}{W_N(G)}$, in which case a price of anarchy (or stability) of 1 is best possible, whereas a price of ∞ is worst possible.

Proposition 1. For every finite 2-person game G , every $n \geq 3$ and every non-adaptive scheduler S , the following holds for the invitation games.

1. $PoA(G_{n,S}) \geq PoA(G)$.
2. $PoS(G_{n,S}) \leq PoS(G)$.

For simplicity, the following definition, and consequently some of our results, are stated for a round robin schedule. Extensions of these results to other nonadaptive schedules (including randomized ones) appear in Section B.1.

Definition 2. We say that a finite 2-player game G is invitation-resilient if for every $n \geq 3$, for every $T \geq 1$, in every subgame perfect equilibrium of the invitation game $G_{n,T}$ with a round robin schedule, in every match the corresponding two players play an equilibrium strategy for G .

Definition 3. We say that a finite 2-player game G is asymptotically-efficient if for every $n \geq 3$, the price of stability of the invitation game $G_{n,T}$ with a round robin schedule tends to 0 as T grows. Namely, $PoS(G_{n,T}) \rightarrow 0$ as $T \rightarrow \infty$.

Our main theorem provides a dichotomy for invitation games.

Theorem 2. Let G be an arbitrary finite 2-player game, and let $G_{n,T}$ be its corresponding full information invitation game with a round robin schedule and $n \geq 3$. Then either G is invitation resilient and then $PoS(G_{n,T}) = PoS(G)$ for every T , or G is asymptotically efficient ($PoS(G_{n,T})$ tends to 0 as T grows).

The proof of Theorem 2 implies the following somewhat paradoxical corollary.

Corollary 3. Given a basic game G , let G^- be the basic game obtained by reducing the payoff of every player in each outcome by a fixed universal constant (say, by 1). There is such a pair of basic games G and G^- for which the corresponding invitation games have the following property: $G_{n,T}^-$ has a subgame perfect equilibrium in which every player has higher payoff than the payoff of every player in every subgame perfect equilibrium of $G_{n,T}$.

Given the payoff matrices of a two player game G that is asymptotically-efficient, the proof of Theorem 2 shows that there is an NP-witness (a proof polynomial in the size of the payoff matrices of G) certifying that G is asymptotically-efficient. However, finding such a witness need not be computationally easy.

Proposition 4. *Given the payoff matrices of the two players in a game G , the computational problem of determining that G is asymptotically-efficient is NP-complete.*

For games that are not invitation resilient, the price of anarchy offers a wide range of behaviors, unlike the price of stability (that necessarily tends to 0 as T grows).

Proposition 5. *Consider arbitrary $0 \leq \epsilon < \delta \leq 1$. Then there is a finite 2-player game G for which $PoA(G) = \epsilon$, and $PoA(G_{n,S})$ tends to δ as the number of rounds grows. ($G_{n,S}$ is the full information invitation game with an arbitrary nonadaptive schedule.)*

1.4 Related work

In our setting there is a basic game that is repeatedly played in the context of an invitation game, and as a result rational players might not play the equilibrium strategy in particular instances of the basic game. A similar phenomenon was previously illustrated in the context of repeated games, in which a basic game is repeated a fixed number of times, each time with the same set of players. We refer to this phenomenon as the repeated games principle, and it states that for certain classes of games, any feasible and individually rational payoff vector of the one-shot basic game can be approximated by the average payoff in a perfect equilibrium of a repeated game with a sufficiently long (but finite) horizon. Some special cases of the repeated game principle were presented in [7] using so called *trigger strategies*, and a characterization of the class of 2-player basic games for which the repeated game principle holds is presented in [4]. Our Proposition 12 in Section B.1 is implied by this characterization. We note however that the repeated games principle does not apply to all games, and in particular not to the game of prisoner’s dilemma. The class of 2-player games for which the price of stability tends to 0 in the invitation version strictly contains the class of games for which the repeated games principle holds.

Despite the failure of the repeated games principle for the game of prisoners dilemma, various other explanations were suggested as to why players might choose to cooperate (hence play out of equilibrium) in repeated versions of this game. One aspect that may lead to cooperation is an unknown horizon – the players do not know how many times the game will be repeated. Technically, this allows to circumvent the issue of backward induction. See [12] for a complete treatment of this issue. An alternative explanation is offered (in a different context, but can easily be adapted to the prisoners dilemma) in [13]: players are willing to settle for a perfect ϵ -equilibrium (no player can gain more than ϵ by unilaterally deviating) instead of a perfect equilibrium. In repeated prisoners dilemma, cooperating in every game and defecting only if the other player previously deviated is an ϵ -equilibrium.

