Renegotiation and Coordination with Private Values*

Yuval Heller[†] and Christoph Kuzmics[‡]

April 8, 2019

Abstract

We study coordination games with pre-play communication in which agents have private preferences over the feasible coordinated outcomes. We present a novel intuitive equilibrium strategy with the following properties: each agent reports his preferred outcome (and nothing else); agents never miscoordinate; if the agents have the same preferred outcome, then they coordinate on this outcome; and otherwise, there is a "fallback norm" that determines the coordinated outcome. We show that this behavior is essentially the unique renegotiation-proof strategy, and that it satisfies appealing properties: independence of the distribution of private preferences, Pareto optimality, high ex-ante expected payoff, and evolutionary stability.

Keywords: coordination games, renegotiation-proof, equilibrium entrants, secret handshake, incomplete information, evolutionary robustness.

JEL codes: C72, C73, D82

Very preliminary and incomplete

*We have benefited greatly from discussions with Bill Sandholm and Tilman Borgers. We would like to express our deep gratitude to participants of LEG2018 conference (Lund University), Bielefeld Game Theory 2018 workshop, Israeli Game Theory 2018 conference in IDC, and seminar audiences in Bar Ilan University and Tel Aviv University for many useful comments. Yuval Heller is grateful to the European Research Council for its financial support (Starting Grant #677057).

[†] Department of Economics, Bar Ilan University, Israel. Email: yuval.heller@biu.ac.il.

[‡] Department of Economics, University of Graz, Austria. Email: christoph.kuzmics@uni-graz.at.

1 Introduction

Various strategic interactions have the following three key properties: agents gain from coordinating their behavior, agents have private preferences about the feasible coordinated outcomes, and agents can communicate before taking their actions.

One common example for these interactions is two pedestrians coming from opposite sides to a narrow pass. In order to avoid bumping into each other, both pedestrians should slightly turn to the same direction (e.g., both slightly turning to the left). In many cases, each agent has private preference regarding the preferred direction (e.g., an agent turning to the left immediately after the narrow pass, would prefer to walk on the left-hand side of the narrow pass). Finally, agents can use non-verbal communication to signal their preferred direction (e.g., tilting their head toward the preferred direction). Casual observation suggests the following stylised facts: agents succeed in not bumping into each other; unlike car traffic, there is no uniform norm such as "always pass on the right"; the coordinated outcome respects the ordinal preferences of the pedestrians (if both prefer left over right, they coordinate on left); and the resulting behavior is simple and relies on very brief communication.[1]

In what follows we describe two important economic applications that have the above three key properties. The first application is market sharing agreements in oligopolistic markets by which a firm sells in a certain region (or serves customers of a certain type), whereas the rivals sell in other regions (see, e.g., Belleflamme and Bloch, 2004 and Motta, 2004, Section 4.1, p. 141). It seems plausible to assume in such situations that the firms prefer coordinating over a sharing scheme over a situation in which they compete in the same regions, each firm has private preferences over the various regions, and firms can communicate (though, each additional communication may incur some risk of anti-trust investigation). A specific example is the 1997 FCC ascending auctions allocating licenses for particular slices of the electromagnetic spectrum, in which the the firms used the trailing digits of their bids to reveal information on their preferred geographic areas and frequencies, and the firms used this information to collude (Cramton and Schwartz, 2000).

The second application is a research joint venture in which some firms agree to share the costs and benefits associated with a given research project (see, e.g., Katz, 1986; Vonortas, 2012). In many cases the firms have private preferences regarding the goals, methods and the extent of knowledge transfer in the joint venture. For a concrete related example, consider co-authors from different academic disciplines working on a joint paper, where each researcher has private preferences regarding issues such as which methodology to apply in the research, and to which journal to submit the paper to.

[1] The example is motivated by Goffman (1971, Chapter 1, p. 6), in which Goffman writes:

> "Take, for example, techniques that pedestrians employ in order to avoid bumping into one another. ... There are an appreciable number of such devices; they are constantly in use and they cast a pattern on street behavior. Street traffic would be a shambles without them."

Our key results are as follows. We present a novel intuitive family of equilibrium strategies with the following properties: each agent reports his preferred outcome (and nothing else); agents never miscoordinate; if the agents have the same preferred outcome, then they coordinate on this outcome; and otherwise, there is a "fallback norm" that determines the coordinated outcome. We show that this behavior is essentially the unique equilibrium strategy that satisfies a mild renegotiation-proof refinement, which requires that there does not exists a Pareto-better equilibrium, where we allow the players to rely on additional communication to coordinate on the alternative equilibrium. Finally, we show that these equilibrium strategy satisfies various appealing properties: independence of the distribution of private preferences, Pareto optimality (with respect to any strategy profile, not only equilibrium strategies), high ex-ante expected payoff, and robustness against various perturbations (in the spirit of evolutionary stability à la Maynard Smith, 1974).

Overview of the Baseline Model We consider a setup in which two agents with private idiosyncratic preferences are randomly matched to play a coordination game, and this game is proceeded by pre-play cheap-talk. Each player can choose one of two actions, L and R. Each player has a privately known value (or type) that is independently drawn from a common atomless distribution F on the unit interval. An agent with value $u \in (0,1)$ obtains a payoff u when both players coordinate on R, a payoff of $1-u$ when both players coordinate on L, and a payoff of zero when the players choose different actions (henceforth, miscoordinate) (see Table 1 in Section 2). After learning her type, but before playing this coordination game, each player simultaneously sends a message to her partner from a finite set of messages. A strategy of a player describes which (possibly mixed) message the agent sends as a function of her own type in the first stage, and which action the player chooses in the second stage as a function of her own type and the messages sent by both players.

Equilibrium Strategies Strategy σ is a symmetric Bayes-Nash equilibrium strategy, abbreviated to *equilibrium strategy*, if each type best replies to the aggregate behavior in the population given by σ itself.

The game admits many equilibria. In particular, it includes *babbling equilibria*, in which the agents' behavior is independent of the pre-play communication; for example, the equilibrium in which agents always choose L. In addition, it can include equilibria in which agents communicate some information about their cardinal preferences, and these reports influence the agents' behavior, see Example 1.

Renegotiation-proofness We define a relatively mild form of renegotiation-proofness. Our motivation for this refinement is that if agents can communicate prior to playing the game, it seems plausible to assume that they can further communicate after observing the realized pair of messages, and that if there exists an alternative Pareto-better equilibrium, the agents will use the additional communication to agree on replacing the original equilibrium with the Pareto-better equilibrium, and as result, the original equilibrium would not be played in the first place. Refinements of renegotiation-proofness have been applied to various setups of contracts with incomplete information (see, e.g., Holmström and Myerson, 1983) and repeated games (Farrell and Maskin, 1989; Benoit and Krishna, 1993). A detailed discussion of the related literature on renegotiation-proofness is presented in Section 9.

Our definition of renegotiation-proofness is as follows. Fix an equilibrium strategy σ. Observe that each message sent by an agent on the equilibrium path induces a posterior belief of the partner about the agent's type. Thus, each pair of realized messages m, m' induce a new coordination game, which differs from the original one by each agent now holding her posterior distribution of types (rather than the prior distribution F). Consider a variant of this induced game, in which the agents have another round of communication in which they can simultaneously send messages from a new set of messages $\hat{M}$ before choosing an action. We say that the equilibrium σ is renegotiation-proof if for each new set of messages $\hat{M}$, and each pair of messages $m, m' \in M$ that are sent on the equilibrium path according to σ, there does not exist an alternative equilibrium $\hat{\sigma}$ of the induced game that Pareto-dominates the original equilibrium induced by σ: that is, we require that it will not be the case that all possible types of both players (in the support of their respective posterior distributions) weakly gain from replacing the original equilibrium induced by σ with the alternative equilibrium, and some types of one of the players gain from this replacement. As mentioned above, the motivation for our definition is that if the Pareto-better equilibrium $\hat{\sigma}$ exists, then the agents will renegotiate to play $\hat{\sigma}$ after observing (m, m'), and foreseeing this, the original equilibrium σ would not be played in the first place.

Summary of Main Results Theorem 1 provides a full characterization of renegotiation proof equilibria in coordination games with independent private values and pre-play communication.

A strategy is a renegotiation proof equilibrium strategy if and only if it is coordinated (the two players always use identical actions), it is ordinal preference revealing (each player reveals her ordinal preference), it is mutual preference consistent (when both players have the same ordinal preference they coordinate on their mutually preferred outcome), and it is balanced (each message that players send leads, conditional on the opponent sending a message indicating an opposite cardinal preference, to the same probability of coordination on action L.

Prime examples of such renegotiation proof equilibria (showing also existence) are strategies that we denote as σ_L and σ_R. In each of these strategies really only two messages are used: players either indicate that they prefer to go "left" or "right". That is all the communication achieves. When both indicate the same preference, they both coordinate on that. If they send "mixed signals" (one indicates a preference for left, the other for right), they coordinate on a fallback norm, which is L for strategy σ_L and R for σ_R. The only way other renegotiation proof equilibria can differ from these two, is that they induce more complicated coordination after "mixed signals". This necessitates sending more than two messages.

As a simple intuition for the main result we sketch why a babbling equilibrium, such as always play L, is not renegotiation-proof because agents can renegotiate to the Pareto-better equilibrium σ_L: each agent communicates his preferred outcome, and the agents change their behavior to playing R if and only if both agents prefer R over L.

Appealing Properties Observe that the strategies σ_L and σ_R have two appealing properties. First, they require very little communication: each agent sends a single bit - L or R, and second, they are completely independent of the distributions of types (both strategies remain equilibria for any distribution of types F, and, moreover, an agent who follows σ_L or σ_R does not have to know anything about F.

In Section 6 we show that σ_L and σ_R satisfy four additional appealing properties: Pareto optimality: there does not exist any strategy profile (even when allowing non-equilibrium profiles) that weakly improves the payoff of all types, and strictly improves the payoff of some types relative to the payoff of either σ_L or σ_R; the ex-ante payoff of either σ_L or σ_R maximizes the agents' ex-ante payoff among all equilibria without mis-coordination;[2] either σ_L or σ_R strictly improves the ex-ante payoff with respect to all babbling equilibria; and σ_L are σ_R are robust against various perturbations (i.e., a population slightly perturbing away from the behavior induced by σ_L or σ_R would converge back to the equilibrium behavior): they are neutrally stable (à la Maynard Smith, 1974), the first-stage behavior induced by σ_L and σ_R is weakly dominant, and the second-stage behavior is a neighborhood invader (Apaloo, 1997; Cressman, 2010).

Variants and Extensions Next, we show that our results are robust to various variants and extensions. Specifically, we show that all our results remain essentially the same if one allows: multiple rounds of communication (instead of a single round of communication as in the baseline model); more than two players (where agents get a payoff of zero, unless all agents coordinate on the same outcome; and allowing a larger set of feasible types, namely all matrix payoffs that satisfy the constraint of always preferring coordination over miscoordination.

Finally, we show that some of our key results can be extended to two additional setups. First, we allow extreme types, for which one of the actions (L or R) is dominant, and we show that as long as these extreme types are not too frequent, then a particular "convex combination" of σ_L and σ_R remains a renegotiation-proof equilibrium strategy. Second, we consider coordination games with more than two actions, and we show that strategy again a particular "convex combination" of σ_L and σ_R remains a renegotiation-proof equilibrium strategy in this setup, and that most aspects of the main result hold also in this setup.

[2] Example 1 in Section 6 shows that given some distributions of types, there exists an equilibrium with miscoordination that induces a higher ex-ante payoff. Note, however, that such an equilibrium does not satisfy renegotiation-proofness, and, thus, we do not consider it as a plausible stable prediction of the players' behavior.

Main Insights Our model induces a few key insights on real life interactions that can be modeled as coordination games with private values (such as the motivated examples presented above: collusion in oligopolistic markets and joint research ventures).

The first insight is that a single bit of communication is sufficient to rule out both mis-coordination,[3] as well as uniform norms; in any renegotiation-proof equilibrium agents always coordinate, and the equilibrium behavior respects the ordinal preferences of the agents (i.e, agents coordinate on a jointly preferred outcome, whenever such an outcome exist). By contrast, adding more options to communicate (i.e., having a larger set of messages or/and longer pre-play communication) does not have any influence on the predicted behavior.

Another insight is the contrast between revealing ordinal and cardinal preferences. Our results imply that agents credibly reveal their ordinal preferences among the feasible outcomes in *any* renegotiation-proof equilibrium; by contrast it is impossible for the agents to credibly reveal information about the cardinal preferences (namely, how much the agent prefers L over R).

Structure The baseline model is presented in section 2. Section 3 presents various equilibrium strategies, and defines the four key properties that we show renegotiation proof equilibria have. In Section 4 present our refinement of renegotiation-proofness. Section 5 presents our main result. In Section 6 we present various appealing properties of renegotiation proof equilibria. Section 7 shows the evolutionary stability of renegotiation-proof equilibria. Section 8 shows the robustness of our results in various variants and extensions of our baseline model. In Section 9 we discuss the related literature. We conclude in Section 10.

2 Model

We consider a setup in which two agents with private idiosyncratic preferences are randomly matched to play a coordination game, and this game is preceded by pre-play cheap-talk.

Players and Types There are two players in ex-ante symmetric positions. Players can choose one of two actions, L and R. Each player has a privately known "value" or "type". The two players' values are independently drawn from a common atomless distribution with cumulative distribution function F on the unit interval $U = [0, 1]$ and with density f. To make it interesting we make the assumption, throughout the paper, that $F(1/2) \in (0, 1)$. This means that not all player types agree on their preferred outcome in the coordination game below.

[3] Remark 2 in Section 3 shows that without communication it might be that the unique robust equilibrium induces a substantial probability of miscoordination.

Payoff Matrix For any realized pair of types, u and v, the players play a coordination game given by the following payoff matrix, where the first entry is the payoff of the player of type u (choosing row) and the second entry the payoff of the player of type v (choosing column).

Table 1: Payoff Matrix of the Coordination Game

	L	R
L	1-u, 1-v	0, 0
R	0, 0	u, v

We call this game the *coordination game without communication* and denote it by Γ.

Pre-Play Communication After learning their type, but before playing this coordination game, the two players each simultaneously send a publicly observable message from a finite set of messages M (satisfying $2 \leq |M| < \infty$), with $\Delta(M)$ the set of all probability distributions over messages in M. We assume that messages are costless, i.e., cheap talk. We call the game, so amended, the *coordination game with communication* and denote it by $\langle \Gamma, M \rangle$.

Strategies A player's (ex-ante) strategy in the coordination game with communication is then a pair $\sigma = (\mu, \xi)$, where $\mu : U \to \Delta(M)$ is a (Lebesgue measurable) *message function* that describes which (possibly random) message is sent for each possible realization of the agent's type, and $\xi : M \times M \to U$ is an *action function* that describes the maximal type that chooses L as a function of the observed message profile; that is, when an agent who follows strategy (μ, ξ) observes a message profile (m, m') (message m sent by the agent, and message m' sent by the opponent), then the agent plays L if her type u is at most $\xi(m, m')$ (i.e., if $u \leq \xi(m, m')$), and she plays R if $u > \xi(m, m')$. Let Σ be the set of all strategies in the game $\langle \Gamma, M \rangle$.

