Negotiations with Limited Specifiability*

Satoshi Fukuda† Yuichiro Kamada‡

December 19, 2018

Abstract

We study negotiations with *limited specifiability*—each party may not be able to fully specify a negotiation outcome. We construct a class of negotiation protocols to conduct comparative statics on specifiability as well as move structures. We find that asynchronicity of proposal announcements narrows down the equilibrium payoff set, in particular leading to a unique prediction in negotiations with a "common interest" alternative. The equilibrium payoff set is not a singleton in general, and depends on the fine details of how limitation on specifiability is imposed. The equilibrium payoff set is weakly larger under limited specifiability than under unlimited specifiability.

Keywords: asynchronous moves, commitment, limited specifiability, negotiations

*We thank Elchanan Ben-Porath, Gabriel Carroll, Rohan Dutta, Jeff Ely, Juan Escobar, Brett Green, Mai Inoue, Ehud Kalai, Yasuaki Kambe, Sebastian Kodritsch, Fuhito Kojima, Nenad Kos, Stephen Morris, Bora Park, Philipp Strack, Bruno Strulovici, Noriyuki Yanagawa, and seminar/conference participants at the University of Chile and the 25th International Conference on Game Theory at Stony Brook University. We thank Omair Butt, Rohan Tripathi, and Lingxuan Wu for their excellent research assistantship.

†Department of Decision Sciences and IGIER, Bocconi University, Via Roentgen 1, Milan 20136, Italy, e-mail: satoshi.fukuda@unibocconi.it

‡Haas School of Business, University of California Berkeley, 2220 Piedmont Avenue, Berkeley, CA 94720-1900, USA, e-mail: y.cam.24@gmail.com

1 Introduction

Negotiations pervade our social, economic, legal and political lives. They take place in the contexts of labor union, legislature, mergers and acquisitions, international trade, climate change, disarmament, and so forth. This paper introduces a novel concept *limited specifiability* in negotiations, and analyzes its effect on negotiation outcomes. Specifiability refers to the degree to which each participant of a negotiation can specify an outcome in their proposals. We say it is limited if one cannot fully specify any of the exact outcomes and unlimited if one can do so for each of the possible outcomes.

Specifiability can be limited for various reasons. Specifying an exact outcome may be prohibitively costly, or the rule of a negotiation may forbid full specification. For example, consider negotiations among different countries, say the Conference of the Parties (COP) meetings for climate change. The representative of each country may not be able to make a proposal that goes against the benefit of a certain influential interest group in her own country, as doing so would result in a loss of support from the interest group. It may be the case that each of the negotiating parties has their own exclusive right to change a certain aspect of negotiation outcomes. Again taking an example from COP, under the new framework adopted for the 2015 Paris Agreement, each country was able to report their target emission level, while they were not able to specify other countries' emission levels.[1] Examples of limited specifiability are also abundant in business: Two firms may be quoting prices of their products until they settle down. Communication between the firms is not usually allowed, but in practice there would be various ways to imperfectly convey reactions to the opponent firm, and such a situation resembles negotiations with limited specifiability (firms cannot specify a price profile but only their own price). Another example would be an online market, whose design may be such that sellers (buyers) can only specify their minimum (maximum) acceptable prices.

Some of these situations are more complicated than others, and sometimes specifiability may vary across time or histories of past proposals and responses. Moreover,

[1] Negotiating parties were allowed to specify an explicit emission allowance of the Annex I countries in 2009 Copenhagen Summit. In the Paris Agreement, some countries proposed "unconditional INDC (intended nationally determined contribution)" in which they only specify their own emission level, while other countries proposed "conditional INDC" in which they specify an emission level conditional on other countries' actions such as monetary transfers (although they were still unable to specify other countries' emission levels). Thus, both conditional and unconditional INDC entailed limited specifiability, but the exact degrees of their specifiability were different from each other.

the specifiability condition may be chosen by a designer of a given negotiation or a market. As the first study of considering various cases like those, we focus on analyzing the effect of specifiability by fixing the degree of specifiability constant over time and treating it as given (as opposed to making it a choice variable). Our objective is to examine when and how different specifiability conditions lead to different outcomes.

In our model, there is a set of alternatives X and each player i is associated with a set of proposals $\mathcal{P}_i$, which is a collection of subsets of X. When a player moves, she expresses a response of "Yes" or "No" to past proposals and makes a counter-proposal from $\mathcal{P}_i$. We say that player i's specifiability is limited if $\mathcal{P}_i$ does not contain a singleton set (of an alternative), and it is unlimited if $\mathcal{P}_i$ contains all the singleton sets. Once the players reach a consensus on an alternative (we will explicitly define the meaning of consensus), they obtain corresponding payoffs from that alternative. If the negotiation continues indefinitely without reaching a consensus, then the players receive pre-determined disagreement payoffs. We analyze a subgame-perfect equilibrium, which we show exists, of this negotiation game.

We find that the timing of making proposals affects the comparison of possible outcomes under different specifiability conditions. When players' proposals are made in a synchronous manner, we obtain a "folk theorem"—any payoff profile no worse than the disagreement payoffs is achievable in a subgame-perfect equilibrium under arbitrary specifiability conditions.

If moves are asynchronous, which we believe to be more natural in many negotiation settings,[2] the equilibrium payoff sets are smaller under both specifiability conditions, and depend on the payoff structure and the rule of the negotiation in an interesting way. In particular, if there exists an alternative that Pareto-dominates all other alternatives no worse than the disagreement payoffs, then it is the unique outcome of the negotiation game. In general, the equilibrium payoff set is not a singleton, and is smaller under unlimited specifiability than under limited specifiability.[3] The main reasons for these results are that asynchronicity helps players commit to final outcomes, and the commitment power varies across different specifiability conditions. In particular, limited specifiability implies that each proposal entails a smaller degree

[2]For example, in the 2015 Paris Agreement, each country reported asynchronously.

[3]In order to contrast the distinct implications of limited and unlimited specifiability, except for one result (Proposition 4), we only consider situations in which every player's specifiability is limited or every player's specifiability is unlimited.

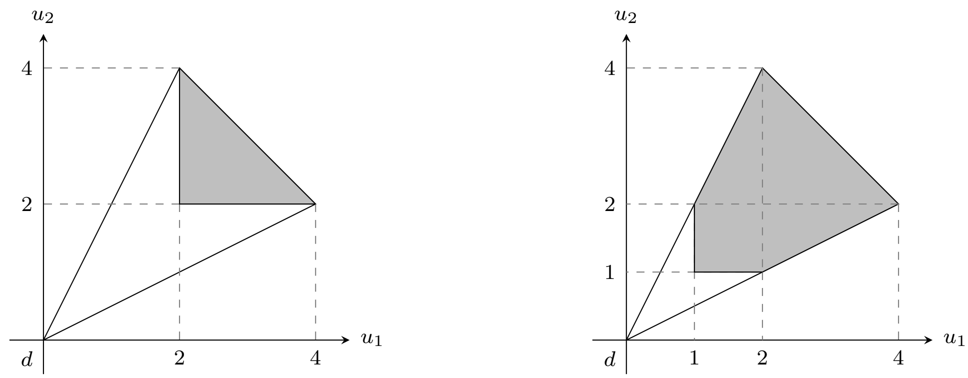

Figure 1: The equilibrium payoff sets under unlimited specifiability (the shaded area of the left panel) and limited specifiability (the shaded area of the right panel).

of commitment to a final outcome, and hence leaves greater scope for punishment conditional on deviations. This leads to a larger set of equilibrium payoffs than in the case under unlimited specifiability.

As an example, Figure 1 depicts the equilibrium payoff sets of two-player asynchronous negotiations under unlimited specifiability (left panel) and limited specifiability (right panel). The disagreement payoffs are $d = (0, 0)$ in both cases. In the game associated with the left panel, each player can announce any feasible payoff profile in the convex hull of $(0, 0)$, $(4, 2)$, and $(2, 4)$. In the game associated with the right panel, in contrast, each player can only announce their own payoff. If the announced payoff profile does not fall into the convex hull, then they get the payoffs d. As the figure shows, any equilibrium payoff profile in the left panel is also an equilibrium payoff profile in the right panel.[4]

Negotiations under limited specifiability and the ones under unlimited specifiability are quite different. If we further add variations of negotiation protocols in terms of the timing of making proposals and try to examine the effect of such variations on the difference of the negotiation outcomes under different specifiability conditions, we need to consider quite a large class of negotiations that at least superficially look very

[4]Note that any additional equilibrium payoff profile in the right panel is Pareto-dominated by some equilibrium payoff profile in the left panel. One may question the relevance of the payoff comparison that relies on such payoff profiles, but the comparison turns out to be robust. It is straightforward to see that, under unlimited specifiability with two players, a player's minimum achievable payoff is still the same as the minimum equilibrium payoff even when the player only knows rationality of the opponent, given an additional assumption that the opponent lexicographically prefers an earlier agreement than later ones in the case of indifference (cf., the intuition explained after Proposition 4).

different from each other. Thus, in order for our comparison of different negotiation protocols to make sense, we need a single coherent framework with which we can study a sufficiently wide class of negotiation protocols. For this purpose, we define a negotiation protocol as a collection of three rules—a *proposer rule*, a *specification rule*, and a *termination rule*. Roughly, a proposer rule determines who speaks when, a specification rule designates a collection of proposals each player can announce, and a termination rule determines the histories under which a negotiation terminates and the outcome associated with each termination. The idea is to vary only one rule in conducting comparative statics, holding fixed the other two.

The heart of this exercise is to define a termination rule solely as a function of histories. This in particular enables us to isolate the termination rule from the specification rules: In other words, we can meaningfully compare two specification rules under a truly single termination rule. In order to define a sensible termination rule, we first define what it means for player i to be *ok* with an alternative x given a history of proposals and responses. We then consider a termination rule such that the negotiation ends with x at a history once all players are ok with x at that history. Briefly, player i is ok with x under a history if her announcement at that history gives rise to the unique intersection $\{x\}$ with the latest proposals by the opponents after which no player announces "No." We call such a termination rule *consensual*. The comparison of outcomes under different specifiability conditions are conducted under the consensual termination rule. As we will show, a two-player negotiation with the consensual termination rule, an asynchronous proposer rule and an unlimited specification rule reduces to the bargaining protocol of Ståhl (1972) and Rubinstein (1982).

The paper is organized as follows. The rest of this section discusses the related literature. Section 2 formulates our model of negotiations by defining proposer, specification, and termination rules. In particular, we define the consensual termination rule. We start with benchmark analyses in Section 3 including the "folk theorem" under synchronous moves. Section 4 analyzes the properties of the equilibrium outcomes with asynchronous moves which are independent of specifiability conditions. Section 5 discusses different predictions under limited and unlimited specification rules. Section 6 provides concluding remarks. The proofs of results that are not proved in the main text can be found in the Appendix. The Online Appendix contains additional discussions.

1.1 Literature Review

Bargaining

Various models of bargaining have been proposed in the literature, including ones with synchronous proposals and others with asynchronous proposals. Bargaining models that originate from Ståhl (1972) and Rubinstein (1982) usually assume asynchronicity, while some other bargaining models such as the Nash demand game (Nash, 1953) assume synchronicity.

Commitment Power under Synchronous vs. Asynchronous Moves

Our paper is related to various strands of the dynamic-game literature which examine the idea that asynchronicity narrows down the equilibrium payoff set. Maskin and Tirole (1987, 1988a,b) study the effect of the timing structure in their oligopoly competition model. In the repeated-games literature, Lagunoff and Matsui (1997) show that if all players have the same payoff function then there is a unique equilibrium outcome.[5] Caruana and Einav (2008) show a uniqueness result under asynchronicity and switching costs in their finite-horizon model. Dutta (2012) and Calcagno et al. (2014) also examine the effect of asynchronicity in their respective finite-horizon models and provide selection results. The general idea behind these results is that player i's choice of action a_i at time t automatically determines her action at $t+1$ under asynchronicity, so i can guarantee the payoff from (a_i, a_{-i}) such that a_{-i} is part of the supergame strategy satisfying $-i$'s best response condition. Our point is that the power of such commitment may be nuanced by the possibility of punishments in negotiation games, and may change depending on specifiability.

Pre-game Cheap-talk Communication—how we should model

The literature of pre-game cheap-talk communication with complete information (e.g., Farrell (1987, 1988) and Rabin (1994)) studies the situation where players convey their intentions for their decisions.[6] One problem endemic in the literature is that it is not clear how to model a negotiation/communication process. Farrell (1988) puts it: "there are no obviously 'right' rules about who speaks when, what he may say, and when discussion ends." Our formulation of negotiation protocols using three

[5]See also Yoon (2001), Lagunoff and Matsui (2001) and Dutta (1995) for conditions on folk theorems in asynchronous repeated games.

[6]See also a survey by Farrell and Rabin (1996).

rules makes it possible to compare equilibrium outcomes under different negotiation protocols.

Specifically, consider the case in which the set of alternatives corresponds to the set of action profiles of a normal-form game. Some models in the literature such as Farrell (1988) and Santos (2000) assume that players can announce action profiles of the underlying game, while others such as Farrell (1987) and Kalai (1981) assume that they can only announce their own actions. The move structure also varies across models. For example, Farrell (1987), Kalai (1981), and Rabin (1994) assume synchronous moves while Santos (2000) assumes asynchronous moves. These models are quite different from each other and hence it is difficult to meaningfully compare their results to understand the effects of limited specifiability and/or move structures. Our model deals with unlimited/limited specifiability and synchronous/asynchronous moves in a unified framework. This would enable one to compare different negotiations.

Pre-game Cheap-talk Communication—what players achieve

We study relationships between the outcome of a negotiation and the structure of an underlying game. In cheap-talk models, however, there exists a babbling equilibrium unless an assumption is imposed on the relationship between talk and choice of actions. So, certain relationships need to be assumed from the outset. For example, Farrell (1987), Rabin (1994), and Santos (2000) assume that any agreement of the communication phase must induce a Nash equilibrium of the underlying game, and thus these models do not address whether we should expect Nash equilibrium after communication. The situations we study, in contrast, are the ones where players can bind their actions of the underlying game during the course of the negotiation. Hence, we can study how the set of equilibrium outcomes of the negotiation game is related to that of Nash equilibria in the given underlying game. It turns out that these two sets can be disjoint in general, while there are certain relationships in special cases (e.g., if the underlying game has a unique Pareto-efficient action profile, then it is a unique equilibrium outcome and is in the set of Nash equilibria).

Also, whether outcomes of bargaining/communication games are restricted to Pareto-efficient outcomes has been studied in the literature (see Crawford (1998) and Farrell and Rabin (1996)). This question has attracted considerable attention especially when the underlying normal-form game has a unique Pareto-efficient action profile (see Farrell (1988), Rabin (1994), and Santos (2000)). Although these papers analyze quite different sets of questions than ours, their results and ours are similar in that the equilibrium outcomes may not be Pareto efficient in general, while if there is a unique Pareto-efficient alternative then it is a unique equilibrium outcome.

Pareto inefficiency also arises in Safronov and Strulovici (2017). They study a repeated game, each period of which may have a renegotiation stage. While the structures of dynamic games are quite different between their and our models, there is some similarity in the results and the intuition: When the Pareto frontier of the feasible payoff set consists of multiple points, inefficient payoff profiles may be sustained as an equilibrium outcome in both models by using as a punishment some payoff profile that is undesirable for the deviator (which in our model exists on the edge of the Pareto frontier). If there is a unique common interest alternative, however, such a punishment is not available, so both models predict a unique efficient outcome.

Revocable Pricing and Asynchronicity in Oligopoly

One important feature of our negotiation game is that announcements of proposals are revocable before the parties ultimately agree on a certain alternative. The industrial-organizations literature (e.g., Bhaskar (1989), Farm and Weibull (1987), and Stahl (1986)) studies revocable pricing to explain "kinked demand curves." In particular, Bhaskar (1989) studies a game which he calls the quick-response game where (i) two firms in a Bertrand duopoly sequentially announce their prices; (ii) they can change their price announcements in reply to those of their opponents; and (iii) they are bounded to take their announced prices once one firm repeats the same price in a row.[7] He shows that the two firms can sustain the monopoly price in a unique equilibrium, which is reminiscent of our uniqueness result for common-interest negotiations.

Bhaskar (1989)'s model has asynchronous moves, while Stahl (1986)'s model has synchronous moves. In Online Appendix D.4, we formulate a "quick response" termination rule within our framework, and show that combining it with an asynchronous proposer rule reduces to Bhaskar (1989)'s model, while combining it with the synchronous proposer rule leads to Stahl (1986)'s model.

Communication Constraints in Mechanism Design

[7]There are no Yes/No responses in his model. Muto (1993) studies quick-response games of two-player normal-form 2×2 games such as the Prisoners' Dilemma and the Battle of the Sexes.