A different kind of explanation of how cooperation emerges in games of prisoners dilemma is demonstrated in [3, 2]: in a mixed population in which different players come up with

their own strategies (not necessarily rational in the game theoretic sense), there is empirical evidence showing that some strategies that allow for cooperation perform better on average (when a player is matched for a sequence of games with a random player from the population) than the strategy of always defecting. If preference of either cooperation or defection is viewed as a hereditary property, then those strategies that do well on average have an evolutionary advantage, and hence become common.

The notion of price of anarchy was formally defined in [9, 14], and price of stability was formally defined in [1]. We note that we slightly change the standard definitions of these concepts so that the value ranges between 0 (no price) and 1 (as bad as it can get).

Our NP-hardness result (Proposition 4) has several predecessors with related results [10, 6], and implicitly follows from them. For completeness, we present a short proof of this proposition, using a characterization of Motzkin and Straus [11] of the size of the maximum independent set in a graph. This characterization has been previously used in NP-hardness proofs for so called *evolutionary stable strategies*, see [8].

2. SUBGAME PERFECT EQUILIBRIA OF INVITATION GAMES

Let G be an arbitrary finite 2-player game played between the row player R and the column player C . In a given match of G , we say that the players play an equilibrium strategy if their (mixed) strategies are in equilibrium. Proposition 1 is an immediate corollary of the following proposition, whose proof appears in Section A.

Proposition 6. *Let G be an arbitrary finite 2-player game, and let $G_{n,T}$ be its corresponding invitation game with full information and a non-adaptive scheduler. Consider an arbitrary equilibrium for G . Then there is a subgame perfect equilibrium for $G_{n,T}$ for which in every match both players (participating in the match) play the given equilibrium for the respective basic game G .*

Note that with an adaptive scheduler and $r \neq 0$, Proposition 6 need not hold. Suppose for example that R has a single dominating strategy in G that gives R payoff $r > 0$ and gives C payoff 0. Suppose that the adaptive scheduler makes P_1 the inviting player in round 1, and depending on P_1 ’s action in round 1, decides whether P_1 will continue to be the inviting player in all future rounds (this happens if P_1 does not play his dominating strategy in round 1), or will never again be the inviting player (this happens if P_1 does play his dominating strategy in round 1). Then if the number of rounds is sufficiently large, P_1 will not play his dominating strategy in round 1.

Theorem 7. *Let G be an arbitrary finite 2-player game, and suppose that there is some value r such that in every equilibrium of G the expected payoff for R is r and the expected payoff for C is 0. Let $G_{n,T}$ be its corresponding invitation game with full information and a non-adaptive scheduler. Then in every subgame perfect equilibrium for $G_{n,T}$, in every match both players (participating in the match) play an equilibrium strategy for the respective basic game G .*

PROOF. Consider an arbitrary equilibrium for $G_{n,T}$. We use backwards induction to show that in every match both

players play an equilibrium strategy for G . The base case is that of the last round. Let P be the inviting player in round T . The first component of the action of P is a choice of player to invite. Let Q be this player. Regardless of the identity of P and Q , the subgame that one is left with is a match of G . By subgame perfection, in this match the players need to play an equilibrium strategy.

Consider now an arbitrary round t , and let P be the inviting player in that round. By our induction hypothesis, in every future round the respective matched players play equilibrium strategies. This implies that in future rounds the expected payoff for inviting players is always r , and the expected payoff for invited players is always 0. Hence players are indifferent to being invited in future rounds. Likewise, in future rounds inviting players are indifferent as to who they invite. Hence nothing that P does in round t effects his expected payoff in future rounds. This means that the match in round t is played without concern of the future, and subgame perfection implies that the match is played without concern of the past. Hence regardless of the identity of player Q invited by P in round T , the in a subgame perfect equilibrium of $G_{n,T}$ the players P and Q play an equilibrium strategy for G in round t . \square

The maximum welfare of a game G is defined as $W^+(G) = \max_{ij}[R_{ij} + C_{ij}]$. In a given match of G , we say that the players play a maximum welfare strategy if R plays i and C plays j satisfying the equality above.

Theorem 8. *Let G be an arbitrary finite 2-player game, and let $G_{n,T}$ be its corresponding full information invitation game with a round robin schedule. Suppose that either one of the following holds:*

1. *There is some value $c \neq 0$ and an equilibrium of the basic game G in which the expected payoff for C is c .*
2. *There are two values $r \neq r'$ and two equilibria for the basic game G , in one of which the expected payoff for R is r and in the other it is r' . (It suffices to consider the case that in both equilibria for G the expected payoff for C is 0, as other cases are handled by 1 above.)*

Then there is some positive integer ℓ such that for every T there is a subgame perfect equilibrium for $G_{n,T}$ in which in all but ℓn matches the players play a maximum welfare strategy (and in other matches players play an equilibrium strategy).

