

Non-cooperative games with preplay negotiations

Valentin Goranko¹ and Paolo Turrini²

¹Department of Applied Mathematics and Computer Science, Technical University of Denmark

²Department of Computing, Imperial College London

April 12, 2013

Abstract

We consider an extension of strategic normal form games with a phase of negotiations before the actual play of the game, where players can make binding offers for transfer of utilities to other players after the play of the game, in order to provide additional incentives for each other to play designated strategies. Such offers are conditional on the recipients playing the specified strategies and they effect transformations of the payoff matrix of the game by accordingly transferring payoffs between players. We introduce and analyze solution concepts for 2-player normal form games with such preplay offers under various assumptions for the preplay negotiation phase and obtain results for existence of efficient negotiation strategies of the players. Then we extend the framework to coalitional preplay offers in N -player games, as well as to extensive form games with inter-play offers for side payments.

1 Introduction

It is well known that some normal form games have no pure strategy Nash equilibria, while others, like the Prisoner's Dilemma, have rather unsatisfactory – e.g., strongly Pareto dominated – ones. These inefficiencies are often attributed to the lack of communication between the players and the impossibility for them to agree on a joint course of action before the play of the game. Indeed, mutually undesirable outcomes could often be avoided if players were able to communicate and *make binding agreements on the strategy to play before the game starts*. However, even if players could freely communicate before the game, enforcing of such contracts is often not possible in practice and, furthermore, it would change the nature of the game from non-cooperative to essentially cooperative.

Here we consider a somewhat weaker and generally more realistic assumption, viz.:

Before the actual game is played any player, say A , can make a binding offer to any other player, say B , to pay him¹, after the game is played, an explicitly declared amount of utility δ if B plays a strategy s specified in the offer by A .

¹We will refer to player A as a female, while to B as a male. This choice is not for the sake of political correctness but to make it easier to distinguish the players from the context.

Once players are endowed with the capacity of performing such moves, a whole negotiation phase emerges before a normal form game is actually played. In other words, we can think of the normal form game that is eventually played as an outcome of another game, played beforehand, in which players engage in exchanging offers on strategies of other players until an agreement is reached on the game to play. Such scenario arises in a wide spectrum of economic, social and political situations, such as collusions, compensations, incentives, concessions, compromises and other kinds of deals in economic and political negotiations, out-of-court settlements of legal cases, or even corruption schemes.

The early literature in economic theory abounds with examples of actors entering negotiations to overcome inefficient resource allocation [Coa60, Pig20, Mea52, Mas94]. Nowadays game theory has developed elaborated models to analyze bargaining among rational decision-makers [OR90, OR94]. Surprisingly enough, in spite of a few proposals of pre-play contracting in games [JW05, EP11], a systematic study of the negotiation process preceding the actual game play seems to be still missing in the literature.

With this paper we initiate such systematic study purporting to fill this gap, by formalizing and studying the negotiation process preceding the actual game play as a bargaining process among the players on the game to play, drawing connections with modern bargaining theory [OR90, OR94]. The paper is intended as a rather non-technical ‘manifesto’ in which we introduce and discuss conceptually our framework and outline a long term research agenda on it. In particular, we discuss our framework in more detail in Section 2, illustrate and discuss preplay offers and offer-induced game transformations in Section 3 and introduce normal form games with preplay negotiations phase in Section 4. Then we analyze solution concepts for 2-player normal form games with preplay offers under various assumptions for the preplay negotiation phase and obtain results for existence of efficient negotiation strategies of both players in Section 5. We extend and briefly discuss the framework to coalitional preplay offers and negotiations in N -player games, as well as to extensive form games with inter-play offers for side payments in Section 6. We end the paper with discussion of related work in Section 7 and concluding remarks and directions for further study in Section 8.

The ideas presented here will be developed in detail in a series of technical companions starting with [Gor12] and [GT].

2 Non-cooperative games with preplay offers: the conceptual framework

In this section we provide a more detailed description of preplay offers, discuss some motivating examples, and lay down several extra conditions that play a role in determining the outcome of the negotiation phase.

2.1 Preplay offers in more detail

We assume that any preplay offer by A to B is *binding for A*, conditional on B playing the strategy s specified by A^2 . However, such offer *does not* create any obligation for B and therefore it does not transform the game into a cooperative one, for B is still at liberty to choose his strategy when the game is actually played. In particular, after her offer A does not

²We will not discuss here the mechanism securing the payments of the preplay offers after the play if the conditions are met. That can be done by a legal contract, or by using a trusted third party, etc.

know before the game is played whether B will play the desired by A strategy s , and thus make use of the offer or not. Furthermore, several such offers can be made, possibly by different players, so the possible rational behaviours of the payers game maintain, in principle, all their complexity. The key observation applying to this assumption, is that *after any binding preplay offer is made, the game remains a standard non-cooperative normal form game, only the payoff matrix changes according to the offer.*

Several previous studies have considered similar or related frameworks, incl. (chronologically) [Ros75, Gut78, Kal81, Gut87, DS91, Var94b, Var94a, aMT06] and most notably the more recent [JW05, EP11]. We discuss and compare related work in more detail in Section 7.

2.2 Motivating examples

First, we introduce the following notation: $A \xrightarrow{\delta/\sigma_B} B$ denotes an offer made by player A to pay an amount δ to player B after the play of the game if player B plays strategy σ_B .

Prisoners' Dilemma Consider a standard version of the Prisoner's Dilemma (PD) game in Figure 1. The only Nash Equilibrium (NE) of the game is (D, D) , yielding a payoff of

	C	D
C	4, 4	0, 5
D	5, 0	1, 1

Figure 1: A Prisoner's Dilemma

$(1, 1)$. Now, suppose $Row \xrightarrow{2/C} Column$, that is, player Row makes to the player $Column$ a *binding offer* to pay her 2 units of utility (hereafter, *utils*) after the game if $Column$ plays C . That offer transforms the game by transferring 2 utils from the payoff of Row to the payoff of $Column$ in every entry of the column where $Column$ plays C , as pictured in Figure 2.

	C	D
C	2, 6	0, 5
D	3, 2	1, 1

Figure 2: An offer to cooperate by player Row.

In this game player Row still has the incentive³ to play D , which strictly dominates C for him, but the dominant strategy for $Column$ now is C , and thus the only Nash equilibrium is (D, C) with payoff $(3, 2)$ – strictly dominating the original payoff $(1, 1)$.

Thus, even though player Row will still defect, the offer he has made to player $Column$ makes is strictly better for $Column$ to cooperate.

Of course, $Column$ can now realize that if player Row is to cooperate, an extra incentive is needed. That incentive *can be created* by an offer $Column \xrightarrow{2/C} Row$, that is, if $Column$, too, makes an offer to Row to pay him 2 utils after the game, if player Row cooperates. Then the game transforms, as in Figure 3.

³Intuitively, *having the incentive* to play a strategy should be understood as realizing that that strategy is not dominated. Later on we will provide a formal and abstract notion of equilibrium, which will rule out dominated strategies to be part of the solution of a game.

	<i>C</i>	<i>D</i>
<i>C</i>	4, 4	2, 3
<i>D</i>	3, 2	1, 1

Figure 3: A second offer, by player Column.

In this game, the only Nash equilibrium is (C, C) with payoff $(4, 4)$, which is also Pareto optimal. Note that this is the same payoff for (C, C) as in the original PD game, but now both players have created incentives for their opponents to cooperate, and have thus escaped from the trap of the original Nash equilibrium (D, D) .

Remark 1 *Clearly, preplay offers can only work in case when at least part of the received payoff can actually be transferred from a player to another. They obviously cannot apply to scenarios such as the original PD, where one prisoner cannot offer to the other to stay in prison for him, even if they could communicate before the play.*

Battle of the Sexes Consider now a typical instance of the Battle of the Sexes (Figure 4), where we call the column player *Him* and the row player *Her*. The game has two NE: one

Her/Him	<i>Ballet</i>	<i>Soccer</i>
<i>Ballet</i>	4, 2	1, 1
<i>Soccer</i>	0, 0	2, 4

Figure 4: The Battle of the Sexes.

preferred by Her: $(Ballet, Ballet)$, and the other – by *Him*: $(Soccer, Soccer)$.

An offer $Him \xrightarrow{1/Soccer} Her$ from Him to Her transforms the game to the one in Figure 5 which is biased in favour of Her.

	<i>Ballet</i>	<i>Soccer</i>
<i>Ballet</i>	4, 2	1, 1
<i>Soccer</i>	1, -1	3, 3

Figure 5: The Battle of the Sexes, transformed by an offer by *Him* favouring *Her*.

That bias, however, can be neutralized by a ‘matching counter-offer’ $Her \xrightarrow{1/Ballet} Him$, which transforms the game to the one in Figure 6.

Both NE profiles yield the same payoffs here and, besides, they are both Pareto optimal and ‘fair’ for both parties. Yet, because of the symmetry, the question of *which Nash equilibria to choose* remains. That symmetry could be broken if a player is able to *signal* to the other player the strategy he would be actually playing. In this setting a signal from *Him* to *Her* can be realized as a further *vacuous offer* for payment of 0 made by any of the players, e.g. $Him \xrightarrow{0/Soccer} Her$ or $Him \xrightarrow{0/Ballet} Her$, which does not change the payoff matrix but only serves to indicate to the other player for which of the two equivalent Nash equilibria to play.

Note that the unilaterally made initial offer by *Him* to *Her* was a self-sacrificing move that put *Him* in a relatively disadvantaged position. As we will see further, that situation can make it non-beneficial for either of the players to make a first offer, even though they

	<i>Ballet</i>	<i>Soccer</i>
<i>Ballet</i>	3, 3	1, 1
<i>Soccer</i>	0, 0	3, 3

Figure 6: The Battle of the Sexes, further transformed by a matching ‘counter-offer’.

together could eventually achieve an improvement of both payoffs compensating the initial sacrifice. This problem can be avoided if we allow *conditional offers* as follows: Him can make an offer $Him \xrightarrow{1/Soccer} Her$, but now, *conditional on Her making to Him the matching counter-offer* $Her \xrightarrow{1/Ballet} Him$, which we hereafter denote as $Him \xrightarrow{1/Soccer \mid 1/Ballet} Her$. The idea is that, unlike the so far considered unconditional offers, *Him’s conditional offer* is only confirmed and enforced if *Her* does make the required counter-offer, else it is cancelled and nullified before the play of the game. We will introduce formally and discuss conditional offers in detail further.

2.3 Additional optional assumptions

There are several important additional assumptions that, depending on the particular scenarios under investigation may, or may not, be realistically made. We therefore do not commit to any of them, but we acknowledge that each of them can make a significant difference in the behaviour and abilities of players to steer the game in the best possible direction for them. So, we consider the possible options for each of them separately and study the consequences under the various combinations of assumptions.

Revocability of offers. Once made, offers may, or may not, be withdrawn during the negotiations phase. Both cases are reasonable and realistic, and we consider each of them separately.

Value of time. Time, measured discretely as the number of explicitly defined steps/rounds of the negotiations, may or may not have value, i.e. players may, or may not, strictly prefer a reward in the present to the same reward in the future. Moreover, time may have the same value for all players, or may be more, or less, valuable for each of them depending on their *patience*.

- In the case when time is of no value, players can keep making and withdrawing offers (if allowed to do so) at no extra cost. Intuitively the effect should be the same as if withdrawn offers were never made.
- In the case when time is of value, making unacceptable or suboptimal offers or withdrawing offers that were made earlier should intuitively lead to inefficient negotiation and, consequently, strategies involving such offers or withdrawing offers would not be subgame perfect equilibrium strategies.

The order of making offers. The order in which offers are made by the different players can be essential, especially in case of irrevocable offers. In such cases we assume that the order in which players can make offers is set by a separate, exogenous protocol which is an added component of the preplay negotiations game; for instance, it can be random. Alternatively, the offers may be required to be made simultaneously by all players, as in [JW05] and [EP11].

Rejection of offers. Once made, offers may, or may not, be officially rejected before the play. A rejection by a player B of an offer made to her by a player A has the same practical effect as a withdrawal of the offer by A , but the choice to withdraw or not is now in the hands of B . Both options can be reasonable in different scenarios.

Conditionality of offers. Offers may be *unconditional*, i.e., not depending on the acceptance or rejection by the player to whom the offer is made, or *conditional* upon an expected (suggested or demanded) counter-offer by the player to whom the offer was made. Acceptance of a conditional offer means both acceptance of the offer and making the expected counter-offer. We emphasize that *after acceptance, a conditional offer results into a pair of unconditional offers, and it therefore transforms the current game into another non-cooperative game*. Rejection/withdrawal means cancellation of both of these unconditional offers. This is a special form of rejection/withdrawal of an offer that can be reasonably assumed under some circumstances (e.g. possibility for extended communication and for a low-cost negotiations), but not in other. We will consider both cases separately.

3 Preplay offers and induced game transformations

In this section we describe the game transformations induced by preplay offers in a general and more technical fashion (Subsection 3.1). We discuss some of their basic properties and then formally extend the framework of preplay offers with conditional offers and briefly discuss withdrawals of offers, punishments and offers contingent on strategy profiles in Subsection 3.2.