Broadly speaking, our paper sheds light on how limitations on communications change an outcome of a negotiation, and thus it relates to the literature studying how constraints on communication affect an allocation in a mechanism design problem. One of the seminal papers in this literature is Green and Laffont (1987). In their model, an agent's message space is restricted to having a smaller number of dimensions than the space of her private information. Kos (2012) studies an auction problem where each bidder's report is restricted to a finite message space while her valuation can take any value in an interval. Battigalli and Maggi (2002) introduce the cost of specifying contingencies in a contract, and Mookherjee and Tsumagari (2014) study the effect of communication costs in mechanism design problems. The ways in which the communication constraints affect equilibrium allocations differ among those papers and ours. While we demonstrate that the set of equilibrium outcomes may expand if the specification rule is limited, the set of achievable outcomes in the models of those papers does not expand in the presence of communication constraints. Roughly, this is because an allocation under communication constraints can be replicated as an allocation without such constraints.

Issue Linkage, Agenda Restrictions, and Pledge-and-Review Bargaining

The literature on multi-issue negotiations with issue linkage or agenda restrictions (see, e.g., Fershtman (1990), In and Serrano (2004), Maggi (2016)) considers settings in which not all issues may be negotiated at once, and motivates such situations by arguing that they come from, for example, cognitive constraints or bounded rationality. In the language of our framework, a history-dependent (limited) specification rule, together with a non-consensual termination rule, can accommodate a situation in which negotiating parties agree on an outcome issue-by-issue. These papers consider a specific class of sets of alternatives and specification rules, and analyze how the outcome of a negotiation varies with the way the specification rule depends on histories. Our paper, in contrast, considers general sets of alternatives and general (but non-history dependent) specification rules, and studies how specification rules affect the set of negotiation outcomes. A similarity is that a restriction on specifiability may enlarge negotiation outcomes in a way that each additional outcome is Pareto-dominated by some negotiation outcome sustained under unlimited specifiability.

In a recent working paper, Harstad (2018) considers what he calls a "pledge-and-review bargaining" model in the context of the COP meetings for climate change.

In the model, each party simultaneously proposes unconditional INDC periodically.[8] The crucial differences from our model are the timing assumption (synchronicity is assumed), time preferences (discounting plays a key role in his analysis), and utility functions (some form of externalities is assumed). The results are similar in that both his and our models predict inefficiency under limited specifiability, but the logic is different because his argument depends on different countries having different time preferences and CO2 emissions having negative externalities.

2 The Model

Environments. An *environment* is a triple $G = \langle N, X, (u_i)_{i \in N} \rangle$. The set of players $N := \{1, 2, \ldots, n\}$ is finite with $n \geq 2$. Let X be the non-empty set of *alternatives*. Player i's payoff function is $u_i : X \to \mathbb{R}$. Throughout, we treat a generic player i as female.

Negotiation Games. The players in N engage in rounds of negotiations, which we call a *negotiation game* (or simply a *negotiation*) of a given G. They make announcements in a given order, where each player's announcement comprises of a subset of alternatives, referred to as a *proposal*, and a *response* to the opponents' previous proposals. The payoffs from an agreed-upon alternative in the negotiation are those specified in the description of the environment. If the players do not agree on any alternative (i.e., the negotiation lasts indefinitely), they obtain the *disagreement payoffs* $d \in \mathbb{R}^n$.

Formally, the negotiation of G is an extensive-form game $\Gamma = \langle G, d, \rho, (\mathcal{P}_i)_{i \in N}, \varphi \rangle$, where ρ is the *proposer rule*, $(\mathcal{P}_i)_{i \in N}$ is the *specification rule*, and φ is the *termination rule*. The proposer rule determines who can speak when. The specification rule designates what each player can potentially announce at each time. The termination rule determines when the players conclude their negotiation. We will formally explain these components in what follows.

Histories. A negotiation takes place in a discrete-time setting, with each time indexed by $t \in \mathbb{N}_0 := \mathbb{N} \cup \{0\}$. At each time $t \in \mathbb{N}$, there is a set of proposers $I^t \subseteq N$, and the proposers in I^t simultaneously announce responses to past proposals as well as

[8] As discussed in footnote 1, INDC stands for "intended nationally determined contribution" in which each negotiating country specifies its own emission level.

make new proposals. Each proposer i's response is $R_i^t \in \{\text{Yes}, \text{No}\}$, and her proposal is $P_i^t \in 2^X$. Together, an event at time t is expressed by a list of the set of proposers, their responses and their proposals $(I^t, ((R_i^t, P_i^t))_{i \in I^t})$.[9] For simplicity of notation in later defining a termination rule, we assume that the situation is as if the negotiation starts after everyone says (No, X) at time 0. Thus, the event at time 0 is exogenously fixed and is expressed as $h^0 := \big((N, ((\text{No}, X))_{i \in N})\big)$. That is, $I^0 = N$, $R_i^0 = \text{No}$, and $P_i^0 = X$ for each $i \in N$. Let $\mathcal{H}^0 := \{h^0\}$.

A history is a (finite or infinite) enumeration of events at consecutive times from time 0. That is, a history h is expressed as $h = \big((I^{t'}, ((R_i^{t'}, P_i^{t'}))_{i \in I^{t'}})\big)_{t'=0}^{t}$ for some $t \in \mathbb{N}_0 \cup \{\infty\}$, where we often omit to write the first element corresponding to $t' = 0$, i.e., $\big(N, ((\text{No}, X))_{i \in N}\big)$, as part of a history. For each history h with finite length, we denote by $t(h) \in \mathbb{N}_0$ the length associated with h. Let $\mathcal{H}^t$ be the set of all histories h with length $t = t(h) \in \mathbb{N}_0$, and let $\mathcal{H}^\infty$ be the set of infinite histories. Let $\mathcal{H}^* := \bigcup_{t \in \mathbb{N}_0 \cup \{\infty\}} \mathcal{H}^t$.

For each history $h \in \mathcal{H} := \bigcup_{t \in \mathbb{N}_0} \mathcal{H}^t$ with finite length, we write h as

$$h = \left(\left(I^{t'}(h), ((R_i^{t'}(h), P_i^{t'}(h)))_{i \in I^{t'}(h)}\right)\right)_{t'=0}^{t(h)}.$$

For such h and $t' \in \{0, \ldots, t(h)\}$, we denote by $h^{t'}$ the subhistory of the form

$$h^{t'} := \left(\left(I^{t''}(h), ((R_i^{t''}(h), P_i^{t''}(h)))_{i \in I^{t''}(h)}\right)\right)_{t''=0}^{t'}.$$

To simplify notation, we sometimes omit parentheses and/or proposers from a given history. We denote $h' \sqsubseteq h$ ($h \sqsupseteq h'$) if h' is a subhistory of h, i.e., $h' = h^{t'}$ for some $t' \in \{0, \ldots, t(h)\}$. We denote $h' \sqsubset h$ ($h \sqsupset h'$) if h' is a proper subhistory of h, i.e., $h' = h^{t'}$ for some $t' \in \{0, \ldots, t(h) - 1\}$.

Proposer Rules. A proposer rule is a function that deterministically assigns the *proposers* who can make announcements after each possible history: $\rho : \mathcal{H} \to 2^N \setminus \{\emptyset\}$. We assume that for any $h \in \mathcal{H}$, there exists $t' \in \mathbb{N}$ such that, for any $h' \in \mathcal{H}$ such that $t(h') = t'$ and $h' \sqsupset h$, we have $N = \bigcup_{h'' \in \mathcal{H} : h \sqsubseteq h'' \sqsubseteq h'} \rho(h'')$. In words, for any history $h \in \mathcal{H}$ with finite length, there is a time t' such that, for any proper superhistory

[9] The assumption that a player can announce $(\text{Yes}, \cdot)$ even when no one has spoken yet may not look natural. However, under the consensual termination rule that we define below, the Yes/No response at that time has no consequence on the set of equilibrium payoffs.

h' of h with length $t' = t(h')$, every player has an opportunity to speak between the histories h and h'. A proposer rule is said to be *synchronous* if $\rho(h) = N$ for all $h \in \mathcal{H}$. The proposer rule is said to be *asynchronous* if $|\rho(h)| = 1$ for all $h \in \mathcal{H}$ and $\rho(h) \neq \rho(h')$ for any $h, h' \in \mathcal{H}$ with $h \sqsubset h'$ and $t(h') = t(h) + 1$. This latter condition means that no player speaks in two adjacent periods. By abusing notation, we often write $\rho(h) = i$ when $\rho(h) = \{i\}$.

Specification Rules. A specification rule $(\mathcal{P}_i)_{i \in N}$ is a profile of subsets of alternatives that each player can potentially propose after each history. For each player i, we fix a history-independent set $\mathcal{P}_i \subseteq 2^X$ from which she chooses her announcement at each history at which she moves.[10] We sometimes refer to $\mathcal{P}_i$ as player i's specification rule.

We assume that, for each $x \in X$, there exists a profile $(P_i)_{i \in N}$ of proposals such that $\bigcap_{i \in N} P_i = \{x\}$. It says that the players can collectively pin down any alternative. Note that we allow each player i to announce (No, P_i) at a history h in reply to the opponents' previous announcements, in which her last proposal under h coincides with P_i. In this case, her announcement (No, P_i) can be interpreted as the message that she is not satisfied with the opponents' previous announcements and yet her proposal is P_i.

We say that $\mathcal{P}_i$ is *unlimited* if $\{x\} \in \mathcal{P}_i$ for all $x \in X$. In words, player i can specify any single alternative in her proposal. In contrast, $\mathcal{P}_i$ is said to be *limited* if $\{x\} \notin \mathcal{P}_i$ for any $x \in X$. That is, player i can specify no single alternative.[11] The specification rule $(\mathcal{P}_i)_{i \in N}$ is said to be *unlimited* if every $\mathcal{P}_i$ is unlimited, and it is *limited* if every $\mathcal{P}_i$ is unlimited.

Termination Rules. A termination rule is a function $\varphi : \mathcal{H} \to X \cup \{\text{Continue}\}$ which determines, conditional on the negotiation having continued up to a history $h \in \mathcal{H}$, whether the negotiation ends with a certain alternative or it continues at h. If the negotiation ends at h with an agreed-upon alternative x, we write $\varphi(h) = x$. If it continues at h, we write $\varphi(h) = \text{Continue}$.

[10] The announcement of the empty set can be interpreted as withholding making proposals. In real negotiations, one may do so until other parties make proposals.

[11] Note that i's specification rule could be neither unlimited nor limited. However, in a spirit similar to that of footnote 3, we only consider limited and unlimited specifiability in order to highlight the effect of changes in specifiability conditions.

We define a particular termination rule which we call the *consensual* termination rule φ^{con}.[12] In order for $\varphi^{\text{con}}(h) = x$ to hold, it requires that every player i is *ok* with the alternative x at the end of the history h. So as to define what it means by a player being ok with an alternative, we first define, for any given history h, two critical times of a negotiation under h. First, for each $h \in \mathcal{H}$ and $j \in N$, denote by $t_j^{\text{sp}}(h) := \max\{\tau \in \mathbb{N}_0 \mid \tau \le t(h) \text{ and } j \in I^\tau(h)\}$ the latest time at which j has spoken by period $t(h)$. Second, for each $h \in \mathcal{H}$ and $i \in I^{t(h)}(h)$, define $t_i^{\text{No}}(h)$ to be the time at which i sees the latest reply of "No" no later than period $t(h)$. That is,

$$t_i^{\text{No}}(h) := \max\left\{\tau \in \mathbb{N}_0 \mid R_j^\tau(h) = \text{No for some } (j,\tau) \in N \times \{0,\dots,t(h)-1\} \cup \{(i,t(h))\}\right\}.^{13}$$

Now, we define, for each $h \in \mathcal{H}$ and $i \in I^{t(h)}(h)$, player $j(\neq i)$'s *relevant proposal* $P_j^{i\text{-rel}}(h)$ for i at h by

$$P_j^{i\text{-rel}}(h) := \begin{cases} P_j^{t_j^{\text{sp}}(h)}(h) & \text{if } t_i^{\text{No}}(h) \le t_j^{\text{sp}}(h^{t(h)-1}) \\ X & \text{otherwise} \end{cases}.$$

In words, for each player i who speaks at h, her opponent j's relevant proposal for i at h is the most recent announcement by j after the most recent "No" in terms of player i's observation at h.

Player i, who speaks at history h, is ok with an alternative x at h if $\{x\}$ is the intersection of her current proposal and all the relevant proposals for her. Formally, player $i \in I^{t(h)}(h)$ is *ok* with $x \in X$ at $h \in \mathcal{H}$ if $\{x\} = P_i^{t(h)}(h) \cap \left(\bigcap_{j\in N\setminus\{i\}} P_j^{i\text{-rel}}(h)\right)$. The right-hand side is an intersection of two sets. The first is i's proposal in the current period. The second is the intersection of her opponents' latest proposals after i observes the latest "No" (i.e., the intersection of her opponents' relevant proposals). The condition states that the intersection of these two sets is a singleton $\{x\}$.

We view this definition as capturing the idea of "ok" for the following two reasons. First, the second term in the intersection is what the other players have left as pos-

[12] Although we think of the consensual termination rule as the most natural in a wide range of applications, we acknowledge that there would be other sensible termination rules in some other applications. We discuss other possible termination rules in Online Appendix D. Except for Online Appendix D, we only consider the consensual termination rule in order to focus on comparing proposer and specification rules.

[13] Note that $t_j^{\text{sp}}(h)$ and $t_i^{\text{No}}(h)$ are well defined because the sets from which the maximums are taken are non-empty by the assumption that all players speak and say "No" at time 0.

sibilities after someone has expressed dissatisfaction by announcing "No." Proposing the first term such that the intersection becomes $\{x\}$ means that x is the only possibility that can be interpreted as what i has left as a possibility. If, in contrast, the intersection consists of multiple alternatives containing x, it is unclear if i is satisfied with x, or simply wants to wait and see the opponents' responses to determine her future responses by not specifying a single alternative but restricting the set of possibilities. Second, player i could have said "No" as her current response, and she would have been able to make the entire intersection a non-singleton unless her proposal itself is a singleton set. That is, she would have been in the situation where she could safely be interpreted as being *not* ok with x (note that, by assumption, for any alternative $y \in X \setminus \{x\}$, i is able to announce P_i with $y \in P_i$).

The consensual termination rule terminates a negotiation with an alternative at a given history, once each player becomes ok with the same alternative when she speaks after the most recent response of "No." That is,

$$\varphi^{\text{con}}(h) := \begin{cases} x & \text{if each } j \in N \text{ is ok with } x \text{ at } h^{t_j^{\text{sp}}(h)} \text{ and } t^{\text{No}}(h) \leq t_j^{\text{sp}}(h) \\ \text{Continue} & \text{otherwise} \end{cases},$$

where $t^{\text{No}}(h)$ is the latest time at which some player says "No" under $h \in \mathcal{H}$:

$$t^{\text{No}}(h) := \max \left\{ t' \in \mathbb{N}_0 \,\middle|\, R_j^{t'}(h) = \text{No for some } (j, t') \in N \times \{0, \cdots, t(h)\} \right\}.$$

Given an environment, the three rules (proposer, specification, and termination rules) generate the set $H(\subseteq \mathcal{H}^*)$ of histories of the negotiation Γ and the set $Z(\subseteq H)$ of terminal histories of Γ. That is, $Z = \{h \in H \mid h \sqsubset h' \text{ implies } h' \notin H\}$. Note that Z includes all histories in H that have infinite length. These three rules are defined independently from each other, in order to examine how a change in a certain rule affects the outcome of the negotiation.

Strategies. A strategy of player $i \in N$ is a plan of what to announce at each history at which she speaks. Letting $H_i := \{h \in H \setminus Z \mid i \in \rho(h)\}$ be the set of non-terminal histories at which player i speaks, her *pure strategy* is a mapping $s_i : H_i \to \{\text{Yes}, \text{No}\} \times \mathcal{P}_i$.

Denote by S_i the set of player i's strategies.[14] Denote by $S := \bigtimes_{i \in N} S_i$ the set of strategy profiles. Each strategy profile $s = (s_i)_{i \in N}$ induces a terminal history sometimes denoted by $h(s) \in Z$.

Outcomes and Payoffs. The *outcome* of the negotiation induced by a strategy profile $s \in S$ under a termination rule φ is defined as follows. If the history $h = h(s) \in Z$ induced by s has finite length and $\varphi(h) = x$, then x is the outcome of the negotiation induced by s. If $h \in Z$ has infinite length, the outcome of the negotiation induced by s is defined as the *disagreement outcome* associated with h.

We abuse notation to write each player i's payoff function in the extensive-form game by $u_i : Z \to \mathbb{R}$, where $u_i(h) = d_i$ for any terminal history h with infinite length (i.e., when the outcome is the disagreement outcome associated with h) and $u_i(h) = u_i(x)$ for any terminal history h which corresponds with an agreed-upon alternative $x \in X$ (i.e., when $\varphi(h) = x$). Note that we assume that the payoff from the disagreement outcome is independent of histories. Note also that there is no discounting.[15] We also abuse notation to define i's payoff function $u_i : S \to \mathbb{R}$, where $u_i(s) = u_i(h(s))$.

Individual Rationality and Pareto Efficiency. We denote by U the *feasible* payoff set: $U := \{u(x) \in \mathbb{R}^n \mid x \in X\}$, where $u(x) := (u_i(x))_{i \in N}$. We say that a payoff profile $v \in U$ is individually rational (IR) if $v \geq d$.[16] A payoff profile $v \in U$ is Pareto

[14]Throughout the paper, we drop the reference to "pure"ness of strategies unless there is room for confusion. In Online Appendix C.2, we give an example in which the equilibrium payoff set expands by allowing for behavioral strategies. Specifically, the equilibrium payoff set includes the di-convex span (Aumann et al., 1968) of the pure-strategy equilibrium payoff profiles.