A complete proof of Theorem 8 appears in Section B. Here we only sketch the main principles used in the proof. The set of all rounds is broken into the main set of rounds which includes the first $T - \ell n$ rounds, and the auxiliary set of rounds which includes the last ℓn rounds. The intention is that in the subgame perfect equilibrium players will play a maximum welfare strategy in the main set of rounds. Players have no incentive to deviate because if they do they will get “punished” in the auxiliary set of rounds. The nature of the punishment depends on properties of equilibria of the basic games.

1. If there is an equilibrium of the basic game G in which the expected payoff for C is $c > 0$, then deviating players are punished by not being invited by other players in the auxiliary rounds, and hence not collecting a payoff of c .

2. If there is an equilibrium of the basic game G in which the expected payoff for C is $c < 0$, then deviating players are punished by being invited by other players in the auxiliary rounds, and consequently suffering a negative payoff of c .
3. If there is an “inferior” equilibrium for the basic game G in which the expected payoff for R is r and a “superior” equilibrium in which the expected payoff for R is $r' > r$, then a deviating player is punished by other players playing the inferior equilibrium whenever invited by the deviating player, rather than the superior equilibrium.

For more details regarding how the above principles can be implemented within a subgame perfect equilibrium, see the full proof in Section B.

It is convenient to think of a strategy for a player in an invitation game as composed of two components. One is the *invitation strategy*, specifying who a player invites when it is his turn to invite, and the other is the *G -strategy* specifying what to play in matches of G . We say that an invitation strategy in *nonadaptive* if it depends only on the round number but not of the history of the matches that the player played. In our proof of item 1 of Theorem 8 players use adaptive invitation strategies, whereas in the proof of item 2 players use nonadaptive invitation strategies.

We remark that even though Theorem 8 is stated for the full information setting, its proof applies without change also to the local information setting.

We now prove Theorem 2.

PROOF. Theorem 7 provides a sufficient condition for a finite 2-player game to be invitation resilient. This condition states that there is some value r such that in every equilibrium of G the expected payoff for R is r and the expected payoff for C is 0. Theorem 8 handles all those games for which this sufficient condition does not hold and claims that in all cases the game is asymptotically efficient. \square

The price of anarchy does not provide the same sort of dichotomy provided by the price of stability. We now prove Proposition 5.

PROOF. Consider a basic game G with two pure strategies per player with payoff matrices $R = \begin{pmatrix} 0 & \frac{1-\delta}{2} \\ 0 & 0 \end{pmatrix}$ and $C = \begin{pmatrix} 1-\epsilon & \frac{1-\delta}{2} \\ 1-\epsilon & 1 \end{pmatrix}$. We have $W^+(G) = 1$ and $PoA(G) = \epsilon$ (in an equilibrium C must not play the second column as then R plays only the first row). For the invitation version of G , consider the equilibrium in which every player has a designated partner that he invites, as long as the designated partner agrees to play the $(\frac{1-\delta}{2}; \frac{1-\delta}{2})$ cell, and switches to an alternative partner once the designated partner deviates. As the number of rounds grows, all but a constant number of rounds (the constant depending on ϵ, δ) are played with $1 - \delta$ welfare. \square

2.1 Prisoners dilemma as the basic game

Prisoners dilemma (PD) is a symmetric 2-player game. A player has two actions: *defect* (denoted here as 0) and *cooperate* (denoted here as 1). Defecting is a dominant action, but if both players cooperate they get a higher payoff than if they both defect. Hence we have $R_{10} < R_{00} < R_{11} < R_{01}$.

Corollary 9. *Let PD be the prisoners dilemma game and consider subgame perfect equilibria of the corresponding invitation game $PD_{n,T}$ with a round robin scheduler.*

1. *When the dominant payoff satisfies $R_{00} = 0$ then in every subgame perfect equilibria of the invitation game $PD_{n,T}$, in all matches both players defect.*
2. *When the dominant payoff satisfies $R_{00} \neq 0$ then there is a value of ℓ that depends only on PD such that for every T , the invitation game $PD_{n,T}$ has a subgame perfect equilibria for which in all but ℓn matches both players cooperate.*

PROOF. Item 1 is a special case of Theorem 7. If $2R_{00} \geq R_{01} + R_{10}$ then item 2 is a special case of (item 1 of) Theorem 8. If $2R_{00} < R_{01} + R_{10}$ then item 2 is technically not a special case of (item 1 of) Theorem 8, but it can be proved in exactly the same way in which item 1 of Theorem 8 is proved. \square

2.2 A paradox for invitation games

One readily observes that Corollary 9 implies Corollary 3.