Remark 1. In principle we should allow more general action functions $\xi : U \times M \times M \to \triangle\{L, R\}$, which specify the probability an agent chooses L as a function of the observed message profile and the agent's type. It is simple to see, however, and proven in Lemma 1 in Appendix B.1, that any "generalized" strategy is (weakly) dominated by a strategy that uses a cut-off action function in the second stage (i.e., choosing L for any type $u \leq u_{m,m'}$ and choosing R for any type $u > u_{m,m'}$). The intuition, is that following the observation of any pair of messages, lower types always gain weakly more (less) than higher types from choosing L (R). We, thus, simplify our notation by considering only cut-off action functions of the form $\xi : M \times M \to U$.[4]

Let $\mu_u(m)$ denote the probability, given message function μ, that a player sends message m if she is of type u. Let $\mu(m) = \mathbb{E}_u[\mu_u(m)]$ be the mean probability that a player of a random type sends message m. Let $\text{supp}(\mu) = \{m \in M | \mu(m) > 0\}$ denote the support (carrier) of μ.

[4]The arbitrary choice that the threshold type plays L does not play any role in our analysis, given the assumption that the distribution of types F is without atoms. All results remain the same under any assumption on the behavior of the threshold types.

With a slight abuse of notation we write $\xi(m,m')=L$ when all types play L (i.e., when $\xi(m,m')=1$), and we write $\xi(m,m')=R$ when all types play R (i.e., when $\xi(m,m')=0$).

3 Equilibrium Strategies

In this section we define the standard notion of (Bayes Nash) equilibrium strategies, present properties that renegotiation proof equilibria turn out to have, and present examples of equilibria in the coordination game with communication with and without these properties.

Definition Given a strategy profile (σ,σ') and a type profile $u,v\in U$, let $\pi_{u,v}(\sigma,\sigma')$ denote the interim (pre-communication) expected payoff of a player of type u who follows strategy σ and faces an opponent of type v who follows strategy σ'. Formally, for $\sigma=(\mu,\xi)$ and $\sigma'=(\mu',\xi')$,

$$\pi_{u,v}(\sigma,\sigma')=\sum_{m\in M}\sum_{m'\in M}\mu_u(m)\cdot\mu_v(m')\cdot\Big((1-u)\mathbf{1}_{\{u\leq\xi(m,m')\}}\mathbf{1}_{\{v\leq\xi'(m',m)\}}+u\mathbf{1}_{\{u>\xi(m,m')\}}\mathbf{1}_{\{v>\xi'(m',m)\}}\Big),$$

where generally $\mathbf{1}_{\{x\}}$ is the indicator function equal to 1 if statement x is true and equal to zero otherwise. Let

$$\pi_u(\sigma,\sigma')=\mathbb{E}_v\left[\pi_{u,v}(\sigma,\sigma')\right]\equiv\int_{v=0}^{1}\pi_{u,v}(\sigma,\sigma')\cdot f(v)\,dv$$

denote the expected interim payoff of a player of type u who follows strategy σ and faces an opponent with a random type who follows strategy σ'. Finally, let,

$$\pi(\sigma,\sigma')=\mathbb{E}_u\left[\pi_u(\sigma,\sigma')\right]=\int_{u=0}^{1}\pi_u(\sigma,\sigma')\cdot f(u)\,du$$

denote the ex-ante expected payoff of an agent who follows strategy σ and faces an opponent who follows strategy σ'.

A strategy σ is a *(symmetric Bayes-Nash) equilibrium strategy* if $\pi_u(\sigma,\sigma)\geq\pi_u(\sigma',\sigma)$ for each $u\in[0,1]$ and each strategy $\sigma'\in\Sigma$. Let $\mathcal{E}\subseteq\Sigma$ denote the set of all equilibrium strategies of $\langle\Gamma,M\rangle$.

Properties We call a strategy $\sigma=(\mu,\xi)\in\Sigma$ *ordinal preference revealing* if there exist two non-empty, disjoint, and exhaustive subsets of supp(μ) denoted by M_L and M_R (i.e., supp(μ) $=M_L\dot{\bigcup}M_R$) such that if $u<1/2$, then $\mu_u(m)=0$ for each $m\in M_R$, and if $u>1/2$, then $\mu_u(m)=0$ for each $m\in M_L$. With an ordinal preference revealing strategy a player indicates her ordinal preferences.

We call it *mutual preference consistent* if whenever $u,v<1/2$ then $\xi(m,m')=\xi(m',m)=L$ for all $m\in$ supp(μ_u) and all $m'\in$ supp(μ_v) and whenever $u,v>1/2$ then $\xi(m,m')=\xi(m',m)=R$ for all $m\in$ supp(μ_u) and all $m'\in$ supp(μ_v). In a mutual preference consistent strategy players with the same ordinal preference coordinate on their mutually preferred choice.

We call it *coordinated* if $\xi(m, m') = \xi(m', m) \in \{L, R\}$ for any pair of messages $m, m' \in \text{supp}(\mu)$. A coordinated strategy never leads to mis-coordination after any (used) message pair.

Given strategy $\sigma = (\mu, \xi) \in \Sigma$ denote by

$$\alpha_m(\sigma) = \int_{v>\frac{1}{2}} \sum_{m' \in \text{supp}(\mu_v)} \mu_v(m') \mathbf{1}_{\{v \leq \xi(m,m')\}} f(v) dv$$

the expected probability of a u-type's opponent playing L conditional on the opponent having a type in $[1/2, 1]$ and conditional on the u-type sending message $m \in M$. Let analogously denote by

$$\beta_m(\sigma) = \int_{v<\frac{1}{2}} \sum_{m' \in \text{supp}(\mu_v)} \mu_v(m') \mathbf{1}_{\{v \leq \xi(m,m')\}} f(v) dv$$

the expected probability of a u-type's opponent playing L conditional on the opponent having a type in $[0, 1/2]$ and conditional on the u-type sending message $m \in M$.

We call a strategy $\sigma = (\mu, \xi) \in \Sigma$ *balanced* if there exists $\alpha \in [0, 1]$ such that $\alpha_m(\sigma) \leq \alpha$ for all $m \in M$ with equality for all $m \in \text{supp}(\mu_u)$ with $u \leq 1/2$, and $\beta_m(\sigma) \geq \alpha$ for all $m \in M$ with equality for all $m \in \text{supp}(\mu_u)$ with $u > 1/2$. We call the above α the *left-tendency* of a balanced strategy σ.

A strategy is balanced if any type who prefers (to coordinate on) one action, say L, and who knows only that she faces an opponent who prefers (to coordinate on) the other action R expects the same probability of her opponent playing action L for all messages that she uses and expects a lower or equal probability of her opponent playing her favorite action for all messages that she could use.

Note that a strategy that is balanced, coordinated, and ordinal preference revealing induces behavior that is independent of the player's cardinal preferences beyond the ordinal content of her preferences, $u \leq 1/2$ or $u > 1/2$.

Examples The following strategies, denoted by σ_L, σ_R, and σ_C, are prime examples (that play a special role in later sections) of strategies that are all ordinal preference revealing, mutual preference consistent, coordinated, and balanced.

The strategies σ_L and σ_R are given by the pairs (μ^*, ξ_L) and (μ^*, ξ_R), respectively. The message strategy function μ^* has the property that there are messages $m_L, m_R \in M$ such that message m_L indicates a preference for L and m_R a preference for R, that is

$$\mu^*(u) = \begin{cases} m_L & u \leq \frac{1}{2} \\ m_R & u > \frac{1}{2}. \end{cases}$$

The action strategy parts ξ_L and ξ_R are defined as follows:

$$\xi_L(m,m') = \begin{cases} R & m = m' = m_R \\ L & \text{otherwise,} \end{cases} \qquad \xi_R(m,m') = \begin{cases} L & m = m' = m_L \\ R & \text{otherwise.} \end{cases}$$

This means that the "fallback norm" of σ_L (which is applied when the agents have different preferred outcomes) is to coordinate on L, while the fallback norm of σ_R is to coordinate on R. In other words the left-tendency of σ_L is one and the left-tendency of σ_R is zero.

Next, we present the strategy $\sigma_C = (\mu_C, \xi_C)$, in which the "fallback norm" is to use a joint lottery to randomly choose the coordinated outcome. The agents implement this joint lottery through each agent simultaneously sending a random bit, and let the coordinated outcome depend on whether the random bits are equal or not.

Assume that $|M| \geq 4$. We denote four distinct messages as $m_{L,0}, m_{L,1}, m_{R,0}, m_{R,1} \in M$, where we interpret the first subscript (R or L) as the agent's preferred direction, and the second subscript (0 or 1) as a random binary number chosen with probability 1/2 each by the agent. Formally, the message function μ_C is defined as follows:

$$\mu_C(u) = \begin{cases} \frac{1}{2} m_{L,0} \oplus \frac{1}{2} m_{L,1} & u \leq \frac{1}{2} \\ \frac{1}{2} m_{R,0} \oplus \frac{1}{2} m_{R,1} & u > \frac{1}{2}, \end{cases}$$

where $\alpha m \oplus (1-\alpha) m'$ denotes the lottery that attaches a probability of α on message m and $1-\alpha$ on message m'.

In the second stage, if both agents share the same preferred outcome they play it. Otherwise, they coordinate on L if their random numbers differ, and coordinate on R if they have chosen the same random number. Formally:

$$\xi_C(m,m') = \begin{cases} R & (m,m') \in \{(m_{R,0}, m_{R,0}), (m_{R,0}, m_{R,1}), (m_{R,0}, m_{L,0}), (m_{R,1}, m_{L,1}) \\ & \qquad (m_{R,1}, m_{R,1}), (m_{R,1}, m_{R,0}), (m_{L,0}, m_{R,0}), (m_{L,1}, m_{R,1})\} \\ L & \text{otherwise.} \end{cases}$$

Note that among all strategies that satisfy the four properties, strategies σ_L and σ_R are the *simplest* in terms of the number of "bits" needed to implement the message strategy. Note that strategy σ_C is in a certain sense *fairest*: conditional on a coordination conflict, i.e., conditional on one agent having type between 0 and 1/2 and the other between 1/2 and 1 both agents expect the same payoff. Strategy σ_L in contrast favors types between 0 and 1/2 and σ_R favors types between 1/2 and 1.

The coordination game with communication $\langle \Gamma, M \rangle$ admits many more equilibria that satisfy only some or even none of the four properties defined above.

The game $\langle \Gamma, M \rangle$, for instance, has two simple babbling equilibria, in which agents ignore the communication and apply a uniform norm of always playing L or of always playing R. These equilibria are coordinated and (trivially) balanced, but not ordinal preference revealing nor mutually preference consistent.

Depending on the distribution of types, the game can also have more inefficient babbling equilibria in which agents sometimes mis-coordinate. Specifically, if there exists a type $x \in (0,1)$ satisfying $x = F(x)$, then it is an equilibrium for the agents to babble, and then for each player to ignore whatever is said and to choose L if and only if her type is below x. Such a babbling equilibrium is (trivially) balanced but satisfies none of the other three properties defined above. Note that, as by assumption $F(0) = 0$ and $F(1) = 1$, all babbling equilibria can be identified with an $x \in [0,1]$ that has the property that $F(x) = x$.

Remark 2 (Robustness of equilibria without communication). When there is no pre-play communication (i.e., $|M| = 1$), then these babbling equilibria of course constitute all equilibria. Arguably, a plausible equilibrium refinement in setups without communication is robustness to small perturbations in the behavior of the population (e.g., requiring Lyapunov stability of the best-reply dynamics, or continuous stability à la Eshel, 1983). Adapting the analysis of Sandholm (2007) to the current setup implies that an equilibrium is robust in this sense if and only if the density of the distribution of types at the relevant threshold x (with $x = F(x)$) is less than one. In particular, if the distribution of types satisfies $f(0), f(1) > 1$, then there exists $x \in (0,1)$ satisfying $x = F(x)$ and $f(x) < 1$. The corresponding equilibrium, which entails inefficient miscoordination is then robust to small perturbations. Thus, coordination games without communication are likely to induce substantial miscoordination if the density of extreme types is high (i.e., if $f(0), f(1) > 1$).

The game also admits equilibria in which agents reveal some information about the cardinality of their preferences (i.e., some information beyond only stating if $u \leq 1/2$ or $u > 1/2$). For example, it is straightforward to show that for any symmetric distribution of types (i.e., $F(1/2) = 1/2$), there exists a type $x \in (0, 1/2)$, such that there is an equilibrium strategy in which each agent communicates to which one of four possible intervals her type belongs to: below x (interpreted as strong preference for L), between x and $1/2$ (mild preference for L), between $1/2$ and $1-x$ (mild preference for R), or above $1-x$ (strong preference for R). Agents with the same preferred outcome, coordinate on this outcome. Agents with different preferred outcomes play as follows: if both agents report a mild preference, they choose a random coordinated outcome based on a joint lottery; if one of the agents reports a strong preference, then they coordinate on this agent's preferred outcome; and if both agents report strong preferences, then they play an inefficient equilibrium with positive probability of mis-coordination. Note, that playing an inefficient outcome in this latter case is necessary for this strategy to be an equilibrium, as otherwise, all agents would prefer reporting a strong preference over a mild preference, contradicting this strategy being an equilibrium. Such an equilibrium is ordinal preference revealing, but not mutually preference consistent, nor coordinated, nor balanced. See also Example 1 in Section 6.

4 Definition of Renegotiation-Proofness

In order to define our notion of renegotiation-proofness it is useful to introduce a bit more notation. For any given strategy in Σ, employed by both players, in the game $\langle \Gamma, M \rangle$, communication and knowledge of this strategy leads to updated and possibly different and asymmetric information about the two agents' types. Suppose the updated distributions of types are given by some distribution functions G and H. The two agents then face a (possibly asymmetric) game of coordination without communication, which we shall denote by $\Gamma(G, H)$. Note that the original game (without communication) Γ is then given by $\Gamma(F, F)$.

Let f_m be the type density conditional on the agent following a given strategy in the game Γ and sending a message m that is sent with positive probability given this strategy.[5]

That is,

$$f_m(u) = \frac{f(u) \cdot \mu_u(m)}{\mu(m)},$$

and let F_m be the cumulative distribution function associated with density f_m.

We allow players to renegotiate (only) after communication. Players in their renegotiation can use any new finite message set, $\tilde{M}$. Given a strategy of the game Γ, employed by both players, we denote the induced "renegotiation" game after a positive probability message pair $m, m' \in M$ by $\langle \Gamma(F_m, F_{m'}), \tilde{M} \rangle$. For a pair of strategies σ, σ' of such a renegotiation game $\langle \Gamma(G, H), \tilde{M} \rangle$ define the *post communication* expected payoffs for a type u agent by

$$\pi_u^H(\sigma, \sigma') = \mathbb{E}_{v \sim H}\left[\pi_{u,v}(\sigma, \sigma')\right] \equiv \int_{v=0}^{1} \pi_{u,v}(\sigma, \sigma') \cdot h(v)\, dv.$$

Define $\mathcal{E}(G, H)$ as the set of all (possibly asymmetric) equilibrium profiles of the coordination game with communication $\langle \Gamma(G, H), \tilde{M} \rangle$ for some finite message set $\tilde{M}$. Furthermore let $\mathcal{S}(G)$ denote the set of all symmetric equilibrium strategies of the coordination game with communication $\langle \Gamma(G, G), \tilde{M} \rangle$ for some finite message set $\tilde{M}$. With a slight abuse of notation for any strategy σ of the game $\langle \Gamma, M \rangle$ we denote its prescription after message pair $m, m' \in M$, i.e., in the game $\langle \Gamma(F_m, F_{m'}), \tilde{M} \rangle$ by σ as well.