3.1 Transformations of normal form games by preplay offers

	B_1	\dots	B_j	\dots
A_1	\dots	\dots	a_{1j}, b_{1j}	\dots
A_2	\dots	\dots	a_{2j}, b_{2j}	\dots
\dots	\dots	\dots	\dots	\dots
A_i	\dots	\dots	a_{ij}, b_{ij}	\dots
\dots	\dots	\dots	\dots	\dots

Figure 7: A general 2-player game

Here we formally define the notion of transformation induced by a preplay offers. For technical convenience we consider general 2-player game with a payoff matrix given in Figure 7; the case of N-player games is a straightforward generalization, but see Section 6.1 for more general types of offers in that case.

Suppose player A makes a preplay offer to player B to pay her additional utility⁴ $\alpha \geq 0$ if B plays B_j . Recall that we denote such offer by $A \xrightarrow{\alpha/B_j} B$. It transforms the payoff matrix of the game as indicated in Figure 8.

⁴The reason we allow vacuous offers with $\alpha = 0$ is not only to have an identity transformation at hand, but also because such offers can be used by players as signaling, to enable coordination, as in the transformed BoS game in Example 2.2.

	B_1	\dots	B_j	\dots
A_1	\dots	\dots	$a_{1j} - \alpha, b_{1j} + \alpha$	\dots
A_2	\dots	\dots	$a_{2j} - \alpha, b_{2j} + \alpha$	\dots
\dots	\dots	\dots	\dots	\dots
A_i	\dots	\dots	$a_{ij} - \alpha, b_{ij} + \alpha$	\dots
\dots	\dots	\dots	\dots	\dots

Figure 8: A general 2-player game with an offer.

We will call such transformation of a payoff matrix a **primitive offer-induced transformation**, or a POI-transformation, for short.

Several preplay offers can be made by each players. Clearly, the transformation of a payoff matrix induced by several preplay offers can be obtained by applying the POI-transformations corresponding to each of the offers consecutively, in any order. We will call such transformations **offer-induced transformations**, or OI-transformations, for short. Thus, every OI-transformation corresponds to a *set* of preplay offers, respectively a set of POI-transformations.

Note that the set generating a given OI-transformation need not be unique. For instance, A can make two independent offers $A \xrightarrow{\alpha_1/B_j} B$ and $A \xrightarrow{\alpha_2/B_j} B$ equivalent to the single offer $A \xrightarrow{\alpha_1+\alpha_2/B_j} B$.

The general mathematical theory of OI-transformations is studied in more detail separately, in [Gor12]. Here we only mention some observations about the game-theoretic effects of OI-transformations, which will be useful later on.

1. An OI-transformation does not change the sum of the payoffs of all players in any outcome, only redistributes it. In particular, OI-transformations preserve the class of zero-sum games.
2. An OI-transformation induced by a preplay offer by player A does not change the preferences of A regarding her own strategies. In particular, (weak or strict) dominance between strategies of player A is invariant under OI-transformations induced by preplay offers of A , i.e.: a strategy A_i dominates (weakly, resp. strongly) a strategy A_j before a transformation induced by a preplay offer made by A if and only if A_i dominates (weakly, resp. strongly) A_j after the transformation.
3. The players can collude to make *any* designated outcome, with any redistribution of its payoffs, a dominant strategy equilibrium, by exchanging sufficiently high offers to make the strategies generating that outcome with that redistribution of the payoffs, strictly dominant.

As we have seen, preplay offers can transform the game matrix radically. However, we note that not every matrix transformation that preserves the sums of the payoffs in every outcome can be induced by preplay offers. In particular, this is the case if the transformed matrix differs from the original one in only one payoff. For general necessary and sufficient condition for a normal form game to be obtained from another by preplay offers see [Gor12].

A major question arising is what should be regarded as a *solution* of a strategic game allowing binding preplay offers. The possible answers to that question crucially depend on the

additional assumptions discussed in the introduction, and on the mechanism and procedure of 'preplay negotiations', and will be discussed further.

3.2 Extending preplay offers and OI-transformations

3.2.1 Conditional offers

A property of unconditional offers is that they always decrease the proponent's payoff at some outcomes, and hence making an unconditional offer comes with a cost. As in the Battle of the Sexes (Example 2.2), this can be a hindrance for making mutually beneficial offers and we will discuss this problem in more details in Section 5.1.2. Furthermore, often in real life situations players who make such preplay offers expect some form of reciprocity from their fellow players and make their offers conditional on an expected 'return of favour'.

For these reasons, we also enable players to *suggest* a transformation of the starting game, by making a *conditional offer*, i.e., making an offer to an opponent for playing a certain strategy, in exchange for a similar 'counter-offer' from that opponent. More precisely, every conditional offer, denoted as $A \xrightarrow{\alpha/\sigma_B \mid \beta/\rho_A} B$ is associated with a **suggested transformation** of the starting game \mathcal{G} into a game $\mathcal{G}(X)$ where $X = \{A \xrightarrow{\alpha/\sigma_B} B, B \xrightarrow{\beta/\rho_A} A\}$.

Two responses of the recipient of a conditional offer $A \xrightarrow{\alpha/\sigma_B \mid \beta/\rho_A} B$ are possible: it can be *accepted* or *rejected* by the player receiving it. If rejected, the offer is immediately cancelled and does not commit any of the players to any payment, and therefore it does not induce any transformation of the game matrix. If accepted, the **actual transformation** induced by the offer is the suggested transformation defined above.

Two important observations:

- an unconditional offer has the same effect as an accepted conditional offer with a trivial counter-offer where $\beta = 0$.
- a conditional offer can be seen as the proposal of two separate unconditional offers *that can only be enforced together*.

Conditional offers can be made to different players. Multiple conditional offers can be made to the same player, contingent upon same or different strategies of the recipient and the proposer, too.

3.2.2 Withdrawals of offers and transformations induced by them

Withdrawal of an offer, i.e. a 'change of mind' by the player who makes the offer, can be simulated in a sense by matching the amount α offered by A contingent on a given strategy σ_B of B by offers from A to B for the same amount α , contingent on *every other strategy* of B . However, his simulated offer withdrawal is costly for A and, while the preferences on all outcomes remain the same for both players, the game is no longer the same. A *proper* withdrawal of A 's offer can only be achieved if A extends his offer to cover all possible strategies of B and B offers in return to pay back the amount α to A unconditionally, that is, makes offers of amount α to A , contingent on all strategies of A .

If a player A withdraws an unconditional offer $A \xrightarrow{\alpha/\sigma} B$ made earlier by her, the transformation $\mathcal{G}(X_A)$ of the game \mathcal{G} induced by that offer must be reverted. The **withdrawal** of a transformation $\mathcal{G}(X_A)$ is again a transformation, induced by the (fictitious) negative offer

$A \xrightarrow{-\alpha/\sigma} B$. Likewise, the transformation associated with withdrawal of an earlier made and accepted conditional offer $A \xrightarrow{\alpha/\sigma|\beta/\rho} B$ consists of the reversal transformations of the two constituent unconditional offers.

A withdrawal can thus be seen as a sort of unconditional reversal of payments that have already been enforced in a previous accepted offer. Thereby the possibility of performing a withdrawal strictly depends on the previous history of the negotiation and this feature will be of fundamental importance when treating preplay negotiations as full-fledged extensive games.

3.2.3 Punishments

Offers, be them conditional or not, can be modeled as future rewards uniformly associated with specific moves. Likewise, one can think of negative offers, or **punishments** as dual operations associated with opponents' moves that transfer payments in the opposite direction. Formally, punishments induce transformations of the type $A \xrightarrow{-\alpha/\sigma_B} B$, where $\alpha \in \mathbb{R}_+$.

	L	R		L	R
U	$a - \alpha, b + \alpha$	c, d		U	$a + \alpha, b - \alpha$
D	$e - \alpha, f + \alpha$	g, h		D	$e + \alpha, f - \alpha$

Figure 9: On the left: transformation of a game by an unconditional offer for playing L ; on the right: transformation of the game by an unconditional punishment for the same strategy. Notice the similarity with making and withdrawing an offer.

Even though preplay offers are assumed to be non-negative, just like withdrawals, negative offers/punishments can be *simulated* by non-negative ones in the sense of effecting 'equivalent' transformed games, in terms of the players' preferences over the outcomes. Namely, note that an outcome is dominant strategy equilibrium in the game resulting from A punishing player B for playing s if and only if it is dominant strategy equilibrium in the game resulting from A rewarding B for playing any strategy other than s (of the same) amount. So, a player A may offer a payment $\alpha > 0$ to player B for *every strategy of B , except a designated, undesirable for A , strategy σ_B* . The net effect of such offer is that, in the transformed game player B would be *punished* by not receiving the offered amount α if he plays the strategy σ_B . Again, such simulation is strictly beneficial for player B , whereas player A is paying a price for the ability to penalize B , so we do not adopt this simulation further.

Also, note the following:

- as in the case of unconditional offer, a punishment has no effect on the relative dominance relation among the punisher's strategies.
- however, unlike unconditional offers, punishments are always rewarding for the punisher, as each outcome after the punishment makes the punisher at least as better off as the same outcome before the punishment.

These observations show that allowing players to punish each other technically amounts to empowering them with the capacity of withdrawing offers that they have never made. What is more, a punishment transforms the game into a one that is more beneficial for the punisher, independently of what will actually be played; think of a player punishing all his opponents for

playing all their strategies. Allowing such form of unrestrained punishments goes against our fundamental principle that preplay offers commit only the proposer and not the receiver and. Moreover, allowing punishments can have detrimental effects on understanding the course of the game with preplay offers so we will not consider punishments further in the paper. Still, we note that more controlled forms of punishment have been considered in the literature, for instance players sacrificing part of their payoffs to punish their fellow players [FG02], or threatening them by playing strategies minimizing their payoff, independently of the cost for the player himself [Col90]. We believe that allowing these milder forms of punishment can help understanding several scenarios of preplay negotiations. Yet, we leave their treatment to future investigation.

3.2.4 Offers contingent on strategy profiles

More complex offers were considered in [JW05] and [EP11], contingent not just on the recipient's strategy but on an entire strategy profile (i.e., on an outcome). It is not difficult to show that such an offer, that redistributes the payoffs of only one outcome, cannot be effected by OI-transformations of the type we consider here. This also follows from a more general result in [Gor12].

For conceptual reasons we do not adopt offers contingent on actions of the offerer in our study. Yet, we compare in detail the framework of [JW05] and [EP11] with ours in Section 7.

4 Normal form games with preplay negotiations phase

Clearly, players would only be interested in making preplay offers inducing payoffs that are optimal for them. Therefore, rational players are expected to 'negotiate' in the preplay phase the play of Pareto optimal outcomes. In particular, if the game has a unique strictly Pareto dominant outcome then the players can negotiate a transformation of the game to make it the (unique) dominant strategy equilibrium. Yet, players that are getting lesser shares of the total payoff may still want to negotiate a redistribution, so even in this case the outcome of the preplay negotiations is not a priori obvious.

Similarly to [JW05], our setting for normal form games with preplay offers begins with a given 'starting' normal form game \mathcal{G} and consists of two phases:

- A *preplay negotiation phase*, where players negotiate on how to transform the game \mathcal{G} by making offers, accepting or rejecting conditional offers they receive, and possibly withdrawing old ones.
- An *actual play* phase where, after having agreed on some OI-transformation X in the previous phase, the players play the game $\mathcal{G}(X)$.

Henceforth we use the acronym PNG for 'preplay negotiation game'.

Major questions that we set out to study are:

- What constitutes an optimal/rational/efficient negotiation strategy?
- When can players agree upon Pareto optimal outcomes in their preplay negotiations if playing rationally?

- What can, or should, players agree upon in the preplay negotiations phase when the original game has several Pareto optimal outcomes?

In this section we first briefly discuss rationality assumptions and solution concepts for the normal form games. Then we model preplay negotiations as extensive games of perfect or almost perfect (if simultaneous offers are allowed) information and define generic notions of rationality and solution concepts for these preplay negotiations phase. In particular, we will define solutions of the preplay negotiations phase as outcomes of SPE strategy profiles. Finally, we discuss the combined solution concepts for the entire games outlined above.

4.1 Solutions and values of normal form games

We will be using i, j, \dots for variables ranging over players, while A, B, \dots will denote individual players.

4.1.1 Normal form games

Let $\mathcal{G} = (N, \{\Sigma_i\}_{i \in N}, u)$ be a normal form game (hereafter abbreviated as NFG), where $N = \{1, \dots, n\}$ a finite set of players, $\{\Sigma_i\}_{i \in N}$ a family of strategies for each player and $u : N \times \prod_{i \in N} \Sigma_i \rightarrow \mathbb{R}$ is a **payoff function** assigning to each player a utility for each strategy profile. The game is played by each player i choosing a strategy from Σ_i . The resulting strategy profile σ is the **outcome** of the play and $u_i(\sigma) = u(i, \sigma)$ is the associated payoff for i . An outcome of a play of the game \mathcal{G} is called **maximal** if it is a Pareto optimal outcome with the highest sum of the payoffs of all players.