[15]In the applications that we have in mind, it would be unrealistic to relate the key determinants of negotiation outcomes solely to impatience. For example, a COP conference continues during a fixed short period of time, while the stake of the negotiation is large and long-term. Thus, it would be unlikely that the discounting that would take place during the course of the negotiation period would affect the outcome. Validity of such an argument relies on how sensitive the results are to the introduction of small discounting. As we discuss in Online Appendix C.1, for each statement of formal results for which we construct an equilibrium, if the number of alternatives is finite and there is no tie in payoffs, then the equilibrium that we construct continues to be an equilibrium under any sufficiently high discount factor strictly less than 1. If, on the other hand, the discount factor is close to zero, an alternative may be an equilibrium outcome even if it is not an equilibrium outcome under no discounting if such an alternative requires less time to reach a consensus (Online Appendix C.1 provides an example).

[16]For any $x = (x_i)_{i \in N}, y = (y_i)_{i \in N} \in \mathbb{R}^n$, we write $x \geq y$ if and only if $x_i \geq y_i$ for all $i \in N$; and $x > y$ if and only if $x_i > y_i$ for all $i \in N$.

efficient if there is no $v' \in U \setminus \{v\}$ such that $v' > v$. We say that an alternative $x \in X$ is individually rational (resp. Pareto efficient) if $u(x)$ is individually rational (resp. Pareto efficient). Denote by $\mathrm{IR}(U, d) := \{v \in U \mid v \geq d\}$ the set of individually-rational payoffs. Denote by $\mathrm{PE}(U) := \{v \in U \mid v \text{ is Pareto efficient}\}$ the set of Pareto-efficient payoffs. We assume that $\mathrm{IR}(U, d)$ is compact and that, for each $i \in N$, there is $v \in \mathrm{IR}(U, d)$ with $v_i > d_i$. The interpretation of the latter part of this assumption is that each player can potentially benefit from the negotiation. In particular, $\mathrm{IR}(U, d)$ is not empty.

We define the *IR-Pareto-meet* (of (U, d)) by $U^M(U, d) := \{v \in \mathrm{IR}(U, d) \mid v_i \geq w_i \text{ for some } (w, i) \in \mathrm{PE}(U) \times N\}$. That is, $U^M(U, d)$ (often shorthanded by U^M) is the set of individually-rational payoff profiles which give each player a payoff no less than her worst Pareto-efficient payoff.[17] We also call the set of alternatives whose payoff profiles lie in the IR-Pareto-meet, $X^M := \{x \in X \mid u(x) \in U^M(U, d)\}$, the *IR-Pareto-meet* (of (X, d)). Figure 2 depicts three examples that illustrate the concepts introduced here.

Equilibrium Concept. Our solution concept is subgame-perfect equilibrium (henceforth "SPE") of the negotiation $\Gamma = \langle G, d, \rho, (\mathcal{P}_i)_{i \in N}, \varphi \rangle$.[18] Let X^{SPE} be the set of SPE outcomes, i.e., the set of outcomes of the negotiation Γ induced by some SPE.[19]

2.1 Illustration of the Consensual Termination Rule

Here, we illustrate the consensual termination rule through two examples. First, consider the bargaining protocol of Ståhl (1972) and Rubinstein (1982). Two players bargain over the share of a pie of size 1: $X = \{(s, 1-s) \in \mathbb{R}^2 \mid s \in [0, 1]\}$. They specify the exact share, so that the specification rule is unlimited: $\mathcal{P}_i = \{\{(s, 1-s)\} \in 2^X \mid s \in [0, 1]\}$ for each $i \in N = \{1, 2\}$. The proposer rule is asynchronous. The

[17] Rabin (1994) calls the "Pareto meet" the set of payoff profiles which give each player no less than her worst Pareto-efficient Nash-equilibrium payoff. Notice that we do not require the "Nash" restriction, and we added the modifier "IR-."

[18] Our negotiation games do not satisfy "continuity at infinity" (Fudenberg and Levine, 1983) due to lack of discounting, and hence the "one-stage deviation" principle cannot be applied. Hence, our proofs demonstrate that a specific strategy profile s^* constitutes a SPE by showing that each s^*_i is a best response to s^*_{-i} in all subgames.

[19] In all of our formal results, we do not consider whether players can achieve the disagreement payoffs d in a SPE. Whether it is possible can be incorporated with minor (but sometimes cumbersome) modifications of the statements.

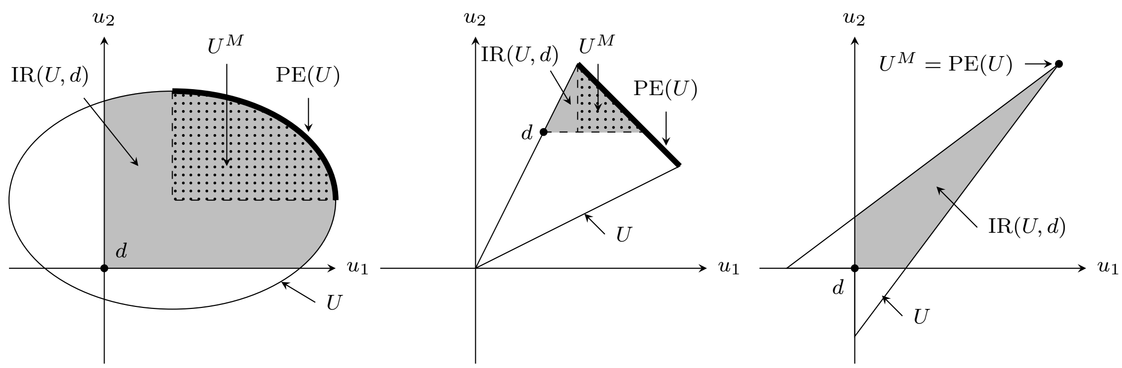

Figure 2: The feasible payoff set (U), the disagreement payoffs (d), the set of individually-rational payoff profiles (IR(U, d); the gray region), the set of Pareto-efficient payoff profiles (PE(U)), and the IR-Pareto-meet (U^M; the dotted region).

consensual termination rule terminates the negotiation with outcome $x = (s, 1 - s)$ at any history of the form $h = (h^{t(h)-2}, (\cdot, \{x\}), (\text{Yes}, \{x\}))$, where we omit the set of proposers from the description of a history.[20]

To see how the negotiation works, notice that, at the end of the first period, proposer i is ok with x after proposing $(R_i, \{x\})$. In general, when responding to j's proposal $(R_j, \{x\})$, we regard player i as accepting the offer x if and only if i's proposal is (Yes, $\{x\}$). In this case, the game ends. Otherwise, i's proposal $(R_i, \{y\})$ is i's counteroffer, and i is always ok with y after such a proposal.

In the second example, we consider an environment $G = \langle N, X, (u_i)_{i \in N} \rangle$ with $N = \{1, 2\}$ where the set of alternatives is the set of action profiles in the following two-player normal-form game: $\overline{G} = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, where A_i is the set of player i's actions and $X = A_1 \times A_2$.[21] In this case we say G is associated with $\overline{G}$. We assume that for the case when specifiability is unlimited, each player can announce an action profile as her proposal. Under limited specifiability, we assume that each player can only specify her own action, and cannot specify the opponent's action.[22] We consider synchronous and asynchronous proposer rules.

Table 1 shows, for each pair of specification and proposer rules, whether each

[20]We frequently use this abuse of notation.

[21]Dutta and Takahashi (2012) give a bargaining interpretation to their finite-horizon game in which two players alternate and choose actions in a $K \times K$ normal-form game, where diagonal action profiles are interpreted as agreements. They consider two models: one in which players' actions are fixed once an agreement is reached, and another in which their actions can be changed until the deadline. In either model, players receive flow payoffs.

[22]Formally, the case of unlimited specifiability corresponds to $\mathcal{P}_i = \{\{a\} \mid a \in A\}$, while that of limited specifiability corresponds to $\mathcal{P}_i = \{\{a_i\} \times A_{-i} \mid a_i \in A_i\}$.

	Specifiability	Proposal	History h	Termination
1	Unlimited	Asynchronous	$(h', (\text{No}, a), (\text{Yes}, a))$	Terminal
2	Unlimited	Synchronous	$(h', ((\text{No}, a), (\text{No}, a)))$	Terminal
3	Unlimited	Synchronous	$(h', ((\text{No}, a), (\text{No}, a')), ((\text{Yes}, a'), (\text{Yes}, a)))$	Non-terminal
4	Limited	Asynchronous	$(h', (\cdot, a_i), (\text{Yes}, a_{-i}), (\text{Yes}, a_i))$	Terminal
5	Limited	Asynchronous	$(h', (\cdot, a_i), (\text{Yes}, a_{-i}), (\text{Yes}, a'_i))$	Non-terminal
6	Limited	Synchronous	$(h', ((\cdot, a_1), (\cdot, a_2)), ((\text{Yes}, a_1), (\text{Yes}, a_2)))$	Terminal
7	Limited	Synchronous	$(h', ((\cdot, a_1), (\cdot, a_2)), ((\text{No}, a_1), (\text{No}, a_2)))$	Non-terminal

Table 1: List of terminal/non-terminal histories. To simplify the notation, we drop reference to the set of proposers. We assume $a \neq a'$ and $a_i \neq a'_i$.

history h is terminal or not under the consensual termination rule, where we abuse notation to denote by a the proposal $\{a\}$ under unlimited specifiability and by a_i the proposal $\{a_i\} \times A_{-i}$ under limited specifiability. We note that Row 1 also corresponds to the bargaining protocol of Ståhl (1972) and Rubinstein (1982). Rows 3 and 5 are especially worth explaining. For Row 3, h is not terminal because two players are ok with different action profiles a and a', and in such a circumstance they would need more conversations to reach a consensus.[23] For Row 5, h is not terminal because, despite the fact that both players' latest responses are "Yes," player $-i$ is ok with (a_i, a_{-i}) in period $t-1$ while player i is ok with (a'_i, a_{-i}) in period t.

3 Benchmark Cases

We start with benchmark observations. First, for any proposer and specification rules, every SPE outcome has to be individually rational under the consensual termination rule.

Proposition 1 (Indiviual rationality)**.** *Any SPE outcome of a negotiation* $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ *is individually rational.*

If there existed a SPE outcome that makes some player worse than the disagreement outcome, then any such player i would be able to profitably deviate by an-

[23]As an example, consider Ann and Bob, who exchange emails about their plan for the next day. Suppose that, after Ann expresses her willingness to go to the movies and Bob expresses his willingness to go to a museum, they simultaneously send replies to each other in which Ann writes "Yes, let's go to the museum" and Bob writes "Yes, let's go to the movies." Unless there is some predetermined rule, their email exchanges would need to continue to settle on a single plan for the day.

nouncing (No, P_i) such that P_i contains an individually-rational alternative y at every history at which she speaks (by assumption, there are an individually-rational alternative $y \in X$ and a proposal $P_i \in \mathcal{P}_i$ with $y \in P_i$). When player i announces (No, P_i), she is not ok with any alternative when P_i is not a singleton and she is ok with y if $P_i = \{y\}$. Thus, any player i can guarantee herself at least the disagreement payoff by repeatedly announcing (No, P_i).

The next observation is a "folk theorem" under synchronous proposer rules: Any individually-rational alternative can be supported as a SPE outcome under the consensual termination rule, irrespective of specification rules.[24]

Proposition 2 ("Folk theorem")**.** *For any negotiation $\langle G, d, \rho, (\mathcal{P}_i)_{i \in N}, \varphi^{\text{con}} \rangle$ such that ρ is synchronous, every $x \in X$ with $u(x) \geq d$ is a SPE outcome.*

To see this, fix x and a profile of proposals $(P_i)_{i \in N}$ such that $\{x\} = \bigcap_{i \in N} P_i$ (such a profile exists by assumption). The following is a SPE which supports x as its outcome. Each player i announces (Yes, P_i) if the announcement profile in the last period entails no deviation (in particular, the players announce (Yes, P_i) at the first period at which they move), while they announce (No, P_i) otherwise. No player has an incentive to deviate because, given the opponents' strategies, after any history, it is only x or the disagreement that can be an outcome.

We can sustain such a strategy profile as a SPE because no single player can influence the opponents' future actions by committing to a proposal, just as in repeated coordination games where a repetition of an inefficient Nash equilibrium is a subgame-perfect equilibrium. Such lack of commitment is partially overcome when a proposer rule is asynchronous.

Propositions 1 and 2 imply the full characterization of the set of SPE outcomes under the synchronous proposer rule:

Corollary 1 (Synchronous proposals)**.** *For any negotiation $\langle G, d, \rho, (\mathcal{P}_i)_{i \in N}, \varphi^{\text{con}} \rangle$ such that ρ is synchronous, $x \in X$ is a SPE outcome if and only if $u(x) \geq d$.*

[24]Kalai et al. (2010) show a "folk theorem" with synchronous moves in the context of what they call the (two-player) "commitment games."

4 Specifiability-Free Results with Asynchronous Moves

The previous section shows lack of prediction under synchronous-move negotiations. The rest of the paper focuses on asynchronous proposer rules. We will see that asynchronicity helps us narrow down our prediction, and the way in which asynchronicity helps players commit to outcomes depends on the structures of environments and specification rules. This section provides predictions that are free from specification rules. Section 5 then discusses how the limitations on specification rules change the predictions.

4.1 Negotiations with a Common-Interest Alternative

We say that a negotiation $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ has a *common-interest alternative* $x^* \in X$ if $u(x^*) > v$ for all $v \in \text{IR}(U, d) \setminus \{u(x^*)\}$.[25] In other words, the negotiation has a common-interest alternative if and only if X^M is a singleton.[26]

Theorem 1 (Unique selection)**.** *Any negotiation $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ such that ρ is asynchronous has a unique SPE outcome x^* if and only if it has a common-interest alternative x^*.*

Note that the theorem applies to general n-player cases. Before explaining the intuition behind the result, let us provide two examples to discuss what this theorem implies.

Example 1. The environment is given by $X = A_1 \times A_2$, where A_i and u_i are those of the "tacit coordination game" studied by Bryant (1983) and Huyck et al. (1990): Each player chooses an effort level $a_i \in A_i = [0, 1]$, and her payoff is $u_i(a_1, a_2) = 2\min\{a_1, a_2\} - \frac{1}{2}a_i$. We let $d = (0, 0)$. The feasible payoff set and the unique SPE payoff profile are depicted in the left panel of Figure 3. The normal-form game has a continuum of Nash equilibria $\{(a_1, a_2) \in A \mid a_1 = a_2 \in [0, 1]\}$, among which the action profile $(1, 1)$ Pareto-dominates all other Nash equilibria. Our result shows that $(1, 1)$ is the *unique* SPE outcome of a negotiation game when the proposer rule is asynchronous. Put differently, any Pareto-dominated Nash equilibrium (i.e., (a_1, a_2) with $a_1 = a_2 \in [0, 1)$) is not sustained as a SPE of a negotiation game with an asynchronous proper rule. □

[25] Note that $u(x^*) \in \text{IR}(U, d)$ holds because $\text{IR}(U, d) \neq \emptyset$.

[26] This and Theorem 2 that we present later imply the "only if" part of Theorem 1.

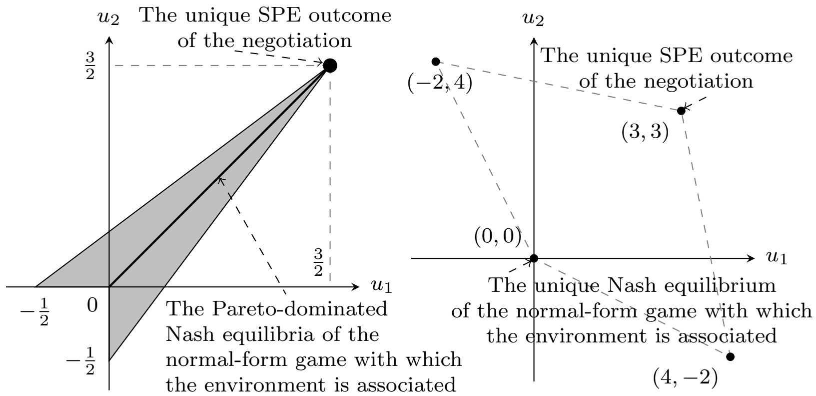

	C_2	D_2
C_1	3, 3	−2, 4
D_1	4, −2	0, 0

Figure 3: The feasible, Nash, and SPE payoff sets for the tacit coordination game (left). A Prisoners' Dilemma game: the payoff matrix (right) and the feasible, Nash, and the SPE payoff sets with $d_1 = d_2 \in (-2, 0]$ (center).