Definition 1. We say that an equilibrium strategy $\sigma = (\mu, \xi) \in \mathcal{E}$ is *post communication equilibrium Pareto dominated* if either there is a message $m \in \text{supp}(\mu)$ and an equilibrium $\tilde{\sigma} \in \mathcal{S}(F_m)$ such

[5] The density f_m depends on the given strategy in the game Γ. For aesthetic reasons we refrain from giving this strategy a name and omit to indicate this obvious dependence in our notation.

that $\pi_u^{F_m}(\sigma,\sigma) \leq \pi_u^{F_m}(\tilde{\sigma},\tilde{\sigma})$ for all $u \in \text{supp}(F_m)$ with strict inequality for some $u \in \text{supp}(F_m)$, or there is a pair of messages $m, m' \in \text{supp}(\mu)$ and an equilibrium profile $\tilde{\sigma} \in \mathcal{E}(F_m, F_{m'})$ such that $\pi_u^{F_{m'}}(\sigma,\sigma) \leq \pi_u^{F_{m'}}(\tilde{\sigma})$ and $\pi_v^{F_m}(\sigma,\sigma) \leq \pi_v^{F_m}(\tilde{\sigma})$ for all $u \in \text{supp}(F_m)$ and all $v \in \text{supp}(F_{m'})$ with strict inequality for some $u \in \text{supp}(F_m)$ or some $v \in \text{supp}(F_{m'})$.

Definition 2. An equilibrium strategy $\sigma = (\mu, \xi) \in \mathcal{E}$ is *renegotiation-proof* if it is not post communication equilibrium Pareto-dominated.

The motivation for renegotiation-proofness is that if the agents can communicate prior to playing the game, then it seems plausible that they can further communicate after observing the realized messages. If there is an observed pair of messages after which the original equilibrium induces the agents to play a strategy profile with a low payoff, then, arguably, the agents can use an additional round of communication to renegotiate the existing "bad" equilibrium of the current induced game, and to coordinate their play on a Pareto-improving equilibrium (which weakly improves the payoff of all possible types of both players). Our refinement of renegotiation-proofness requires that no such Pareto-improving equilibria exists in any induced game with additional communication.

Refinements of renegotiation-proofness were presented in various setups in the existing literature; though, to the best of our knowledge we are the first to apply this refinement to one-shot games with private values and pre-play cheap-talk. In section 9, we thoroughly discuss the related literature on renegotiation-proofness, including somewhat related notions in the evolutionary game theory, such as secret hand-shakes (see, e.g., Robson, 1990) and robustness to equilibrium entrants (Swinkels, 1992).

We have chosen to define a mild notion of renegotiation-proofness because it already suffices for the sharp characterization given in Theorem 1. Our refinement is mild in the following ways: we allow players to renegotiate only after observing their realized messages (but not before), and when players play a symmetric induced game, we allow them only to implement an alternative symmetric equilibrium (rather, then also allowing them to play asymmetric equilibria, in which an agent's behavior may explicitly depend on its role in the game).

In Section 6 we show that renegotiation proof equilibria satisfy a much stronger refinement of Pareto optimality, which requires there is no Pareto better feasible strategy profile (also allowing non-equilibrium strategy profiles). As argued also in Section 6, this fact implies that our renegotiation proof equilibria satisfy stronger refinements of renegotiation-proofness, including: a refinement that allows agents to play an asymmetric equilibrium in a symmetric induced game, and a refinement a la Benoit and Krishna, 1993 that allows players to renegotiate also in earlier stages: in the interim stage before observing the realized messages, and, possibly, also in the ex-ante stage, before observing one's own type.

5 Main Result

With all this in place we can state our main result.

Theorem 1. *A strategy σ of the game with communication $\langle \Gamma, M \rangle$ is a renegotiation-proof equilibrium strategy if and only if it is ordinal preference revealing, mutual preference consistent, coordinated, and balanced.*

The complete proof of this theorem involves a series of lemmas and is given in all its details, including lemmas and their proofs, in Appendix B.2. The "if" part, i.e., that any strategy satisfying the four properties must be renegotiation proof, is fairly straightforward. We here provide a sketch of the proof of the "only if" part.

Lemma 3 provides a preliminary property: agents never mis-coordinate after sending the same message. That is, we show that whenever two players send the same message m they must coordinate (either on L or on R) with probability one. Otherwise, due to the agents having a symmetric behavior after sending the same messages, each agent plays L with the same average probability $p \in (0,1)$ in the induced game $\Gamma(F_m, F_m)$. This implies that the expected payoff for each type u of each player is equal to $\max((1-u) \cdot q, u \cdot (1-q))$. This, in turn, implies that the players can renegotiate to a Pareto-dominating equilibrium in which the agents use an additional round of communication to do a joint lottery, in which with probability q they both play L and with probability $1-q$ they both play R.[6] This new equilibrium induces each type of each player with a higher payoff of $(1-u) \cdot q + u \cdot (1-q)$.

Using Lemma 3 we then show that renegotiation-proof equilibrium strategy must be *ordinal preference revealing* in Lemma 4: Assume to the contrary that some message m is sent by two types with different preferred outcomes. Lemma 3 implies that agents always coordinate on the same outcome after observing (m, m). Assume without loss of generality that this outcome is L. Then, players can renegotiate to a Pareto-dominating equilibrium in which each agent communicates her preferred outcome, and the agents play the opposite outcome, R, if and only if both agents prefer it. This contradicts that the original strategy is renegotiation-proof.

Lemma 4 implies that we can indeed partition $\mathrm{supp}(\mu)$ into two disjoint sets M_L and M_R, and that all types $u < 1/2$ send messages in M_L, while all types $u > 1/2$ send messages in M_R.

Lemma 4 is then a key ingredient in the proof that a renegotiation proof equilibrium strategy satisfies the remaining three properties.

To show that a renegotiation proof equilibrium strategy is *mutual preference consistent*, which we do

[6] Specifically, assume that $q = \frac{k}{n}$ is a rational number. The joint lottery can be done as follows (Aumman, Maschler, and Stearns, 1968): $\tilde{M} = \{1, .., n\}$; each player i sends a random message m_i uniformly. Both players play L if $(m_i + m_j) \bmod n \le k$, and they both play R otherwise. A more elaborate argument, which does not assume q to be a rational number, is presented in Lemmas 2 and 3).

in Lemma 5, assume to the contrary that agents sometimes play a different outcome than (L, L) after observing message pair $m, m' \in M_L$. Note that by force of Lemma 4 both agents must have a type below 1/2. The two players can, therefore, renegotiate to a Pareto dominating outcome which is to always play L after observing the pair of messages (m, m'). An analogous argument shows that if $m, m' \in M_R$, then agents must always coordinate on R in a renegotiation proof equilibrium strategy.

In Lemma 6 we show that there is no miscoordination after observing (m, m') where $m \in M_L$ and $m' \in M_R$ (the previous steps imply that there is no miscoordination in the remaining cases). Assume to the contrary that there is mis-coordination in the game induced after observing (m, m'). This implies that there are types of both players playing both actions. In particular, there is a type $u < 1/2$ of the player who sent message m that plays action R and gets a payoff of at most $u < 1/2$ (and would get a payoff of at most 1/2 also if she were playing L). This implies that all types of the this player get a payoff of at most $u < 1/2$, and the same is true for all types of the other player who sent message m' by an analogous argument. We conclude the step by showing that the players can renegotiate to a Pareto-dominating equilibrium in which they do a 1/2-1/2 joint lottery and use this to both play L with probability 1/2 and to both play R with probability 1/2 (this equilibrium induces all types of both players an expected payoff of 1/2).

Lemmas 3, 5, and 6 together imply that a renegotiation proof equilibrium strategy is *coordinated*.

Finally we show, in Lemma 7, that a renegotiation proof equilibrium strategy must be *balanced*. Using Lemma 4 we can, for each $m \in M_L$, redefine $\alpha_m \in [0, 1]$ as the average probability the partner plays L, conditional on the player sending message m and the partner sending a message in M_R. To show balancedness, assume to the contrary, and w.l.o.g., that $m, m' \in M_L$, and $\alpha_m > \alpha_{m'}$. This implies that all types $u < 1/2$ strictly prefer to send message m over sending message m' (as the former induces a higher probability to coordinate on their favored outcome L), and we get a contradiction to $m' \in M_L$ (i.e., to having types $u < 1/2$ sending message m'). This concludes the proof of Theorem 1.

6 On Efficiency

In this section we investigate the efficiency properties of renegotiation proof equilibria. We first argue that ex-ante efficiency (i.e., the first best) cannot be achieved by any equilibrium of any coordination game with communication. We then provide an example of an equilibrium with high ex-ante payoffs that is, however, not renegotiation proof. We then show that all renegotiation proof equilibria, while not necessarily ex-ante payoff optimal among all equilibria, are at least interim (pre communication) Pareto efficient, i.e., not interim (pre communication) Pareto dominated by any other feasible (not necessarily equilibrium) outcome of a coordination game with communication. Finally we show that at least one of the two "extreme" renegotiation proof equilibria, σ_L and σ_R, provides the highest ex-ante payoff among all coordinated equilibria and ex-ante payoff dominates any equilibrium of the coordination game without communication.

First Best The first-best ex-ante payoff can only be induced by a strategy that is coordinated and such that the coordinated outcome depends heavily on the cardinal preferences of the two agents, namely on how strong is the preference of an agent for her favorite outcome. The first best strategy is one that induces coordination on L whenever $u + v \leq 1$ and coordination on R otherwise. In other words, the first-best ex-ante payoff can only be induced by a strategy in which each agent reveals her type, and the two agents then choose the favorite outcome, L or R, of the more extreme type (i.e., the type that is farther away from $1/2$). Note that this strategy is not an equilibrium: Each player has an incentive to present a more extreme type than her real type (e.g., all types $u > 1/2$ would claim to have type 1).[7]

High payoff non-coordinated equilibria Equilibria with mis-coordination (which are not renegotiation-proof due to Theorem 1) may induce agents to credibly reveal some cardinal information about their type. This can happen if there is a message that induces higher probabilities of coordinating on the agent's favored outcome but also higher probabilities of mis-coordination compared with some other available message. Such a message can then be chosen by extreme types with a u far from $1/2$, while moderate types with u closer to $1/2$ choose the other message. Potentially, such equilibria with mis-coordination may induce a higher ex-ante payoff, if the benefit of signaling the extremeness of the type outweighs the loss due to mis-coordination. To see this consider the following concrete example.

Example 1. For simplicity we let the distribution of types F be discrete with four atoms $1/10 + \epsilon$, $1/2 - \epsilon$, $1/2 + \epsilon$, $9/10 - \epsilon$, with a probability of $1/4$ on each atom and $\epsilon > 0$ sufficiently small.[8] The game admits three babbling equilibria: always coordinating on L with an ex-ante payoff of $1/2$, always coordinating on R with an ex-ante payoff of $1/2$, and playing L if and only if the type is less than $1/2$ with an ex-ante payoff of $7/20 < 1/2$, for all ϵ sufficiently small. Theorem 1 (together with the symmetry of the distribution F) implies that when adding a single bit of communication (i.e., $|M| \geq 2$), any renegotiation-proof equilibrium strategy (and, in particular, σ_L and σ_R) induces the same expected ex-ante payoff of $3/5 > 1/2$ for any ϵ sufficiently small.

This game also has a (non-renegotiation-proof) equilibrium strategy with mis-coordination that yields a higher ex-ante payoff than the renegotiation proof payoff of $3/5$ provided the message set M has sufficiently many elements. To simplify the presentation we here allow the players to use public correlation devices to determine their joint play after sending messages, which can be approximately implemented by a sufficiently large message set. See the proof of Lemma 3 and Footnote 12 in the

[7]Note that the ex-ante efficient strategy could only really be implemented in a coordination game with communication with a continuum message set. For finite message sets it could only be approximated. For the same reasons no such approximation could be an equilibrium.

[8]One can easily adapt the example to an atomless distribution of types, in which each atom is replaced with a continuum of nearby types.

appendix. Let $m_L, m_l, m_r, m_R \in M$ and consider strategy $\sigma = (\mu, \xi)$ as follows. Let $\mu(1/10 + \epsilon) = m_L$, $\mu(1/2 - \epsilon) = m_l$, $\mu(1/2 + \epsilon) = m_r$, and $\mu(9/10 - \epsilon) = m_R$, and let $\xi(m_a, m_b) = L$ if $a, b \in \{L, l\}$, $\xi(m_a, m_b) = R$ if $a, b \in \{r, R\}$, $\xi(m_L, m_r) = L$, $\xi(m_l, m_R) = R$, $\xi(m_l, m_r)$ be a joint lottery to coordinate on L or R with probability 1/2 each, and finally let $\xi(m_L, m_R)$ be a joint lottery to coordinate on L or R with probability 3/10 each, and to play the inefficient mixed equilibrium in which each type plays her favored outcome with probability $9/10 - \epsilon$ with probability 4/10. It is straightforward to verify that, for e.g., $\epsilon = 1/100$, this strategy is indeed an equilibrium strategy with an ex-ante payoff of around 0.627 which is higher than the ex-ante payoff of 3/5 of all the renegotiation proof equilibria. This equilibrium strategy is not coordinated (also not balanced and not mutually preference consistent) and, hence, by Theorem 1 is not renegotiation proof.

Interim (pre-communication) Pareto Optimality An (ex-ante) symmetric *(type-dependent) outcome function* is a function $\phi : [0,1]^2 \to \Delta\left(\{L, R\}^2\right)$ assigning to each pair of types a possibly correlated action profile with the condition that $\phi_{u,v}(a, b) = \phi_{v,u}(b, a)$ for any $a, b \in \{L, R\}$, where $\phi_{u,v} \equiv \phi(u, v)$.[9] We thus interpret $\phi_{u,v}$ as the correlated action profile played by the two players when a player of type u meets a player of type v. Let Φ be the set of all symmetric type-dependent outcome functions.

Note that this set Φ is defined without reference to the set of available messages of the coordination game with communication. Indeed any strategy of any coordination game with communication (with any finite message set) induces an outcome function in Φ, but not all outcome functions in Φ can be generated by a strategy of a given coordination game with communication. Note also that an outcome function in Φ does not need to constitute equilibrium behavior. One can interpret the set of outcome functions Φ as the outcome functions that can be implemented by a social planner who perfectly observes the types of both players, and as a function of that can force them to play arbitrarily.

For each type $u \in [0, 1]$, let $\pi_u(\phi)$ denote the expected payoff of a player of type u under outcome function ϕ, i.e.,

$$\pi_u(\phi) = \boldsymbol{E}_v\left[(1-u)\cdot\phi_{u,v}(L, L) + u\cdot\phi_{u,v}(R, R)\right].$$

Definition 3. A strategy $\sigma \in \Sigma$ is *interim (pre communication) Pareto dominated* by a type-dependent outcome function $\phi \in \Phi$ if $\pi_u(\sigma, \sigma) \le \pi_u(\phi)$ for each type $u \in [0, 1]$ with this holding as a strict inequality for all u in a positive probability subset of $[0, 1]$. A strategy $\sigma \in \Sigma$ is *interim (pre communication) Pareto optimal* if it is not interim (pre communication) Pareto dominated by any $\phi \in \Phi$.

Proposition 1. *Every renegotiation proof strategy of a coordination game with communication (with a message set with at least two elements) is interim (pre communication) Pareto optimal.*

[9] We restrict attention to symmetric outcome functions here for two mostly aesthetic reasons. First, it makes the paper conceptually consistent, given that the subject of the paper, coordination games with communication, is a class of (ex-ante) symmetric games. Second, this prevents us from having to here introduce player subscripts which we do not need anywhere else in the paper. Proposition 1 below, however, also holds even if we allow asymmetric outcome functions.

The proof of Proposition 1, as well as all other proofs from this section, are given in Appendix B.3.

Above we have given an example of an equilibrium strategy that provides a higher ex-ante payoff than any renegotiation proof equilibrium. This strategy involved a certain degree of mis-coordination. In the following proposition we show that any equilibrium without mis-coordination, i.e., any coordinated equilibrium must have an ex-ante expected payoff that is less than or equal to the maximal ex-ante payoff of the two “extreme” renegotiation proof strategies σ_L and σ_R.