4.1.2 Solution concepts and solutions

Let \mathbf{G}_N be the set of all normal form games for a set of players N . By **solution concept for \mathbf{G}_N** we mean a map \mathfrak{S} that associates with each $\mathcal{G} \in \mathbf{G}_N$ a non-empty set $\mathfrak{S}(\mathcal{G})$ of outcomes of \mathcal{G} , called the **\mathfrak{S} -solution of the game**. At times we will talk about players' strategies that are consistent with some solution concept. For a player i , we denote \mathfrak{S}_i to be the restriction of the mapping \mathfrak{S} to i returning, instead of full outcomes, only strategies of player i consistent with \mathfrak{S} in the sense that $\mathfrak{S}_i(\mathcal{G}) = \{\sigma_i \in \Sigma_i \mid \sigma \in \mathfrak{S}(\mathcal{G})\}$. Slightly abusing notation we will also consider mappings of the form \mathfrak{S}_{-i} to indicate the mapping $\mathfrak{S}(\mathcal{G})$ restricted to player i 's opponents. Solution concepts formalize the concepts of rationality of the players in the strategic games. A \mathfrak{S} -solution of a strategic game \mathcal{G} basically tells us what outcomes of the game the players could, or should, select in an actual play of that game, if they adopt the solution concept \mathfrak{S} .

In this work we do not commit to a specific solution concept for the normal form games but we assume that the one adopted by the players satisfies the necessary condition that *every outcome in any solution prescribed by that solution concept must survive iterated elimination of strictly dominated strategies*. We will call such solution concepts **acceptable**. This condition reflects the assumption that players would never play strategies that are dominated, and that this exclusion is a common knowledge amongst them and can be used in their strategic reasoning. Thus, the weakest acceptable solution concept is the one that returns all outcomes surviving iterated elimination of strictly dominated strategies.

Games for which the solution concept \mathfrak{S} returns a single outcome will be called **\mathfrak{S} -solved**. For instance, every game with a strongly dominating strategy profile is \mathfrak{S} -solved for any

acceptable solution concept \mathfrak{S} . Games for which \mathfrak{S} returns only maximal outcomes will be called **optimally \mathfrak{S} -solvable**. If for every player all these maximal outcomes provide the same payoffs, we call the game **perfectly \mathfrak{S} -solvable**. Games that are \mathfrak{S} -solved and perfectly \mathfrak{S} -solvable (i.e., \mathfrak{S} returns one maximal outcome) will be called **\mathfrak{S} -perfectly solved**.

A desirable objective of a preplay negotiation is to transform the starting NFG into a perfectly \mathfrak{S} -solvable one. Ideally, it should moreover be a \mathfrak{S} -perfectly solved one, but this is not always possible: cf. any symmetric coordination game like, e.g., the transformed Battle of the Sexes game on Figure 6.

4.1.3 Players' expected values of a game

It is necessary for the preplay negotiation phase for each player to have an **expected value** of any NFG that can be played. Naturally, that expected value would depend not only on the game but also on the adopted solution concept and on the player's level of risk tolerance. A risk-averse player would assign as expected value the minimum of his payoffs over all outcomes in the respective solution, while a risk-neutral player could take the probabilistic expected value of these payoffs, etc. Note that the expected value of any \mathfrak{S} -solved game for any player i naturally should equal the payoff for i from the only outcome in the solution.

For sake of definiteness, unless otherwise specified further, we adopt here the conservative, risk-averse approach and will define for every acceptable solution concept \mathfrak{S} , game \mathcal{G} and a player i , the expected value of \mathcal{G} for i relative to the solution concept \mathfrak{S} to be:

$$v_i^{\mathfrak{S}}(\mathcal{G}) = \max_{\sigma_i \in \mathfrak{S}_i(\mathcal{G})} \min_{\sigma_{-i} \in \mathfrak{S}_{-i}(\mathcal{G})} u_i(\sigma)$$

4.2 Preplay negotiation games

The purpose of the preplay negotiations game is to reach a 'best possible' agreement between all players on an OI transformation of the original game \mathcal{G} . Before we define these concepts more explicitly, we need to introduce and discuss some preliminary notions. This subsection introduces main features of preplay negotiation games such as the concept of history, the order of moves, the possibility of players come to a disagreement, and finally a notion of solution for these games.

4.2.1 Moves and histories in preplay negotiations games

Depending on some of the optional assumptions, the players can have several possible moves in the preplay negotiations game. Let us consider the most general case, where both conditional offers and withdrawals of offers are allowed. Then the moves available to the player whose turn is to play depend on whether or not he/she has received since his/her previous move any conditional offers. If so, we say that the player has **pending conditional offers**. The possible moves of the player in turn are as follows.

1. If the player has no pending conditional offers, he/she can:
 - (a) *Make an offer* (conditional or not).
 - (b) *Pass*.
 - (c) (Optional) *Withdraw an offer* he/she has made at a previous move.

(d) (Optional) *Opt out* (see Section 4.2.3).

2. If the player has pending conditional offers, for each of them he/she can:

- (a) *Accept the pending offer* by making the requested counter-offer to the player who has made the conditional offer, and then make an offer of his/her own or pass or opt out (when available).
- (b) *Reject the pending offer*, and then make an offer of his/her own or pass or opt out (when available).

If *all players have passed* at their last move, or *any player has opted out*, the preplay negotiations game is over.

Note that while conditional offers always require a response by the player receiving them (acceptance or rejection), this is not the case for unconditional offers. The effect of the latter ones is to immediately update the normal form game as specified by the offer.

We say that an offer of the game is **passing** if its acceptance by the opponents is followed by a pass of the proponent. In other words, the one making the offer would be happy to end the game with the suggested transformation. Likewise, an acceptance is passing if, once declared, it is followed by a pass move of the same player. In other words, with a passing acceptance a player declares agreement to terminate the game with the proposed transformation. As we will see, passing moves (i.e. offers or acceptances that are passing), are the only way for players to terminate the game in agreement and *the only way to effectively deviate from undesired outcomes*.

We now define the notion of a **history** in a preplay negotiations games as a finite or infinite sequence of admissible moves by the players who take their turns according to an externally set protocol (see further). Every finite history in such a game is associated with the current NFG: the result of the OI-transformation of the starting game by all offers that are so far made, accepted (if conditional) and currently not withdrawn. The current NFG of the empty history is the input NFG of the preplay negotiations game.

A **play** of a preplay negotiations game is any finite history at the end of which the preplay negotiations game is over, or any infinite history.

In order to eventually define realistic solution concepts for preplay negotiations games we need to endow every history in such games with value for every player. Intuitively, **the value of a history** is the value for the player of the current NFG associated with that history, in the case of non-valuable time, and the same value accordingly discounted in the case of valuable time.

These notions can be defined formally, e.g., as in Osborne and Rubinstein's bargaining model [Rub82, OR90, OR94]. This is done in the technical companion [GT] of the present paper.

4.2.2 The order of making moves

Depending on some of the optional assumptions, the players can have several possible moves in the preplay negotiations game, which they can make simultaneously, in several rounds, or by taking turns according to some externally set protocol or by a randomized procedure. We will focus on turn-based negotiations games but will not discuss here how the order of making moves is determined. We only state that when time is not valuable the order in which the players take turns to make their offers is irrelevant for the eventual solution, and multi-round

negotiations with simultaneous offers are reducible to multi-round turn-based negotiations with any order of making moves.

4.2.3 Disagreements

The PNG may terminate if all players pass at some stage, in which case we say that the players have reached agreement, or may go on forever, in which case the players have failed to reach agreement; we call such situation a (*passive*) *disagreement* and we denote any such infinite history with D . We will not discuss disagreements and their consequences here, but will make the explicit assumption that *any agreement is better for every player than disagreement* in terms of the payoffs, e.g. by assigning payoffs of $-\infty$ in the entire game for each player if the PNG evolves as a disagreement. However, we also outline a more flexible and possibly more realistic alternative, whereby players can explicitly express tentative agreements with the status quo before every move they make, essentially by saying “So far so good, but let me try to improve the game further by offering ...”, or express disagreements, by essentially saying “No, I am not happy with the way the negotiations have developed since the last time I agreed, so I’d like to improve the game by offering instead ...”. This type of negotiations involves, besides the other moves listed above, also formal *statements of acceptance or non-acceptance of the current NFG*, where the input NFG is automatically accepted by all players and at every stage of the negotiations, the **current NFG** is the one on which they are currently negotiating by making offers, whereas the **currently accepted NFG** is the last current one for which all players have explicitly stated acceptance. Then if at any stage of the PNG any player is currently unhappy and realizes that he cannot improve further because of the other players not willing to accept his best conditional offers, then he can terminate the negotiations by explicitly **opting out**, which would revert to the current game to the currently accepted NFG.

4.2.4 Preplay negotiations games and their solutions

Preplay negotiation games (PNG) can be defined generically, in terms of the notion of history, as turn-based, possibly infinite, extensive form games. A preplay negotiation game starts with an input NFG \mathcal{G} and either ends with a transformed game \mathcal{G}' or goes on forever, which we discuss further. The **outcome of a play of the PNG** is the resulting transformed game \mathcal{G}' in the former case and ‘Disagreement’ (briefly D) in the latter case.

The notion of PNG and its outcomes can be defined formally, e.g., like a bargaining game as in Osborne and Rubinstein’s bargaining model [Rub82, OR90, OR94]. This is done in the technical paper [GT].

Solution of PNG By a solution of a PNG we mean *the set of all transformed normal form games that can be obtained as outcomes of plays effected by subgame perfect equilibrium strategy profiles in the PNG.*

4.3 Feasible moves and efficient negotiation strategies

In order to understand how solutions of PNG look and to make statements about existence of ‘good’ solutions, we need to discuss the notions of ‘feasibility of moves’ and efficiency of negotiation strategies’. We emphasize that in this context we will use the term ‘efficient’ not

in its standard game-theoretic sense, i.e., by applying it to outcomes, but to the way outcomes are reached.

4.3.1 Feasible offers and moves

In principle, players can make offers that would induce transformations decreasing their expected value of the game. However, generally such offers would not be realistic to expect⁵. We say that a player's offer is **weakly feasible** if it does not decrease that player's expected value of the game in the transformed by that offer game; the offer is **feasible** if it strictly increases that expected value.

This is a generic notion of feasibility of offers, which needs to be extended further to the notion of feasible moves in the PNG. The latter is specific to some of the optional additional assumptions which we will discuss in more detail for the 2-player case in the next section. We will show that, in order for a player's strategy in the preplay negotiation phase to be a part of a solution, i.e., subgame perfect equilibrium, it must only involve weakly feasible moves.

4.3.2 Minimal offers

With their preplay offers players want to create incentives for the other players to play desired strategies. So, feasibility is a necessary condition for an offer to be made in an actually played PGM, but it is not sufficient for it to be a part of a subgame perfect equilibrium strategy. Clearly, an optimal offer from a player to another would be a *minimal feasible one* providing a sufficient incentive for the recipient of the offer to play the desired transformation, but not more than that. The question of what is a minimal offer that achieves such objective crucially depends on the adopted solution concept and, in particular, on the rationality assumptions and reasoning skills of the recipient. For instance, if the players know the solution, induced by the adopted solution concept, of the starting normal form game \mathcal{G} then they also know which outcomes can be selected among the ones surviving the iterated elimination process. Thereafter, if a player A wants to induce with a preplay offer another player B to play a given strategy σ_B then, for any acceptable solution concept, it would suffice for A to make any sufficiently large offer that would turn σ into a strictly dominant strategy for B . But, such offer may be prohibitively costly or, depending on the solution concept and the rationality assumptions for B , unnecessarily generous. For instance, when a player B receives an offer $A \xrightarrow{\delta/\sigma_B} B$, he should naturally expect that A considers playing A 's best response to σ_B , so B can anticipate the outcome of the transformed game, and if B considers that outcome better than his current expected value, that should suffice for A .

A technical detail: it is often the case that no minimal offer exists that achieves the objective, e.g., to turn the desired strategy into a *strictly* dominant one. For instance, if it suffices for A to pay to B any amount that is *greater than* d for that purpose, then any offer of $d + \epsilon$, for $\epsilon > 0$, should do. Clearly, however, there is a practical minimum beyond which a player in question would not bother optimizing any further, so we will often refer to offers of payments d^+ meaning $d + \epsilon$ for 'sufficiently small $\epsilon > 0$ ' without specifying the value of ϵ , but still allowing its further reduction, as long as it remains strictly positive.

⁵But, they may be necessary in some circumstances, e.g., when conditional offers are not allowed but withdrawals of offers are.

4.3.3 Efficient negotiation strategies

Definition 2 (Efficient negotiation strategies) *Intuitively, a strategy in the preplay negotiation game is an efficient negotiation strategy if it only involves making (minimal) feasible offers, it passes once they are accepted, and – in the case when conditional offers are allowed – at no point prescribes withdrawal of earlier made offers. It is strongly efficient if the vector of payoffs of the outcome it attains is a redistribution of the vector of payoffs of a maximal outcome.*

A number of important relevant questions arise:

- Is it the case that every subgame perfect equilibrium strategy of a PNG is an efficient negotiation strategy?
- If not, can the inefficient ones be replaced by efficient ones generating the same, or at least as good solution?
- Is it the case that, when generalized offers are allowed there are subgame perfect equilibria where each player need not make more than one offer (combining all intended subsequent offers by that player)? Does this depend on the number of players?
- Under what conditions can a given (maximal) Pareto optimal outcome in the starting NFG become the unique outcome of the final NFG?