Example 2. The environment is given by $X = A_1 \times A_2$, where $A_i = \{C_i, D_i\}$ and u_i are those of the Prisoners' Dilemma game depicted in the middle and right panels of Figure 3. We set disagreement payoffs $d_1 = d_2 \in (-2, 0]$. The normal-form game has a unique Nash equilibrium (D_1, D_2). The Pareto frontier consists of three points, corresponding to action profiles (C_1, C_2), (C_1, D_2), and (D_1, C_2). However, (C_1, C_2) is a unique individually-rational and Pareto-efficient action profile, and thus Theorem 1 applies. Our result shows that (C_1, C_2) is a *unique* SPE outcome of a negotiation game when the proposer rule is asynchronous.[27] Thus, the set of Nash equilibria of the normal-form game with which the environment is associated and the set of SPE outcomes in the negotiation game can be disjoint, even if every element in the former set is individually rational. □

Equilibrium existence will be shown in the next subsection (Corollary 2) for the general model presented in Section 2.[28] Here we explain the intuition behind why only the common-interest alternative can be an equilibrium outcome, using Example 2. Consider the case with limited specifiability in which each player can only announce her own action (a similar argument holds under any specifiability condition).

[27] Kalai (1981) studies a pre-play communication model in which players take actions in the Prisoner's Dilemma game, and shows that (C_1, C_2) is a unique equilibrium outcome.

[28] The proof for existence is constructive. For example, consider Example 1 with the specification rule limited in that each player can only announce her own action. One can construct a SPE in which each player announces (Yes, 1) as long as the opponent has announced $(\cdot, 1)$ in the previous period; otherwise, she announces (No, 1). This construction is generalized in the proof of Theorem 2 in the next subsection to deal with non-common interest negotiations.

Assume also that player 1 moves in odd periods while player 2 does in even periods. First, notice that the individually-rational payoff profiles are $(0,0)$ and $(3,3)$. So, by Proposition 1, the two players' payoffs have to be equal in any SPE.

After any history of the form $h = (h^{t(h)-2}, (\text{Yes}, C_1), (\text{Yes}, C_2))$, player 1 can terminate the negotiation with (C_1, C_2) by announcing (Yes, C_1). Since she can attain her best possible equilibrium payoff conditional on h, (C_1, C_2) is the unique SPE outcome of the subgame starting at h.

Given any history of the form $h^{t(h)-1} = (h^{t(h)-2}, (\text{Yes}, C_1))$, player 2 announcing (Yes, C_2) guarantees player 1 a payoff of 3, which, in turn, guarantees himself a payoff of 3. In other words, (C_1, C_2) is the unique SPE outcome of the subgame starting at $h^{t(h)-1}$.

At the start of the negotiation, player 1's announcement (Yes, C_1) induces the history $(h^0, (\text{Yes}, C_1))$, which is of the form $h^{t(h)-1}$ as above. Hence, player 1 can guarantee herself a payoff of 3, which means player 2's minimum SPE payoff is also 3. This implies that, in any SPE, the payoff must be 3 for each player. Thus, the unique SPE outcome is the unique individually-rational Pareto-efficient action profile (C_1, C_2).

In the above argument, we used the fact that two players must receive equal payoffs in any SPE. The argument actually depends only on the fact that there is a common action profile whose payoffs strictly dominate all other payoff profiles (including the disagreement payoffs), and it is why the result can be extended to any negotiation with a common-interest alternative as in Theorem 1. In other words, the theorem depends on the fact that if i receives the best individually-rational payoff under a given strategy profile, it fully pins down $-i$'s payoff under that strategy profile. The next subsection deals with the case in which there is no common interest alternative, and shows that multiple equilibrium outcomes exist.

In standard repeated Prisoners' Dilemma games, each player i unconditionally choosing D_i is a subgame-perfect equilibrium, and one may wonder why such a strategy profile cannot constitute a SPE in our negotiation game. The reason is that a termination of a negotiation is endogenously determined, and hence, under history h, announcing C_1 and announcing D_1 have asymmetric implications on the process of negotiation. While C_1 leads to a termination with the outcome (C_1, C_2), D_1 does not lead to a termination. Proposing D_1 cannot terminate the negotiation because player 2 would not be ok with (D_1, C_2).

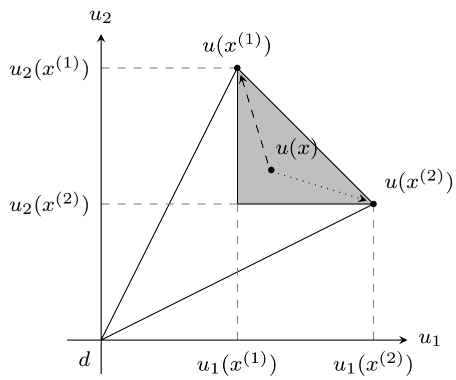

Figure 4: Illustration of the proof of Theorem 2: The shaded area is $\{u(x) \mid x \in X^M\}$. The dashed and dotted arrows indicate punishments for players 1 and 2, respectively.

4.2 Negotiations without a Common-Interest Alternative

We now turn to negotiations that do not (necessarily) have a common interest alternative and seek predictions free from specification rules. The first result is also independent of proposer rules.

Theorem 2 (IR-Pareto-meet)**.** *For any negotiation* $\langle G, d, \rho, (\mathcal{P}_i)_{i \in N}, \varphi^{\text{con}} \rangle$*, every* $x \in X^M$ *is a SPE outcome.*

If the individually-rational Pareto frontier consists of multiple points, then the fact that player i receives the best individually-rational payoff does not pin down $-i$'s payoff. This leads to a possibility of punishment. As a consequence, the set of equilibrium payoffs consists of multiple points. As an extreme case, in the case of two players, suppose that a given environment is strictly competitive (Osborne and Rubinstein, 1994) in that $u_1(x) \geq u_1(x')$ if and only if $u_2(x) \leq u_2(x')$ for all $x, x' \in X$. Then, Theorem 2 implies that any individually-rational alternatives are SPE outcomes since they are in the IR-Pareto-meet.

The idea of the proof is as follows. For each player i, let $x^{(i)}$ be her worst individually-rational and Pareto-efficient alternative.[29] We show that every point in the IR-Pareto-meet X^M can be attained by using $x^{(i)}$ to punish i's deviations. Figure 4 depicts the intuition for a two-player case.

We note that showing that $x^{(i)}$ can be sustained in the subgame after a deviation is by no means trivial, as we need to make sure that there exists a strategy profile in

[29] $x^{(i)}$ exists for each i because $\text{IR}(U, d)$ is non-empty and compact.

which players take best responses even off the equilibrium path: Off the equilibrium path, some players may have already been ok with an alternative x, and it may be of the remaining players' best interest to agree on x even if x is not a SPE outcome. Checking if such an agreement is of "best interest" of these remaining players is complicated because each player off the path needs to correctly forecast the opponents' future actions.

Under our assumption that $\mathrm{IR}(U, d)$ is a non-empty compact set, the IR-Pareto-meet is always non-empty. Hence, any negotiation has a SPE.

Corollary 2 (Existence of SPE)**.** *Any negotiation* $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\mathrm{con}}\rangle$ *has a SPE.*

Theorem 2 shows that the set of SPE alternatives always includes the IR-Pareto meet X^M. The question arises as to how large the SPE set can be beyond X^M. Hereafter, we restrict attention to two-player negotiations for which sharp contrasts between the SPE outcomes under unlimited and limited specifiability are obtained.

We start with providing the players' SPE payoff lower bounds under asynchronous proposer rules. For this purpose, let

$$\begin{aligned} v^{[i,M]} &:= \max\{v_i \mid v \in \mathrm{IR}(U,d)\}; \\ v^{[i,m]} &:= \min\left\{v_i \,\middle|\, v \in \mathrm{IR}(U,d) \text{ and } v_{-i} \geq v^{[-i,M]}\right\}; \quad \text{and} \\ \underline{u}_i &:= \min\left\{v_i \,\middle|\, v \in \mathrm{IR}(U,d) \text{ and } v_{-i} \geq v^{[-i,m]}\right\}. \end{aligned}$$

Call $\underline{u}_i$ the *worst Pareto-guaranteeing payoff* for player $i \in N$. In what follows, it turns out to be her SPE payoff lower bound. Let $\underline{u} := (\underline{u}_1, \underline{u}_2)$. We sometimes denote $\underline{u}$ by $\underline{u}(U, d)$ to make clear its dependence on (U, d). The payoffs $\underline{u}$ exist because $\mathrm{IR}(U, d)$ is non-empty and compact. Note also that, again by the non-emptiness and compactness of $\mathrm{IR}(U, d)$, for each $i \in N = \{1, 2\}$, there is $x^{(i)} \in X^M$ such that $v^{[i,m]} = u_i(x^{(i)})$ and $v^{[-i,M]} = u_{-i}(x^{(i)})$. By construction, $x^{(i)}$ is i's worst alternative in X^M. The next example illustrates the computation of $\underline{u}$.

Example 3. Let $U = \mathrm{conv}(\{(0, 0), (4, 2), (2, 4)\})$. In the left panel of Figure 5, we set $d = (0, 0)$. By inspection, we obtain $(v^{[1,M]}, v^{[2,m]}) = (4, 2)$, $(v^{[1,m]}, v^{[2,M]}) = (2, 4)$, and $(\underline{u}_1, \underline{u}_2) = (1, 1)$.

In the middle panel of Figure 5, we change d to $d = (\frac{3}{2}, 3)$. We have $(v^{[1,M]}, v^{[2,m]}) = (3, 3)$, $(v^{[1,m]}, v^{[2,M]}) = (2, 4)$, and $(\underline{u}_1, \underline{u}_2) = d$.

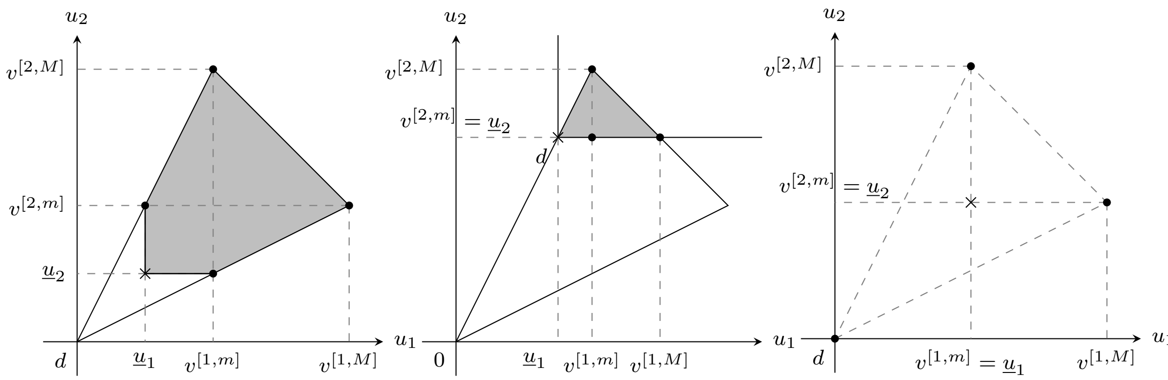

Figure 5: The computation of the worst Pareto-guaranteeing payoffs: the case with $d = (0,0)$ (left); the case with $d = (\frac{3}{2}, 3)$ (middle); and the case with finite X and $d = (0,0)$ (right).

Suppose now that $U = \{(0,0),(4,2),(2,4)\}$ as in the right panel of Figure 5. Let $d = (0,0)$. We have $(v^{[1,M]}, v^{[2,m]}) = (4,2)$, $(v^{[1,m]}, v^{[2,M]}) = (2,4)$, and $(\underline{u}_1, \underline{u}_2) = (2,2)$. □

Proposition 3 (The SPE lower bounds)**.** *Fix a two-player negotiation $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ such that ρ is asynchronous. If $x \in X$ is a SPE outcome, then $u(x) \geq \underline{u}$.*

To prove the proposition, consider a history at which player 1 has proposed P_1 such that $P_1 \cap P_2 = \{x^{(2)}\}$ for some P_2 (such P_1 exists by assumption), where $u_1(x^{(2)}) = v^{[1,M]}$ and $u_2(x^{(2)}) = v^{[2,m]}$. Since $v^{[1,M]}$ is 1's maximum possible SPE payoff and $v^{[1,M]} > d_1$ by assumption, player 2 can guarantee himself a payoff of $v^{[2,m]}$ at that history by announcing (Yes, P_2). Observe that $v^{[2,m]}$ is player 2's minimum individually-rational and Pareto-efficient payoff.

Solving backwards, in a similar vein, at any history, player 1 can guarantee herself the minimum payoff $\underline{u}_1$ such that her opponent 2 can obtain 2's minimum individually-rational and Pareto-efficient payoff. Hence, $\underline{u}_1$ is 1's SPE payoff lower bound. In words, player 1 can guarantee herself the minimum payoff which gives player 2 at least as high as player 2's minimum individually-rational and Pareto-efficient payoff.

Note that if there is a common-interest alternative x^*, we have $v^{[i,M]} = v^{[i,m]} = u_i(x^*)$ for each $i \in \{1,2\}$. By replacing $v^{[1,M]}$, $v^{[2,m]}$ and $x^{(2)}$ with $u_1(x^*)$, $u_2(x^*)$ and x^*, respectively, the above intuition echoes with the intuition explained for Theorem 1 showing unique selection in the case of common-interest negotiations. In particular, we have $\underline{u} = u(x^*)$.

Recall the folk theorem under synchronous proposer rules (Proposition 2). The implication of Proposition 3 is that certain payoffs may not be achievable in SPE under asynchronous proposer rules. That is, asynchronicity helps narrow down SPE payoffs.

5 Unlimited vs. Limited Specifiability with Asynchronous Moves

This section compares negotiations with different specifiability conditions under asynchronous moves. Section 5.1 analyzes negotiations under unlimited specifiability. The main result is that the set of SPE payoffs is completely characterized by the IR-Pareto-meet under unlimited specifiability. We then move on to the case of limited specifiability in Section 5.2. We show that the set of SPE payoff profiles under limited specifiability always includes the set of SPE payoff profiles under unlimited specifiability, and provide conditions under which extra payoff profiles are achieved.

5.1 Unlimited Specifiability

We first show that a lower bound of player i's SPE payoffs is given by her worst individually-rational and Pareto-efficient payoff when i's specifiability is unlimited.

Proposition 4 (Player with unlimited specifiability)**.** *Fix a two-player negotiation* $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ *where* ρ *is asynchronous. If* $\mathcal{P}_i$ *is unlimited and* $x \in X$ *is a SPE outcome, then* $u_i(x) \geq v^{[i,m]}$.

The result follows because a player whose specifiability is unlimited can commit to choosing her opponent's best individually-rational alternative. Recall that $x^{(i)}$ satisfies $u_{-i}(x^{(i)}) = v^{[-i,M]}$ and $u_i(x^{(i)}) = v^{[i,m]}$. That is, it gives player $-i$ the best individually-rational payoff and gives player i her worst Pareto-efficient and individually-rational payoff. Suppose that player i's specifiability is unlimited so that $\{x^{(i)}\} \in \mathcal{P}_i$. Given any non-terminal history at which i speaks, if she announces $(\text{No}, \{x^{(i)}\})$, then (i) player $-i$ can guarantee herself a payoff no less than $u_{-i}(x^{(i)})$ because she can terminate the negotiation by announcing (Yes, P_{-i}) such that $x^{(i)} \in P_{-i}$, and (ii) player i can guarantee herself an individually-rational payoff because she can keep announcing $(\text{No}, \{x^{(i)}\})$. These two properties imply that, after any

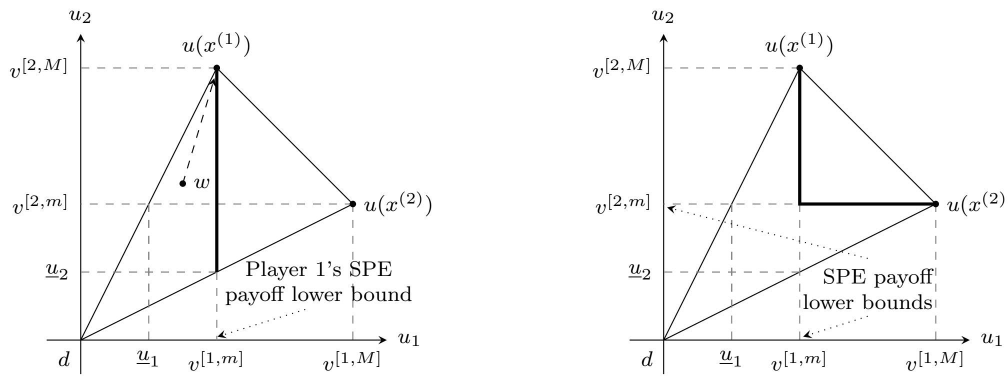

Figure 6: The left panel illustrates the proof of Proposition 4 with $i = 1$. The dashed arrow indicates player 1's deviation to $x^{(1)}$ under a strategy profile inducing w. Player 1's SPE payoff lower bound is $v^{[1,m]}$. The right panel illustrates Corollary 3, where $X^{\text{SPE}} = X^M$.

non-terminal history at which i speaks, she can guarantee herself her worst Pareto-efficient and individually-rational payoff $v^{[i,m]}$. This logic is illustrated in the left panel of Figure 6. For any strategy profile sustaining the payoff profile w in the panel, under any non-terminal history, player 1 can announce $(\text{No}, x^{(1)})$. Player 2 could respond by "Yes" so he can secure the payoff $v^{[2,M]}$, and hence player 1 can secure a payoff of $v^{[1,m]}$, which is larger than w_1.