PROOF. Let G be the prisoners dilemma game with payoffs $R_{10} = -2$, $R_{00} = 0$, $R_{11} = 2$, $R_{01} = 4$ (where 0 signifies *defect* and 1 signifies *cooperate*). By item 1 of Corollary 9, in every subgame perfect equilibrium players always defect, and hence every player receives total payoff of 0. The game G^- is the prisoners dilemma with payoffs smaller by 1, namely $R_{10} = -3$, $R_{00} = -1$, $R_{11} = 1$, $R_{01} = 3$. Taking $\ell \geq 2$, the proof of item 2 of Corollary 9 shows (when n is even) a subgame perfect equilibrium in which every player P_i gets payoff of $2(T/n - \ell)R_{11} + 2\ell R_{00}$. Taking the number of blocks T/n to be at least five, the corollary is proved. (With extra work the proof can be extended also to the case that n is odd, but details are omitted.) \square

Corollary 3 has a paradoxical flavor. The basic games G and G^- are the same game up to a fixed additive shift of payoffs. For the corresponding invitation games, if the number of matches that a player P_i participates in is fixed in advance to m_i (e.g., because players use nonadaptive invitation strategies) then the shift in payoffs for P_i between $G_{n,T}$ and $G^-_{n,T}$ is just the fixed constant m_i . The aspect that makes the proof of Corollary 3 possible is the use of adaptive invitation strategies, by which one player threatens to change the value of m_i if P_i deviates from the proposed equilibrium strategies.

2.3 Constant sum basic games

It is instructive to see what our results imply for the case that the basic game G is a constant sum game. By the minimax theorem, there are values r and c such that in all equilibria for G the expected payoff for R is r and the expected payoff for C is c (in a 0-sum game $r = -c$). In the special case that $c = 0$, Theorem 7 implies that in every subgame perfect equilibrium, in all matches both players play their minimax strategies. However, if $c \neq 0$, Theorem 7 no longer applies. Indeed, arguments similar to those appearing in the proof of Theorem 8 can be used in order to design a 0-sum game G and a subgame perfect equilibrium for the corresponding invitation games $G_{n,T}$ in which in all but the last few matches players do not play their minimax strategies.

2.4 Computational complexity

It is well known that for many natural restrictions on Nash equilibria, finding a Nash equilibrium satisfying these restrictions is NP-hard [10, 6]. In Lemma 10 we present yet another proof of such a result, that we shall use in order to prove Proposition 4.

Given a graph G , consider the following 2-player game that we refer to as the k -clique game. The payoff matrix A for the row player has $n + 1$ rows and n columns. The top n by n submatrix is the adjacency matrix of a graph G . All the entries of the bottom row are $\frac{k-1}{k}$. For the column player, the top n by n submatrix is the identity matrix, and the bottom row is all 0.

Lemma 10. *If G has a clique of size k , then there is a Nash equilibrium for the k -clique game in which the support for the row player has at least k rows, and the expected payoff for the column player is positive. If G does not have a clique of size k , then in all Nash equilibria the support for the row player is the last row, his payoff is $\frac{k-1}{k}$, and the payoff for the column player is 0.*

PROOF. If S is a maximal clique of size at least k , then both players playing uniformly over S gives a Nash equilibrium with payoff $1 - 1/|S|$ for the row player and $1/|S|$ for the column player. This proves the first part of the theorem.

We shall refer to a Nash equilibrium in which the row player plays row $n + 1$, the payoff for the row player is $\frac{k-1}{k}$ and the payoff for the column player is 0, as a *standard* equilibrium. If S is a maximal clique of size smaller than k , then the row player playing row $n + 1$ and the column player playing uniformly over S gives a standard equilibrium. To prove the second part of the theorem it remains to show that if G has no clique of size k , there is no non-standard equilibrium. Suppose for the sake of contradiction that there is a non-standard equilibrium. Let x_i be the probability that the i th column is played in this equilibrium, and let S be the support for the mixed strategy for the column player (namely, $S = \{i | x_i > 0\}$). S must also be in the support for the row player (otherwise the payoff for one of the strategies of the column player is 0, implying that this is the payoff for all the column player's strategy, implying that only row $n + 1$ is played, and hence this is a standard equilibrium). Being in the support of an equilibrium, all rows of S give the row player the same expected payoff, and it must be at least $\frac{k-1}{k}$ (otherwise the row player plays row $n + 1$ instead). Let y range over nonnegative vectors whose entries sum to 1, and let M be the adjacency matrix of G . We have thus established that $\max_y y^T M y \geq \frac{k-1}{k}$, by taking $y = (x_1, \dots, x_n)$. The well known Motzkin-Straus theorem [11] implies that G contains a k -clique. \square

We now prove Proposition 4.