To answer these questions we need an analysis of the solutions of the PNG game. The next section will provide such partial analysis for the case of two players.

5 Analyzing two-players normal form games with preplay negotiations

Let the players be A and B . We will only sketch some technical results but will defer the proof details to the technical paper [GT].

As discussed earlier, there are several important optional assumptions that can be adopted or not, and that choice would affect essentially the nature and the solutions of the preplay negotiations games. The most important ones are:

1. Whether time is valuable.
2. Whether conditional offers are allowed.
3. Whether offers can be withdrawn later.
4. Whether generalized offers (i.e., simultaneous multiple offers to several other players) are allowed.

Depending on the choices for each of these, 16 possible cases arise. In the case of 2 player games, the last item is of no particular importance, so we will not discuss it in detail here, but will simply assume that generalized offers are allowed. We analyze only the more interesting of the resulting 8 cases and only briefly discuss the others.

Before doing this, though, let us state a useful general result, also valid in the case of many players PNG. An extensive form game is said to have the **One Deviation Property**

(ODP) [OR94, Lemma 98.2] if, in order to check that a strategy profile is a Nash equilibrium in (some subgame of) that game, it suffices to consider the possible profitable deviations of each player not amongst *all* of its strategies (in that subgame), but only amongst the ones differing from the considered profile in the first subsequent move (in that subgame).

Lemma 3 (ODP for PNG) *Every PNG has the one deviation property.*

Proof sketch. It is easy to check that a strategy profile of a PNG is a subgame perfect equilibrium if and only if it is a subgame perfect equilibrium of the same PNG without disagreement histories as, notice, strategies leading to disagreement are cannot be used as credible threats. But, a PNG without disagreement histories is an extensive game of perfect information and finite horizon. By [OR94, Lemma 98.2] the PNG has the one deviation property. ■

5.1 Preplay negotiations with no conditional offers

We begin with the case of more restricted preplay negotiations, where conditional offers are not possible, or not allowed. As we will see further, the strategic reasoning in such preplay negotiations games is rather different from the case with conditional offers, because any player who makes an unconditional offer puts himself in a disadvantaged position by offering unilaterally a payment to the other player and thus transforming the payoff matrix to the other player's advantage. Therefore, generally, players are more interested in receiving, rather than in making, unconditional offers and this affects essentially the preplay negotiations phase.

Here we focus on the *locally rational* behavior of players exchanging unconditional offers, by first determining the *best* (for the offerer) rational unconditional offer that a player can make on a given 2-player normal form game. Then we illustrate with some examples possible evolutions and outcomes of the preplay negotiation phase consisting of exchanging such best offers and draw some conclusions. Thus, here we analyze and illustrate the rationality of *moves*, rather than full-blown strategies, suggesting that every good notion of equilibrium used to analyze PNGs without conditional offers should take this rationality into account, while leaving so far open the question of how the subgame perfect equilibria in such PNGs are composed. We also leave untreated for now the question of how the value of time affects the outcomes of the preplay negotiations games in this case, by tacitly assuming that time is not valuable.

5.1.1 The effect of allowing withdrawals of unconditional offers

We first argue that when withdrawals of unconditional offers are allowed, conditional offers can be simulated, too, even though at the cost of some time delay. Indeed, if player A wants to make a conditional offer $A \xrightarrow{\alpha/B_j \mid \beta/A_i} B$ she can make the unconditional offer $A \xrightarrow{\alpha/B_j} B$ expecting the matching (or better) unconditional offer $B \xrightarrow{\beta/A_i} A$ from B . How can the receiver B guess the expected matching offer, if side communication is not possible or not allowed? Note that the offer $A \xrightarrow{\alpha/B_j} B$ has 2 effects: it changes the payoff table in a way beneficial for B and indicates that player A wants player B to play B_j . Therefore, B can naturally expect that (disregarding for a moment all other offers) A intends to play her best response to B_j . However, an offer from B to A may change A 's best response to B_j in a way, that would make it more beneficial for A , and at least as beneficial for B , if A plays another

strategy, say A_i . By inspecting the possibilities B can identify his options for matching offers that would make A 's unconditional offer worth her while. If B has more than one such options, he can guess and try. If the expected matching offer is not received in the next round of the preplay negotiations, A can subsequently withdraw her offer, but later can make it again, possibly repeating this 'ritual' until B eventually realizes what is expected from him and delivers it, or until A gives up expecting.

Thus, the case of unconditional offers with withdrawals is essentially reducible to the case where conditional offers are allowed, which will be treated further. We only note here that when time is valuable the simulation suggested above may be costly and leading to side effects, and leave the details for a further work.

5.1.2 Preplay negotiations with unconditional offers and no withdrawals

The case when no withdrawals of offers are allowed is essentially different. As we will see further, in this case the players can be genuinely obstructed from making the first offer without disadvantaging themselves, and this can be crucial for the outcome of the negotiations. We can distinguish 3 types of unconditional offers:

1. **vacuous offers**, of the kind $A \xrightarrow{0/\sigma} B$ for payment of 0. These can be used instead of passing, but also, more importantly, as a kind of *signaling*, i.e., indication that A expects B to play σ , for breaking the symmetry in case of symmetric games with several equivalent optimal equilibria (as in the BoS game 2.2).
2. **ϵ -offers**, of the kind $A \xrightarrow{\epsilon/\sigma} B$ for a small enough $\epsilon > 0$. These can be used similarly, for breaking the symmetry, when B has more than one best for him moves which, however, yield different payoffs for A . Using such a move, A can make any of these strictly preferable for B and, thus, can turn a weak equilibrium into a strict one, with minimal cost.
3. **effective offers**, of the kind $A \xrightarrow{d/\sigma} B$ for a (large enough) $d > 0$. These are the standard offers used to change the recipient's preferences and influence his choice of strategy in the 2nd phase.

It is easy to observe that in ideal negotiations between two players none of them needs to make two or more consecutive offers, in between which the other player has passed or made a vacuous offer. Indeed, no player would be better off by making offers in the same game contingent on two or more different strategies of the opponent; in fact, such multiple offers send to the opponent confusing signals. Furthermore, two or more offers by the same player that are contingent on the same strategy of the opponent can be combined into one.

So, leaving aside the question of who starts the preplay negotiations game, in the case where only unconditional and irrevocable offers are allowed, the PNG consists of a sequence of alternating offers made by the two players until both of them pass.

Thus, in order to capture the notion of efficient negotiations in this case, we need to analyze the question of *what are the best unconditional and irrevocable offers that a player can make on a given NFG?*

5.1.3 Computing the best unconditional offers of a player

What is a rational player's reasoning when considering making an unconditional and irrevocable offer to another player in a given NFG \mathcal{G} ? Suppose, player A considers making such an offer to player B . Then, for each strategy B_j of B , player A considers making an offer contingent on B playing B_j . To make sure that B will play B_j in the resulting game, it suffices to make the latter a strictly dominant strategy for B . The necessary payment for that, however, can be prohibitively high for A because after that payment A 's best response to B_j may yield a worse payoff than the current (e.g., maxmin) expected value for A of the original game. So, a more subtle reasoning is needed, as follows.

1. For each strategy B_j of B , player A looks at her best response to B_j . Suppose for now that it is unique, say A_{i_j} . Then, this is what B would expect A to play if B knows that A expects B to play B_j . In this case, A computes the minimal payment needed to make B_j not necessarily a strictly dominant strategy, but a *best response to A_{i_j}* , i.e., the minimal payment that would make the strategy profile $\sigma_{i_j,j} = (A_{i_j}, B_j)$ a Nash equilibrium. That payment is

$$\delta_{i_j,j}^A = \max_k (u_B(A_{i_j}, B_k) - u_B(\sigma_{i_j,j})).$$

If it is positive, or is 0 but reached not only for $k = j$ but also for other values of k , then, in order to break B 's indifference and make $\sigma_{i_j,j}$ a *strict* Nash equilibrium, A has to add to $\delta_{i_j,j}^A$ a small enough $\epsilon > 0$, thus eventually producing the minimal necessary payment δ_j^A .

2. If A 's best response to B_j is not unique, then A should compute the minimal payment δ_j^A needed to make B_j the best response of B to *each* of A 's best responses to B_j . Clearly, that should be the maximum of all $\delta_{i_j,j}^A$ computed above, possibly plus a small enough $\epsilon > 0$.
3. Once δ_j^A is computed, A computes her expected payoff in the transformed game $\widehat{\mathcal{G}}_{B_j}$ after an offer $A \xrightarrow{\delta_j^A/B_j} B$, which is:

$$v^A(\widehat{\mathcal{G}}_{B_j}) = u_A(\sigma_{i_j,j}) - \delta_j^A.$$

4. Finally, A maximizes over j :

$$v^A(\widehat{\mathcal{G}}) = \max_j v^A(\widehat{\mathcal{G}}_{B_j}).$$

If the maximum is achieved for more than one j , then A can choose any of them, or –better– the one yielding the least payoff for B , thus stimulating B to make her a further offer.

If this maximum is 0 and reached for only one value of j , then there is no need for A to make any offer, because in this case there is a unique Nash equilibrium in the game and A cannot make any offer that would improve on her payoff yielded by that Nash equilibrium. If the maximum is 0, but reached for more than one values of j , then A must still make a vacuous offer $A \xrightarrow{0/B_j} B$ in order to indicate to B for which Nash equilibrium she will play.

The reasoning for B is completely symmetric, eventually producing the value $v^B(\widehat{\mathcal{G}})$.
The definition of $v^A(\widehat{\mathcal{G}})$ implies the following:

Proposition 4 *Given the NFG \mathcal{G} , the value $v^A(\widehat{\mathcal{G}})$ is the best payoff that player A can guarantee as a result of the players playing any Nash equilibrium induced by an unconditional offer from A to B in the respective transformed game.*

It is now up to player A to decide whether to make the respective offer leading to the value $v^A(\widehat{\mathcal{G}})$ – if that offer would improve her current expected value – or to pass, possibly by making only a vacuous offer, for the sake of indicating to B on which of the several equivalent Nash equilibria to coordinate (as in the symmetric coordination game), when appropriate.

Example 5 (Solving a game by exchange of unconditional offers) *Consider the following NFG \mathcal{G} between players R (row) and C (column), which has no pure strategy NE:*

	C1	C2	C3
R1	2, 10	10, 4	5, 1
R2	6, 0	4, 4	6, 3

The maximin solution is (R2, C2) with payoffs (4, 4), which is not Pareto optimal.

Suppose, player R is to make the first offer and let us see what is the best offer that R make to C . (Hereafter we will often denote $d + \epsilon$ by d^+ and $d - \epsilon$ by d^- .)

- *The best response of R to C1 is R2.
So, $\delta_{2,1}^R = 4 - 0 + \epsilon = 4^+$ and $v^R(\widehat{\mathcal{G}}_{C1}) = 6 - 4^+ = 2^-$*
- *The best response of R to C2 is R1.
So, $\delta_{1,2}^R = 10 - 4 + \epsilon = 6^+$ and $v^R(\widehat{\mathcal{G}}_{C2}) = 10 - 6^+ = 4^-$.*
- *The best response of R to C3 is R2.
So, $\delta_{2,3}^R = 4 - 3 + \epsilon = 1^+$ and $v^R(\widehat{\mathcal{G}}_{C3}) = 6 - 1^+ = 5^-$.*

Thus, $v^R(\widehat{\mathcal{G}}) = v^R(\widehat{\mathcal{G}}_{C3}) = 5^-$, meaning that R 's best offer to C is $R \xrightarrow{1^+ / C3} C$. The resulting transformed game is

	C1	C2	C3
R1	2, 10	10, 4	$4^-, 2^+$
R2	6, 0	4, 4	$5^-, 4^+$

It has one Nash equilibrium (R2, C3) yielding payoffs ($5^-, 4^+$) which are strictly better than the players maximin values, but not yet Pareto optimal.

Now, let us compute the best offer of C to R in the transformed game.

- *The best response of C to R1 is C1 and $\delta_{1,1}^C = 4^+$. So, $v^C(\widehat{\mathcal{G}}_{R1}) = 10 - 4^+ = 6^-$.*
- *The best response of C to R2 is C3 and $\delta_{2,3}^C = 0$. Thus, $v^C(\widehat{\mathcal{G}}_{R2}) = 5^-$.*

So, $v^C(\widehat{\mathcal{G}}) = 6^-$, which is better than C 's current value of 4^+ . Thus, C can improve his value by making the offer $C \xrightarrow{4^+ / R1} R$. The resulting transformed game is

	C1	C2	C3
R1	$6^+, 6^-$	$14^+, 0^-$	$8, -2$
R2	$6, 0$	$4, 4$	$5^-, 4^+$

It has one Nash equilibrium $(R1, C1)$, where the strategy $R1$ is strictly dominant for R , yielding payoffs $(6^+, 6^-)$ which are strictly better than the previous ones of $(5^-, 4^+)$, but not yet Pareto optimal. So, let us see whether R can improve any further the resulting game, given that the strategy $R1$ is already his best response to all strategies of C :

- For $C1$: $\delta_{1,1}^R = 0$ and $v^R(\widehat{\mathcal{G}}_{C1}) = 6^-$
- For $C2$: $\delta_{1,2}^R = 6^- - 0^- + \epsilon = 6^+$ and $v^R(\widehat{\mathcal{G}}_{C2}) = 14^+ - 6^+ = 8$.
- For $C3$: $\delta_{1,3}^R = 6^- - 2 + \epsilon = 8$ and $v^R(\widehat{\mathcal{G}}_{C3}) = 8 - 8 = 0$.