Combining this result with that of Theorem 2, the set of SPE outcomes is the IR-Pareto-meet under unlimited specifiability. This corollary is illustrated in the right panel of Figure 6.

Corollary 3 (Unlimited specifiability)**.** *Fix a two-player negotiation* $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ *such that* ρ *is asynchronous and* $(\mathcal{P}_i)_{i\in N}$ *is unlimited. Any* $x \in X$ *is a SPE outcome if and only if it is in the IR-Pareto-meet* X^M.

5.2 Limited Specifiability

5.2.1 A Benchmark Result

While Theorem 2 states that the set of SPE alternatives of any negotiation must include the IR-Pareto-meet, Corollary 3 asserts that the set of SPE alternatives under unlimited specifiability is the IR-Pareto-meet. Thus, the set of SPE alternatives under limited specifiability is at least as large as that under unlimited specifiability.

Corollary 4 (SPE payoff sets under limited and unilimited specifiability)**.** *Fix two-player negotiations* $\Gamma = \langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ *and* $\Gamma' = \langle G, d, \rho', (\mathcal{P}'_i)_{i\in N}, \varphi^{\text{con}}\rangle$ *such that* ρ *and* ρ' *are asynchronous,* $(\mathcal{P}_i)_{i\in N}$ *is unlimited and* $(\mathcal{P}'_i)_{i\in N}$ *is limited. If an alternative* $x \in X$ *is a SPE outcome of* Γ*, then it is a SPE outcome of* Γ'*.*

In general, it is possible that some alternative that is not a SPE outcome under unlimited specifiability is a SPE outcome under limited specifiability (examples appear shortly). The reason why this is possible under limited specifiability is that, unlike unlimited specifiability, player i saying (No, P_i) is not a commitment to agreeing on an alternative that is an element of P_i. This is because, even if player $-i$ responds with "Yes" to (No, P_i), player i can still revise her proposal by not announcing (Yes, P_i).

As Proposition 3 shows, however, this does not mean that there is no commitment under limited specifiability: Player i's announcing (No, P_i) is a commitment to agreeing on an alternative no worse than x for i after $-i$ subsequently announces (Yes, P_{-i}) such that $\{x\} = P_i \cap P_{-i}$, because i then has an option to announce (Yes, P_i). This is why the SPE payoff set can still be strictly smaller than the entire set of individually-rational payoff profiles even under limited specifiability. In the extreme case, for any negotiation with a common-interest alternative, players' commitment power is so strong that the SPE outcome is always the unique Pareto-efficient alternative. We now see that the exact degree to which limited specifiability enlarges the SPE payoff set depends critically on the fine details of the payoff structures and on how a restriction on specifiability is imposed.

5.2.2 Unilateral Improvability and SPE Payoffs

We have seen that the SPE payoff set under limited specifiability is weakly larger than the one under unlimited specifiability. Now we show that whether an additional alternative can be supported as a SPE outcome depends on the payoff structure and the way in which specifiability is limited.

To see this point, we say that $x \in X$ is *unilaterally improvable for player* i if for all (P_1, P_2) such that $P_1 \cap P_2 = \{x\}$, there exists $(P'_i, x') \in \mathcal{P}_i \times X$ such that $P'_i \cap P_{-i} = \{x'\}$ and $u(x') > u(x)$. That is, player i can unilaterally deviate and create an intersection that Pareto-improves upon x.

Theorem 3 (Unilateral improvability and limited specifiability)**.** *Fix a two-player negotiation* $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ *such that* X *is finite,* ρ *is asynchronous, and* $(\mathcal{P}_i)_{i\in N}$

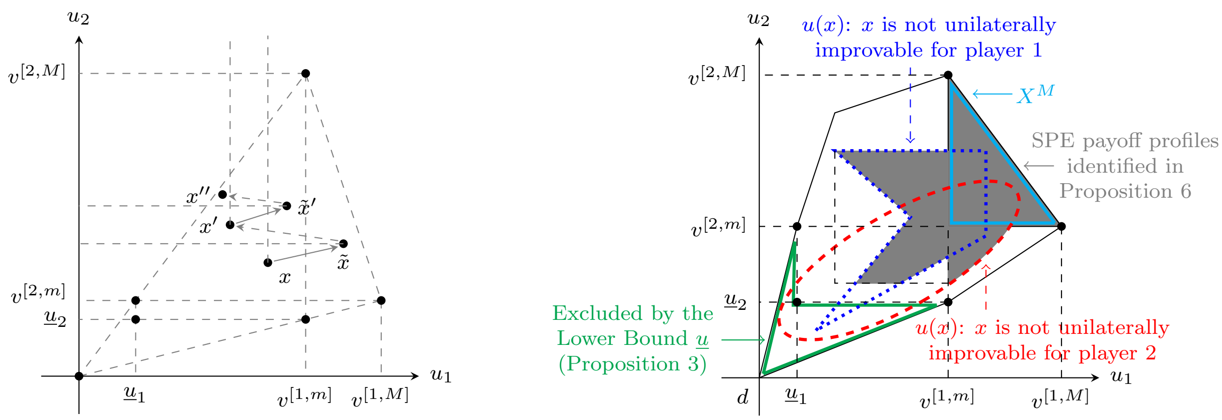

Figure 7: The left panel illustrates the intuition for the proof of Proposition 5 (the "if" part of Theorem 3). The solid arrows indicate deviations by player 1 and the dashed arrows indicate punishments against such deviations. The right panel depicts the set of payoff profiles that satisfy the conditions of Proposition 6 (the "only if" part of Theorem 3). The region enclosed by the dotted line segments and the one with the dashed ellipse correspond, respectively, to the set of x that is not unilaterally improvable for player 1 and the one that is not unilaterally improvable for player 2. Thus, their intersections with $\{u(x) \mid u_i(x) < v^{[i,m]} \text{ and } u_{-i}(x) \geq v^{[-i,m]}\}$ with $i = 1$ and $i = 2$, respectively, determine the set of SPE payoff profiles. The shaded area corresponds to the set of $u(x)$ such that x satisfies one of the three conditions in Proposition 6.

is limited. Then, $X^{\text{SPE}} = X^M$ *if and only if, for each* $i \in N$*, any* $x \in X$ *with* $u_i(x) < v^{[i,m]}$ *and* $u_{-i}(x) \geq v^{[-i,m]}$ *is unilaterally improvable for* i*.*

The "if" part of this theorem follows from the following proposition providing a sufficient condition for $v^{[i,m]}$ to be a lower bound of player i's SPE payoff.

Proposition 5 (Implication of unilateral improvability)**.** *Fix a two-player negotiation* $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ *such that* X *is finite,* ρ *is asynchronous, and* $(\mathcal{P}_i)_{i\in N}$ *is limited. If there is* $i \in N$ *such that any* $x \in \{y \in X \mid u_i(y) < v^{[i,m]} \text{ and } u_{-i}(y) \geq v^{[-i,m]}\}$ *is unilaterally improvable for* i*, then there is no SPE outcome in* $\{y \in X \mid u_i(y) < v^{[i,m]}\}$*. In other words,* i*'s SPE payoff lower bound is* $v^{[i,m]}$*.*

The proof of this proposition consists of two steps. First we prove that, under the given condition, no alternative x with $u_i(x) < v^{[i,m]}$ and $u_{-i}(x) \geq v^{[-i,m]}$ is a SPE outcome. Second, we use this result to show that no alternative x with $\underline{u} \leq u(x) < (v^{[1,m]}, v^{[2,m]})$ is a SPE outcome.

The proof iteratively constructs payoff bounds using unilateral improvability and equilibrium conditions. We explain the intuition using the left panel of Figure 7. Suppose to the contrary that an alternative x with $u_1(x) < v^{[1,m]}$ and $u_2(x) \geq v^{[2,m]}$ is a SPE outcome. Then, unilateral improvability implies that player 1 has a deviation to some $\tilde{x}$ (making a proposal that makes $\{\tilde{x}\}$ the intersection of the proposals) that Pareto-dominates x. By the equilibrium condition, such a deviation has to be punished by an off-equilibrium outcome x' that gives player 1 a payoff no more than $u_1(x)$. Hence, $u_1(x') \leq u_1(x)$ and $u_2(x') > u_2(x)$. Now, since x' is unilaterally improvable for player 1, there exist $\tilde{x}' \in X$ and $P_2 \in \mathcal{P}_2$ such that $u(\tilde{x}') > u(x')$ and $x', \tilde{x}' \in P_2$. Again, a deviation to $\tilde{x}'$ has to be punished by an off-equilibrium outcome x'' that gives player 1 a payoff no more than $u_1(x')$. Going forward, we need to be able to find an infinite sequence of distinct alternatives that goes to the north-west direction given by alternations of deviations and punishments, but this contradicts the assumption imposed in the proposition that X is finite.

We have two remarks. First, the argument holds as long as $\{x \in X \mid u_i(x) < v^{[i,m]} \text{ and } u_{-i}(x) \geq v^{[i,m]}\}$ is finite even if X is infinite. Second, Online Appendix B.1 provides a counterexample to this "if" part (Proposition 5) for the case when X is infinite (or more precisely, $\{x \in X \mid u_i(x) < v^{[i,m]} \text{ and } u_{-i}(x) \geq v^{[i,m]}\}$ is infinite).

For the second step, we consider player 1's deviation to announce a proposal that contains an alternative with payoff $(v^{[1,M]}, v^{[2,m]})$. Such a deviation has to be punished by a continuation strategy that leads to an outcome y giving player 1 a payoff no greater than $v^{[1,m]}$ and player 2 a payoff no less than $v^{[2,m]}$. The bound $v^{[2,m]}$ follows because one option player 2 has after the deviation is to become ok with the alternative with the payoff profile $(v^{[1,M]}, v^{[2,m]})$, which guarantees player 1 a payoff of $v^{[1,M]}$, which in turn guarantees player 2 a payoff of $v^{[2,m]}$. By assumption, however, the alternative y is unilaterally improvable for player 1, and by a similar argument as for the first step, such an alternative cannot be supported in the continuation play.

The "only if" part of Theorem 3 is a consequence of the following stronger result that provides a sufficient condition for an alternative to be a SPE outcome.

Proposition 6 (SPE outcomes under limited specifiability). *Fix a two-player negotiation $\langle G, d, \rho, (\mathcal{P}_i)_{i\in N}, \varphi^{\text{con}}\rangle$ such that ρ is asynchronous and $(\mathcal{P}_i)_{i\in N}$ is limited. Any $x \in X$ is a SPE outcome if one of the following three conditions holds:*

1. $x \in X^M$.

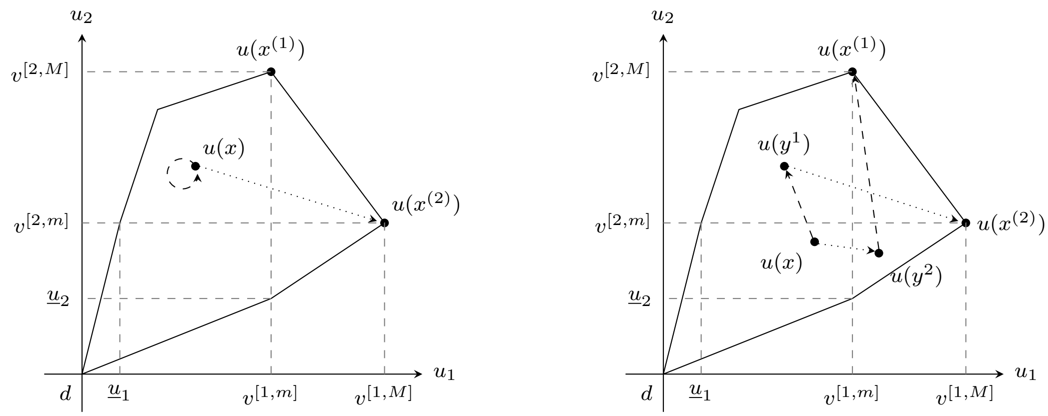

Figure 8: Description of SPE strategies: The left panel illustrates condition 2 of Proposition 6. The dashed arrow indicates a punishment for player 1 while the dotted arrow does for player 2. The alternative x is not unilaterally improvable for player 1. The right panel illustrates condition 3 of Proposition 6. The dashed arrows indicate punishments for player 1 while the dotted arrows do for player 2. The alternative x is not unilaterally improvable for either player 1 or 2, and each y^j is not unilaterally improvable for player $j \in N$.

2. *$d_i \leq u_i(x) < v^{[i,m]}$ and $u_{-i}(x) \geq v^{[-i,m]}$ for some $i \in N$, and x is not unilaterally improvable for i.*

3. *$\underline{u} \leq u(x) < (v^{[1,m]}, v^{[2,m]})$, and the following two conditions hold.*

   (a) x is not unilaterally improvable for each player $j \in N$.

   (b) There exists $(y^1, y^2) \in X^2$ such that, for each $j \in N$, $u_j(y^j) \leq u_j(x)$ and $u_{-j}(y^j) \geq v^{[-j,m]}$ hold and y^j is not unilaterally improvable for j.

The right panel of Figure 7 illustrates the set of SPE payoff profiles that satisfy the conditions stated in Proposition 6. Any payoff profile in the shaded area is a SPE payoff profile.

The proof of Proposition 6 is constructive. Figure 8 illustrates the construction in two cases. In each panel, x is an alternative that satisfies a condition given in the proposition. First, consider the left panel of Figure 8, which illustrates condition 2 of Proposition 6. The following is a punishment strategy that can be used to sustain $u(x)$. If player 2 deviates at some history, then, from that history on, the players switch to another SPE that supports $x^{(2)}$. The dotted arrow in the figure shows such a punishment for player 2. If player 1 deviates, then the players do not switch but continue sustaining $u(x)$. The dashed arrow in the figure shows such a punishment for player 1. To sustain such a punishment, after 1's deviation, player 2 responds with "No" and sticks to the original alternative x unless he can terminate the negotiation with a better outcome for him. However, the assumptions that player 1's specifiability is limited and x is not unilaterally improvable for player 1 imply that no profitable deviation by player 1 would induce player 2 to terminate the negotiation to get a better payoff.

Recall that, under unlimited specifiability, $u(x)$ is not a SPE payoff profile. In particular, a deviation to agreeing on $x^{(1)}$ is profitable. The reason that player 2 does not have an incentive to switch to $x^{(1)}$ under limited specifiability is that if he does so, then the players switch to a SPE which supports $x^{(2)}$ and it gives a lower payoff to player 2 than $u_2(x)$. Here, it is crucial that player 2 cannot be ok with $x^{(1)}$ by himself when he deviates. That is, *player 2 has low commitment power due to his limited specifiability, which gives player 1 greater scope for punishment: Player 1 has an opportunity to punish player 2 conditional on 2's deviation.*

Second, the right panel illustrates condition 3 of Proposition 6. Here, the payoff profile $u(x)$ can be sustained under limited specifiability by using y^1 and y^2 as punishments. That is, we construct a SPE that induces x by a threat to punish player i by an off-path outcome y^i, where y^i is a SPE outcome by the above argument.

We make the following two remarks on Proposition 6. First, it is possible that an alternative $x \in X$ with $u_i(x) < v^{[i,m]}$ and $u_{-i}(x) \geq v^{[-i,m]}$ that is unilaterally improvable for i can be supported as a SPE outcome if there exists another alternative y^i with $u_i(y^i) \leq u_i(x)$ and $u_{-i}(y^i) \geq u_{-i}(x)$ that is not unilaterally improvable for i and that is used as a punishment strategy for i's deviation. Online Appendix B.2 provides such an example. This can happen because the sequence of deviations and punishment described in the left panel of Figure 7 can terminate at such y^i.

Second, Proposition 6 shows that it is possible for negotiations under limited specifiability to lead to more SPE outcomes, but the extent to which this happens depends on the given problem, i.e., the payoff structure of a given environment and how the limitation on the specification rule is imposed. The next example illustrates this point.

Example 4. First, consider the environment associated with the normal-form game in the left panel of Figure 9 and a negotiation with an asynchronous proposer rule and limited specifiability. Specifically, we assume that each player can only announce their action. The right panel of Figure 9 shows, in particular, that a payoff profile

	L	R
U	0, 0	2, 4
D	4, 2	1, 3

	L	R
U	1, 3	2, 4
D	4, 2	0, 0

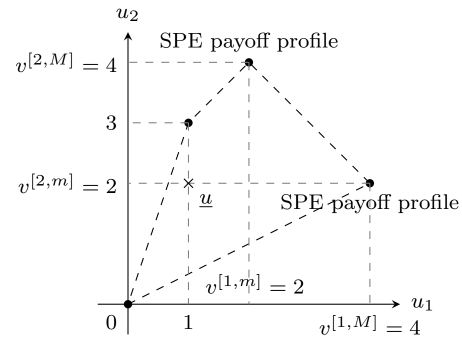

Figure 9: Illustration of Example 4: For a negotiation associated with the normal-form game in the left panel, $(1, 3)$ *cannot* be sustained in a SPE, and the worst Pareto-guaranteeing payoff is a loose lower bound. For a negotiation associated with the normal-form game in the middle panel, $(1, 3)$ is sustained in a SPE. The right panel depicts the feasible and SPE payoff sets.