PROOF. As implied by the proof of Theorem 2, a game is asymptotically efficient if and only if at list one of the two conditions listed in Theorem 8 holds. Each of these condition has a short witness (exhibiting the respective Nash equilibria) certifying that it holds (if it indeed holds). Hence the problem of determining that a game is asymptotically efficient is in NP. Lemma 10 implies that determining that a game is asymptotically efficient is NP-hard, because for the k -clique game is equivalent to determining whether the underlying graph has a k -clique. \square

Acknowledgements

We thank Reshef Meir and an anonymous reviewer for directing us to related references.

3. REFERENCES

- [1] Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., and Roughgarden, T. (2004). The Price of Stability for Network Design with Fair Cost Allocation. In Proceedings of the 45th IEEE Symposium on Foundations of Computer Science, (FOCS-04), pp. 59-73.
- [2] Robert Axelrod. The Evolution of Cooperation, Basic Books, 1984.
- [3] Robert Axelrod, William D. Hamilton. The Evolution of Cooperation. Science, New Series, Vol. 211, No. 4489. (March 27, 1981), pp. 1390-1396.
- [4] Jean-Pierre Benoit and Vijay Krishna. Finitely Repeated Games. Econometrica, Vol. 53, No. 4 (July 1985), pp. 905-922.
- [5] Eric Berne. Games People Play. New York: Ballantine Books. 1964.
- [6] Vincent Conitzer, Tuomas Sandholm. Complexity Results about Nash Equilibria. IJCAI 2003: 765-771.
- [7] James W. Friedman. Cooperative equilibria in finite horizon noncooperative supergames. Journal of Economic Theory, Volume 35, Issue 2, August 1985, Pages 390-398.
- [8] Kousha Etesami, Andreas Lochbihler. The computational complexity of evolutionarily stable strategies. Int. J. Game Theory 37(1): 93-113 (2008).
- [9] Koutsoupias, E., and Papadimitriou, C. (1999). Worst-Case Equilibria. In Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 404-413.
- [10] I. Gilboa and E. Zemel. Nash and Correlated Equilibria: Some Complexity Results. Games and Economic Behavior, 1 (1989), 80-93.
- [11] T.S. Motzkin and E.G. Straus. Maxima for graphs and a new proof of a theorem of Turan. Can. J. Math. 17, 533-540 (1965).
- [12] Abraham Neyman, Cooperation in Repeated Games when the Number of Stages is not Commonly Known. Econometrica, 67 (1999), pp. 45-64.
- [13] Roy Radner. Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives. Journal of Economic theory 22, 136-154 (1980).
- [14] Roughgarden, T., and Tardos, E. (2002). How Bad is Selfish Routing?. Journal of the ACM, 49(2), 236-259.
- [15] R. Selten. Re-examination of the perfectness concept for equilibrium points in extensive games. Int. J. Game Theory 4 (1975), 25-55.

APPENDIX

A. PROOF OF PROPOSITION 6

In this section we prove Proposition 6.

PROOF. Fix an arbitrary equilibrium for G . Let r (c , respectively) be the expected payoff of R (C , respectively) in this equilibrium. We say that a player plays his equilibrium strategy in G if he plays the strategy that corresponds to this given fixed equilibrium.

Consider the following profile of strategies for the players. For every i , when it is player's P_i turn to invite, he always invites player P_{i+} (regardless of the history up to this round), where i^+ signifies addition in cyclic order, and its inverse i^- signifies subtraction in cyclic order. (This choice of using the function i^+ is arbitrary and was made for concreteness.) In every match (regardless of the history), both players play their equilibrium strategies.

Let m_i be the number of rounds in which P_i is the inviting player under the nonadaptive schedule. Then under the profile of strategies specified above, for each player P_i his expected payoff is $rm_i + cm_{i-}$. Unilaterally deviating from the above strategy will not increase the expected payoff of P_i . Hence the profile of strategies is an equilibrium. It is subgame perfect, as can be shown by backward induction. \square

B. PROOF OF THEOREM 8

For the proof of Theorem 8, we define various parameters of G :

The maximum welfare. $W^+(G) = \max_{ij}[R_{ij} + C_{ij}]$. For this optimal choice of ij that results in $W^+(G)$ (if there are several possible pairs of strategies i, j that lead to maximum welfare, fix one pair arbitrarily) we denote the respective payoffs by $r^* = R_{ij}$ and $c^* = C_{ij}$. In a given match of G , we say that the players play a maximum welfare strategy if R plays i and C plays j as above.

The maximum difference. $D_R(G) = \max_{ii'jj'}[R_{ij} - R_{i'j}]$. $D_C(G) = \max_{ijjj'}[C_{ij} - C_{ij'}]$. $D(G) = \max[D_R(G), D_C(G)]$. The parameter $D(G)$ provides an upper bound on how much a player can gain (in a single game of G) by replacing one strategy by another.