Thus, $v^R(\widehat{\mathcal{G}}) = v^R(\widehat{\mathcal{G}}_{C2}) = 8$, which is better than R 's current value of 6^+ , hence R 's best offer to C now is $R \xrightarrow{6^+ / C2} C$. The resulting transformed game is

	C1	C2	C3
R1	$6^+, 6^-$	$8, 6$	$8, -2$
R2	$6, 0$	$-2^-, 10^+$	$5^-, 4^+$

It has a strictly dominant strategies equilibrium $(R1, C2)$ yielding payoffs $(8, 6)$ which are strictly better than the previous ones $(6^+, 6^-)$. In fact, this is the only Pareto maximal outcome in the game, and one can now check that none of the players can make any further improving offers. Thus, this is the end of the negotiation phase.

We leave to the reader to check that if C makes the first offer, the negotiation phase will end with a slightly different game but with the same solution, and after each player making only one offer. As we will see further, such confluence is not always the case.

5.1.4 Weakness of unconditional offers

The example above demonstrates the potential power of unconditional offers to solve normal form games. The next ones demonstrate their weakness showing that in preplay negotiation games where no conditional offers and no withdrawals are allowed the players may not be able to reach any Pareto optimal outcome. Moreover, the expected value of the game that a player can achieve by making an effective unconditional offer in such a preplay negotiations game, need not be better than the original expected value of the game for that player. In fact, it can be strictly worse, for *every* player, than the value yielded by the maxmin strategy profile in the original game, as shown by the following example.

Example 6 (No player benefits by making an effective offer) Consider the following game \mathcal{G} between players R (row) and C (column):

	C1	C2
R1	$3, 3$	$1, 1$
R2	$9, 1$	$0, 8$
R3	$0, 7$	$8, 1$

Note that the maxmin solution is $(R1, C1)$ with payoffs $(3, 3)$, but it is not Pareto optimal. The players have the potential to negotiate a mutually better deal in any of the outcomes in rows 2 and 3. However, it turns out that none of them can make a first unconditional offer that would improve his expected payoff. Indeed, computing their best offers according to the procedure outlined above produces the following:

- The best response of R to $C1$ is $R2$. Then, $\delta_1^R = 8 - 1 + \epsilon = 7^+$ and $v^R(\widehat{\mathcal{G}}_{C1}) = 9 - 7^+ = 2^-$. Respectively $\delta_2^R = 6^+$ and $v^R(\widehat{\mathcal{G}}_{C2}) = 8 - 6^+ = 2^-$.
- Likewise, $v^C(\widehat{\mathcal{G}}_{R1}) = 3 - 6^+ = -3^-$; $v^C(\widehat{\mathcal{G}}_{R2}) = 8 - 8^+ = 0^-$; $v^C(\widehat{\mathcal{G}}_{R3}) = 7 - 9^+ = -2^-$.

Thus, $v^R(\widehat{\mathcal{G}}) = 2^-$ and $v^C(\widehat{\mathcal{G}}) = 0^-$. Both values are less than the respective maxmin values of 3. Therefore, no player is interested in making a first offer and the negotiation phase ends at start.

5.1.5 The disadvantage of making the first unconditional offer

Even when each of the players can start an effective negotiation ending with a solved game, the solution may essentially depend on who makes the first effective offer.

Example 7 (Making the first offer can be disadvantageous) We leave it to the reader to check that in the following game between R and C

	$C1$	$C2$
$R1$	1, 8	10, 4
$R2$	4, 10	1, 11
$R3$	4, 0	2, 2

if the first offer is made by R the preplay negotiation game ends with

	$C1$	$C2$
$R1$	1, 8	$6^-, 8^+$
$R2$	4, 10	$-3^-, 15^+$
$R3$	4, 0	$-2^-, 6^+$

where the only acceptable (surviving iterated elimination of strictly dominated strategies) outcome is $(R1, C2)$ yielding payoffs $(6^-, 8^+)$, whereas if the first offer is made by C the preplay negotiation game ends with

	$C1$	$C2$
$R1$	$4^+, 5^-$	9, 5
$R2$	4, 10	$-3^-, 15^+$
$R3$	4, 0	$-2^-, 6^+$

where the only acceptable outcome is again $(R1, C2)$, but now yielding payoffs $(9, 5)$. Note that in both cases the disadvantaged player is the one who has made the first offer.

The example above also indicates that, in the case under consideration, the greedy approach, where a player always makes the best effective offer he can, may not be his best

strategy, but passing the turn to the other player – that is, making a vacuous offer – could be strategically more beneficial. On the other hand, if both players keep exchanging only vacuous offers or passing, then they will never improve their expected values of the starting game. Yet, one can check that any pair of strategies in the example above, whereby one of the player takes the initiative by making the first effective move with his best first offer and thereafter always responding with his currently best effective offers until possible and then passing, while the other player remains passive until that happens and thereafter keeps responding with her best offers until possible and then passing, is a subgame perfect equilibrium strategy in the preplay negotiation phase for that game.

Stocktaking We have demonstrated that, on the one hand, by exchanging only unconditional and irrevocable offers players can often achieve mutually better outcomes of normal form games, but on the other hand their bargaining powers to achieve their *best* outcomes in such games can be substantially affected by the potential disadvantage of making the first effective offer in such games. We therefore believe that the analysis of PNG with unconditional offers warrants the use of equilibria that go beyond SPE. In particular, features such as *vacuous offers* used for *signaling* a future intention on the strategy to be played become essential. Such analysis should take into account equilibria generated both by backward induction on the negotiation game as well as by forward induction, where past moves can be used to justify rational behavior in the future (see [OR94]), may become relevant. We leave the overall analysis of this case, and in particular the further investigation of the best negotiation strategies and the analysis of the effect of valuable time, to future work. We now proceed with the study of more powerful preplay negotiation games where players can make conditional offers.

5.2 Conditional offers with no withdrawals and non-valuable time

First, a general observation: in the case of non-valuable time players assign the same value to the NFG associated with the current moment and the same game associated with any other moment in the future, which means that players can afford composing generalized offers over time by a sequence of simple offers as well as making and withdrawing offers at no extra cost. Therefore, the issue of allowing or not generalized offers and their withdrawals is of marginal importance. For this reason we will assume that generalized offers and withdrawals of offers are allowed. However, we will argue further that in efficient negotiations withdrawals are avoided.

To analyze equilibrium strategies of PNG when time is not valuable it is useful to consider an interesting class of games where players display some coherence in playing. We say that players have **stationary acceptance strategies** when they have a minimal acceptance threshold (each player accepts any conditional offer that guarantees him least some amount d , which may vary among the players) and a minimal passing threshold (each player, when allowed, passes at histories associated to games guaranteeing them at least some amount $d' \geq d$).

Proposition 8 *Every subgame perfect equilibrium strategy profile of a 2-player PNG with conditional offers and non-valuable time consisting of stationary acceptance strategies is strongly efficient.*

Proof sketch. Suppose not. Let d^- be a vector of expected utilities that is not the redistribution of a maximal outcome of the starting game associated to some subgame perfect equilibrium strategy. Without loss of generality we can assume that such strategy yields a history h that ends with: 1) the proposal of d^- ; 2) acceptance of that proposal; 3) pass; 4) pass. Consider now some redistribution d^* of a maximal outcome where both players get more than in d^- and the history h where the last four steps are substituted by the following ones: 1) the proposal of d^* ; 2) the acceptance of that proposal; 3) pass; 4) pass. By stationarity of strategies and the ODP, the player moving at step 1) is better off deviating from d^- and instead proposing d^* . Contradiction. ■

The condition of stationarity of acceptance strategies is needed if we want to talk about SPE that do not lead to inefficiency. The general idea is that if players were not sticking to stationary acceptance strategies there could be a *suboptimal* outcome guaranteeing for both players an expected utility of respectively d_A and d_B . It is then enough that off the equilibrium path player A threatens player B with a *stubborn* but maximal stationary acceptance strategy giving him less than d_B and that player B threatens A with an expected payoff of strictly less than d_A . As players are not obliged to be consistent in their acceptance policies, d_A and d_B can be the result of a subgame perfect equilibrium strategy. The following example will provide a detailed instance of such games.

Example 9 (Attaining inefficiency) Consider the following starting game (Figure 10). As there are no dominant strategy equilibria, there are acceptable solution concepts assigning 2 to each player.

	L	R
U	2, 2	4, 3
D	3, 3	2, 2

Figure 10: Attaining inefficient divisions

We now construct a strategy profile of the PNG starting from the game in Figure 10 that (i) is a SPE strategy profile and (ii) it attains an inefficient outcome. In what follows we say that a player "proposes a given outcome with a given payoff distribution" to mean that the player makes a conditional offer which, when accepted, would make that specific outcome, with that specific distribution of the payoffs, the unique (dominant strategy equilibrium) outcome in the solution of the transformed game.

1. At the root node player A proposes outcome (D, L) with payoff distribution $(3, 3)$.
2. After such proposal player B accepts. However, if A had made a different offer (so, off the equilibrium path) B would reject and keep proposing a distribution of 5 for him and 2 for A and accepting (and passing on) maximal outcomes guaranteeing him at least 5. A , on the other hand, would not have better option than proposing the same distribution (5 for B and 2 for her) and accepting only maximal outcomes guaranteeing her at least 2. Notice that once they enter this subgame neither A nor B can profitably deviate from such distribution.
3. If, however, B did not accept the $(3, 3)$ deal then A would keep proposing a redistribution of $(5, 2)$ (5 for her, 2 for him) and accepting at least that much. B on the other hand

would also stick to the same distribution, accepting at least 2. Once again, no player can profitably deviate from this stationary strategy profile is adopted.

4. After player B has accepted the deal (3,3), then A passes. If A did not pass, player B would go back to his (2,5) redistribution threat.

Likewise with the next round. That eventually leads to the inefficient outcome (3,3).

It is easy to check that the strategy profile described above is a subgame perfect equilibrium. No player can at any point deviate profitably by proposing the outcome (U,L) with payoff distribution (3.5,3.5).

We now show that SPE strategy profiles do exist in PNG and that in fact every redistribution of a maximal outcome can be attained.

Proposition 10 *Let \mathcal{E} be PNG with non-valuable time starting from a NFG \mathcal{G} and let $d = (x_A, x_B)$ be any redistribution of a maximal outcome of the starting NFG. The following strategy profile $\sigma = (\sigma_A, \sigma_B)$ is a subgame perfect equilibrium:*

For each player $i \in \{A, B\}$:

- if i is the first player to move, he proposes a transformation of \mathcal{G} where the vector of expected values in the transformed game is d ;
- when i can make an offer and the previously made offer has not been accepted, he proposes a transformation of the current NFG where the vector of expected values in the transformed game is d ;
- when i can make an offer and the previously made offer has been accepted, he passes;
- when i has a pending offer of a suggested transformation where the vector of expected values in the transformed game is d' , he accepts it if and only if $x'_i \geq x_i$, and rejects it otherwise;
- i never withdraws any previously made offer;
- when i can pass and the other player has just passed, he passes;
- when i can pass and the other player has not just passed, he proposes d ;
- when i has just accepted a proposal he passes;

Proof sketch.

We have to show that in every subgame this strategy is a Nash equilibrium. Consider any history of the PNG and the player moving at that history. By Lemma 3 we can restrict ourselves to considering only first move deviations to the above described strategy. Suppose in that history the player has a pending offer d^* . If he accepts it then the outcome will be d^* , if he rejects it, it will be the starting offer d . And he will accept if and only if he will get more from d^* than from d . So the acceptance component is optimal. Similar arguments hold for passing and withdrawing. Notice also that deviating from proposing d is never strictly beneficial. ■

As a straightforward consequence of the previous proposition we obtain:

Corollary 11 *The game associated to the outcome of a subgame perfect equilibrium strategy profile consisting of stationary acceptance strategies in a 2-player PNG with conditional offers and non-valuable time is optimally solvable.*

Stocktaking The analysis of 2-player PNG with non-valuable time shows that efficiency can be attained when conditional offers are allowed. In particular, there are SPE strategies where any redistribution of the vector of payoffs of a maximal outcome can be made the unique solution of the *final* normal form game (Proposition 10). The result can be even stronger in the case of SPE consisting only of stationary acceptance strategies, where players display a consistency in accepting and passing — i.e. if anywhere in their negotiation strategy they accept (pass upon) an offer guaranteeing them a certain amount then they everywhere accept (pass upon) an offer guaranteeing them at least that amount (Proposition 8).