$(1, 3)$ is at least as high as the worst Pareto-guaranteeing payoff profile $\underline{u} = (1, 2)$.

The payoff profile $(1, 3)$ *cannot* be sustained in a SPE. To see this, suppose that it is sustained under some SPE. In order to sustain this payoff profile, player 2 must announce $(\cdot, R)$ at some point on the equilibrium path at which the negotiation does not terminate. After such an announcement, however, player 1 can announce (Yes, U), after which player 2 has an option to say (Yes, R) that ends the negotiation. This means that, after player 2's announcement of $(\cdot, R)$ that does not end the negotiation which would be a necessary step for $(1, 3)$ to be sustained, player 1 can guarantee a payoff of 2. Hence, $(1, 3)$ cannot be sustained in a SPE.[30] Notice that this argument hinges on the fact that (D, R) is unilaterally improvable for player 1.

Next, consider the environment associated with the normal-form game in the middle panel of Figure 9, where we replace the payoff profile under (U, L) with $(1, 3)$ and the one under (D, R) with $(0, 0)$. Under asynchronous moves and limited specifiability, the payoff profile $(1, 3)$ is sustainable under SPE. This is because if player 1 deviates by proposing D, then player 2 announces (No, L) in equilibrium. If player 2 deviates by not announcing (No, L), the players switch to the Pareto-efficient outcome (D, L) (with payoffs $(4, 2)$). Formally, the result follows from Proposition 6 as (U, L) is not unilaterally improvable for player 1.

The reason for the difference is that the alternative sustaining the payoff profile

[30] We can generally establish that an alternative $x \in X$ with $u(x) \geq \underline{u}$ is not a SPE outcome if the following hold: For any $(i, x) \in N \times X$ with $u_i(x) < v^{[i,m]}$ and for any (P_1, P_2) with $P_1 \cap P_2 = \{x\}$, there is $P'_i \in \mathcal{P}_i$ such that $P'_i \cap P_{-i} = \{x^{(i)}\}$.

$(1,3)$ is unilaterally improvable for player 1 in the first game but it is not in the second game. Under the game in the left panel of Figure 9, an alternative that (i) strictly Pareto-dominates $(1,3)$ and (ii) gives player 2 the best feasible payoff is contained in player 2's imprecise proposal that would sustain the given payoff profile. In contrast, no such alternative can be found under the modified game in the proposal sustaining the given payoff profile (player 1's proposal U contains a strictly Pareto-dominating alternative with a payoff profile $(2,4)$, but player 2's payoff (which is 3) is higher than his worst Pareto-efficient payoff (which is 2)). $\square$

5.2.3 Tightness of Payoff Bounds

Section 5.2.2 shows that, under limited specifiability, whether an alternative x with $u_i(x) \leq v^{[i,m]}$ for some $i \in N$ is sustained as a SPE outcome depends on the structure of the environment and on the way in which the specification rule is imposed. Here, we demonstrate the tightness of payoff lower bounds. First, we show that for any feasible payoff set U, there exists a negotiation whose feasible payoff set is U in which the SPE lower bounds $\underline{u}$ are tight. Second, we show that there exists a negotiation whose feasible payoff set is U in which the SPE payoff set is the IR-Pareto-meet.

Corollary 5 (Existence of games with tight bounds)**.** *Let $U \subseteq \mathbb{R}^2$ and $d \in U$ be such that $\{v \in U \mid v \geq d\}$ is compact and, for each $i \in N$, there is $w \in \{v \in U \mid v \geq d\}$ with $w_i > d_i$. Let ρ be an asynchronous proposer rule.*

1. *Suppose $|U| \geq 2$. There is a two-player negotiation $\Gamma^L = \langle G^L, d, \rho, (\mathcal{P}_i^L)_{i\in N}, \varphi^{\text{con}}\rangle$ with the following properties: (i) $G^L = \langle N, X^L, (u_i^L)_{i\in N}\rangle$ satisfies $U = \{u^L(x) \in \mathbb{R}^2 \mid x \in X^L\}$; (ii) $(\mathcal{P}_i^L)_{i\in N}$ is limited; and (iii) $x \in X^L$ is a SPE outcome of Γ^L if and only if $u^L(x) \geq \underline{u}(U,d)$.*

2. *Suppose $3 \leq |U| < \infty$. There is a two-player negotiation $\Gamma^H = \langle G^H, d, \rho, (\mathcal{P}_i^H)_{i\in N}, \varphi^{\text{con}}\rangle$ with the following properties: (i) $G^H = \langle N, X^H, (u_i^H)_{i\in N}\rangle$ satisfies $U = \{u^H(x) \in \mathbb{R}^2 \mid x \in X^H\}$; (ii) $(\mathcal{P}_i^H)_{i\in N}$ is limited; and (iii) $x \in X^H$ is a SPE outcome of Γ^H if and only if it is in the IR-Pareto-meet.*

To show part 1 of Corollary 5, we first construct an environment associated with a normal-form game in which each player's action corresponds to her own payoff. Let

$X^L = \mathbb{R}^2$. The payoff functions are defined as follows:

$$u^L(v_1, v_2) = \begin{cases} (v_1, v_2) & \text{if } (v_1, v_2) \in U \\ d & \text{otherwise} \end{cases}.^{31}$$

Let the specification rule be such that each player i announces her own payoff: $\mathcal{P}_i^L = \{\{v_i\} \times \mathbb{R} \mid v_i \in \mathbb{R}\}$. No alternative $x \in X^L$ with $u^L(x) \geq \underline{u}(U, d)$ is unilaterally improvable, and thus Propositions 3 and 6 establish the result.

To obtain part 2 of Corollary 5, let $X^H = U$ and $u^H(x) = x$. Suppose first that $|\{v \in U \mid v \geq d\}| \geq 2$. Then, let the specification rule be $\mathcal{P}_i^H = \{\{v, (v^{[i,M]}, v^{[-i,m]})\} \mid v \in U \setminus \{(v^{[i,M]}, v^{[-i,m]})\}\}$ for each $i \in N$.[32] Theorem 3 establishes the result because, for any $v \in U$ such that $v_{-i} \geq v^{[-i,m]}$ and $v_i < v^{[i,m]}$, the alternative v is not unilaterally improvable for i. Suppose next that $|\{v \in U \mid v \geq d\}| = 1$, that is, $\{v \in U \mid v \geq d\} = \{v^*\}$. Then, for any specification rule, the negotiation has a unique SPE payoff profile v^*. Choosing distinct $v_1, v_2, v^* \in U$, we can define the following limited specification rule: For each $i \in N$, let $\mathcal{P}_i^H = \bigcup_{w \in \{v^1, v^2, v^*\}} \{\{w, v\} \mid v \in U \setminus \{w\}\}$. This establishes the desired result.

To conclude the discussion on part 2 of Corollary 5, let us go back to the example in Section 1, where U is the convex hull of $\{(0,0), (2,4), (4,2)\}$ and $d = (0,0)$. The arguments made here show that the set of SPE payoffs under unlimited specifiability is as depicted in the left panel of Figure 1, while that under limited specifiability is as depicted in the right panel of Figure 1.

6 Conclusion

This paper introduced a novel concept that we called limited specifiability in negotiations, and examined its effect on SPE outcomes. We showed that the effect of limitation on specifiability depends on the move structure. Although there is no difference in the SPE payoff sets under synchronicity, a difference exists when moves are asynchronous because asynchronicity helps players make a commitment. The extent to which such a difference arises can be explained by a tradeoff between commitment

[31] Note that $d \in U$ implies that the feasible payoff set of the environment G^L coincides with U.

[32] We assume $|U| \geq 3$ to ensure that the specification rule satisfies the requirement that for each $x \in X^H$, there is $(P_1, P_2) \in \mathcal{P}_1^H \times \mathcal{P}_2^H$ such that $P_1 \cap P_2 = \{x\}$.

and punishment: Limited specifiability implies lower commitment power, which provides greater scope for punishment. In one extreme, the power of commitment is so strong when the negotiation has a common-interest alternative that there is a unique SPE outcome under arbitrary specifiability conditions. In the other extreme, the set of SPE payoff profiles can be quite different, and we provided lower bounds of SPE payoffs which we showed to be tight in a certain sense, as well as a condition that we called unilateral improvability whose absence guarantees that a given alternative is a SPE outcome under limited specifiability.

In order to have all these comparative statics make sense, we defined a negotiation protocol as a collection of three rules: proposer, specification and termination rules. The generality of the model enables us to nest many possible negotiation protocols as special cases of our model, which we believe would facilitate meaningful comparison among different models.

The paper suggests a number of avenues for future research. First, our paper is merely a first step of the study of limited specifiability. One could depart from our perfect-information assumption to allow for imperfect or incomplete information. If the limitation on specifiability originated from imperfect information in such a model, then it would be of interest to see how potential resolution of the limitation interacts with incentives of making proposals. Also, we mainly focused on the situations where either every player's specification rule is limited or it is unlimited. The future research may investigate more complicated cases that lie in between those two cases, which would widen the applicability of our results. Finally, one could imagine endogenizing specifiability, by perhaps introducing costs associated to the degree of specifiability as in the "writing cost" of contracts in Battigalli and Maggi (2002).

Second, studying the effect of limited specifiability in different contexts may prove fruitful. For example, it may be studied in such contexts as cheap talk, delegation, and contracts. We hope that the concept of limited specifiability will shed new lights on those applications.

Finally, our framework of negotiation protocols may facilitate unification of the literature. Online Appendix D enlists possible variations of proposer, specification and termination rules, demonstrating wide applicability of our framework to various negotiation models. This suggests that one could use the idea of dividing negotiation protocols into three rules to formally compare existing models in the literature with each other. Such an exercise would lead to an understanding of the effect of negotiation protocols on the outcomes in a unified framework.

A Appendix

A.1 Proofs for Section 4

Proof of Theorem 1

We use Theorem 2 to prove Theorem 1. The proof of Theorem 2 does not depend on Theorem 1.

"If" part: Since Theorem 2 in Section 4.2 implies that $x^* \in X^M$ is a SPE outcome, we show that x^* is a unique SPE outcome. Consider the shortest terminal history h under which each player i announces (Yes, P_i), where $\bigcap_{i \in N} P_i = \{x^*\}$ (such a history h exists by assumption). At the history $h^{t(h)-1}$, player $i_1 = \rho(h^{t(h)-1})$ can guarantee herself a payoff of $u_{i_1}(x^*)$, her maximum possible SPE payoff (note that any $y \in X$ with $u_i(y) > u_i(x^*)$, if it exists, is not individually rational for some other player). Since $u_{i_1}(x^*) > d_{i_1}$ by assumption, x^* is the unique outcome in the subgame starting at $h^{t(h)-1}$ in any SPE. Next, at $h^{t(h)-2}$, player $i_2 = \rho(h^{t(h)-2})$ can guarantee herself a payoff of $u_{i_2}(x^*)$, her maximum possible SPE payoff. Since $u_{i_2}(x^*) > d_{i_2}$ by assumption, x^* is the unique outcome in the subgame starting at $h^{t(h)-2}$ in any SPE. Solving backwards in this way, for each $j \in \{1, \dots, t(h)\}$, at any $h^{t(h)-j}$, player $i_j = \rho(h^{t(h)-j})$ can guarantee herself a payoff of $u_{i_j}(x^*)$, her maximum possible SPE payoff. Hence, x^* is the unique SPE outcome in the subgame starting at the initial history h^0.

"Only if" part: Suppose that x^* is a unique SPE outcome. First, note that $X^M = \{x^*\}$ because Theorem 2 implies that if $X^M (\neq \emptyset)$ is not a singleton set then the negotiation has multiple SPE outcomes. Next, if there is $x \in X \setminus \{x^*\}$ such that $u_i(x) = \max_{v \in \mathrm{IR}(U,d)} v_i$ for some $i \in N$, then $x \in X^M$, a contradiction. Hence, for each $i \in N$, x^* is the unique alternative that generates the maximum individually-rational payoff. Thus, $u(x^*) > v$ for all $v \in \mathrm{IR}(U,d) \setminus \{u(x^*)\}$, as desired. $\square$

Proof of Theorem 2

Let $x^{(0)} := x \in X^M$. For each $j \in N$, let $x^{(j)} \in X$ be player j's worst individually-rational and Pareto-efficient alternative. For each $j \in \{0\} \cup N$, fix $(P_i^{(j)})_{i \in N}$ with $\{x^{(j)}\} = \bigcap_{i \in N} P_i^{(j)}$ (such a profile exists by assumption). We denote $P_i = P_i^{(0)}$ for each $i \in N$. Note that it is possible that $x = x^{(j)}$ for some $j \in N$.

Let h^* be the shortest terminal history under which every player i announces (Yes, P_i) at any subhistory at which she speaks (such a history h^* exists by assumption). Denote by $Q_0 := \{h \in H \setminus Z \mid h \sqsubseteq h^*\}$ the set of non-terminal subhistories of h^*. Let Q_j be the set of $h \in (H \setminus Z) \setminus Q_0$ such that $j = \min I^\tau(h)$ with $\tau = \min\{t' \in \mathbb{N} \mid (R_i^{t'}(h), P_i^{t'}(h)) \neq (\text{Yes}, P_i) \text{ for some } i \in N\}$. That is, Q_j is the set of non-terminal histories under which player $j \in N$ deviates from announcing (Yes, P_j) first.

We define the following strategy profile s^*. For each $i \in N$ and $h \in H_i$, let

$$s_i^*(h) := \begin{cases} (\text{Yes}, P_i) & \text{if } h \in H_i \cap Q_0 \\ s_i^{(j)}(h) & \text{if } h \in H_i \cap Q_j \text{ for some } j \in N \end{cases},$$

where $s_i^{(j)}(h)$ is defined as

$$s_i^{(j)}(h) := \begin{cases} (\text{Yes}, \tilde{P}_i(h)) & \text{if } h \in H_i \cap Q_{j,1} \\ (\text{Yes}, P_i^{(j)}) & \text{if } h \in H_i \cap Q_{j,2} \\ (\text{No}, P_i^{(j)}) & \text{if } h \in H_i \cap (Q_j \setminus (Q_{j,1} \cup Q_{j,2})) \end{cases}$$

with the following properties. Intuitively, the set $Q_{j,1}$ contains any non-terminal history $h \in Q_j$ such that there are a group of players who can collectively terminate the negotiation with an outcome $\tilde{x}(h) \in X \setminus \{x^{(j)}\}$ at h, which gives a strictly greater payoff to each of them. Note that $Q_{j,1}$ could be empty (e.g., it is empty in negotiations with a common-interest alternative). Formally, $h \in Q_{j,1}$ if and only if (i) $h \in Q_j$ and (ii) there exists a sequence $\tilde{h}(k^*) := ((N_\ell, ((\text{Yes}, \tilde{P}_m(h)))_{m \in N_\ell}))_{\ell=1}^{k^*}$ with $k^* \in \mathbb{N}$ such that, denoting $\tilde{h}(k) := ((N_\ell, ((\text{Yes}, \tilde{P}_m(h)))_{m \in N_\ell}))_{\ell=1}^{k}$ for each $k \in \{1, \ldots, k^*\}$, $N_1 = \rho(h)$, $N_{\ell+1} = \rho(h, \tilde{h}(\ell))$ for all $\ell \in \{1, \ldots, k^* - 1\}$, $\varphi^{\text{con}}(h, \tilde{h}(k^*)) = \tilde{x}(h) \in X \setminus \{x^{(j)}\}$, and that $u_\ell(\tilde{x}(h)) > u_\ell(x^{(j)})$ for all $\ell \in \bigcup_{k=1}^{k^*} N_k$. Note that $\tilde{x}(h)$ is well defined as a single alternative. This is because of the following: Since $x^{(j)}$ is Pareto efficient, $\bigcup_{k=1}^{k^*} N_k \neq N$ for any choice of $\tilde{h}(k^*)$. Also, any player in $N \setminus (\bigcup_{k=1}^{k^*} N_k)$ has to be ok with $\tilde{x}(h)$ at h. Hence, for any choice of $\tilde{h}(k^*)$, $\tilde{x}(h)$ is uniquely determined.

The set $Q_{j,2}$ has any non-terminal history $h \in Q_j \setminus Q_{j,1}$ with the following properties: Either (i) every player ℓ who spoke at the end of h announced $(\text{No}, P_\ell^{(j)})$, i.e., $h = \left(h^{t(h)-1}, (I^{t(h)}(h), ((\text{No}, P_\ell^{(j)}))_{\ell \in I^{t(h)}(h)})\right)$; or (ii) every player ℓ has been announcing $(\text{Yes}, P_\ell^{(j)})$ since the most recent announcement of "No" at time $t^{\text{No}}(h) \leq t(h)-1$, i.e., $h = \left(h^{t^*}, ((I^k(h), ((\text{Yes}, P_\ell^{(j)}))_{\ell \in I^k(h)}))_{k=t^*+1}^{t(h)}\right)$ with $t(h) - 1 \geq t^* := t^{\text{No}}(h)$.