We now prove Theorem 8.

PROOF. The proof involves a case analysis, but there is a common thread to all cases. Assume for simplicity that n is even. (After the proof we sketch how this assumption can be removed.) Arrange all players in disjoint pairs. Two players of the same pair will be called *partners*. We shall also have another arrangement of players in pairs, for which two players in the same pair are *fallback partners*. (For every player, his partner and fallback partners are different.)

Suppose that there is an equilibrium of G in which the expected payoff of C is $c > 0$, and let r denote the expected payoff of R in this equilibrium. Let ℓ be such that $c\ell \geq \max[D(G), r - r^*]$, and choose T such that $T/n > \ell$. We propose the following subgame perfect equilibrium for $G_{n,T}$. In equilibrium:

1. Every inviting player always invites his partner.
2. In every match in the first $T/n - \ell$ blocks the players play the maximum welfare strategy.
3. In every match in the last ℓ blocks the players play the equilibrium strategy.

We now describe how players respond to situations in which other players deviate from equilibrium:

1. If a player is invited by any player other than his partner, on that match he plays the equilibrium strategy.
2. If a player P observes a deviation by his partner Q in any of the matches that are (or should have been) held between them (either Q inviting a player other than P , or Q not playing the deterministic maximum welfare strategy in the first blocks, or Q not playing the

equilibrium strategy in the last blocks), then P never invites Q again. Instead P invites his fallback partner and plays the equilibrium strategy. Moreover, if Q invites P after Q has deviated, P plays the equilibrium strategy.

For subgame perfect equilibrium, we need to also describe how a player reacts to his own deviations from equilibrium. This represents what other players “believe” that this player will do, though of course once deviating, the player might deviate again. In our subgame perfect equilibrium strategy:

- If a player P previously deviated, then whenever it is P 's turn to invite he invites his partner, and in every match P plays the equilibrium strategy for G .

If all players follow the above equilibrium strategy, every player receives expected payoff $(T/n - \ell)W^+(G) + \ell(c + r)$. A player has nothing to gain from unilaterally deviating. In the last ℓ blocks only equilibrium strategies for G are played. If a player deviates in any of the first $T/n - \ell$ blocks, he gains at most $D(G)$ (if the deviation was in a play of G against his partner) or at most expected $r - r^*$ (if the deviation was by inviting a player other than his partner), and loses an expected payoff of at least $c\ell$, by not being invited in the last ℓ blocks. The net gain is not positive. The equilibrium is subgame perfect because once a player deviates, all his games are played in an equilibrium for G , and backward induction applies.

Suppose now that there is no equilibrium of G in which the expected payoff of C is strictly positive, but there is an equilibrium of G in which the expected payoff of C is $c < 0$. Let r denote the expected payoff of R in this equilibrium. Let ℓ be such that $c\ell \geq \max[D(G), r - r^*]$, and choose T such that $T/n > \ell$. We propose the following equilibrium for $G_{n,T}$. In equilibrium:

1. In the first $T/n - \ell$ blocks every inviting player always invites his partner.
2. In the last ℓ blocks every inviting player always invites his fallback partner.
3. In every match in the first $T/n - \ell$ blocks the players play the maximum welfare strategy.
4. In every match in the last ℓ blocks the players play the equilibrium strategy.

We now describe how players respond to situations in which other players deviate from equilibrium:

1. If a player is invited by any player other than his partner, on that match he plays the equilibrium strategy for G .
2. If a player P observes a deviation by his partner Q in any of the matches that are (or should have been) held between them in the first $T/n - \ell$ blocks (either Q inviting a player other than P , or Q not playing the deterministic maximum welfare strategy), then P invites Q in all his remaining turns (including in the last ℓ blocks) and plays the equilibrium strategy for G . Moreover, if Q invites P after Q has deviated, P plays the equilibrium strategy for G .

We need to also describe how a player reacts to his own deviations from equilibrium. In our subgame perfect equilibrium strategy:

1. If a player P previously deviated then whenever it is P 's turn to invite he invites his fallback partner, and in every match P plays the equilibrium strategy for G .

If all players follow the above equilibrium strategy, every player receives expected payoff $(T/n - \ell)W^+(G) + \ell(c + r)$. A player has nothing to gain from unilaterally deviating. In the last ℓ blocks only equilibrium strategies for G are played. If a player deviates in any of the first $T/n - \ell$ blocks, he gains at most $D(G)$ (if the deviation was in a play of G against his partner) or at most expected $r - r^*$ (if the deviation was by inviting a player other than his partner), and loses an expected payoff of at least $c\ell$, by being invited by his partner in the last ℓ blocks. The net gain is not positive. The above equilibrium is subgame perfect because once a player deviates, all his games are played in an equilibrium for G , and backward induction applies.