However not all equilibria display desirable properties, as Example 9 and Proposition 10 clearly show: there exist subgame perfect equilibrium strategy profile of a 2-player PNG with conditional offers and non-valuable time where (i) offers are made that are not feasible, (ii) the vector of payoffs of the outcome it attains is not a redistribution of the vector of payoffs of a maximal outcome, i.e. it is not strongly efficient.

To partially address the issues related to inefficiency we now introduce the possibility for players to unilaterally put an end to the negotiations.

5.2.1 Negotiations with ‘opt out’ moves

We now analyze the consequences of allowing players the special *opt out* move, with which they can end the negotiations unilaterally and make the currently accepted NFG the outcome of the whole PNG. Formally, the **currently accepted NFG** at history h of a PNG starting from a NFG \mathcal{G} is defined as follows:

- if $h = \emptyset$, it is \mathcal{G} ;
- if $h = (h', a)$ for some atomic move a that is not an unconditional offer or an acceptance of an offer, then it is the currently accepted NFG at the history h' ;
- if $h = (h', a)$ and a is either an unconditional offer or an acceptance of an offer, then it is the updated game according to the atomic move a .

Proposition 12 *Let σ be a subgame perfect equilibrium strategy of a PNG with opt out move and let h be the resulting history. Then σ guarantees to all players at least as much as they had in the currently accepted NFG; in particular, at least as much as in the original game.*

Proof sketch. It follows from the fact that, starting with the original game which is automatically agreed upon, each currently accepted NFG must make each player better off than in the previous one, otherwise opting out is a profitable deviation. ■

By introducing the possibility of opting out the set of subgame perfect equilibria reduces further. Strategies, such as the one described in Proposition 10, contemplating an unreasonably high or unreasonably low reward for the proponent, will not be equilibria anymore. However this extra option does not solve the problem of attaining inefficiency, as the game in Figure 10 again shows. It has, however, several other advantages: first of all, the equilibrium strategies of the PNG will guarantee for both players at least the expected payoff of the starting NFG; and second, it gives the players the possibility of making a more effective use of unconditional offers: players could alone put an end to the negotiations without requiring their opponents’ permissions, by first making an unconditional offer and then opting out.

Remaining issues Summing up, while SPE strategies in a 2-player PNG can attain efficiency, three important issues are still remaining:

- some SPE strategies are not strongly efficient, i.e., players can find it rational to display a form of inconsistency in their accepting and passing policies (non-stationary acceptance strategies);
- players can keep making unfeasible moves as a part of a SPE strategy, i.e., there are forms of equilibria where some players strictly decrease their expected payoff with respect to the original game;
- even strongly efficient strategies do not yield games that are perfectly solved, i.e., there is no notion of *most fair* redistribution of the payoff vectors in the solution of the original game.

Clearly, when time is of no value, conditional offers alone are not sufficient to guarantee that fair and efficient outcomes are reached in any reasonable amount of time. The next subsection will study PNG games where time is of value. We will see that, under an appropriate discounting factor, all the problems mentioned above will be solved at once.

5.3 Conditional offers and withdrawals, with valuable time

We now assume time to be valuable, by introducing, for each player i , a *payoff discounting* factor $\delta_i \in (0, 1)$. The discounting factor is a measure of the players' impatience, i.e., how much they value time, and it modifies the payoffs accordingly as time goes by. Thus, the players have no interest in delaying the negotiations by making redundant moves or sub-optimal offers. In particular, the issue of whether and when it would be beneficial for a player to withdraw an earlier offer becomes essential. The general intuition in this case, which we will justify further, is that the only SPE strategy profiles for the preplay negotiation games with valuable time would consist of just 2 moves:

1. The first player to move makes his best conditional offer which the other player would eventually accept (by adding for himself a small premium for the time saving)
2. The other player accepts the conditional offer.

The reasoning behind this intuition is that in a SPE strategy profile:

1. If any player is ever going to make an offer, he/she would never make any earlier offer that gives him, if accepted, a lesser value of the resulting game. (That would be a costly waste of time.)
2. If any player is ever going to accept a given offer (or any other offer which would give him at least the same expected value of the resulting NFG), he/she should do it the very first time when such offer is made to him/her. (Again, otherwise he would have wasted valuable time.)

To analyze PNG with valuable time we will impose some additional restrictions, mainly for technical reasons. Namely, we will assume that:

	L	R
U	a, b	c, d
D	e, f	g, h

	L	R
U	$a\delta_1, b\delta_2$	$c\delta_1, d\delta_2$
D	$e\delta_1, f\delta_2$	$g\delta_1, h\delta_2$

	L	R
U	$a\delta_1^2, b\delta_2^2$	$c\delta_1^2, d\delta_2^2$
D	$e\delta_1^2, f\delta_2^2$	$g\delta_1^2, h\delta_2^2$

Figure 11: The game with valuable time at time instants 0, 1 and 2. The payoffs of players get discounted according to their individual impatience.

- every game associated with a history of a PNG does not have outcomes in the solution that assign negative utility to players. Notice, that we do allow payoff vectors consisting of negative reals to be present in the game matrix, only we do not allow such vectors to be associated to outcomes in the solution. The constraint we impose has several practical consequences:
 - players' expected payoffs *decrease* in time, i.e. the discounting factor δ has always a negative effect on the expected payoff.
 - players are allowed to make offers that redistribute the payoff vectors associated to outcomes in the solution, leaving some nonnegative amount to each player.
- players are not allowed to opt out, i.e., the only possible disagreement is obtained when players keep bargaining forever;
- players will not have the possibility of withdrawing a previously made offer.
- the expected payoff of each player at any disagreement history is 0.

We will discuss the impact of such restrictions and the scenarios in which they do not hold in the technical companion [GT].

Henceforth, we employ the following notational conventions:

- we denote by (x, t) the payoff vector x at time t , i.e., the payoff vector where each component x_i is discounted by δ_i^t .
- we denote \mathcal{G}_X the set of all possible redistributions of payoffs of outcomes in \mathcal{G} that assign nonnegative payoffs to all players.

Some observations. In every 2-person PNG with valuable time starting from a normal form game \mathcal{G} where the set of players we can observe the following:

1. For each $x, y \in \mathcal{G}_X$ such that $x \neq y$, if $u_i(x, 0) = u_i(y, 0)$ then $u_j(x, 0) \neq u_j(y, 0)$. This holds because the set \mathcal{G}_X is made by maximal outcomes and subtracting some payoff to a player means adding it to the other.
2. $u_{-i}(b^i, 1) = u_{-i}(b^i, 0) = u_{-i}(D)$, where b^i is the highest payoff that i obtains in \mathcal{G}_X and D is any disagreement history. As b^i is the best agreement for player i it is also the worst one for player $-i$. However, due to the constraints we have imposed of the offers, the payoff for player $-i$ is 0, the same as the expected payoff at any disagreement history.

3. If x is Pareto optimal amongst the payoff vectors in \mathcal{G}_X then, by definition of \mathcal{G}_X , there is no y with $u_i(x, 0) \geq u_i(y, 0)$ for each $i \in N$. Moreover, x is a redistribution of a maximal outcome in \mathcal{G} .
4. There is a unique pair (x^*, y^*) with $x^*, y^* \in \mathcal{G}_X$ such that $u_A(x^*, 1) = u_A(y^*, 0)$ and $u_B(y^*, 1) = u_B(x^*, 0)$ and both x^*, y^* are Pareto optimal amongst the payoff vectors in \mathcal{G}_X .

The importance of the properties listed above lies in the fact that they are the four fundamental assumptions of Rubinstein's perfect equilibrium solution of the bargaining problems in [Rub82], see also Osborne and Rubinstein's bargaining model [OR94, p.122]. The first 3 statements above are quite straightforward. To see the last one, let $x^* = (x_A^*, x_B^*)$ and $y^* = (y_A^*, y_B^*)$ and let the sum of the payoffs in any maximal outcome in \mathcal{G} be d . Then $(x_A^*, x_B^*, y_A^*, y_B^*)$ is the unique solution of the following, clearly consistent and determined system of equations:

$$\begin{aligned} y_A &= \delta_A x_A, \\ x_B &= \delta_B y_B, \\ x_A + x_B &= d, \\ y_A + y_B &= d. \end{aligned}$$

The solution (see also [Rub82, OR94]) is:

$$x_A = d \frac{1 - \delta_B}{1 - \delta_A \delta_B}; \quad y_A = \delta_A d \frac{1 - \delta_B}{1 - \delta_A \delta_B}; \quad x_B = \delta_B d \frac{1 - \delta_A}{1 - \delta_A \delta_B}; \quad y_B = d \frac{1 - \delta_A}{1 - \delta_A \delta_B}.$$

In the remaining part of the section we will explicitly view preplay negotiation as a bargaining process on how to play the starting normal form game. Also, due to the previous observations and assumptions, we can adapt the results from [Rub82, OR94] to show that when time is valuable not only all equilibria attain efficiency but they also do it by redistributing the payoff vector in relation to players' impatience.

Proposition 13 *Let (x^*, y^*) be the unique pair of payoff vectors defined above. Then, in a PNG with valuable time starting from a NFG \mathcal{G} with a unique discounting factor δ for both players, the strategy of player A in every subgame perfect equilibrium consisting of stationary acceptance strategies satisfies the following (to obtain the strategy for B simply swap x^* and y^*):*

- if A is the first player to move she 'proposes' outcome x^* , i.e., makes a conditional offer that, if accepted, would update the game into one with its dominant strategy equilibrium yielding the Pareto maximal outcome x^* as payoff vector;
- when A has a pending offer y' , she accepts it if and only if the payoff she gets in y' is at least as much as in y^* ;
- when A can pass, she passes if and only if the expected value associated to the proposed game y'_A is at least as much as y_A^* ; otherwise she proposes x^* .

Proof sketch. The argument is a variant of the proof given in [Rub82, OR94] and it basically follows the same construction. It shows first that the best subgame perfect equilibrium payoff for player A in any subgame G_A starting with her proposal — let us denote it by $M_A(G_A)$ — yields the same utility as the worst one — $m_A(G_A)$ — which, in turn, is

the payoff of A at x^* . Similar considerations can be made for player B . This is obtained by comparing $M_A(G_A)$ with the (appropriately discounted) worst SPE for player 2 in case A 's proposal is rejected. Then it shows that in every SPE the initial proposal is x^* , which is immediately accepted by the other player. Finally, it shows that the acceptance and the passing conditions given cannot be improved upon. ■

Notice that in each subgame perfect equilibrium history, players agree as soon as possible and divide (almost) evenly any of the maximal outcomes in the game. Thus, introducing value of time solves both problems of efficiency and fairness at once.

Stocktaking When time is valuable and players' impatience is measured by a discount vector δ that is multiplied to players' payoff vectors at each time step, the SPE are essentially unique, efficient and redistribute a maximal payoff vector in *fair* way, depending on players' impatience. Moreover, similar to what is observed in [OR94], stationary strategies emerge in all equilibria: even though we did not impose on players to display consistency in accepting and passing they do in equilibrium.

It must be noted that the result we have obtained is strictly dependent on our modeling assumption concerning the discounting factors. Different ways of discounting time, for instance by subtracting a fixed cost δ' at each time step, would substantially change the solution predictions, like in [Rub82], and these are subject of further study.

6 Extended frameworks with offer-induced game transformations

The framework with offer-induced game transformations of non-cooperative games that can be extended in various ways. Here we discuss briefly two of the most important cases.

6.1 Coalitional preplay negotiations in multi-player normal form games

	Y	N		Y	N
Y	3, 3, 3	-1, 8, -1	Y	-1, -1, 8	-5, 4, 4
N	8, -1, -1	4, 4, -5	N	4, -5, 4	0, 0, 0
	Y			N	

Figure 12: The three-person common project: a player may either contribute €9 to a common project, or contribute nothing. Each €9 contributed produces an additional €3. The total amount is divided equally among the players, independently of their contribution. Note that the money contributed by a player is subtracted from his final payoff.

The analysis of N -player normal form games with preplay negotiations phase, for $N > 2$, is much more complicated than the 2-players case. To begin with, the benefit for a player A of player B playing a strategy induced by an offer from A to B crucially depend on the strategies that the remaining players choose to play, so an offer from a player to another player does not have the clear effect that it has in the 2-player case. Thus, a player may have to make a *collective offer* to several (possibly all) other players in order to orchestrate their plays in the best possible for him way. Furthermore, a player may be able to benefit in different ways by making offers for side payments to different players or groups of players,

and the accumulated benefit from these different offers may or may not be worth the total price paid for it. Lastly, when all players make their offers pursuing their individual interests only, the total effect may be completely unpredictable, or even detrimental for all players. It is therefore natural that groups of players get to collaborate in coordinating their offers.

Thus, a coalitional behaviour naturally emerges here, and the preplay negotiation phase incorporates playing a coalitional game to determine the partition of all players into coalitions that will coordinate their offers and moves in the negotiation phase. However, we emphasize again that the transformed normal form game played after the preplay negotiation phase should remain a non-cooperative game where every player eventually plays for himself.