We now show that each player $i \in N$ following s_i^* is a best response to s_{-i}^* in any subgame, which implies that the SPE s^* induces the history h^* and the outcome x. To show this, first consider a subgame starting at $h \in H_i \cap (Q_j \setminus Q_{j,1})$ for some $j \in N$. The continuation strategy profile $s^*|_h := (s_i^*|_h)_{i \in N}$ induces $x^{(j)}$, where we define $s_i^*|_h$ to be a mapping from $\{h' \in H_i \mid h' \sqsupseteq h\}$ to $\{\text{Yes}, \text{No}\} \times \mathcal{P}_i$. If player i announces "No" at $h' \in H_i$ with $h' \sqsupseteq h$, then it is impossible for any $x' \in X$ with $u_i(x') > u_i(x^{(j)})$ to be an outcome under s_{-i}^*. Suppose to the contrary that some alternative $x' = \varphi^{\text{con}}(h'')$ with $u_i(x') > u_i(x^{(j)})$ and $h'' \sqsupseteq h'$ is an outcome under s_{-i}^*. Since player i cannot terminate the negotiation with outcome x' at h' by saying "No," $h'' \neq h'$. Let $h'' \sqsupset h'$. Since $u(x^{(j)}) \in \text{PE}(U)$, some player k with $u_k(x') \leq u_k(x^{(j)})$ is ok with x' at h''. This is impossible because such player k, who follows s_k^*, must not be okay with x' at h''.

For any strategy of player i such that she announces "Yes" after each history at which it is her turn to move in the the subgame starting at h, either every player k keeps announcing $(\text{Yes}, P_k^{(j)})$ to agree upon $x^{(j)}$ or some player announces "No" in the subgame starting at h. In the subgame starting at history h at which some player announces "No," no alternative x' with $u_i(x') > u_i(x^{(j)})$ can be an outcome under $s_{-i}^*|_h$.

Second, consider $h \in H_i \cap Q_{j,1}$. If player i follows $s_i^*|_h$, then $s^*|_h$ induces $\tilde{x}(h)$. Suppose, on the other hand, player i deviates at a history $h' \in H_i$ with $h' \sqsupseteq h$. If she announces "No" at h', then any alternative x' with $u_i(x') > u_i(x^{(j)})$ (which, of course, includes $\tilde{x}(h)$) cannot be an outcome. If her announcement is "Yes," then either every player k keeps announcing $(\text{Yes}, \tilde{P}_k)$ to agree upon $\tilde{x}(h)$ or some player announces "No" at some point. In the latter case, any induced history $h \in H_i$ no longer belongs to $Q_{j,1}$.

Third, consider the subgame starting at $h \in H_i \cap Q_0$. Player i gets a payoff of $u_i(x)$ by following $s_i^*|_h$, as $s^*|_h$ induces h^*. Her deviation induces a non-terminal history $h' \in Q_i \setminus Q_{i,1}$, as any other player k follows s_k^*. Thus, player i's maximum possible payoff in the subgame starting at h' is $u_i(x^{(i)})$, which is no greater than $u_i(x)$ by the definition of $x^{(i)}$. $\square$

A.2 Proofs for Section 5

Proof of Proposition 5

As discussed in the main text, the proof consists of two steps.

First Step: Fix $i \in N$ satisfying the condition in the statement of Proposition 5. Let $Y_i := \{x \in X \mid u_i(x) < v^{[i,m]} \text{ and } u_{-i}(x) \geq v^{[-i,m]}\}$. Suppose to the contrary that there is $y \in X^{\text{SPE}} \cap Y_i$. Let s^* be a SPE that induces y, and let h^* be the finite terminal history induced by s^*. Let h be a subhistory of h^* such that $i \in \rho(h)$ and that if $h' \sqsubset h^*$ satisfies $i \in \rho(h')$ then $h' \sqsubseteq h$. Letting $(R_{-i}, P_{-i}) = (R_{-i}^{t(h)}(h), P_{-i}^{t(h)}(h))$ and $(R_i, P_i) = s_i^*(h)$, we have $P_1 \cap P_2 = \{y\}$. Since y is unilaterally improvable for player i, there exists $(y', P_i') \in X \times \mathcal{P}_i$ such that $u(y') > u(y)$ and $\{y'\} = P_i' \cap P_{-i}$. Choose one such (y', P_i'), and consider i's deviation to announce (Yes, P_i') at h. On the one hand, in the subgame starting at $(h, (\text{Yes}, P_i))$, one strategy player $-i$ can take is to keep announcing (Yes, P_{-i}). The consensual termination rule terminates the negotiation under such a strategy profile at $(h, (\text{Yes}, P_i'), (\text{Yes}, P_{-i}))$ with the outcome y'. Thus, player $-i$'s payoff conditional on the history $(h, (\text{Yes}, P_i'))$ under $s^*|_{(h,(\text{Yes},P_i'))}$ is at least $u_{-i}(y')$. On the other hand, since s^* is a SPE, player i's deviation to announce (Yes, P_i') cannot lead to a payoff strictly higher than $u_i(y)$. These facts imply that $s^*|_{(h,(\text{Yes},P_i'))}$ leads to an outcome in $\{y' \in X \mid u_i(y') \leq u_i(y) \text{ and } u_{-i}(y') > u_{-i}(y)\}$.

Thus, there is an infinite sequence $(y^k)_{k\in\mathbb{N}}$ such that $y^{k+1} \in \{y' \in X \mid u_i(y') \leq u_i(y^k) < v^{[i,m]} \text{ and } u_{-i}(y') > u_{-i}(y^k) \geq v^{[-i,m]}\}$ for each $k \in \mathbb{N}$. By construction, $y^{k+1} \neq y^\ell$ for all $\ell \leq k$. This contradicts the assumption that X is finite.

Second Step: Pick $x \in X$ with $\underline{u} \leq u(x) < (v^{[1,m]}, v^{[2,m]})$. Suppose to the contrary that x is sustained by a SPE s^*. Let h^* be the finite terminal history induced by s^*. Let h be a subhistory of h^* such that $i \in \rho(h)$ and that if $h' \sqsubset h^*$ satisfies $i \in \rho(h')$ then $h' \sqsubseteq h$. Let $(R_i, P_i) = s_i^*(h)$ and $(R_{-i}, P_{-i}) = (R_{-i}^{t(h)}(h), P_{-i}^{t(h)}(h))$. We have $P_1 \cap P_2 = \{x\}$. Let $(P_1^{(-i)}, P_2^{(-i)})$ be such that $P_1^{(-i)} \cap P_2^{(-i)} = \{x^{(-i)}\}$, where $x^{(-i)} \in X$ satisfies $u(x^{(-i)}) = (v^{[i,M]}, v^{[-i,m]})$.

At the history $h' = (h, (\text{No}, P_i^{(-i)}))$, if player $-i$ announces $(\text{Yes}, P_{-i}^{(-i)})$, then player i can receive the best SPE payoff $v^{[i,M]}$ by announcing $(\text{Yes}, P_i^{(-i)})$. Thus, player $-i$ can secure herself a payoff of $v^{[-i,m]}$ at h'. By the equilibrium condition, letting $y \in X$ be the outcome induced by $s^*|_{h'}$, we have $u_i(y) \leq u_i(x) < v^{[i,m]}$ and $u_{-i}(y) \geq v^{[-i,m]} (> u_{-i}(x))$. Now, one must be able to construct an infinite sequence defined in the first step, which leads to a contradiction. $\square$

Proof of Proposition 6

Part 1: This part follows from Theorem 2.

Part 2: Fix $(i,x) \in N \times X$ such that $d_i \leq u_i(x) < v^{[i,m]}$ and $u_{-i}(x) \geq v^{[-i,m]}$. Fix (P_1, P_2) with $\{x\} = P_1 \cap P_2$. Let $(y^i, y^{-i}) = (x, x^{(-i)})$. That is, $u_{-i}(y^{-i}) = v^{[-i,m]}$ and $u_i(y^{-i}) = v^{[i,M]}$. Choose $(P_1^{(-i)}, P_2^{(-i)})$ such that $P_1^{(-i)} \cap P_2^{(-i)} = \{y^{-i}\}$. Note that the profiles of proposals (P_1, P_2) and (P_1^{-i}, P_2^{-i}) exist by assumption, and that for each $j \in N$, y^j is not unilaterally improvable for j.

Let $Q_i \subseteq H \setminus Z$ be the set of non-terminal histories with the following two properties: First, $h^0 \in Q_i$. Second, any $h \in (H \setminus Z) \setminus \{h^0\}$ is in Q_i if and only if, for any $h' \in H \setminus Z$ with $h' \sqsubset h$ and $-i \in \rho(h')$, the history $h^{t(h')+1}$ satisfies the following:

$$h^{t(h')+1} = \begin{cases} (h', (\text{Yes}, P_{-i})) & \text{if } h' = h^0 \\ (h', (\text{Yes}, P_{-i})) & \text{if } h' \neq h^0 \text{ and } (R_i^{t(h')}(h'), P_i^{t(h')}(h')) = (\text{Yes}, P_i) \\ (h', (\text{No}, P_{-i})) & \text{if } h' \neq h^0 \text{ and } (R_i^{t(h')}(h'), P_i^{t(h')}(h')) \neq (\text{Yes}, P_i) \end{cases}.$$

We let $Q_{-i} := (H \setminus Z) \setminus Q_i$.

Consider the following strategy profile s^*. For player $i \in N$ and $h \in H_i$, let

$$s_i^*(h) := \begin{cases} (\text{Yes}, P_i) & \text{if } h \in H_i \cap Q_i \\ (\text{Yes}, P_i^{(-i)}) & \text{if } h \in H_i \cap Q_{-i} \text{ and } \varphi^{\text{con}}(h, (\text{Yes}, P_i^{(-i)})) = y^{-i} \\ (\text{No}, P_i^{(-i)}) & \text{if } h \in H_i \cap Q_{-i} \text{ and } \varphi^{\text{con}}(h, (\text{Yes}, P_i^{(-i)})) \neq y^{-i} \end{cases}.$$

Note that, at $h \in H_i \cap Q_i$, since i's specifiability is limited and x is not unilaterally improvable for i, there is no $(y, P_i') \in X \times \mathcal{P}_i$ such that $\varphi^{\text{con}}(h, (\text{Yes}, P_i'), (\text{Yes}, P_{-i})) = y$ and $u(y) > u(x)$.

For player $-i$, let $h \in H_{-i}$. If $h = h^0 \in H_{-i}$, then let $s_{-i}^*(h) = (\text{Yes}, P_{-i})$. If

$h \in Q_i \setminus \{h^0\}$, then let

$$s^*_{-i}(h) := \begin{cases} (\text{Yes}, P_{-i}) & \text{if } (R_i^{t(h)}(h), P_i^{t(h)}(h)) = (\text{Yes}, P_i) \\ (\text{Yes}, P_{-i}) & \text{if } h = (h^{t(h)-2}, (\cdot, P_{-i}), (\text{Yes}, \tilde{P}_i)), t(h) \geq 2, \\ & \tilde{P}_i \cap P_{-i} = \{y\}, \text{ and } u_{-i}(y) > u_{-i}(x) \\ (\text{No}, P_{-i}) & \text{otherwise} \end{cases}.$$

Next, we define $s^*_{-i}(h)$ for $h \in Q_{-i}$. First, define $Q^*_{-i} \subseteq Q_{-i}$ by $h \in Q^*_{-i}$ if and only if there is $\tilde{P}_{-i}$ such that $\varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_{-i})) = \tilde{x}(h)$ and $u_{-i}(\tilde{x}(h)) > u_{-i}(y^{-i})$. If $h \in Q^*_{-i}$, then $s^*_{-i}(h) = (\text{Yes}, \tilde{P}_{-i})$. Second, $h \in Q_{-i} \setminus Q^*_{-i}$, then let

$$s^*_{-i}(h) := \begin{cases} (\text{Yes}, P_{-i}^{(-i)}) & \text{if } P_i^{t(h)}(h) = P_i^{(-i)} \\ (\text{No}, P_{-i}^{(-i)}) & \text{otherwise} \end{cases}.$$

We show that s^* is a SPE, i.e., for each $j \in N$, following s^*_j is a best response to s^*_{-j} in any subgame. The strategy profile s^* induces the history $((\text{Yes}, P_1), (\text{Yes}, P_2), (\text{Yes}, P_1))$ or $((\text{Yes}, P_2), (\text{Yes}, P_1), (\text{Yes}, P_2))$, and the outcome x in both cases. We show i's best-response condition first, and then $-i$'s best-response condition.

Player i's best-response condition: Take $h \in H_i$. If $h \in Q_{-i}$, then $s^*|_h$ induces player i's best SPE outcome y^{-i}. Suppose that $h \in Q_i$. The continuation strategy profile $s^*|_h$ induces x. Notice that any non-terminal history in H_i induced by a continuation strategy profile $(s_i, s^*_{-i})|_h$ is in $H_i \cap Q_i$. If player i proposes $(\text{Yes}, \tilde{P}_i)$ so that player $-i$ can terminate the negotiation with $y \in X$ such that $\{y\} = \tilde{P}_i \cap P_{-i}$ and $u_{-i}(y) > u_{-i}(x)$, then player i receives a payoff $u_i(y) \leq u_i(x)$. Otherwise, a possible outcome is either x or the disagreement outcome. Hence, s^*_i is a best response to s^*_{-i} in the subgame starting at $h \in H_i$.

Player $-i$'s best-response condition: Take $h \in H_{-i}$. First, suppose that $h \in Q_{-i} \setminus Q^*_{-i}$. The continuation strategy profile $s^*|_h$ induces y^{-i}. Any continuation strategy profile $(s_{-i}, s^*_i)|_h$ induces either y^{-i} or the disagreement outcome.

Second, suppose that $h \in Q^*_{-i}$. The continuation strategy profile $s^*|_h$ induces $\tilde{x}(h)$, and player $-i$ gets a payoff of $u_{-i}(\tilde{x}(h)) > u_{-i}(y^{-i})$. If player $-i$ does not terminate the negotiation with $\tilde{x}(h)$ at h, then the outcome following the continuation play is either y^{-i} or the disagreement outcome.

Finally, suppose that $h \in Q_i$. Assume that $h = h^0 \in H_{-i}$. The continuation strategy profile $s^*|_h$ induces the outcome x. If $-i$ announces $(R'_{-i}, P'_{-i}) \neq (\text{Yes}, P_{-i})$ at h^0, then $s_i^*(R'_{-i}, P'_{-i}) = (\text{No}, P_i^{(-i)})$. At the history $((R'_{-i}, P'_{-i}), (\text{No}, P_i^{(-i)})) \in H_{-i} \cap Q_{-i}$, player $-i$ can obtain at most $u_{-i}(y^{-i})$.

Now, assume that $h \neq h^0$. We consider the following three cases: (i) $(R_i^{t(h)}(h), P_i^{t(h)}(h)) = (\text{Yes}, P_i)$; (ii) $h = (h^{t(h)-2}, (\cdot, P_{-i}), (\text{Yes}, \tilde{P}_i))$ with $t(h) \geq 2$, $\tilde{P}_i \cap P_{-i} = \{y\}$, and $u_{-i}(y) > u_{-i}(x)$; and (iii) otherwise. In case (i), the continuation strategy profile $s^*|_h$ induces x. If player $-i$ uses s_{-i} and announces $s_{-i}(h) \neq (\text{Yes}, P_{-i})$ at h, then player i announces $s_i^*(h, s_{-i}(h)) = (\cdot, P_i^{(-i)})$. If this announcement terminates the negotiation then the outcome is y^{-i}. If not, then in the subgame starting at the resulting history $h' = (h, s_{-i}(h), s_i^*(h, s_{-i}(h))) \in H_{-i} \cap (Q_{-i} \setminus Q_{-i}^*)$, player $-i$ can obtain at most $u_{-i}(y^{-i})$.

In case (ii), player $-i$ can obtain a payoff of $u_{-i}(y)$ by terminating the negotiation. In fact, following s_{-i}^*, player $-i$ can indeed terminate the negotiation. For the case in which she does not terminate the negotiation, consider $h' = (h, (R_{-i}^{t(h)+1}(h'), P_{-i}^{t(h)+1}(h')))$. If $(R_{-i}^{t(h)+1}(h'), P_{-i}^{t(h)+1}(h')) = (\text{No}, P_{-i})$ then, at $h'' = (h', s_i^*(h'))$, player $-i$ can obtain a payoff of at most $u_{-i}(x)$. If $(R_{-i}^{t(h)+1}(h'), P_{-i}^{t(h)+1}(h')) \neq (\text{No}, P_{-i})$, then in the subgame starting at $h'' = (h', s_i^*(h')) = (h', (\text{No}, P_i^{(-i)}))$, player $-i$ can obtain a payoff of at most $u_{-i}(y^{-i})$.