We now turn to prove item 2 of Theorem 8. Namely, the expected payoff of C is 0 in all equilibria for G , and there are two different equilibria, the *inferior* one in which the expected payoff of R is r , and the *superior* one in which the expected payoff of R is $r' > r$. Let ℓ be such that $(r' - r)\ell \geq \max[D(G), r - r^*]$, and choose T such that $T/n > \ell$. We propose the following subgame perfect equilibrium for $G_{n,T}$. In equilibrium:

1. Every inviting player always invites his partner.
2. In every match in the first $T/n - \ell$ blocks the players play the maximum welfare strategy.
3. In every match in the last ℓ blocks the players play the *superior* equilibrium strategy.

We now describe how players respond to situations in which other players deviate from equilibrium:

1. If a player is invited by any player other than his partner, on that match he plays the *inferior* equilibrium strategy.
2. If a player P observes a deviation by his partner Q in any of the matches that are (or should have been) held between them (either Q inviting a player other than P , or Q not playing the deterministic maximum welfare strategy in the first blocks, or Q not playing the superior equilibrium strategy when playing as C in the last blocks), then in future rounds P continues to invite Q . In matches with Q : if P is inviting he plays the superior equilibrium, and if Q is inviting P plays the inferior equilibrium.

We need to also describe how a player reacts to his own deviations from equilibrium. In our subgame perfect equilibrium strategy, a player P who previously deviated:

1. Always invites his partner and plays the inferior equilibrium.
2. When invited by his partner, P plays the superior equilibrium.
3. When invited by a player other than his partner, P plays the inferior equilibrium.

If all players follow the above strategy, every player receives expected payoff $(T/n - \ell)W^+(G) + \ell r'$. A player has nothing to gain from unilaterally deviating. In the last ℓ blocks invited players get 0 payoff, whereas and inviting player P gets payoff r , except for payoff of $r' > r$ when

inviting their partner (conditioned on P not deviating earlier). If a player deviates in any of the first $T/n - \ell$ blocks, he gains at most $D(G)$ (if the deviation was in a play of G against his partner) or at most expected $r - r^*$ (if the deviation was by inviting a player other than his partner), and loses an expected payoff of at least $(r' - r)\ell$, by getting payoff r instead of r' in the last ℓ blocks. The net gain is not positive. The above equilibrium is subgame perfect because once a player deviates, all his games are played in an equilibrium for G , and backward induction applies. \square

We presented the proof of Theorem 8 only for the case that n is even. We sketch how the proof of Theorem 8 can be extended to the case that n is odd. Because n is odd, we cannot partition all players into pairs, and instead we partition all but three of the players into pairs, and the remaining three players form a triple (P_1, P_2, P_3) . In odd-numbered blocks, every player in the triple considers the player among the triple that follows him (in cyclic order) to be his partner, and the player that precedes him to be his auxiliary partner. In even-numbered blocks, every player in the triple considers the player that precedes him to be his partner, and the player that follows him to be his auxiliary partner. In two consecutive blocks (an odd block followed by an even block) we can identify three pairs of rounds, where in each such pair two players of the triple consider each other to be their partner, and the remaining player to be their auxiliary partner. Hence each player in the triple is in a situation equivalent to having one partner in one set of $T/2$ pairs of rounds (each player in the partnership invites in one round within a pair), and a different partner in a different set of $T/2$ pairs of round. The proof of Theorem 8 extends to this setting (with the modification that each player in the triplet is involved in 2ℓ blocks of an equilibrium strategy for G , instead of only ℓ blocks).

B.1 Extensions to other schedules

Theorem 8 shows how invitation games naturally lead to situations in which individual matches are not played in an equilibrium for the basic game. The statement of the theorem assumes a round robin schedule. However, variants of the theorem easily extend to other schedules. One natural schedule to consider is what we call the random schedule: in every round the inviting player is chosen uniformly at random, independently of all other events. Theorem 8 extends as is to this schedule.

Corollary 11. *Let G be an arbitrary finite 2-player game, and let $G_{n,T}$ be its corresponding full information invitation game with the random schedule. Suppose that either one of the following holds:*

1. *There is some value $c \neq 0$ and an equilibrium of the basic game G in which the expected payoff for C is c .*
2. *There are two values $r \neq r'$ and two equilibria for the basic game G , in one of which the expected payoff for R is r and in the other it is r' .*

Then there is some positive integer ℓ such that for every T there is a subgame perfect equilibrium for $G_{n,T}$ in which in all but ℓn matches the players play the maximum welfare strategy (and in other matches players play an equilibrium strategy).