Proposition 2. *Let $\sigma \in \mathcal{E}$ be a coordinated equilibrium strategy. Then*

$$\pi(\sigma, \sigma) \leq \max\{\pi(\sigma_L, \sigma_L), \pi(\sigma_R, \sigma_R)\}.$$

Remark 3 (All-stage renegotiation-proofness à la Benoit and Krishna (1993)). One could refine the notion of renegotiation-proofness to allow agents to renegotiate for a Pareto-improving equilibrium additionally also in earlier stages: in the interim stage before observing the realized messages induced by the original equilibrium, and in the ex-ante stage before each agent observes his own type. This more restrictive definition of renegotiation-proofness à la Benoit and Krishna (1993) would: call the strategies satisfying our definition of renegotiation proofness *ex-post renegotiation-proof* strategies; say that an ex-post renegotiation-proof strategy is *interim renegotiation-proof* if it is not Pareto dominated (in the original game after each agent observes his own type, yet before observing the realized message profile) by any ex-post renegotiation-proof strategy; and say that an interim renegotiation-proof strategy is *all-stage renegotiation-proof* if there is no other interim renegotiate-proof strategy that induces a higher ex-ante expected payoff to both players (before each player knows his own type).

It is immediate that any ex-post renegotiation-proof strategy that satisfies interim Pareto optimality is interim renegotiation-proof strategy. Proposition 2 implies that either σ_L or σ_R maximizes the ex-ante payoff among all interim renegotiation-proof strategies, which immediately implies that either σ_L or σ_R is an all-stage renegotiation-proof strategy. Moreover, if $\pi(\sigma_R, \sigma_R) \neq \pi(\sigma_L, \sigma_L)$, then one can show that either σ_L or σ_R is the unique coordinated strategy that maximizes the ex-ante payoff, which implies that it is the unique all-stage renegotiation-proof strategy.

Next, we show that σ_L or σ_R provides a strictly higher ex-ante expected payoff than any equilibrium of the game without communication (and thus any babbling equilibrium of the game with communication).

Recall from Remark 2 and the text preceding it that in the coordination game without communication any equilibrium is characterized by a cut-off value $x \in [0, 1]$ such that $x = F(x)$ with the interpretation that types $u \leq x$ play L and types $u > x$ play R.

Let $\pi_u(x, x')$ denote the payoff of an agent with type u who follows a strategy with cut-off x and faces a partner of unknown type who follows a strategy with cut-off x', which is given by

$$\pi_u(x, x') = \mathbf{1}_{u \leq x} \cdot F(x') \cdot (1-u) + \mathbf{1}_{u > x} \cdot F(x') \cdot u,$$

and let $\pi(x, x') = E_u(\pi_u(x, x'))$ be the ex-ante expected payoff of an agent who follows strategy x and faces a partner who follows x'.

The following result shows that either σ_L or σ_R strictly improves the ex-ante payoff.

Proposition 3. *Let x be an equilibrium strategy cutoff in the coordination game without communication. Then*

$$\pi(x, x) < \max\{\pi(\sigma_L, \sigma_L), \pi(\sigma_R, \sigma_R)\}.$$

7 Evolutionary Stability

A common interpretation of a Nash equilibrium is a convention that that is reached as a result of a process of social learning when similar games are repeatedly played within a large population. Specifically, consider a population in which a pair of agents from a large population are occasionally randomly matched and play the coordination game with communication $\langle \Gamma, M \rangle$. The agents can observe past behavior of other agents playing similar games in the past. It seems plausible that the aggregate behavior of the population would gradually converge into a self-enforcing convention, which is a symmetric Nash equilibrium of $\langle \Gamma, M \rangle$ (see, e.g., Weibull, 1995; Sandholm, 2010 for textbook introduction).

Arguably, for such a convention to be a reasonable long run prediction of the underlying game, one should require that the convention is robust to small perturbations in the behavior of the population. In this section we show that both σ_L and σ_R satisfy three properties that imply robust to various perturbations (and the results can be extended to renegotiation proof equilibria such as σ_C):

1. Neutral stability (à la Maynard Smith and Price, 1973, and evolutionary stability if $|M| = 2$). This implies that σ_L and σ_R are robust to the presence of a few experimenting agents who behave differently than the rest of the population.

2. The action function μ^* is weakly dominant (taking as given the second-stage behavior induced by σ_L and σ_R), and it is strictly dominant if $|M| = 2$). This implies that σ_L and σ_R are robust to any perturbation (including, large perturbations) that changes the first-stage behavior in the population.

3. The message function is a neighborhood invader (Apaloo, 1997; Cressman, 2010, which refines the notion of continuous stability à la Eshel and Motro, 1981) in any second-stage induced game. This implies that σ_L and σ_R are robust to any sufficiently small perturbation that changes the

second-stage behavior of agents.[10]

As the formal analysis of these robustness properties is somewhat lengthy, we relegate it to Appendix A.

8 Extensions

In this section we show that our results can be extended to various setups that allow having, multiple rounds of communication, more than two players, a multi-dimensional set of types, extreme types with dominant actions, or more than two actions. The first three of these five extensions are straightforward and leave the main results and their proofs essentially unchanged. We indicate where this is not the case. These extensions are described verbally only. For the latter two extensions results and their proofs change more substantially and are therefore provided.

8.1 Multiple Rounds of Communication

Consider a variant of the coordination game with communication in which agents have fixed and finite number of rounds of communication. In each such round of this communication phase they simultaneously send messages from the set of messages M. Players observe messages after each round and can, thus, condition their message choice and then their final action choice on the history of observed message pairs up to the point in time where they take their message or action decision. Renegotiation then possibly takes place once at the end of this communication phase but before the final action choices are made. Straightforward adaptations of Theorem 1, as well as Proposition 1, 3, and 2 then hold in this setting.

Thus, regardless of the length of the pre-play communication, agents only reveal their preferred outcome (but nothing about the strength of this preference), and additional rounds of communication cannot improve the ex-ante expected payoff relative to the payoff induced by a single round of communication with a binary message.

8.2 More Than Two Players

Consider a variant of the coordination game in which there are (a finite number of) three or more players who play a symmetric coordination game (with private values) with pre-play communication. Before players choose actions, they simultaneously send a message each from a finite set of messages.

[10] We further conjecture that σ_L and σ_R satisfy a slightly weaker form of evolutionary robustness à la Oechssler and Riedel (2002).

In such a setting the appropriate version of Theorem 1 still holds: renegotiation proof equilibrium strategies are those that are ordinal preference revealing, mutual preference consistent, coordinated, and balanced. Also appropriate versions of Propositions 1 and 3 hold: renegotiation proof equilibrium strategies are interim (pre communication) Pareto undominated and are ex-ante payoff improving over all babbling equilibria (equilibria of the game without communication).

Proposition 2 does not extend to this setting: with three players, for instance, for some distributions of values F, the strategy that determines the fallback option by majority vote (in case of messages that indicate different preferred actions) is an ex-ante payoff improvement over a simple (and in the multi-player setting more radical) fallback norm of choosing say action L in every case of disagreement.

8.3 Multi-Dimensional Set of Types

In our model we made the simplifying assumption that mis-coordination provides the same (normalized to zero) payoff to both players. Now consider the following multi-dimensional set of types. Let $U \subseteq [0,1]^4$ be the set of payoff matrices, with u_{ab} the payoff if a player chooses action a while her opponent chooses action b, for a two-dimensional two-action coordination game:

$$U = \{(u_{LL}, u_{LR}, u_{RL}, u_{RR}) \,|\, \min\{u_{LL}, u_{RR}\} > \max\{u_{RL}, u_{LR}\}\}.$$

That is, each payoff is between zero and one, and all types strictly prefer coordination on the same action as the partner (and obtaining at least $\min\{u_{LL}, u_{RR}\}$) over mis-coordination (which yields at most $\max\{u_{RL}, u_{LR}\}$). Given a type $u = (u_{LL}, u_{LR}, u_{RL}, u_{RR})$, let $\phi_u \in [0,1]$ denote type u's *left-tendency*, which is the probability of the opponent playing L that induces an agent of type u to be indifferent between the two actions:

$$\phi_u = \frac{u_{RR} - u_{RL}}{u_{LL} + u_{RR} - u_{RL} - u_{LR}}.$$

Observe that an agent with left-tendency ϕ_u prefers to play L (R) if her partner plays L with probability larger (smaller) than ϕ_u. Thus, the left-tendency ϕ_u replaces what we denoted by u in the main model.

Straightforward adaptations of the proofs of the baseline model show that our main results, Theorem 1 as well as Propositions 1, 2, and 3 extend to this setup.

8.4 Extreme Types with Dominant Actions

In this subsection we show how to extend our analysis to a setup in which some types have an extreme preference in favor of one of the actions, such that it becomes a dominant action for them.

Let $a < 0$ and $b > 1$. We extend the set of types to be the interval $[a, b]$. Observe that action L (R) is a dominant action for any type $u < 0$ ($u > 1$) as coordinating on R (L) yields such a type a negative payoff of $u < 0$ ($1 - u < 0$). We call types with a dominant action (i.e., $u < 0$ or $u > 1$) *extreme*, and types that do not have a strictly dominant action (i.e., $u \in [0, 1]$) *moderate*. We assume that the cumulative distribution of types F is continuous (atom-less) and has a support $[a, b]$.

To simplify the presentation of the results we assume that

$$\alpha \equiv \frac{F(0)}{F(0) + (1 - F(1))} = \frac{k}{n}.$$

is a rational number (i.e., k and n are natural numbers).[11] Observe that $\alpha = 1/2$ in the case of a symmetric distribution F (i.e., $F(0) = 1 - F(1)$). We further assume that the extreme types are a minority both among the agents who prefer action R and among the agents who prefer action L, i.e.,

$$F(0) < \frac{1}{2} \cdot F\left(\frac{1}{2}\right), \quad \text{and} \quad 1 - F(1) < \frac{1}{2} \cdot \left(1 - F\left(\frac{1}{2}\right)\right). \tag{1}$$

A simple way to adapt $\sigma_L = (\mu^*, \xi_L)$ to the setup with extreme types is to adapt ξ_L by having extreme types following their dominant action in the second stage (and moderate types play in the same way as in the baseline model). In what follows we show that σ_L is no longer an equilibrium strategy with extreme types (the argument why σ_R is no longer an equilibrium strategy is analogous). Observe that sending message m_R by a moderate type leads to coordination (on R or L depending on the opponent message) with probability one, while sending message m_L leads to coordination (on L) only with probability $(F(1)) < 1$.

This implies that types $u < 1/2$ sufficiently close to $1/2$ strictly prefer to send message m_R over m_L (as the former induces a higher probability to coordinate the same action as the partner), which contradicts σ_L being an equilibrium strategy.

In what follows we show that strategy σ_α, according to which moderate players with different preferred outcomes coordinate on L with probability α, is a renegotiation-proof equilibrium strategy. Formally, assume that $|M| \geq 2 \cdot n$. Denote $2n$ distinct messages as $\{m_{L,1}, ..., m_{L,n}, m_{R,1}, ..., m_{R,n}\} \in M$, where we interpret $m_{L,i}$ ($m_{R,i}$) as expressing a preference for L (R) and choosing at random the number i from the set of numbers $\{1, ..., n\}$ in the joint lottery described below. We arbitrarily interpret any message $m \in M \backslash \{m_{L,1}, ..., m_{L,n}, m_{R,1}, ..., m_{R,n}\}$ as equivalent to $m_{L,1}$ (i.e., any such message is interpreted as having a preference for L and randomly choosing 1). Given message $m \in M$, let $i(m)$ denote its associated random number, e.g., $i(m_{L,j}) = j$. Let $M_R = \{m_{R,1}, ..., m_{R,n}\}$ and $M_L = M \backslash M_R$.

[11] In order to deal with irrational α-s one needs either to slightly weaken the results below to show that there exists a renegotiation-proof ϵ-equilibrium strategy (in which each type of each player gains at most ϵ from deviating) for any $\epsilon > 0$, or to allow an infinitive set of messages or a continuous “sunspot”.

Let $\sigma_\alpha = (\mu_\alpha, \xi_\alpha)$ be defined as follows:

$$\mu_\alpha (u) = \begin{cases} \frac{1}{n} \cdot m_{L,1} + ... + \frac{1}{n} \cdot m_{L,n} & u \leq \frac{1}{2} \\ \frac{1}{n} \cdot m_{R,1} + ... + \frac{1}{n} \cdot m_{R,n} & u > \frac{1}{2}. \end{cases}$$

Thus, the first-stage strategy μ_α induces each agent to reveal whether his preferred outcome is L or R, and to uniformly choose a number between 1 and n. In the second stage, if both agents share the same preferred outcome they play it. Otherwise, moderate types coordinate on L if the sum of their random numbers modulo n is at most k, and coordinate on R otherwise. Extreme types play their strictly dominant action. Formally:

$$\xi_\alpha (m, m') = \begin{cases} 0 & (m, m') \in M_R \times M_R \text{ or } ((m, m') \notin M_L \times M_L \text{ and } ((i(m) + i(m')) \, mod \, n) > k) \\ 1 & \text{otherwise.} \end{cases} \tag{2}$$

Next we show that σ_α is a renegotiation-proof equilibrium strategy. Formally,

Proposition 4. *σ_α is a renegotiation-proof equilibrium strategy.*

Proof. In what follows, we present new arguments in the proof: (1) the argument why μ_α is a best reply against σ_α in the first stage, and (2) the argument why ξ_α is a best reply against σ_α in the second stage. All other arguments in the proof are very similar to the arguments in the proof of Theorem 1, are omitted for brevity.

We begin by showing that μ_α is a best reply against σ_α in the first stage. The probability of mis-coordination conditional on the agent sending message $m \in M_R$ is equal to the probability of $F(0)$ that the partner has an extreme "left" type times the probability of $(1 - \alpha)$ that the joint lottery induces moderate players to play R. Similarly, the probability of mis-coordination conditional on the agent sending message $m \in M_L$ is equal to the probability of $1 - F(1)$ that the partner has an extreme "right" type times the probability of α that the joint lottery induces moderate players to play L. These two probabilities of mis-coordination coincide due to the definition of α in (2) above:

$$F(0) \cdot (1 - \alpha) = (1 - F(1)) \cdot \alpha \tag{3}$$

$$\Leftrightarrow \alpha = \frac{F(0)}{F(0) + (1 - F(1))}$$

This implies that (1) type 1/2 is indifferent between sending message in M_R and sending a message in M_L as both messages induce the same probability of mis-coordination, and (2) any type $u < 1/2$ ($u > 1/2$) strictly prefers to send a message in M_L (M_R) as sending such a message induces a higher probability of coordination on his preferred outcome L (R). This implies that μ_α is a best reply against σ_α in the first stage.