In case (iii), the continuation strategy profile $s^*|_h$ induces x. Suppose that player $-i$ does not announce (No, P_{-i}) at h. If she terminates the negotiation, then she can get a payoff of at most $u_{-i}(x)$. If she uses a strategy s_{-i} so as not to terminate the negotiation, then the resulting history $h' = (h, s_{-i}(h))$ is in $H_i \cap Q_{-i}$. If player i terminates the negotiation at h', player $-i$ gets $u_{-i}(y^{-i})$. If not, the resulting history $h'' = (h', s_i^*(h'))$ is in $H_{-i} \cap (Q_{-i} \setminus Q_{-i}^*)$. Player $-i$ can obtain a payoff of at most $u_{-i}(y^{-i})$ in the subgame starting at h''. Hence, s_{-i}^* is a best response to s_i^* in the subgame starting at $h \in H_{-i}$.

Part 3: Fix $x \in X$ with $(v^{[1,m]}, v^{[2,m]}) > u(x)(\geq \underline{u})$ satisfying the conditions of the statement. Fix (P_1, P_2) such that $\{x\} = P_1 \cap P_2$. For each $i \in \{1, 2\}$, fix (P_1^i, P_2^i) such that $\{y^i\} = P_1^i \cap P_2^i$. Also, for each $i \in N$, let $(P_1^{(i)}, P_2^{(i)})$ be such that $\{x^{(i)}\} = P_1^{(i)} \cap P_2^{(i)}$. Note that each $x^{(i)} \in X$ satisfies $u_i(x^{(i)}) = v^{[i,m]}$ and $u_{-i}(x^{(i)}) = v^{[-i,M]}$. Our proof of this part consists of the following three steps. In the first step, in order to define a strategy profile s^* that induces the alternative x, we partition the set of non-terminal histories $H \setminus Z$. In the next step, using the partition, we define the strategy profile s^*. In the last step, we show that each s_i^* is a best response to s_{-i}^*.

Partitioning $H \setminus Z$: Fix $j = \rho(h^0)$. We partition $H \setminus Z$ into Q_1, Q_2, and $Q_0 := \{h^0, ((\text{Yes}, P_j)), ((\text{Yes}, P_j), (\text{Yes}, P_{-j}))\}$. Here, we define each Q_i to be the set of non-terminal histories under which player i deviates from announcing (Yes, P_i) first. Formally, $h \in Q_j$ if and only if $h \in H \setminus Z$ satisfies either (i) $h^1 \neq ((\text{Yes}, P_j))$ or (ii) $h^2 = ((\text{Yes}, P_j), (\text{Yes}, P_{-j}))$ and $(R_j^3(h), P_j^3(h)) \neq (\text{Yes}, P_j)$. Likewise, $h \in Q_{-j}$ if and only if $h \in H \setminus Z$ satisfies $h^1 = ((\text{Yes}, P_j))$ and $h^2 \neq ((\text{Yes}, P_j), (\text{Yes}, P_{-j}))$.

For each $i \in N$, we further partition Q_i into Q_i^{on} and $Q_i^{\text{off}} := Q_i \setminus Q_i^{\text{on}}$. We define Q_i^{on} so that $h \in Q_i^{\text{on}}$ if and only if $h \in Q_i$ satisfies either of the following two properties: First, $h^{t(h)-1} \in Q_0$. Second, for any proper subhistory $h' \sqsubset h$ with $h' \in Q_i \cap H_{-i}$,

$$h^{t(h')+1} = \begin{cases} \big(h', (\text{Yes}, P_{-i}^i)\big) & \text{if } h^{t(h')-1} \in Q_i \text{ and } P_i^{t(h')}(h) = P_i^i \\ \big(h', (\text{No}, P_{-i}^i)\big) & \text{if } h^{t(h')-1} \notin Q_i \text{ or } P_i^{t(h')}(h) \neq P_i^i \end{cases}.$$

Defining s^*: We define the following strategy profile s^*. For any $h \in H_i \cap Q_0$, let $s_i^*(h) := (\text{Yes}, P_i)$. For any $h \in H_i \cap Q_i^{\text{off}}$, we let

$$s_i^*(h) := \begin{cases} (\text{Yes}, P_i^{(-i)}) & \text{if } \varphi^{\text{con}}(h, (\text{Yes}, P_i^{(-i)})) = x^{(-i)} \\ (\text{No}, P_i^{(-i)}) & \text{otherwise} \end{cases}.$$

For any $h \in H_i \cap Q_{-i}^{\text{off}}$, we let

$$s_i^*(h) := \begin{cases} (\text{Yes}, \tilde{P}_i) & \text{if there is } \tilde{P}_i \text{ such that } \varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_i)) = \tilde{x}(h) \\ & \text{and } u_i(\tilde{x}(h)) > v^{[i,m]} = u_i(x^{(i)}) \\ (\text{Yes}, P_i^{(i)}) & \text{if } (R_{-i}^{t(h)}(h), P_{-i}^{t(h)}(h)) = (\text{No}, P_{-i}^{(i)}) \text{ or } \varphi^{\text{con}}(h, (\text{Yes}, P_i^{(i)})) = x^{(i)} \\ (\text{No}, P_i^{(i)}) & \text{otherwise} \end{cases}.$$

For any $h \in H_i \cap Q_i^{\text{on}}$, let $s_i^*(h) := (\text{Yes}, P_i^i)$. Let $h \in H_i \cap Q_{-i}^{\text{on}}$. If there is $\tilde{P}_i$ such that $\varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_i)) = \tilde{x}(h)$ and $u_i(\tilde{x}(h)) > u_i(y^{-i})$, then we let $s_i^*(h) := (\text{Yes}, \tilde{P}_i)$. Otherwise, we let

$$s_i^*(h) := \begin{cases} (\text{Yes}, P_i^{-i}) & \text{if } P_{-i}^{t(h)}(h) = P_{-i}^{-i} \\ (\text{No}, P_i^{-i}) & \text{if } P_{-i}^{t(h)}(h) \neq P_{-i}^{-i} \end{cases}.$$

Showing that s_i^* is a best response to s_{-i}^*: We show that player $i \in N$ following s_i^* is a best response to s_{-i}^* in any subgame starting at $h \in H_i$. First, let $h \in Q_i^{\text{off}}$. The continuation strategy profile $s^*|_h$ induces i's best SPE outcome $x^{(-i)}$.

Second, let $h \in Q_{-i}^{\text{off}}$. Suppose that there is no $\tilde{P}_i$ such that $\varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_i)) = \tilde{x}(h)$ and $u_i(\tilde{x}(h)) > u_i(x^{(i)})$. The continuation strategy profile $s^*|_h$ induces the outcome $x^{(i)}$. Any continuation strategy profile $(s_i, s_{-i}^*)|_h$ induces either $x^{(i)}$ or the disagreement outcome. Next, suppose that there is $\tilde{P}_i$ such that $\varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_i)) = \tilde{x}(h)$ and $u_i(\tilde{x}(h)) > u_i(x^{(i)})$. The continuation strategy profile $s^*|_h$ induces $\tilde{x}(h)$. If player i does not terminate the negotiation at h with $\tilde{x}(h)$, then the outcome following the continuation play is either $x^{(i)}$ or the disagreement outcome.

Third, let $h \in Q_i^{\text{on}}$. We start with showing that $s^*|_h$ induces y^i. Since $h \in Q_i^{\text{on}} \cap H_i$, we have

$$h = \left(h^{t(h)-2}, \left(R_i^{t(h)-1}(h), P_i^{t(h)-1}(h)\right), \left(R_{-i}^{t(h)}(h), P_{-i}^i\right)\right).$$

If $R_{-i}^{t(h)}(h) = \text{Yes}$, then the negotiation terminates with y^i at $(h, (\text{Yes}, P_i^i))$. If $R_{-i}^{t(h)}(h) = \text{No}$, then the negotiation terminates with y^i at $(h, (\text{Yes}, P_i^i), (\text{Yes}, P_{-i}^i))$.

We now show that s_i^* is a best response to s_{-i}^* in the subgame starting at $h \in H_i$. Assume that $R_{-i}^{t(h)}(h) = \text{No}$. If player i announces (Yes, P_i) such that $P_i \cap P_{-i}^i = \{y\}$ and $u_{-i}(y) > u_{-i}(y^i)$ for some y, then player $-i$ terminates the negotiation with y by announcing (Yes, P_{-i}^i). However, since y^i is not unilaterally improvable for i, we have $u_i(y) \leq u_i(y^i)$. If player i announces $(\hat{R}_i, \hat{P}_i) \neq (\text{Yes}, P_i)$, then player $-i$'s following s_{-i}^* induces $(h, (\hat{R}_i, \hat{P}_i), (\cdot, P_{-i}^i)) \in H_i \cap Q_i^{\text{on}}$.

Assume that $R_{-i}^{t(h)}(h) = \text{Yes}$. Thus, $P_i^{t(h)-1}(h) = P_i^i$ so that player $-i$ is ok with y^i at $h^{t(h)-1}$. If player i terminates the negotiation, then the outcome must be y^i. In fact, $s_i^*(h) = (\text{Yes}, P_{-i}^i)$ terminates the negotiation. If she uses s_i so as not to terminate the negotiation at h, then $(h, s_i(h), (\cdot, P_{-i}^i)) \in H_i \cap Q_i^{\text{on}}$.

Fourth, let $h \in Q_{-i}^{\text{on}}$. Suppose that there is no $\tilde{P}_i$ such that $\varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_i)) = \tilde{x}(h)$ and $u_i(\tilde{x}(h)) > u_i(y^{-i})$. The continuation strategy profile $s^*|_h$ induces y^{-i}. If player i uses s_i and induces $(h, s_i(h)) \in H_{-i} \cap Q_{-i}^{\text{off}}$, then the outcome following the continuation play is either $x^{(i)}$ or the disagreement outcome. If player i's announcement induces $(h, s_i(h)) \in H_{-i} \cap Q_{-i}^{\text{on}}$, then player $-i$ is ok with y^{-i} at $(h, s_i(h), (\text{Yes}, P_{-i}^i))$. If this history is non-terminal (i.e, in $H_i \cap Q_{-i}^{\text{on}}$), then there is no $\tilde{P}_i$ such that $\varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_i)) = \tilde{x}(h)$ and $u_i(\tilde{x}(h)) > u_i(y^{-i})$.

Suppose that there is $\tilde{P}_i$ such that $\varphi^{\text{con}}(h, (\text{Yes}, \tilde{P}_i)) = \tilde{x}(h)$ and $u_i(\tilde{x}(h)) > u_i(y^{-i})$. Player i can obtain a payoff of $u_i(\tilde{x}(h)) > u_i(y^{-i})$ by following s_i^*. Suppose that player i does not terminate the negotiation with $\tilde{x}(h)$ at h, inducing

$$h' = \left(h^{t(h)-1}, (\text{Yes}, P_{-i}^{t(h)}), (R_i^{t(h)+1}(h'), P_i^{t(h)+1}(h'))\right).$$

If $h' \in H_{-i} \cap Q_{-i}^{\text{off}}$, then player i can obtain a payoff of at most $u_i(x^{(i)})$ in the subgame starting at h'. Suppose that $h' \in H_{-i} \cap Q_{-i}^{\text{on}}$. Then, in the subgame starting at $\left(h, (\cdot, P_i^{-i}), (\text{Yes}, P_{-i}^{-i})\right) \in H_i \cap Q_{-i}^{\text{on}}$, player i can obtain a payoff of at most $u_i(y^i)$.

Fifth, consider $h \in Q_0$. The continuation strategy profile $s^*|_h$ induces x. If player i follows s_i and announces $s_i(h) \neq (\text{Yes}, P_i)$ at h, then, since i's specifiability is limited and x is not unilaterally improvable for i, player i's maximum payoff in the continuation play against s_{-i}^* is $u_i(y^i) \leq u_i(x)$.

Overall, s_i^* is a best response to s_{-i}^* in the subgame starting at any history. Since this holds for each $i \in N$, the strategy profile s^* is a SPE. By construction, it induces the history $((\text{Yes}, P_j), (\text{Yes}, P_{-j}), (\text{Yes}, P_j))$ with $j = \rho(h^0)$ and the outcome x. $\square$

References

Aumann, R. J., M. B. Maschler, and R. E. Stearns (1968): "Repeated Games of Incomplete Information: An Approach to the Non-Zero Sum Case," Report of the U.S. Arms Control and Disarmament Agency, Contract ST-143, Prepared by Mathematica, Inc., Princeton, NJ., Chapter IV, pp. 117-216. Published as Chapter V in Aumann, R. J. and M. B. Maschler (1995): *Repeated Games with Incomplete Information*, MIT Press.

Battigalli, P. and G. Maggi (2002): "Rigidity, Discretion, and the Costs of Writing Contracts," *American Economic Review*, 92, 798–817.

Bhaskar, V. (1989): "Quick Responses in Duopoly Ensure Monopoly Pricing," *Economics Letters*, 29, 103–107.

Bryant, J. (1983): "A Simple Rational Expectations Keynes-Type Model," *Quarterly Journal of Economics*, 98, 525–528.

Calcagno, R., Y. Kamada, S. Lovo, and T. Sugaya (2014): "Asynchronicity and Coordination in Common and Opposing Interest Games," *Theoretical Economics*, 9, 409–434.

Caruana, G. and L. Einav (2008): "A Theory of Endogenous Commitment," *Review of Economic Studies*, 75, 99–116.

Crawford, V. (1998): "A Survey of Experiments on Communication via Cheap Talk," *Journal of Economic Theory*, 78, 286–298.

Dutta, P. K. (1995): "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, 66, 1–32.

——— (2012): "Coordination Need not be a Problem," *Games and Economic Behavior*, 76, 519–534.

Dutta, P. K. and S. Takahashi (2012): "Repeated Bargaining over a Finite Horizon," Mimeo.

Farm, A. and J. W. Weibull (1987): "Perfectly Flexible Pricing in a Homogeneous Market," *Scandinavian Journal of Economics*, 89, 487–495.

Farrell, J. (1987): "Cheap Talk, Coordination, and Entry," *RAND Journal of Economics*, 18, 34–39.

——— (1988): "Communication, Coordination and Nash Equilibrium," *Economics Letters*, 27, 209–214.

Farrell, J. and M. Rabin (1996): "Cheap Talk," *Journal of Economic Perspectives*, 10, 103–118.

Fershtman, C. (1990): "The Importance of the Agenda in Bargaining," *Games and Economic Behavior*, 2, 224–238.

Fudenberg, D. and D. Levine (1983): "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, 31, 251–268.

Green, J. R. and J.-J. Laffont (1987): "Limited Communication and Incentive Compatibility," in *Information, Incentives, and Economic Mechanisms: Essays in Honor of Leonid Hurwicz*, ed. by T. Groves, R. Radner, and S. Reiter, Basil Blackwell, 308–329.

Harstad, B. (2018): "Pledge-and-Review Bargaining," Mimeo.

Huyck, J. B. V., R. C. Battalio, and R. O. Beil (1990): "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure," *American Economic Review*, 80, 234–248.

In, Y. and R. Serrano (2004): "Agenda Restrictions in Multi-Issue Bargaining," *Journal of Economic Behavior & Organization*, 53, 385–399.

Kalai, A. T., E. Kalai, E. Lehrer, and D. Samet (2010): "A Commitment Folk Theorem," *Games and Economic Behavior*, 69, 127–137.

Kalai, E. (1981): "Preplay Negotiations and the Prisoner's Dilemma," *Mathematical Social Sciences*, 1, 375–379.

Kos, N. (2012): "Communication and Efficiency in Auctions," *Games and Economic Behavior*, 75, 233–249.

Lagunoff, R. and A. Matsui (1997): "Asynchronous Choice in Repeated Coordination Games," *Econometrica*, 65, 1467–1477.

——— (2001): "Are "Anti-Folk Theorems" in Repeated Games Nongeneric?" *Review of Economic Design*, 6, 397–412.

Maggi, G. (2016): "Issue Linkage," in *Handbook of Commercial Policy*, ed. by K. Bagwell and R. W. Staiger, Elsevier, vol. 1, 513–564.

Maskin, E. and J. Tirole (1987): "A Theory of Dynamic Oligopoly, III: Cournot Competition," *European Economic Review*, 56, 549–569.

——— (1988a): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56, 571–599.

——— (1988b): "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica*, 31, 947–968.

Mookherjee, D. and M. Tsumagari (2014): "Mechanism Design with Communication Constraints," *Journal of Political Economy*, 122, 1094–1129.

Muto, S. (1993): "Alternating-Move Preplays and vN-M Stable Sets in Two Person Strategic Form Games," Mimeo.

Nash, J. (1953): "Two Person Cooperative Games," *Econometrica*, 21, 128–140.

Osborne, M. J. and A. Rubinstein (1994): *A Course in Game Theory*, MIT Press.

Rabin, M. (1994): "A Model of Pre-game Communication," *Journal of Economic Theory*, 63, 370–391.

Rubinstein, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97–109.

Safronov, M. and B. Strulovici (2017): "From Self-enforcing Agreements to Self-sustaining Norms," Mimeo.

Santos, V. (2000): "Alternating-announcements Cheap Talk," *Journal of Economic Behavior & Organization*, 42, 405–416.

Stahl, D. O. I. (1986): "Revocable Pricing Can Yield Collusive Outcomes," *Economics Letters*, 22, 87–90.

Ståhl, I. (1972): "Bargaining Theory," Economic Research Institute at the Stockholm School of Economics (EFI).

Yoon, K. (2001): "A Folk Theorem for Asynchronously Repeated Games," *Econometrica*, 69, 191–200.