PROOF. The proof follows from that of Theorem 8 essentially with no change, because players are assumed to be expectation maximizers. Within every block of n matches, the round robin schedule lets every player invite exactly once, whereas the random schedule lets every player invite once in expectation. As expectation maximizers are concerned only with expectation, the arguments in the proof of Theorem 8 carry over to the case of a random scheduler. \square

Recall that the proof of Theorem 8 applies not only in the full information setting, but also to the local information setting (both in the synchronous version and the asynchronous version). The proof of Corollary 11 as stated need not apply in the local information setting. Extending Corollary 11 so that it does apply in the local information setting is beyond the scope of this paper.

Returning to deterministic schedules, Theorem 8 need not hold as stated for arbitrary nonadaptive schedules, due to the fact that some players might invite less often than others. However, the main qualitative message of Theorem 8 still holds, namely, some matches may be governed by welfare considerations irrespective of consistency with equilibria for the basic game. We demonstrate this in two propositions for which we only sketch the proofs (as they follow principles similar to the proof of Theorem 8). Proposition 12 concerns the second class of basic games considered in Theorem 8. For that class, the proof applies even in the two player repeated game setting (without the need of an invitation setting), and is sketched mainly for completeness, as similar results already appeared in [7, 4]. This class of basic games does not include the prisoners dilemma.

Proposition 12. *Let G be an arbitrary finite 2-player game, and let G^T be its T -fold repeated setting. Let c denote the maximum possible payoff for C (the largest entry in C 's payoff matrix). Suppose that among the equilibria for the basic game G , there is a superior equilibrium in which the expected payoffs for R and C are (r_1, c_1) , and an inferior equilibrium in which the expected payoffs for R and C are (r_2, c_2) , with $r_1 > r_2$ (there are no restrictions on the values of c_1 and c_2). Then there is some positive integer ℓ depending only on G such that if $T > \ell$, there is a subgame perfect equilibrium for G^T in which there is a match in which the payoff for C is c .*

PROOF. The proposed equilibrium for G^T when no player deviates is as follows. In the first round, the players play the pure strategies that give C payoff c . In all future rounds the players play the superior equilibrium. If R deviates in the first round, then in all rounds the players play the inferior equilibrium. No other deviation affects future play of the players.

We sketch the proof of why this is a subgame perfect equilibrium. A unilateral deviation of C gains nothing for C , because in every round C is playing a best response. A unilateral deviation of R in round 1 may gain in round 1, but causes R to suffer the inferior equilibrium in all future rounds, instead of the superior equilibrium, hence losing $(r_1 - r_2)\ell$. In all other rounds R gains nothing by deviating, because he is playing a best response. \square

Proposition 12 extends to the case when the roles of R and C are reversed. Namely, when r denotes the maximum possible payoff for R and $c_1 > c_2$, there is a subgame perfect

equilibrium for G^T in which there is a match in which the payoff for R is r .

Proposition 13 corresponds to the first class of games considered in Theorem 8 (which does include the prisoners dilemma if the payoff for defect-defect is nonzero). This proposition does require the invitation setting and at least three players.

Proposition 13. *Let G be an arbitrary finite 2-player game, and let $G_{n,S}$ be its corresponding full information invitation game with schedule S . Let r denote the maximum possible payoff for R (the largest entry in R 's payoff matrix). Suppose that there is some value $c \neq 0$ and a distinguished equilibrium of the basic game G in which the expected payoff for C is c . Then there is some positive integer ℓ depending only on G such that if the number of rounds in schedule S is larger than ℓn , there is a subgame perfect equilibrium for $G_{n,S}$ in which in some match the inviting player gets payoff r .*

PROOF. Given S , there is a player (say, P) that invites in more than ℓ rounds. Fix some arbitrary other player Q .

Suppose that $c > 0$. Then we propose the following subgame perfect equilibrium, All players always invite Q , and Q always invites P . In all matches the distinguished equilibrium is played, except for the first match among those in which it is P 's turn to invite. In that match between P and Q they play the pure strategies that give R a payoff of r . If Q deviates from this strategy in this particular match, then P never invites Q again, and instead invites some other default player P' and they play the distinguished equilibrium. No other deviation results in a change of strategy of players.

We sketch the proof of why this is a subgame perfect equilibrium. A unilateral deviation of any player other than Q gains nothing for that player, because in every round the player is playing a best response. A unilateral deviation of Q in the first match in which Q is invited by P may result in Q gaining in that round, but causes Q to not be invited by P in future rounds, losing a payoff of $c\ell$. In all other rounds Q gains nothing by deviating, because he is playing a best response.

If $c < 0$ the same principles as above apply, except for exchanging Q for Q' (and Q' for Q) in the rounds in which P invites, except for P 's first turn. \square