Next we show ξ_α is a best reply against σ_α in the second stage. Let $m, m' \in \text{supp}(\mu_\alpha)$. Consider the induced game $\Gamma(F_m, F_{m'})$. If $(m, m') \in M_R \times M_R$ ($(m, m') \in M_L \times M_L$), then a partner who follows σ_α is going to play R (L) for sure, which implies that following $\xi_\alpha(m, m')$ and playing R (L) is the best reply against σ_α in the induced game $\Gamma(F_m, F_{m'})$. If $(m, m') \in M_R \times M_L$ and $((i(m) + i(m'))\ mod\, n) > k$ ($(m, m') \in M_L \times M_R$ and $((i(m) + i(m'))\ mod\, n) \le k$), then a moderate partner (who sent message m') who follows σ_α is going to play R (L). Eq. (1) implies that the probability that the partner is moderate is strictly more than $1/2$, which, in turn, implies that following $\xi_\alpha(m, m')$ and playing R (L) is the best reply against σ_α in the induced game $\Gamma(F_m, F_{m'})$ (as playing R (L) yields the player an expected payoff of at least $\frac{u}{2} \ge \frac{0.5}{2}$ ($\frac{1-u}{2} \ge \frac{0.5}{2}$), while playing L (R) yields the player a strictly smaller expected payoff that is strictly less than $\frac{1-u}{2} \le \frac{0.5}{2}$ ($\frac{u}{2} \le \frac{0.5}{2}$). □

Minor adaptations to the proof of Theorem 1 show that the result applies also to the current setup with extreme type. Note that in the current setup with extreme types, one should interpret the property of a strategy being *coordinated* (namely, that $F_m(\xi(m, m')) = F_{m'}(\xi(m', m)) \in \{0, 1\}$ for any $m, m' \in \text{supp}(\mu^*)$) as showing that there is no mis-coordination among non-extreme types; there might be mis-coordination only in matches in which one of the players is an extreme type, and the partner has the opposite preferred outcome.

8.5 Coordination Games with More Than 2 Actions

In this subsection we extend our main result to coordination games with more than two actions.

Adaptation of the Model Let Γ_n be a coordination game in which agents simultaneously send messages from a finite set M ($2 \cdot n \le |M| < \infty$), and then simultaneously choose an action from the ordered set $A = (a_1, ..., a_n)$ with $2 < n < \infty$. The type of an agent $u = (u_1, ..., u_n)$ is an element of $[0, 1]^n$, where we interpret the k-th component u_k as the payoff of the agent if both players choose action a_k. If the players choose different actions (mis-coordinate), then they both get a payoff of zero. We assume that the distribution of types is a continuous (atomless) distribution with a support of $[0, 1]^n$. Let f denote its density. For each action a_i,let $p(a_i)$ be the probability that the preferred action of a random agent is a_i (i.e., the probability that $u_i = \max(\{u_1, .., u_n\})$).

A player's (ex-ante) strategy in Γ_n is a pair $\sigma = (\mu, \xi)$, where $\mu : U \to \Delta(M)$ is a *message function* that describes which (possibly random) message is sent for each possible realization of the agent's type, and $\xi : M \times M \times U \to \Delta(A)$ is an *action function* that describes the distribution of actions chosen as a function of the agent's type and the observed message profile; that is, when an agent of type u who follows strategy (μ, ξ) observes a message profile (m, m') (message m sent by the agent, and message m' sent by the opponent), then the agent plays action a_k with probability $\xi_u(m, m')(a_k)$. Let Σ_n denote the set of all strategies of the game Γ_n.

The notions of equilibrium strategy, the induced game $\Gamma_n^{m,m'}$, the induced game with additional communication $\Gamma_n^{m,m',M'}$, and renegotiation-proof equilibrium strategy are adapted to the current setup in a straightforward way.

Strategies σ_L / σ_R are not Equilibrium Strategies We begin by demonstrating that there is no simple way to adapt σ_L and σ_R to the setup with more than 2 actions. Example 2 shows that a simple adaptation of σ_L does not constitute an equilibrium strategy with more than two actions.

Example 2 (Simple adaptation of σ_L / σ_R is not an equilibrium strategy)**.** Consider the following strategy $\tilde{\sigma} = \left(\tilde{\mu}, \tilde{\xi}\right)$, which is a simple adaptation of σ_L / σ_R to the setup with more than two actions. Fix n distinct messages $m_1, m_2, ..., m_n \in M$. We begin by defining the action function $\tilde{\mu}$: each agent sends his preferred outcome (with an arbitrary tie-breaking rule of choosing the action with the smaller index in case of multiple preferred outcomes). Formally,

$$\tilde{\mu}(u_1, ..., u_n) = \begin{cases} m_1 & u_1 = \max(\{u_1, .., u_k\}) \\ ... & ... \\ m_k & u_k = \max(\{u_1, .., u_k\}) > \max(\{u_1, .., u_{k-1}\}) \\ ... & ... \\ m_n & u_n = \max(\{u_1, .., u_n\}) > \max(\{u_1, .., u_{n-1}\}) \end{cases}$$

Thus, the first-stage strategy μ^* induces each agent to reveal his preferred outcome. In the second stage both agents play the preferred action with the smaller index, i.e.,

$$\tilde{\xi}(m_k, m_l) = a_{\min(k,l)}.$$

That is, the |Fallback" norm, in cased the agents have different preferred outcomes, is to play the outcome with the smaller index. We conclude the example by showing that $\tilde{\sigma}$ is not an equilibrium strategy. Consider an agent with type u that satisfies: (1) coordinating on a_n is the agent's preferred outcome, i.e., $u_n > \max(\{u_1, .., u_{n-1}\})$, (2) coordinating on a_1 is only slightly less preferred than coordinating on a_n, i.e., $u_n - u_1 << 1$, and (3) coordinating on any other action yields a very low payoff, i.e., $\max(u_2, ..., u_{n-1}) << 1$. Such an agent obtains a payoff of

$$\sum_{i \leq n} p(a_i) \cdot u_i \approx p(a_1) \cdot u_1 + p(a_n) \cdot u_n \approx (p(a_1) + p(a_n)) \cdot u_1$$

when following strategy $\tilde{\sigma}$ (because strategy $\tilde{\sigma}$ induces both players to coordinate on the partner's preferred outcome). The agent can obtain a strictly higher payoff of u_1 by deviating to sending message m_1,

σ_C is a Renegotiation-Proof Equilibrium Strategy Next we show that a simple adaptation of σ_C remains a renegotiation-proofness equilibrium strategy also with with more than two actions.

Fix $2 \cdot n$ distinct messages $m_1^0, m_2^0, ..., m_n^0, m_1^1, m_2^1, ..., m_n^1 \in M$, where we interpret m_i^j as indicating that the agent's preferred outcome is the i-th outcome, and that the random binary number chosen by the agent is $j \in \{0,1\}$. Let $\sigma_C = (\mu_C, \xi_C)$ be extended to the current setup as follows:

$$\mu_C(u_1,...,u_n) = \begin{cases} \frac{1}{2} \cdot m_1^0 + \frac{1}{2} \cdot m_1^1 & u_1 = \max(\{u_1,..,u_k\}) \\ ... & ... \\ \frac{1}{2} \cdot m_k^0 + \frac{1}{2} \cdot m_k^1 & u_k = \max(\{u_1,..,u_k\}) > \max(\{u_1,..,u_{k-1}\}) \\ ... & ... \\ \frac{1}{2} \cdot m_n^0 + \frac{1}{2} \cdot m_n^1 & u_n = \max(\{u_1,..,u_n\}) > \max(\{u_1,..,u_{n-1}\}) \end{cases}$$

Thus, the first-stage strategy μ_C induces each agent to reveal his preferred outcome, and to uniformly choose a binary number (either, zero or one). In the second stage, if both agents share the same preferred outcome they play it. Otherwise, they coordinate on the preferred action with the smaller index if both agent have chosen the same random number, and they coordinate on the preferred outcome with the larger index it the agents have chosen different random numbers. Formally:

$$\xi_C\left(m_k^i, m_l^j\right) = \begin{cases} a_k & (k \leq l \text{ and } i = j) \;\; OR \;\; (k \geq l \text{ and } i \neq j) \\ a_l & \text{otherwise.} \end{cases}$$

Our first result shows that σ_C is a renegotiation-proof equilibrium strategy.

Proposition 5. *σ_C is a renegotiation-proof equilibrium strategy.*

Sketch of proof. In what follows we show why the first-stage behavior is a best reply against σ_C. The remaining arguments why the second-stage behavior is a best reply, and why σ_C is renegotiation-proof are analogous to the proof of Theorem 1 and are omitted for brevity.

Observe that an agent who sends message m_k^i obtains an expected payoff of $1/2 \cdot u_k + 1/2 \cdot \sum_{l \leq n} p(a_l) \cdot u_l$ when facing a partner who follows strategy σ_C. This implies that the first-stage best reply of an agent for which the preferred outcome is the k-th outcome, is to send either message m_k^0 or m_k^1, which, in turn, implies that following the message function μ_C is indeed a first-stage best reply of the agents. $\square$

Necessary Conditions for Renegotiation-Proofness Our final result adapts Theorem 1 to the current setup with more than two actions. Specifically, we show that any renegotiation-proof equilibrium strategy, satisfies the first the following three properties:

1. Each agent reveals his preferred outcome.

2. Agents always coordinate on a jointly preferred outcome (if such an outcome exist)

3. The agents never mis-coordinate.

   The condition of balancedness is replaced with the following condition:

4. For each subset of actions $A' \subseteq A$ and each action $a_i \in A'$, there is a message $m_i \in M_i$ that maximizes the probability that the coordinated outcome would be in A'. (Observe that this property is satisfied by σ_C, while it is not satisfied by σ_L and σ_R.)

The intuition behind part (4) is as follows. Part (1) implies that all types must reveal their preferred outcome in a renegotiation-proof strategy. A type that is almost indifferent between coordinating on any action in A', while disliking coordinating on any action in $A \backslash A'$, will reveal his most preferred outcome, only if this will not decrease the probability of coordinating on one of the actions in A'. Formally:

Theorem 2. *Let $\sigma = (\mu, \xi)$ be a renegotiation-proof equilibrium strategy. Then the following statements are true.*

1. Each agent reveals his preferred outcome*: There exist n non-empty, disjoint, and exhaustive subsets of supp(μ) denoted by $M_1, ..., M_n$ (i.e., supp(μ) $= \dot{\bigcup} M_i$) such that for each $i \leq n$: $u_i > \max\left(\{u_1, .., u_n\} \setminus \{u_i\}\right)$ implies that $\mu_u(m) = 0$ for each $m \notin M_i$.*

2. Agents always coordinate on a jointly preferred outcome *(if such an outcome exist): If $m, m' \in M_i$ for some $i \leq n$, then $\xi(m, m', u) = a_i$ for each type $u \in U$ satisfying $\mu_u(m) > 0$. In this case we omit the parameter u and write $\xi(m, m') = a_i$.*

3. Agents never miscoordinate*: $\xi(m, m', u) = \xi(m, m', u') \in A$ for any $m, m' \in$ supp(μ) and any pair of types $u, u' \in U$ satisfying $\mu_u(m) > 0$ and $\mu_{u'}(m') > 0$.*

4. *For each subset of actions $A' \subseteq A$ and each action $a_i \in A'$, there is a message $m_i \in M_i$, such that*
$$\sum_{m' \in M} \mu(m') \cdot \mathbf{1}_{\xi(m', m_i) \in A'} \geq \sum_{m' \in M} \mu(m') \cdot \mathbf{1}_{\xi(m', m) \in A'}$$
for each message $m \in M$.

Proof. Let σ be a renegotiation-proof equilibrium strategy. The proofs of parts (1–3) are very similar to the proofs of the analogous parts of Theorem 1 and are omitted for brevity. In what follows we prove part (4). Assume to the contrary that there exists an action $a_i \in A' \subseteq A$ and a message $m^* \notin M_i$ such that

$$\sum_{m' \in M} \mu(m') \cdot \mathbf{1}_{\xi(m', m_i) \in A'} < \sum_{m' \in M} \mu(m') \cdot \mathbf{1}_{\xi(m', m^*) \in A'} \equiv p^*$$

for each $m_i \in M_i$. Let p_i be the maximal probability of coordination on an action in A' conditional on the agent sending a message from M_i:

$$p_i = \max_{m_i \in M_i} \left(\sum_{m' \in M} \mu(m') \cdot \mathbf{1}_{\xi(m', m_i) \in A'} \right).$$

Note that $p_i < p^*$. Let $0 < \epsilon << 1$ be sufficiently small such that $p_i \cdot u_i + \epsilon < p^* \cdot u_i - \epsilon$. Consider a type $\hat{u} = (u_1, ..., u_n)$ satisfying: (1) $u_i > \max(\{u_1, .., u_n\} \setminus \{u_i\})$, (2) $u_i - u_j < \epsilon$ for each action $a_j \in A'$, (3) $u_k < \epsilon$ for each $a_k \notin A'$. Due to the property proved in part (2) an agent with type $\hat{u}$ who follows σ has to send a message from M_i. This implies that his payoff would be at most $p_i \cdot u_i + \epsilon$. By contrast, deviating to sending the message m^* yield the agent a strictly higher payoff of $p^* \cdot u_i - \epsilon$, contradicting σ being a renegotiation-proof equilibrium strategy. □

9 Related Literature

While our notion of renegotiation proofness may not be exactly the same as any notion in the literature as far as we can see, it is however very much inspired by previous notions and, we feel, simply appropriately adapted to the problem at hand. At the heart of renegotiation proofness is the idea that people will find it hard to ignore obvious Pareto-improving alternatives to any suggestions as to how they should behave or what choices they should (collectively) make. When people are forward looking, their anticipation of revisions of plans that are not Pareto-efficient can constrain possible equilibrium behavior. Renegotiation proofness concepts have been developed in the context of infinitely and finitely repeated games with complete information in e.g., Farrell and Maskin (1989), Van Damme (1989), Bernheim and Ray (1989), Evans and Maskin (1989), and in e.g., Benoit and Krishna (1993) and Wen (1996), respectively. There is a sizable literature on the renegotiation proofness of contracts in the presence of asymmetric information possibly starting with Hart and Tirole (1988) and Dewatripont (1989). See also Maestri (2017) and Strulovici (2017) for more recent contributions in this area.

When there is incomplete information analysts have to be careful as to what information people have when they renegotiate. In a seminal article Holmström and Myerson (1983) distinguish between three possible cases: the ex-ante stage before people even have their own private information, the interim stage at which people know their private information but nothing else, and the ex-post stage at which people know everything. They then call a decision rule (a mechanism - one could also think of it as a strategy profile) “durable” if it is immune to renegotiation at any of these stages.

The closest concept to ours may be that of posterior efficiency of Forges (1994). Forges (1994) argues that the final outcome of a mechanism (or here strategy profile) will not necessarily fully reveal all initially privately held information. Posterior efficiency then only demands that the outcome be efficient given the information that the people can infer from the outcome of the mechanism alone. Similarly we here demand that the strategy profile prescribes an action profile after messages are sent that is efficient given the information revealed by the messages sent. Thus, our players have more information than just the prescribed action profile as they in fact observe also the messages sent which in our case typically provide additional information. See also Kawakami (2016, page 897) on this point. Our definition of renegotiation proofness necessarily also differs from Forges (1994) posterior efficiency in that in our strategic setting we impose the additional (sequential rationality) requirement that agents play an equilibrium action profile given the information they have. The domain of problems in Forges (1994) is the domain of Bayesian collective choice problems, in which agents only choose which message to send but do not make any other strategic choices. The contribution of our paper is not the possibly slightly novel solution concept, but the characterization of this, we believe in our context most appropriate solution concept, for our particular problem of communication in coordination game with private values.