Here we only begin to discuss this more general framework, by first classifying the different types of offers that players or coalitions can make to others. For each of them we give an example in terms of the Common Project game in Figure 12, where we call the respective players Row, Column, and Table:

1. **One-to-one offers:** of the type $A \xrightarrow{\delta/\sigma_B} B$ discussed in the previous sections. A player may place several such offers to different players, and each offer is independent from the rest and only conditional on the strategy played by its sole recipient. Figure 13 illustrates them for the case of the three-person common project.

	Y	N		Y	N
Y	$-2^-, 8^+, 3$	$-1, 8, -1$		$-6^-, 4^+, 8$	$-5, 4, 4$
N	$3^-, 4^+, -1$	$4, 4, -5$		$-1^-, 0^+, 4$	$0, 0, 0$
	Y			N	
	Y	N		Y	N
Y	$-7 - 2\epsilon, 8^+, 8^+$	$-6^-, 8, 4^+$		$-6^-, 4^+, 8$	$-5, 4, 4$
N	$-2 - 2\epsilon, 4^+, 4^+$	$-1^-, 4, 0^+$		$-1^-, 0^+, 4$	$0, 0, 0$
	Y			N	

Figure 13: One-to-one offers. Above: Row offers 5^+ to Column for him contributing to the project. This is enough to make him contribute, but it does not make Row better off in the unique dominant strategy equilibrium (N, Y, N) . Below: Row offers 5^+ independently to each Column and Table for contributing. Now (N, Y, Y) is the dominant strategy equilibrium, but again Row does not benefit from the cooperation of the other two.

2. **Many-to-one (collective) offers:** a group (coalition) of players \mathcal{A} makes a collective offer to a single player B for a total payment of bonus, conditional on B playing the strategy specified in the offer.

The additional issue arising in collective offers is how the coalition \mathcal{A} should split amongst themselves the cost of the bonus due to player B if he complies. The distribution of the due contributions generates a standard in cooperative game theory problem, which will analyze in a follow-up work. Here we assume that a reasonable and commonly acceptable solution to that problem is adopted, e.g., using Shapley value based on the expected values of the normal form game for each player and coalition, and that solution computes on the side the distribution of the due contributions. Once determined and agreed upon, that distribution is explicitly specified as a fixed part of

the offer, and accordingly determines the transformation of the payoff matrix of the game. An instance of this is given in Figure 14.

Formally, we will denote such collective offers by $\mathcal{A} \xrightarrow{\delta_{\mathcal{A}}/\sigma_B} B$ where $\delta_{\mathcal{A}} : \mathcal{A} \rightarrow \mathbb{R}^+$ is the function which specifies the due contribution $\delta_{\mathcal{A}}(A_i)$ for each player $A_i \in \mathcal{A}$ to the total bonus payable to B , while σ is the strategy of B on which the offer is conditional.

	Y	N		Y	N
Y	0.5 ⁻ , 8 ⁺ , 0.5	-1, 8, -1	Y	-3.5 ⁻ , 4 ⁺ , 5.5	-5, 4, 4
N	5.5 ⁻ , 4 ⁺ , -3.5	4, 4, -5	N	1.5 ⁻ , 0 ⁺ , 1.5	0, 0, 0
	Y	N		Y	N

Figure 14: A many-to-one offer. Row and Table offer *together* 5⁺ to Column for him contributing to the project. The amount is divided about evenly between the two. Notice that the offer is enough to make Column contribute and makes *all* players better off in the unique resulting dominant strategy equilibrium (N, Y, N).

3. **One-to-many (conjunctive) offers:** a player A offers to a group of other players $\mathcal{B} = \{B_1, \dots, B_k\}$ side payments of bonuses to each of them conditional on *each of them* playing a strategy prescribed in A 's offer⁶. Such offer presumes that the players from \mathcal{B} coordinate their actions and play as a coalition, because if even one of them deviates from the prescribed to him strategy, the entire offer becomes null and void and noone from the group of recipients gets paid. Formally, we will denote such offers by $A \xrightarrow{\delta_{\mathcal{B}}/\sigma_{\mathcal{B}}} \mathcal{B}$, where $\delta_{\mathcal{B}} : \mathcal{B} \rightarrow \mathbb{R}^+$ is the function which specifies the promised bonus $\delta_{\mathcal{B}}(B_i)$ for each player $B_i \in \mathcal{B}$ and $\sigma_{\mathcal{B}}$ is the strategy profile for \mathcal{B} on which the offer is conditional. An illustration of such offer is given in Figure 15.

	Y	N		Y	N
Y	-7 - 2ε, 8 ⁺ , 8 ⁺	-1, 8, -1	Y	-1, -1, 8	-5, 4, 4
N	-2 - 2ε, 4 ⁺ , 4 ⁺	4, 4, -5	N	4, -5, 4	0, 0, 0
	Y	N		Y	N

Figure 15: A one-to-many offer. Row offers 10 + 2ε to Column and Table for them *collectively* contributing to the project, dividing the amount equally among the two. Notice that this does not make their action of contributing part of a dominant strategy equilibrium, even though (N, Y, Y) is a Nash Equilibrium and N is a dominant strategy for Row. In either case where both Column and Table contribute, Row is utterly worse off.

4. **Many-to-many (collective and conjunctive) offers:** where a coalition \mathcal{A} makes a collective offer to a group \mathcal{B} . This combines the previous 2 types of offers in an obvious way. Many further issues arise here, one of them being whether \mathcal{A} and \mathcal{B} may intersect, in which case there could be an obvious conflict of interests for the players in the intersection. Figure 16 illustrates this complex form of offer.

⁶Alternatively, the player A could offer just one total bonus to the entire group \mathcal{B} and leave it to them to distribute amongst themselves, but this is a risky option because A would not have control on that distribution that would ensure that each player in \mathcal{B} would receive a sufficient incentive to play the prescribed by A strategy.

	Y	N		Y	N
Y	0, 1 ⁻ , 8 ⁺	-4, 6 ⁻ , 4 ⁺	Y	-1, -1, 8	-5, 4, 4
N	8, -1, -1	4, 4, -5	N	4, -5, 4	0, 0, 0
	Y			N	

Figure 16: A many-to-many offer: Column and Row offer together 6⁺ collectively to Row and Table to make them contribute. The payments are divided as follows. Column offers to pay 2⁺ while Row offers to pay 4; if cooperating Table will receive 5⁺, while Row will receive 1. Notice that in this case, too, the resulting game has no dominant strategy equilibrium.

Furthermore, each of these types of offers can be made conditional on counter-offers. Thus, generally, every player would be involved on both sides of several, possibly conflicting offers, and would have to decide which ones to accept, commit or withdraw as a proposer, and which ones to reject or ignore as a recipient.

Thus, the preplay negotiations phase here is much more complex and less determined than in the 2-player case. It would involve, for instance, solving (possibly repeatedly) a corresponding coalitional game to determine a stable partition into coalitions and then conducting negotiations between these coalitions. We leave the analysis of the N -player preplay negotiation games to a future work.

6.2 Inter-play offers in turn-based extensive form games

The problem of underperformance is not limited to normal form games, where players cannot observe the outcome of the opponents' actions during the play. It also arises in some extensive form games, such as the Centipede game, where the Backward Induction strategy profile can recommend an utterly inefficient solution. The idea of preplay offers of bonuses to other players can be applied quite effectively in extensive form games by means of **inter-play offers**, where, before every move of a player, the other player(s) can make him individual or coalitional offers conditional on his forthcoming move. The players from both sides can consider these offers through some commonly accepted solution concept, e.g. Backward Induction (BI) which would provide current values for each player of every subgame arising after the possible moves of A .

A good illustration of the potential power of such inter-play offers is the Centipede game which can be easily transformed in a way stimulating a degree of cooperation. Consider the version of the Centipede game on Figure 17, where player I plays at the odd-numbered nodes and player II plays at the even-numbered nodes.

As well known, BI prescribes player I to go down at node 1, yielding the value (2, 1) of the game. If however, before I's move player II has the opportunity to make an offer to I, then II can offer him a bonus payment⁷ of $1^+ = 1 + \epsilon$ for any $\epsilon > 0$ on the condition that I goes right at node 1. This offer transforms the game tree to the one given on Figure 18. In the resulting game it is strictly more beneficial for I to go right at node 1. Note that the BI value of the resulting game is $(2^+, 2^-)$, which is a strict improvement for *both* players. At node 2 the situation is symmetric. Now, player I can make an offer of 1^+ to player II conditional on her going right. That offer transforms the game tree again, to the one given on top of Figure 18. Note that in this game the original payoffs are restored in the subgame rooted at node

⁷Recall our notation: $d^+ = d + \epsilon$; $d^- = d - \epsilon$.

3. Thereafter, the argument recurs producing further transformations, eventually leading to the game shown at the bottom of Figure 19, where player I would again be better off going right and ending the game with the mutually most beneficial payoffs of $(6^+, 6^-)$.

The message is clear: inter-play bonus payments can naturally stimulate cooperation in non-cooperative extensive form games.




Figure 17: A starting Centipede game.




Figure 18: The Centipede game after the offer of II at node 1.

7 Related work and comparisons

The present study has a rich pre-history and we do not purport to provide a comprehensive citation of all related previous work and literature here, but will only mention various links with earlier studies and then will discuss in more detail and compare with the most relevant recent work.

7.1 Pre-history and relevant early references

Here is a selection of related topics and earlier references on relevant studies:

- To begin with, preplay offers technically fall broadly in the scope of *externalities*. There is abundant literature on these, of which we only mention some of the early works: [Mea52], [Mas94], [Var94b]. More specifically, preplay offers are a special type of so called in cooperative game theory *side payments*.

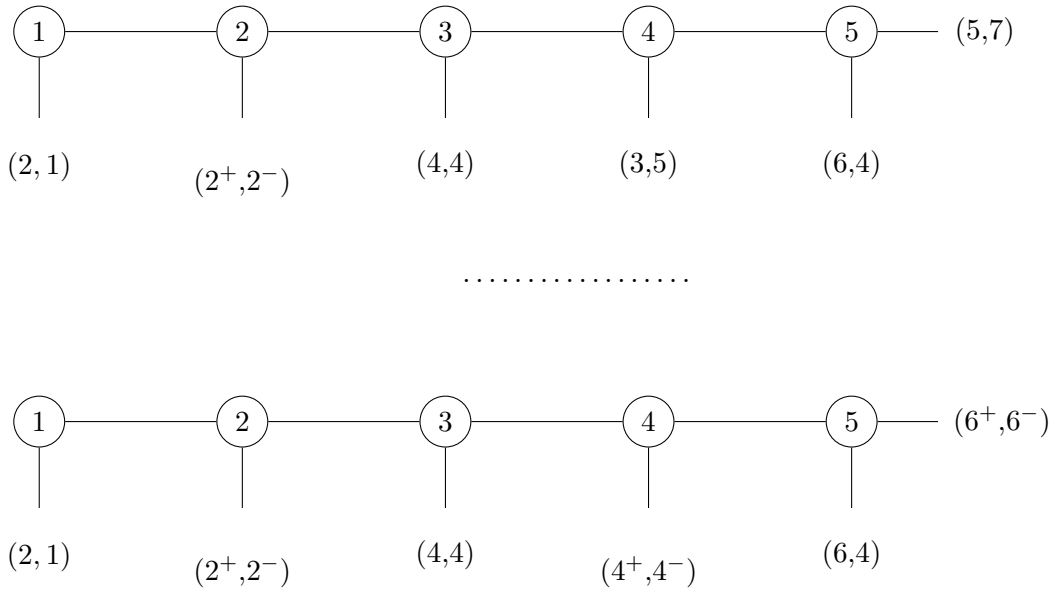


Figure 19: The Centipede games after the subsequent offers at nodes 2 ... 5.

- Coase theorem, [Coa60] describes how efficiency of an allocation of goods or simply an outcome can be obtained in presence of externalities, i.e. when actors' possible decisions affect positively or negatively the payoffs of the other actors involved. The claim, which is usually provided in a rather informal fashion, states that if there are no transaction costs and it is possible to bargain on the effect of the externalities, the process will lead to an efficient outcome regardless of the initial allocation of property rights, i.e. regardless of who is endowed with the capacity of performing the action in question.
- Rosenthal [Ros75] proposes one of the earliest known to us models of preplay negotiations, where 'players successively commit themselves irrevocably, according to a specified exogenous ordering, to coalitional strategies conditionally on the rest of the players in the coalition agreeing to play their parts of the coalitional strategy'. He defines a special solution concept, the induced outcome, and provides some sufficient conditions for its existence and uniqueness are given.
- Various two-stage games with preplay communication have been studied in the literature. They seem to go back to Guttman [Gut78] and [Gut87]. Kalai [Kal81] studies preplay negotiation procedures as sequences of pre-defined length of "preplays", each being a joint strategy of all players. Matthews and Postlewaite [MP89] consider preplay communication in the context of two-person sealed-bid double auctions. Danziger and Schnytzer [DS91] consider a 2-stage game for implementing Lindahl's voluntary-exchange mechanism. In a series of papers, incl. [Far98], Farrell considers two-stage games, with preplay 'cheap talk' followed by actual play, and discusses the role of preplay communication in ensuring Nash equilibrium profile in the actual play. Also, Watson [Wat91] studies two-stage 2-person normal form games with preplay communication and d'Aspremont and Gérard-Varet [dGV80] study Stackelberg-solvable games with preplay

communication.