Another "path" to our paper is the literature on costless pre-play communication, especially when paired with an evolutionary analysis. There are various strands of literature on this topic, but perhaps the most germane is the literature that started with studying costless pre-play communication before players engage in a complete information coordination game by Robson (1990) (see also earlier the related notion of "green beard effect" in Hamilton, 1964; Dawkins, 1976). This has spurred a sizable literature including Sobel (1993), Blume, Kim, and Sobel (1993), Wärneryd (1993), Kim and Sobel (1995), Bhaskar (1998), and Hurkens and Schlag (2003). Simplifying, the general insights are as follows. Suppose that the game has two Pareto-rankable equilibria. Then the Pareto-inferior equilibrium is not evolutionary stable as it can be invaded by mutants who use a previously unused message as a secret handshake: if their opponent does not use the same handshake they simply play the inferior Pareto equilibrium (as do all incumbents), but if their opponent also uses the secret handshake they both play the Pareto-superior equilibrium. Suppose the game has an equilibrium that is not Pareto-dominated by another equilibrium but is Pareto-dominated by some non-equilibrium strategy profile. Then the same argument would suggest that the so Pareto-dominated equilibrium is unstable, yet the mutant strategy profile - by virtue of not being an equilibrium - is itself also unstable. To avoid this one can appeal to the notion of "robustness to equilibrium entrants" introduced by Swinkels (1992) that only those mutants are considered that when they play against each other mutants play an equilibrium. Then, for instance, everyone defecting is the unique strategy in the prisoner's dilemma that is stable with respect to equilibrium entrants (mutants). Our notion of renegotiation proofness has a similar flavor as we request that a renegotiation proof equilibrium is both an equilibrium and efficient among equilibria. Given the incomplete information in our model we believe that the most appropriate place for applying this idea is at the stage after messages are sent, essentially at the posterior stage in the language of Forges (1994) as explained above. Recently, Newton (2017) shows an evolutionary foundation for the agents developing the ability to renegotiate into a Pareto-better outcome ("collaboration" in the terminology of Newton (2017)).

Another related literature deals with stable equilibria in coordination games with private values, but without pre-play communication. Sandholm (2007) (extending earlier results of (Fudenberg and Kreps, 1993; Ellison and Fudenberg, 2000)) shows that mixed Nash equilibria of the game with complete information can be purified in the sense of Harsanyi (1973) in an evolutionary stable way (see also 2). Finally, two related papers analyse stag-hunt games with private values. Baliga and Sjöström (2004) show that introducing pre-play communication induces a new equilibrium in which the Pareto-dominant action profile is played with a high probability. In a recent paper Jelnov, Tauman, and Zhao (2018) show that in some cases a small probability to have another interaction can have a substantial influence of the set of equilibrium outcomes in the Stag Hunt games with private values.

10 Conclusion

TBD

References

Apaloo, J. (1997): "Revisiting strategic models of evolution: The concept of neighborhood invader strategies," *Theoretical Population Biology*, 52(1), 52–71.

Aumman, R. J., M. B. Maschler, and R. Stearns (1968): *Repeated Games of Incomplete Information: An Approach to the Non-zero-sum Case*, Report to the U.S. Arms Control and Disarmament Agency. Mathematica Inc., Princeton (reprinted in Chapter 5 of *Repeated games with incomplete information*, 1995, MIT Press).

Baliga, S., and T. Sjöström (2004): "Arms races and negotiations," *The Review of Economic Studies*, 71(2), 351–369.

Banerjee, A., and J. W. Weibull (2000): "Neutrally stable outcomes in cheap-talk coordination games," *Games and Economic Behavior*, 32, 1–24.

Belleflamme, P., and F. Bloch (2004): "Market sharing agreements and collusive networks," *International Economic Review*, 45(2), 387–411.

Benoit, J.-P., and V. Krishna (1993): "Renegotiation in finitely repeated games," *Econometrica: Journal of the Econometric Society*, pp. 303–323.

Bernheim, B. D., and D. Ray (1989): "Collective dynamic consistency in repeated games," *Games and Economic Behavior*, 1(4), 295–326.

Bhaskar, V. (1998): "Noisy communication and the evolution of cooperation," *Journal of Economic Theory*, 82(1), 110–131.

Blume, A., Y.-G. Kim, and J. Sobel (1993): "Evolutionary stability in games of communication," *Games and Economic Behavior*, 5(4), 547–575.

Cramton, P., and J. A. Schwartz (2000): "Collusive bidding: Lessons from the FCC spectrum auctions," *Journal of regulatory Economics*, 17(3), 229–252.

Cressman, R. (2010): "CSS, NIS and dynamic stability for two-species behavioral models with continuous trait spaces," *Journal of theoretical biology*, 262(1), 80–89.

Dawkins, R. (1976): *The Selfish Gene*. Oxford University Press, New York.

Dewatripont, M. (1989): "Renegotiation and information revelation over time: The case of optimal labor contracts," *The Quarterly Journal of Economics*, 104(3), 589–619.

Ellison, G., and D. Fudenberg (2000): "Learning purified mixed equilibria," *Journal of Economic Theory*, 90(1), 84–115.

Eshel, I. (1983): "Evolutionary and continuous stability," *Journal of theoretical Biology*, 103(1), 99–111.

Eshel, I., and U. Motro (1981): "Kin selection and strong evolutionary stability of mutual help," *Theoretical population biology*, 19(3), 420–433.

Evans, R., and E. Maskin (1989): "Efficient renegotiation-proof equilibria in repeated games," *Games and Economic Behavior*, 1(4), 361–369.

Farrell, J., and E. Maskin (1989): "Renegotiation in repeated games," *Games and economic behavior*, 1(4), 327–360.

Forges, F. (1994): "Posterior efficiency," *Games and Economic Behavior*, 6(2), 238–261.

Fudenberg, D., and D. M. Kreps (1993): "Learning mixed equilibria," *Games and Economic Behavior*, 5(3), 320–367.

Goffman, E. (1971): *Relations in public: Microstudies of the social order*. Basic Books.

Hamilton, W. D. (1964): "The genetical evolution of social behaviour. I," *Journal of Theoretical Biology*, 7(1), 1–16.

Harsanyi, J. C. (1973): "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points," *International journal of game theory*, 2(1), 1–23.

Hart, O. D., and J. Tirole (1988): "Contract renegotiation and Coasian dynamics," *The Review of Economic Studies*, 55(4), 509–540.

Heller, Y. (2010): "Minority-proof cheap-talk protocol," *Games and Economic Behavior*, 69(2), 394–400.

Holmström, B., and R. B. Myerson (1983): "Efficient and durable decision rules with incomplete information," *Econometrica: Journal of the Econometric Society*, pp. 1799–1819.

Hurkens, S., and K. H. Schlag (2003): "Evolutionary insights on the willingness to communicate," *international Journal of game theory*, 31(4), 511–526.

Jelnov, A., Y. Tauman, and c. Zhao (2018): "Stag Hunt with Unknown Outside Options," *mimeo*.

Katz, M. L. (1986): "An analysis of cooperative research and development," *The RAND Journal of Economics*, pp. 527–543.

Kawakami, K. (2016): "Posterior renegotiation-proofness in a two-person decision problem," *International Journal of Game Theory*, 45(4), 893–931.

Kim, Y.-G., and J. Sobel (1995): "An evolutionary approach to pre-play communication," *Econometrica*, 63(5), 1181–1193.

Maestri, L. (2017): "Dynamic contracting under adverse selection and renegotiation," *Journal of Economic Theory*, 171, 136–173.

Maynard Smith, J. (1974): "The theory of games and the evolution of animal conflicts," *Journal of Theoretical Biology*, 47(1), 209–221.

Maynard Smith, J., and G. R. Price (1973): "The logic of animal conflict," *Nature*, 246, 15.

Motta, M. (2004): *Competition policy: theory and practice*. Cambridge University Press.

Newton, J. (2017): "Shared intentions: The evolution of collaboration," *Games and Economic Behavior*, 104, 517–534.

Oechssler, J., and F. Riedel (2002): "On the dynamic foundation of evolutionary stability in continuous models," *Journal of Economic Theory*, 107(2), 223–252.

Robson, A. J. (1990): "Efficiency in evolutionary games: Darwin, Nash, and the secret handshake," *Journal of Theoretical Biology*, 144(3), 379–396.

Sandholm, W. H. (2007): "Evolution in Bayesian games II: Stability of purified equilibria," *Journal of Economic Theory*, 136(1), 641–667.

——— (2010): "Local stability under evolutionary game dynamics," *Theoretical Economics*, 5(1), 27–50.

Sobel, J. (1993): "Evolutionary stability and efficiency," *Economics Letters*, 42(2-3), 301–312.

Strulovici, B. (2017): "Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts," *Econometrica*, 85(2), 585–616.

Swinkels, J. M. (1992): "Evolutionary stability with equilibrium entrants," *Journal of Economic Theory*, 57(2), 306–332.

Van Damme, E. (1989): "Renegotiation-proof equilibria in repeated prisoners dilemma," *Journal of Economic theory*, 47(1), 206–217.

Vonortas, N. S. (2012): *Cooperation in research and development*, vol. 11. Springer Science & Business Media.

Wärneryd, K. (1993): "Cheap talk, coordination, and evolutionary stability," *Games and Economic Behavior*, 5(4), 532–546.

Weibull, J. W. (1995): *Evolutionary Game Theory*. MIT Press.

Wen, Q. (1996): "On renegotiation-proof equilibria in finitely repeated games," *Games and Economic Behavior*, 13(2), 286–300.

A Evolutionary Stability (for online publication?)

In this appendix we analyze the stability properties of strategies σ_L and σ_R (the results can be extended to other renegotiation proof equilibria, but we do not present the details here for brevity). Specifically, we show that σ_L and σ_R satisfy three properties that imply robustness to various perturbations: evolutionary stability (à la Maynard Smith and Price, 1973); behavior in the first stage is weakly dominant (given the second-stage behavior); and behavior in the second stage is a neighborhood invader (Apaloo, 1997; Cressman, 2010).

A.1 Preliminary Definition: Equivalent Strategies

We say that two strategies and are almost-surely realization equivalent (abbr., equivalent) if they induce the same behavior for almost all types (regardless of the opponent's behavior). Formally,

Definition 4. A condition holds *for almost all types* if the set of types that satisfy the condition $\tilde{U} \subseteq U$ has mass one (as measured by the distribution f); i.e.,

$$\int \mathbf{1}_{u\in\tilde{U}} \cdot f(u)\, du = 1.$$

Definition 5. Strategies $\sigma = (\mu, \xi)$ and $\tilde{\sigma} = \left(\tilde{\mu}, \tilde{\xi}\right)$ are *almost-surely realization equivalent* (abbr., *equivalent*) if for almost all types $u \in [0,1]$: $\mu_u(m) = \tilde{\mu}_u(m)$ for every message $m \in M$, and $u \leq \xi(m, m')$ if and only if $u \leq \tilde{\xi}(m, m')$ for every $m \in M$ satisfying $\mu_u(m) > 0$ and every $m' \in M$.

If σ and $\tilde{\sigma}$ are equivalent strategies we denote it by $\sigma \approx \tilde{\sigma}$. It is immediate that equivalent strategies always obtain the same ex-ante expected payoff. Formally,

Fact 1. *Let $\sigma \approx \tilde{\sigma} \in \Sigma$ be two equivalent strategies and let $\sigma' \in \Sigma$ an arbitrary strategy. Then $\pi(\sigma, \sigma') = \pi(\tilde{\sigma}, \sigma')$.*

A.2 Evolutionary/Neutral stability

An equilibrium strategy σ is neutral (evolutionary) stable if it achieves a weakly (strictly) higher ex-ante expected payoff against any (non-equivalent) best-reply strategy, relative to the payoff that the best-reply strategy achieves against itself. Formally

Definition 6 ((adaptation of Maynard Smith and Price, 1973))**.** Equilibrium strategy $\sigma \in \mathcal{E}$ is neutrally (evolutionary) stable if for any non-equivalent strategy $\tilde{\sigma} \not\approx \sigma$,

$$\pi(\tilde{\sigma}, \sigma) \geq \pi(\sigma, \sigma) \Rightarrow \pi(\sigma, \tilde{\sigma}) \geq \pi(\tilde{\sigma}, \tilde{\sigma}) \;\; (\pi(\sigma, \tilde{\sigma}) > \pi(\tilde{\sigma}, \tilde{\sigma}))\,.$$

The refinement of neutral stability is, arguably, a necessary requirement for an equilibrium to correspond to a stable convention in a population (see, e.g., Banerjee and Weibull, 2000). If σ is an equilibrium strategy that is not neutrally stable, then a few experimenting agents who play a best reply strategy σ' can invade the population. These experimenting agents would fare the same against the incumbents, while they outperform the incumbents when being matched with other experimenting agents. This implies that, on average, these experimenting agents would be more successful than the incumbents, and their frequency in the population would increase in any payoff-monotone learning dynamics. this, in turn, implies that the population will move away from σ.

Our first result shows that both σ_L and σ_R are neutrally stable, and, moreover, they are evolutionary stable if there are two feasible messages. Formally,

Proposition 6. *σ_L and σ_R are neutrally stable strategies of the coordination game with communication $\langle \Gamma, M \rangle$. Moreover, if $|M| = 2$, then σ_L and σ_R are evolutionary stable strategies.*

Sketch of proof. Let σ' be a best reply strategy against σ_L (the argument for σ_R is analogous). In what follows we show that, for almost all types σ' induces the same behavior as σ_L, except possibly, that types less than $1/2$ may send different messages $m \neq m_R$. Assume to the contrary that σ' differs from σ_L by either:

1. having a positive mass of types below $1/2$ sending message m_R (resp., having a positive mass of types above $1/2$ sending messages $m \neq m_R$); observe that any such type u obtains a payoff strictly less than $1-u$, while she could obtain a payoff of $1-u$ by following σ_L (resp., a payoff of at most $1-u < 1/2$, while she could obtain a payoff strictly larger than $1-u$ by following σ_L); this implies that $\pi(\sigma', \sigma_L) < \pi(\sigma_L, \sigma_L)$ and we get a contradiction; or

2. having a positive mass of types behaving differently on the equilibrium path after observing message profile (m, m') ; observe that the opponent plays a pure profile after (m, m'), and inducing a different behavior than σ_L in the second stage implies that these types miscoordinate the opponent's action with a positive probability, and, thus, obtain a strictly lower payoff; this implies that $\pi(\sigma', \sigma_L) < \pi(\sigma_L, \sigma_L)$ and we get a contradiction.

The fact that any best reply strategy σ' of σ_L is, essentially, equivalent to σ_L (i.e., it only differs by some types replacing the message m_L with another message $m \neq m_R$) implies that $\pi(\sigma_L, \sigma') = \pi(\sigma', \sigma')$, which implies that σ_L is neutrally stable. Moreover, if $|M| = 2$, then the above argument implies that only equivalent strategies $\sigma' \approx \sigma^*$ are best replies against $\tilde{\sigma}$, which implies that σ_L is evolutionary stable. □

A.3 Message Function is Dominant

In this subsection we show that the behavior in the first stage induced by strategies σ_L (resp., σ_R), namely the action function μ^*, is a weakly dominant message function (and strictly dominant when $|M| = 2$), when taking as given that the behavior in the second stage is according to the action function ξ_L (resp., ξ_R). This suggests that the behavior in the first stage that is induced by σ_L (resp., by σ_R) is robust to any perturbation that keeps the behavior in the second stage unchanged. Specifically, it implies that even if the message function used by the population is perturbed in an arbitrary way (possibly, even a large perturbation), then the original function μ^* yields a weakly higher payoff than any other message function, which suggests that the behavior in the first stage would converge back to play μ^* under any payoff-monotone learning dynamics.

Proposition 7 shows that following message-function μ^* yields a weakly higher payoff relative to following any other message-function when facing a partner who follows in the second stage either ξ_L or ξ_R . Moreover, the inequality is strict whenever the alternative message-function is essentially different than μ^* in the sense of inducing low types to play m_R or inducing high types to play $m \neq m_R$. Formally

Proposition 7. *Let μ' be an arbitrary message function. Then for any type $u \neq 1/2$:*

1. $\pi_u\left(\left(\mu^*, \xi_L\right), \left(\mu', \xi_L\right)\right) \geq \pi_u\left(\left(\mu', \xi\right), \left(\mu', \xi_L\right)\right)$ *with a strict inequality if, either, (I)* $\mu'_u\left(m_R\right) > 0$ *and* $u < 1/2$*, or (II)* $\mu'_u\left(m_R\right) < 1$ *and* $u > 1/2$*; and*
2. $\pi_u\left(\left(\mu^*, \xi_R\right), \left(\mu', \xi_R\right)\right) \geq \pi_u\left(\left(\mu', \xi_R\right), \left(\mu', \xi_R\right)\right)$ *with a strict inequality if, either, (I)* $\mu'_u\left(m_L\right) > 0$ *and* $u > 1/2$*, or (II)* $\mu'_u\left(m_L\right) < 1$ *and* $u < 1/2$.