- Our preplay negotiation games are also closely related to *bargaining* [Rub82, OR90, OR94], [Mye97].
- Another related early work is Varian [Var94b] where he studies variations of ‘compensatory mechanisms’ where, instead of making offers, players declare compensations for which they are prepared to play one or another strategy (in favour of another player who is willing to pay such compensation and makes a binding offer for it). Although the flavour of such variation is somewhat different, technically it reduces to a type of games with preplay offers that we have considered here.
- Kalai et al [CFK91] and more recently, Monderer and Tennenholtz [aMT06] consider the use of ‘agents’ or ‘mediators’ playing on behalf of the players, and show how such mechanisms can be used to achieve more efficient outcomes in non-cooperative games.
- The idea of combining competition and cooperation in non-cooperative games has been considered often since the early times of game theory, and has more recently evolved in theories of *co-opetition* by Brandenburger and Nalebuff [BN97] as well as the more recent [DC]. Also, related in spirit are some theories of coalitional rationality, see Ambrus [Amb09].

7.2 Detailed comparison with most relevant recent work

To our knowledge, Jackson and Wilkie [JW05] have been the first to study arbitrary transfer functions from a player to a player in a normal form game. That work was preceded by earlier relevant literature, such as [Gut78, DS91, Var94b, Qin], where, however, only limited forms of payments were considered, such as payments proportional to the actions taken by the other players or only contingent on own actions. Jackson and Wilkie’s framework bears substantial similarities with ours, as it studies a two-stage transformations on a normal form game where players (i) announce transfers functions which update the initial normal form game; and (ii) play the updated game. Jackson and Wilkie study the subgame perfect equilibria of the two stage game and show under what conditions equilibria of the original game survive in the update game. They focus on the 2-player case, but they also extend their results to the N-player case. Below we describe their framework and emphasize on the difference with ours. In [JW05]:

- Transfers from a player A to a player B are of the form (in our notation) $A \xrightarrow{\delta/\sigma} B$ where $\sigma \in \prod_{i \in N} \Sigma_i$, $\delta \in \mathbb{R}^+$ and $\delta = 0$ whenever $A = B$, i.e. players are allowed to make positive side payments to other players *that are conditional on the entire outcome played*, and not only on the recipient’s individual strategy, as in our framework. Technically, every unconditional offer from player A to player B can be simulated by a set of such transfers from A to B . This is, however, not the case for conditional offers, which would instead require a *set* of transfers from B to A as well, or the possibility for δ to be negative, i.e. the introduction of punishments.
- Players announce their transfer functions *simultaneously*. Canceling an offer by another player is technically possible in such framework, by simply making payments at every

outcome which cancel out, e.g. $A \xrightarrow{\delta/\sigma} B$ and $B \xrightarrow{\delta/\sigma} A$, but cannot be a deliberate response to an anticipated action.

- In [JW05] the authors study strategies that can be *supported*, i.e. that they are subgame perfect equilibria of the two-stage game and Nash-equilibria of the original game that also *survive* — i.e. remain equilibria — in the updated game. In particular they focus on the (interesting) relation between the solo-payoff, i.e. the Nash equilibrium payoff that a player can guarantee by making offers, and the supportability of strategies. For the sake of precision, the solo-payoff u_i^s , starting from a normal form game \mathcal{G} where mix strategies are allowed, is defined as $\sup_{t_i}(\min_{\mu \in NE(t_{-i}^0, t_i)} EU_i(\mu, t_{-i}^0, t_i))$, where:
 - $\mu \in NE(t_{-i}^0, t_i)$ denotes the set of Nash equilibria of the second stage game (the one updating \mathcal{G} with players' simultaneous offer) when player i offers t and the other players offer nothing;
 - $EU_i(\mu, t_{-i}^0, t_i)$ the expected utility for player i associated to the profile μ, t_{-i}^0, t_i (as specified above);
 - \sup and \min simply return the expected values;

Jackson and Wilkie show two important results for the two-player case, the main bulk of their paper: (i) that every Nash equilibrium x of the starting game survives if and only if it yields for every player i a utility that is higher than the one given by i 's solo-payoff; and (ii) that a transfer function together with an outcome are supportable if and only if they yield for every player i a utility that is higher than the one given by i 's *minimal* solo-payoff, the solo-payoff obtained by making *minimal* offers. It is worth noticing that the definition of minimal offer they adopt is essentially the one we have adopted here: the minimal transfer function needed to change the game solution.

- The role of time is not considered and players cannot build upon the game obtained in the second stage, by for instance subsequently making further offers.

	C	D
C	$4 - t_{CC}, 4 + t_{CC}$	$0, 5$
D	$5 - t_{DC}, 0 + t_{DC}$	$1, 1$

Figure 20: Prisoner's dilemma and payments on outcomes.

Jackson and Wilkie's main result consists of showing that side-payments cannot guarantee efficiency, i.e. there are games in which strategies leading to the choice of Pareto optimal outcomes in the second-stage game are not supported. Roughly the argument, also reported in [EP11], is that if the original game is the one in Figure 1 and the Column player believes (remember that the transfers are simultaneous) that the game can be updated into the one of Figure 20 — for instance by the Row player making the specified offers — then he can cancel all the offers out while offering a new transfer $t_{CD} \geq 1$. In this way he will induce the play of CD (it is a Nash equilibrium), which is again inefficient (in the sense that a better outcome could be reached for both players by making different offers, notice that an outcome is Pareto optimal in the big game if it is a redistribution of payoffs of a maximal Pareto optimal outcome of the original game).

Ellingsen and Paltseva [EP11] generalize Jackson and Wilkie's framework in several ways:

- Transfers from a player A to a player B are again of the form $A \xrightarrow{\delta/\sigma} B$ where $\sigma \in \prod_{i \in N} \Sigma_i$, but now $\delta \in \mathbb{R}$ and $\delta = 0$ whenever $A = B$, i.e. players are allowed to propose both rewards and punishments contingent upon entire strategy profiles. This boils down to players not only making offers but also proposing contracts, left for the other players to sign or reject.
- The game played is composed of three stages: (i) the one in which players propose contracts, (ii) the one in which players decide whether to sign a contract, (iii) the one in which players play the game updated by the signed contract.
- Contracts are proposed on mix strategies, and non-deterministic contracts are considered, i.e. it is possible to make randomize offers.

In a nutshell, while in [JW05] each player A specifies the nonnegative transfer to the other players for each pure strategy profile σ , in [EP11] each player specifies a (possibly negative) transfer to the other players for each (possibly mixed) strategy profile σ and, at the same time, specifies a signing decision for each contract of the other players. Ellingsen and Paltseva show that their more general contracting game always has efficient equilibria. In particular they show that all the efficient outcomes guaranteeing to each player at least as much as the worst Nash-equilibrium payoff in the original game can be attained in some equilibrium.

Clearly the message conveyed by this stream of contributions is that efficiency can be reached if the structure of players' offers is complex enough. On the one hand Jackson and Wilkie show that promises are not enough to attain efficient outcomes, while Ellingsen and Paltseva show that contracting is. Our results lie on a rather different axis, as we restrict the type of offers to ones that only commit the proposer, not the recipient, but on the other hand we focus on the effects of additional factors in the preplay negotiation game, e.g. valuable time, conditional offers, and withdrawals, have on attaining outcomes that display desirable properties, such as efficiency and fairness. Moreover we also discuss how equilibrium strategies themselves display desirable properties, i.e. being efficient negotiation strategies.

8 Further agenda and concluding remarks

The main purpose of this paper was to initiate a systematic study of our framework of preplay negotiations in non-cooperative games, and to outline a broad and long-term research agenda for that study. We have indicated a number of conceptual and technical problems and have only sketched some results, but most of the work to be done is left to the future. The focal problems of the study initiated here are:

- to analyze the game-theoretic effects of preplay/interplay offers for payments between individual players and coalitions in strategic and extensive form games, with complete and incomplete information;
- to develop the theory of preplay negotiations and, in particular, to develop the concept of efficient negotiations under various assumptions considered here;
- to analyze the optimality and efficiency of the solutions that can be achieved in preplay negotiation games;

- to expand the study into a systematic theory of cooperation through negotiations in non-cooperative games.
- to apply the developed theory and the obtained results both descriptively and prescriptively to various real-life scenarios where our framework applies.

References

- [Amb09] Attila Ambrus. Theories of coalitional rationality. *Journal of Economic Theor*, 144(2):676–695, 2009.
- [aMT06] Dov Monderer and Moshe Tennenholtz. Strong mediated equilibrium. 2006. in: Proc. of AAAI 2006.
- [BN97] Adam Brandenburger and Barry Nalebuff. *Co-opetition: A Revolutionary Mindset That Combines Competition and Cooperation*. Doubleday, New York, 1997.
- [CFK91] Kenneth Judd Chaim Fershtman and Ehud Kalai. Cooperation through delegation. *International Economic Review*, 32:551–559, 1991.
- [Coa60] Robert H. Coase. The problem of social cost. *Journal of Law and Economics*, pages 1–44, 1960.
- [Col90] James Coleman. *Foundations of Social Theory*. Belknap Harvard, 1990.
- [DC] Daniele Schilirò David Carfi. A model of coopetitive game and the Greek crisis. <http://arxiv.org/abs/1106.3543>.
- [dGV80] Claude d’Aspremont and Louis-André Gérard-Varet. Stackelberg-solvable games and pre-play communication. *Journal of Economic Theory*, 23:201–217, 1980.
- [DS91] Leif Danziger and Adi Schnytzer. Implementing the Lindahl voluntary-exchange mechanism. *European Journal of Political Economy*, 7:55–64, 1991.
- [EP11] Tore Ellingsen and Elena Paltseva. Non-cooperative contracting. Submitted, <http://www2.hhs.se/personal/ellingsen/pdf/Non-cooperativeContracting5.pdf>, 2011.
- [Far98] Joseph Farrell. Communication, coordination and nash equilibrium. *Economics Letters*, 27:209–214, 1998.
- [FG02] Ernst Fehr and Simon Gächter. Altruistic punishment in humans. *Nature*, 415(6868):137–140, 2002.
- [Gor12] Valentin Goranko. Transformations of normal form games by preplay offers for payments among players. Under submission, http://www2.imm.dtu.dk/~vfgo/papers/submitted/Goranko-GameTransformations_120729.pdf, 2012.
- [GT] Valentin Goranko and Paolo Turrini. Two-player preplay negotiation games with conditional offers. Available at <http://arxiv.org/abs/1304.2161>; 2012.

- [Gut78] Joel M. Guttman. Understanding collective action: Matching behavior. *American Economic Review*, 68(2):251–255, 1978.
- [Gut87] Joel M. Guttman. A non-Cournot model of voluntary collective action. *Economica*, 54(213):1–, 1987.
- [JW05] Matthew O. Jackson and Simon Wilkie. Endogenous games and mechanisms: Side payments among players. *Review of Economic Studies*, 72(2):543–566, 2005.
- [Kal81] Ehud Kalai. Preplay negotiations and the prisoner’s dilemma. *Mathematical Social Sciences*, 1:375–379, 1981.
- [Mas94] Eric S. Maskin. The invisible hand and externalities. *American Economic Review*, 84(2):333–337, 1994.
- [Mea52] James Meade. External economies and diseconomies in a competitive situation. *Economic Journal*, 62(245):54–67, 1952.
- [MP89] Steven Matthews and Andrew Postlewaite. Pre-play communication in two- person sealed-bid double auctions. *Journal Of Economic Theory*, 48:238–263, 1989.
- [Mye97] Roger Myerson. *Game Theory: Analysis of Conflict*. Harvard UP, 1997.
- [OR90] Martin Osborne and Ariel Rubinstein. *Bargaining and Markets*. Academic Press, 1990.
- [OR94] Martin Osborne and Ariel Rubinstein. *A course in game theory*. MIT Press, 1994.
- [Pig20] Alfred C. Pigou. *The Economics of Welfare*. Macmillan and Co., 1920.
- [Qin] Cheng-Zhong Qin. Penalties and rewards as inducements to cooperate. Departmental Working Papers, Department of Economics, UCSB, UC Santa Barbara, 2002.
- [Ros75] Robert W. Rosenthal. Induced outcomes in cooperative normal-form games. *Review of Economic Studies*, 975. Discussion Paper No. 178, Center for. Math. Studies, Northwestern University.
- [Rub82] A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50:97–110, 1982.
- [Var94a] Hal R. Varian. Sequential contributions to public goods. *Journal of Public Economic*, 53:165–186, 1994.
- [Var94b] Hal R. Varian. A solution to the problem of externalities when agents are well-informed. *American Economic Review*, 84(5):1278–93, 1994.
- [Wat91] Joel Watson. Communication and superior cooperation in two-player normal form games. *Economics Letters*, 35:267–271, 1991.