Proof. In what follows we prove Part (1) of Proposition 7, which deals with ξ_L. The proof for Part (2), which deals with ξ_R, is analogous and is omitted for brevity. The fact that the opponent plays ξ_L in the second stage, implies that an agent of type $u < 1/2$ ($u > 1/2$) can gain an expected payoff of $1 - u\ \left(\left(1 - F\left(1/2\right)\right) \cdot u + F\left(1/2\right) \cdot \left(1 - u\right)\right)$ by following $\left(\mu^*, \xi_L\right)$, and that no other strategy can yield a higher payoff. Moreover, if $\mu'_u\left(m_R\right) > 0$ and $u < 1/2$ ($\mu'_u\left(m_R\right) < 1$ and $u > 1/2$), then following message-function μ' yields type u a payoff strictly less than $1 - u\ \left(\left(1 - F\left(1/2\right)\right) \cdot u + F\left(1/2\right) \cdot \left(1 - u\right)\right)$ due to having a positive probability (probability lager than $F\left(1/2\right)$) of inducing the partner to play the action, which is less preferred by the agent. This implies part (1) of Proposition 7.

An immediate corollary of Proposition 7 is that the inequalities are strict for any $\mu' \neq \mu^*$ when $|M| = 2$. Formally: □

Corollary 1. *Assume that $|M| = 2$, and that $\mu'_u \neq \mu_u$ for some type $u \neq 1/2$. Then:*

1. $\pi_u\left(\left(\mu^*, \xi_L\right), \left(\mu', \xi_L\right)\right) > \pi_u\left(\left(\mu', \xi\right), \left(\mu', \xi_L\right)\right)$*; and*
2. $\pi_u\left(\left(\mu^*, \xi_R\right), \left(\mu', \xi_R\right)\right) > \pi_u\left(\left(\mu', \xi_R\right), \left(\mu', \xi_R\right)\right)$.

A.4 Action Function is a Neighborhood Invader Strategy

The set of actions of each player in the induced second-stage game $\Gamma(F_m, F_{m'})$ (the game played after observing a pair of messages (m, m')) is, essentially, an interval (i.e., each player has to choose her maximal threshold for playing L). This induced game is asymmetric whenever the message profile is asymmetric, i.e., $m \neq m'$. As argued by Eshel and Motro (1981) and Eshel (1983) when the set of strategies is a continuum, a stable convention should be robust to perturbations that slightly change the strategy played by all agents in the population. Cressman (2010) formalizes this requirement by the notion of neighborhood invader strategy (adapting the related notion of Apaloo, 1997). In what follows we show that the action function induced by σ_L and σ_R is a neighborhood invader strategy in any induced game $\Gamma(F_m, F_{m'})$.

Fix a message strategy μ and a pair of messages $m_1, m_2 \in \text{supp}(\mu)$. We identify a strategy in the induced game $\Gamma_{F_{m_1},F_{m_2}}$ with thresholds x_i, which is interpreted as the maximal type for which player $i \in \{1,2\}$ plays L. We say that strategy x_i of player i is equivalent to x_i' (denoted by $x_i \approx x_i'$) in the induced game $\Gamma_{F_{m_1},F_{m_2}}$, if $F_{m_i}(x_i) = F_{m_i}(x_i')$, which implies that both thresholds induce the same behavior with probability one. Let $\pi^{m_1,m_2}(x_1,x_2)$ denote the expected payoff of an agent with a random type sampled from f_{m_1} who uses threshold x_1 when facing a partner with a random unknown type sampled from f_{m_2} who uses threshold x_2.

Strategy profile (x_1, x_2) is a strict equilibrium in the induced game $\Gamma_{F_{m_1},F_{m_2}}$, if any best reply to x_j is equivalent to x_i, i.e., $\pi^{m_1,m_2}(x_1', x_2) \geq \pi^{m_1,m_2}(x_1, x_2) \Rightarrow x_1' \approx x_1$, and $\pi^{m_2,m_1}(x_2', x_1) \geq \pi^{m_2,m_1}(x_2, x_1)$ $\Rightarrow x_2' \approx x_2$.

We say that the strict equilibrium (x_1, x_2) is a neighborhood invader strategy in the induced game $\Gamma_{F_{m_1},F_{m_2}}$ if the population converges to (x_1, x_2) from any non-equivalent nearby strategy profile (x_1', x_2') in two steps: (1) x_i yields a strictly higher payoff against x_j relative to the payoff of x_i' against x_j (which implies convergence from $\left(x_i', x_j'\right)$ to $\left(x_i, x_j'\right)$, and (2) due to (x_1, x_2) being a strict equilibrium, x_j yields a strictly higher payoff against x_i relative to the payoff of x_j' against x_i (which implies the convergence from $\left(x_i, x_j'\right)$, to (x_i, x_j)). Formally:

Definition 7 (Adaptation of Cressman (2010, Def. 5))**.** Fix a message strategy μ and a pair of messages $m_1, m_2 \in \text{supp}(\mu)$. A strict Nash equilibrium (x_1, x_2) is a *neighborhood invader strategy profile* in the induced game $\Gamma_{F_{m_1},F_{m_2}}$ if there exists $\epsilon > 0$, such that for each (x_1', x_2') satisfying $x_1' \not\approx x_1$, $x_2' \not\approx x_2$, $|x_1' - x_1| < \epsilon$ and $|x_2' - x_2| < \epsilon$, then either $\pi^{m_1,m_2}(x_1, x_2') > \pi^{m_1,m_2}(x_1', x_2')$ or $\pi^{m_2,m_1}(x_2, x_1') > \pi^{m_2,m_1}(x_2', x_1')$.

Proposition 8 shows that the profile of action functions induced by σ_L (or, similarly, by σ_R) is a neighborhood invader strategy in any induced game.

Proposition 8. *Let $m_1, m_2 \in supp(\mu^*)$. Then $(\xi_L(m_1, m_2), \xi_L(m_1, m_2))$ and $(\xi_R(m_1, m_2), \xi_R(m_1, m_2))$ are strict equilibria and neighborhood invader strategy profiles in the induced game $\Gamma_{F_{m_1},F_{m_2}}$.*

Proof. We present the proof for $(\xi_L(m_1,m_2), \xi_L(m_1,m_2))$ (the proof for $(\xi_R(m_1,m_2), \xi_R(m_1,m_2))$ is analogous). Observe that $m_1, m_2 \in \text{supp}(\mu^*)$ imply one of three cases: (1) $m_1 = m_2 = m_L$, (2) $m_1 = m_2 = m_R$, and (3) $m_1 = m_R$, $m_2 = m_L$. We analyse each case as follows: $\square$

1. $m_1 = m_2 = m_L$. This implies that $\xi_L(m_1,m_2) = \xi_L(m_2,m_1) = 1$ and $F_{m_1}(1/2) = F_{m_2}(1/2) = 1$. Let $\bar{x} < 1/2$ be sufficiently close to $1/2$ such that $F_{m_1}(\bar{x}), F_{m_2}(\bar{x}) > 1/2$. Observe that $\pi^{m_1,m_2}(1,x) > \pi^{m_1,m_2}(y,x)$ for any $x > \bar{x}$ and any $y \not\approx 1$. This proves that $(\xi_L(m_1,m_2), \xi_L(m_1,m_2))$ is a strict equilibrium and a neighborhood invader strategy profile.

2. $m_1 = m_2 = m_R$. This implies that $\xi_L(m_1,m_2) = \xi_L(m_2,m_1) = 0$ and $F_{m_1}(1/2) = F_{m_2}(1/2) = 0$. Let $\bar{x} > 1/2$ be sufficiently close to $1/2$ such that $F_{m_1}(\bar{x}), F_{m_2}(\bar{x}) < 1/2$. Observe that $\pi^{m_1,m_2}(0,x) > \pi^{m_1,m_2}(y,x)$ for any $x < \bar{x}$ and any $y \not\approx 0$. This proves that $(\xi_L(m_1,m_2), \xi_L(m_1,m_2))$ is a strict equilibrium and a neighborhood invader strategy profile.

3. $m_1 = m_R$, $m_2 = m_L$. This implies that $\xi_L(m_1,m_2) = \xi_L(m_2,m_1) = 1$, $F_{m_1}(1/2) = 0$, and $F_{m_2}(1/2) = 1$.

   (a) Observe that (1) $\pi^{m_1,m_2}(1,1) > \pi^{m_1,m_2}(x,1)$ for any $x \not\approx 1$ and (2) $\pi^{m_2,m_1}(1,1) > \pi^{m_2,m_1}(x,1)$ for any $x \not\approx 1$, which implies that $(\xi_L(m_1,m_2), \xi_L(m_1,m_2))$ is a strict equilibrium.

   (b) Let $\bar{x} > 1/2$ be sufficiently close to $1/2$ such that $F_{m_1}(\bar{x}) < 1/2$. Observe that $\pi^{m_2,m_1}(1,x) > \pi^{m_1,m_2}(y,x)$ for any $x < \bar{x}$ and any $y \not\approx 1$. This proves that $(\xi_L(m_1,m_2), \xi_L(m_1,m_2))$ is a neighborhood invader strategy profile.

A.5 Remark on Evolutionary Robustness

Oechssler and Riedel (2002) present a strong notion of stability, called evolutionary robustness, which refines both evolutionary stability and being a neighborhood invader. An evolutionary robust strategy σ^* is required to be robust against small perturbation in the strategy played by the population, which may combine both (1) a few experimenting agents who follow arbitrary strategies, and (2) many agents who follow strategies that are only slightly different than σ^*. Specifically, if φ is a distribution of strategies that is sufficiently close to σ^* (in the L_1 norm induced by the weak topology), evolutionary robustness à la Oechssler and Riedel requires that $\pi(\sigma^*, \varphi) > (\varphi, \varphi)$.

One can show that σ_L and σ_R do not satisfy this condition (and, we conjecture, that no strategy can satisfy such a strong condition). However, we conjecture that adaptations to the arguments presented in the results of this appendix can show that σ_L and σ_R satisfy a somewhat weaker notion of evolutionary robustness: for each strategy distribution φ sufficiently close to σ_L (σ_R), there exists a finite sequence of strategy distributions $\varphi_1, \varphi_2, ..., \varphi_k$, such that $\pi(\varphi_1, \varphi) \geq (\varphi, \varphi)$, $\pi(\varphi_2, \varphi_1) \geq (\varphi_1, \varphi_1)$, ..., $\pi(\varphi_k, \varphi_{k-1}) \geq (\varphi_{k-1}, \varphi_{k-1})$, and $\pi(\sigma_L, \varphi_1) \geq (\varphi_1, \varphi_1)$ ($\pi(\sigma_R, \varphi_1) \geq (\varphi_1, \varphi_1)$), with strict inequalities if $|M| = 2$ and φ is not realization equivalent to σ_L (σ_R).

B Proofs (for online publication)

B.1 Undominated action strategies

Let $\Gamma(F, G)$ be a coordination game without communication (possibly played after some communication in the original game Γ).

A generalized action strategy in this game is a measurable function $\xi : U \to \Delta(\{L, R\})$ that describes a mixed action as a function of the player's type. A generalized action strategy is a *cut-off strategy* if there exists a type $x \in [0, 1]$ such that $\xi(u) = L$ for each $u < x$ and $\xi(u) = R$ for each $u > x$.

A generalized action strategy ξ is *(weakly) dominated* by generalized action strategy $\tilde{\xi}$ if $\pi(\xi, \xi') \leq \pi\left(\tilde{\xi}, \xi'\right)$ for any generalized action strategy ξ' of the opponent.

Lemma 1. *Let ξ be a generalized action strategy. Then there exists a cut-off strategy $\tilde{\xi}$, such that ξ is dominated by $\tilde{\xi}$.*

Proof. Let $x \in [0, 1]$ be such that $F(x) = \boldsymbol{E}_{u \sim F}[\xi_u(L)] = \int_u \xi_u(L) f(u) du$. Let $\tilde{\xi}$ then be the cutoff strategy with cutoff x, i.e.,

$$\tilde{\xi}_u(L) = \begin{cases} 1 & u \leq x \\ 0 & u > x. \end{cases}$$

Let ξ' be an arbitrary generalized strategy of the opponent. By construction strategies ξ and $\tilde{\xi}$ induce the same average probability of choosing L. Strategies $\tilde{\xi}$ and ξ differ in that $\tilde{\xi}$ induces lower types to choose L with a higher probability, and higher types to choose L with a lower probability, i.e., $\tilde{\xi}_u(L) \geq \xi_u(L)$ for any type $u \leq x$ and $\tilde{\xi}_u(L) \leq \xi_u(L)$ for any type $u > x$. The fact that lower types always gain weakly more (less) from choosing L (R) relative to higher types implies that $\pi(\xi, \xi') \leq \pi\left(\tilde{\xi}, \xi'\right)$. □

Note also that any best response to some arbitrary generalized action strategy must be a cut-off strategy.

B.2 Proof of Theorem 1

We first prove the "if" part of the theorem. Suppose $\sigma = (\mu, \xi) \in \Sigma$ is ordinal preference revealing, mutual preference consistent, coordinated, and balanced.

As σ is *ordinal preference revealing* $\text{supp}(F_m) \subseteq [0, 1/2]$ or $\text{supp}(F_m) \subseteq [1/2, 1]$ for any message $m \in \text{supp}(\mu)$. Consider any message pair $m, m' \in \text{supp}(\mu)$. There are three cases to consider. Suppose first that both $\text{supp}(F_m), \text{supp}(F_{m'}) \subseteq [0, 1/2]$. Then as σ is *mutual preference consistent* we have that

$\xi(m, m') = \xi(m', m) = L$. Thus ξ describes best response behaviour after this message pair. Moreover this behavior is the best possible outcome for any type in $[0, 1/2]$ and thus for any type in $\text{supp}(F_m)$ and $\text{supp}(F_{m'})$. The second case of both $\text{supp}(F_m), \text{supp}(F_{m'}) \subseteq [1/2, 1]$ is completely analogous. Suppose then third and finally that, w.l.o.g., $\text{supp}(F_m) \subseteq [0, 1/2]$ and $\text{supp}(F_{m'}) \subseteq [1/2, 1]$. As σ is *coordinated* we have that $\xi(m, m') = \xi(m', m) = L$ or $\xi(m, m') = \xi(m', m) = R$. Action strategy ξ, therefore, again describes best response behavior. Moreover, one player (for all her types) obtains her most preferred outcome. The only way to improve the outcome for the other player would be by making the former player deviate from her most preferred outcome. Thus, there is no equilibrium strategy σ' in the game $\Gamma(F_m, F_{m'})$ with any new finite message set that Pareto dominates σ after this message pair.

All this shows that action strategy ξ is a best response to μ and to itself given μ and that moreover it cannot be interim equilibrium Pareto-improved upon. It remains to be shown that the message strategy μ is optimal given the opponent chooses $\sigma = (\mu, \xi)$.

Consider type $u \in [0, 1/2]$ and consider this type's choice of message. As σ is *balanced* and *coordinated* different messages $m \in M$ can only trigger different probabilities of coordinating on L with a highest likelihood of such coordination for any message $m \in \text{supp}(\mu_u)$. Type u is, therefore, indifferent between any message $m \in \text{supp}(\mu_u)$ and weakly prefers any $m \in \text{supp}(\mu_u)$ over any message $m' \notin \text{supp}(\mu_u)$. An analogous statement holds for types $u \in [1/2, 1]$. This concludes the proof of the "if" part of the theorem.

To prove the "only if" part we use a series of lemmas. The first lemma is a technical one.

Lemma 2. *Let $p, q \in (0, 1)$. Then there exists a rational number $\alpha \in (0, 1)$ satisfying $\frac{q-p}{1-p} < \alpha < \frac{q}{p}$.*

Proof. Note that, as $q < 1$,

$$0 \leq (q - p)^2 = q^2 - 2 \cdot q \cdot p + p^2 < q - 2 \cdot q \cdot p + p^2,$$

which implies that

$$p \cdot q - p^2 < q - q \cdot p \Leftrightarrow p \cdot (q - p) < q \cdot (1 - p) \Leftrightarrow \frac{q - p}{1 - p} < \frac{q}{p}.$$

The result then follows from the fact that $\frac{q-p}{1-p} < 1$. □

The following lemma (which relies on Lemma 2) proves that in a renegotiation proof equilibrium strategy agents never miscoordinate after observing identical messages.

Lemma 3. *Let $\sigma = (\mu, \xi)$ be a renegotiation-proof equilibrium strategy. Then either $\xi(m, m) \geq \sup \{u|\mu_u(m) > 0\}$ or $\xi(m, m) \leq \inf \{u|\mu_u(m) > 0\}$ for each $m \in$ supp (μ).*

Proof. Let $m \in \text{supp}\,(\mu)$. Assume to the contrary that $\inf \{u|\mu_u(m) > 0\} < \xi(m, m) < \sup \{u|\mu_u(m) > 0\}$.

This implies that there exist types $u_L, u_R \in [0,1]$ satisfying $\mu_{u_L}(m), \mu_{u_R}(m) > 0$ and $u_L < \xi(m,m) < u_R$. Let $q = F_m(\xi(m,m))$ be the average probability that the partner plays L, conditional on observing the pair of identical messages (m,m). If $q = 0$, then type u_R obtains a payoff of zero in $\Gamma(F_m, F_m)$, which contradicts σ being an equilibrium strategy (as type u_R could play L and obtain a positive payoff). Similarly, if $q = 1$, then type u_L obtains a payoff of zero in $\Gamma(F_m, F_m)$, which contradicts σ being an equilibrium strategy. Thus, we are left with the case in which $q \in (0,1)$.

Note that the expected payoff of a player of type u conditional on observing message profile (m,m) is equal to

$$\max\{(1-u)\cdot q, u\cdot(1-q)\}. \tag{4}$$

Let $p = F_m(1/2)$. If $p = 0$ ($p = 1$) then the strategy in which all agents play R (L) is a symmetric equilibrium of the game $\Gamma(F_m, F_m)$, which Pareto dominates $\xi(m,m)$. We are left with the case in which $p \in (0,1)$. Due to Lemma 2, there exists a rational number $\alpha \equiv \frac{k}{n} \in (0,1)$ satisfying $\frac{q-p}{1-p} < \alpha < \frac{q}{p}$.

Let $\tilde{M} = \{l, r\} \times \{1, ..., n\}$ and consider the following symmetric equilibrium $\tilde{\sigma}$ of the game $\langle \Gamma(F_m, F_m), \tilde{M} \rangle$. We interpret the message of each player as a preferred direction (l or r) and a random number between 1 and n. In $\tilde{\sigma}$ players send message l if and only if their type is smaller than 1/2; send a random number between 1 and n according to the uniform distribution; play L after observing $((l,a),(l,b))$ for any numbers a and b; play R after observing $((r,a),(r,b))$ for any numbers a and b; play L after observing $((l,a),(r,b))$ if $a+b < k \mod n$; and play L after observing $((l,a),(r,b))$ if $a+b \geq k \mod n$.

Observe that $\tilde{\sigma}$ is indeed an equilibrium of the game $\langle \Gamma(F_m, F_m), \tilde{M} \rangle$ because following any pair of messages the players coordinate for sure; each agent with $u < 1/2$ ($u > 1/2$) strictly prefers to report that her preferred direction is l (r), as this induces her to coordinate on L (R) with a high probability of $p + \alpha\cdot(1-p)$ ($1 - p + (1-\alpha)\cdot p$) instead of with a low probability of $\alpha \cdot p$ ($(1-p)\cdot(1-\alpha)$); and each agent is indifferent between sending any random number, as this has no effect on the probability of coordinating on L (which is equal to $\alpha = \frac{k}{n}$), given that the partner chooses his random number uniformly.[12]

The payoff of each type $u \leq 1/2$ in equilibrium $\tilde{\sigma}$ is given by

$$(p + \alpha\cdot(1-p))\cdot(1-u) + (1 - (p + \alpha\cdot(1-p)))\cdot u > \tag{5}$$

$$q\cdot(1-u) + (1-q)\cdot u \geq \max\{(1-u)\cdot q, u\cdot(1-q)\}, \tag{6}$$

where the first inequality is implied by $\frac{q-p}{1-p} < \alpha \Leftrightarrow q < p + \alpha\cdot(1-p)$ and $u \leq 1/2$.

[12] The method to use simultaneous communication to implement a jointly controlled lottery was introduced in Aumman, Maschler, and Stearns (1968). See Heller (2010) for a recent implementation, which is robust to joint deviations of some of the players.

The payoff of each type $u > 1/2$ in equilibrium $\tilde{\sigma}$ is given by

$$((1-p)+(1-\alpha)\cdot p)\cdot u+(1-((1-p)+(1-\alpha)\cdot p))\cdot(1-u) > \tag{7}$$
$$(1-q)\cdot u+q\cdot(1-u) \geq \max\left\{(1-u)\cdot q, u\cdot(1-q)\right\}, \tag{8}$$

where the first inequality is implied by $\frac{p}{q} > \alpha \Leftrightarrow 1-\alpha > \frac{p-q}{p} \Leftrightarrow (1-p)+(1-\alpha)\cdot p > 1-q$ and $u > 1/2$. This implies that all types obtain a strictly larger payoff in $\tilde{\sigma}$ (relative to the expected payoff of σ in the game $\langle\Gamma(F_m, F_m)\rangle$), contradicting σ being renegotiation-proof. □

The following lemma shows that, in a renegotiation proof equilibrium strategy, for each message m that is sent with positive probability any two player types who send the same such message have the same preferred outcome.

Lemma 4. *Let $\sigma = (\mu, \xi)$ be a renegotiation-proof equilibrium strategy. Then for any $m \in supp\,(\mu)$, $F_m\,(1/2) \in \{0, 1\}$.*

Proof. Assume to the contrary that $F_m\,(1/2) \in (0, 1)$. Let $\tilde{\sigma}$ be the following symmetric equilibrium of the symmetric induced game with additional communication $\langle\Gamma(F_m, F_m), \{l, r\}\rangle$. Each player sends message l if her type is at most $1/2$ and sends message r otherwise. If both players send the same message then they both play the respective action. That is they both play L after message pair (l, l) and R after (r, r). If the players send different messages then they play the original equilibrium action prescribed in ξ. The fact that, by Lemma 3, the probability of coordination under ξ is one, implies that strategy $\tilde{\sigma}$ is indeed a symmetric equilibrium of $\langle\Gamma(F_m, F_m), \{l, r\}\rangle$. Note that all types are strictly better off in $\tilde{\sigma}$ than in σ, which contradicts σ being a renegotiation-proof equilibrium strategy. □

Given $\sigma = (\mu, \xi) \in \Sigma$ define $M_L = M_L(\sigma) = \{m \in \text{supp}\,(\mu)\,|F_m\,(1/2) = 1\}$ and $M_R = M_R(\sigma) = \{m \in \text{supp}\,(\mu)\,|F_m\,(1/2) = 0\}$. Lemma 4 implies that $M_L \cap M_R = \emptyset$ and $M_L \cup M_R = \text{supp}(\mu)$ for any renegotiation proof equilibrium strategy σ and, thus, implies that any renegotiation proof equilibrium strategy must be ordinal preference revealing.

The following lemma shows that, in a renegotiation proof equilibrium strategy, agents always coordinate on a jointly preferred outcome if such an outcome exists, i.e., it shows that renegotiation proof equilibrium strategies must be mutual preference consistent.

Lemma 5. *Let $\sigma = (\mu, \xi)$ be a renegotiation-proof equilibrium strategy. If $m, m' \in M_L$, then $F_m\,(\xi(m, m')) = 1$. If $m, m' \in M_R$, then $F_m\,(\xi(m, m')) = 0$.*

Proof. Assume to the contrary that $m, m' \in M_L$ and $F_m\,(\xi(m, m')) < 1$. By Lemma 4 the two player types, sending message m and m', must both have types $u, v \leq 1/2$. Consider the induced game without communication $\Gamma(F_m, F'_m)$. Let $\tilde{\sigma}$ be such that both agents always play L. This is an equilibrium of

$\Gamma(F_m, F'_m)$ in which all types who send either message m or m' are weakly better and some types are strictly better off than in σ. This contradicts σ being renegotiation proof. The remaining case is proven analogously. □

The following lemma shows that, in a renegotiation proof equilibrium strategy, agents never miscoordinate when one players sends a message in M_L and the other in M_R.

Lemma 6. *Let $\sigma = (\mu, \xi)$ be a renegotiation-proof equilibrium strategy. If $m \in M_L$ and $m' \in M_R$, then $F_m(\xi(m, m')) = F_{m'}(\xi(m', m)) \in \{0, 1\}$.*

Proof. Assume to the contrary that $m \in M_L$, $m' \in M_R$, and $F_m(\xi(m, m')) \neq F_{m'}(\xi(m', m))$ or $F_m(\xi(m, m')) \in (0, 1)$. This assumption implies that the strategy σ induces players to mis-coordinate with a positive probability after observing message pair (m, m'). By Lemma 4 we have that $F_m(1/2) = 1$ and $F_{m'}(1/2) = 0$. This means that there are types $u < \frac{1}{2}$ who, in σ, send message m and who play R after observing (m, m'). As σ is an equilibrium this implies that it must be the case that average probability that this type's opponent (a m' message sender) plays L is strictly less than $\frac{1}{2}$. Otherwise this type with $u < 1/2$ would prefer to play L. An analogous argument implies that the average probability that a m' message sender plays R after observing (m, m') is strictly less than $1/2$.

Now consider the induced game with additional communication $\langle \Gamma(F_m, F_{m'}), \{0, 1\} \rangle$. Consider the equilibrium $\tilde{\sigma}$ of this game, in which both players send their messages uniformly, and they play L if they have sent the same message and play R otherwise. Observe that each player is indifferent between the two feasible messages, because the probability of coordinating on L is $1/2$ regardless of the agent's message (as long as the opponent sends her message uniformly). Further observe that all types obtain an expected payoff of $1/2$ in $\tilde{\sigma}$. This implies that all types are then strictly better off in $\tilde{\sigma}$ than in σ, a contradiction to σ being a renegotiation proof equilibrium strategy. □

Finally, Lemma 7 (together with Lemma 5) establishes that renegotiation proof equilibrium strategies are balanced.

Lemma 7. *Let $\sigma = (\mu, \xi)$ be a renegotiation-proof equilibrium strategy. Then σ is balanced.*

Proof. Recall that

$$\alpha_m(\sigma) = \int_{v > \frac{1}{2}} \sum_{m' \in \text{supp}(\mu_v)} \mu_v(m') \mathbf{1}_{\{v \leq \xi(m, m')\}} f(v) dv.$$

Given Lemma 4 and the fact that σ is coordinated by Lemmas 3, 5, and 6, we can rewrite this as $\alpha_m(\sigma) = \sum_{m' \in M_R} \mu(m') F_m(\xi(m, m'))$. This is now the probability of a u-type (with $u \leq 1/2$) ending up coordinating on L conditional on her sending message $m \in M_L$ and her opponent having a type in $[1, 2, 1]$ (i.e., sending a message in M_R). Given that σ is coordinated, $1 - \alpha_m(\sigma)$ is then the probability of such a u-type ending up coordinating on R conditional on her on her sending message $m \in M_L$ and her opponent having a type in $[1, 2, 1]$. Now if there were any other message $m'' \in M$ that were to yield a higher $\alpha_{m''}(\sigma)$ then this would be preferred by any such u type which contradicts $m \in M_L$. If there were any other message $m'' \in M_L$ that gives a lower $\alpha_{m''}(\sigma)$ then m would be preferred by any such u type which contradicts $m'' \in M_L$.

This proves that there exists $\alpha \in [0, 1]$, such that for each $m \in M_L$, $\alpha_m(\sigma) = \alpha$. By an analogous argument we obtain that there exists $\beta \in [0, 1]$, such that for each $m \in M_R$, $\beta_m(\sigma) = \beta$. That $\alpha = \beta$ is immediately implied by the fact that both terms are just two different ways of computing the same quantity. Both compute the overall average probability of players coordinating on action L after observing a "mixed" pair of messages: one message from M_L and the other message from M_R. □

Proof of Theorem 1:

Proof. Lemma 4 (which uses Lemma 3, which in turn uses the technical Lemma 2) proves that a renegotiation proof equilibrium strategy must be ordinal preference revealing. Lemma 5 (which uses Lemma 4) proves that it must be mutual preference consistent. Lemmas 3, 5, and 6 (the latter two lemmas using Lemma 4) together imply that it must be coordinated. Lemma 7 (which also uses Lemma 4 and that just established property that a renegotiation proof equilibrium strategy is coordinated), finally, proves that it must be balanced. □

B.3 Proofs of Section 6

Proof of Proposition 1. By Theorem 1 the equilibrium payoff of a renegotiation proof strategy σ is determined by its left-tendency $\alpha \in [0, 1]$ and, for each type $u \in [0, 1/2]$ given by

$$\pi_u(\sigma, \sigma) = (1 - u) \left[F(1/2) + \alpha \left(1 - F(1/2)\right)\right] + u(1 - \alpha) \left[1 - F(1/2)\right],$$

and for each type $u \in (1/2, 1]$ given by

$$\pi_u(\sigma, \sigma) = (1 - u)\alpha F(1/2) + u \left[(1 - F(1/2)) + F(1/2)(1 - \alpha)\right].$$

The payoff to a u-type for a given outcome function ϕ is given by

$$\pi_u(\phi) = (1 - u)\, \boldsymbol{E}_v \phi_{u,v}(L, L) + u \boldsymbol{E}_v \phi_{u,v}(R, R).$$

Now suppose that ϕ interim (pre communication) Pareto dominates σ. Then $\pi_u(\phi) \geq \pi_u(\sigma, \sigma)$ for all $u \in [0, 1]$ with a strict inequality for a positive measure of u. As $\pi_u(\sigma, \sigma)$ is a convex combination of two payoffs, this implies that for any $u \leq 1/2$

$$\boldsymbol{E}_v \phi_{u,v}(L,L) \geq F(1/2) + \alpha(1 - F(1/2))$$

and for any $u > 1/2$

$$\boldsymbol{E}_v \phi_{u,v}(R,R) \geq (1 - F(1/2)) + F(1/2)(1-\alpha)$$

with at least one of the inequalities holding strictly for a positive probability set of types. We can write

$$\boldsymbol{E}_v \phi_{u,v}(L,L) = F(1/2)\boldsymbol{E}_{\{v\leq 1/2\}}\phi_{u,v}(L,L) + (1 - F(1/2))\,\boldsymbol{E}_{\{v>1/2\}}\phi_{u,v}(L,L),$$

where, for instance, $\boldsymbol{E}_{\{v>1/2\}}$ denotes the expectation conditional on $v > 1/2$. To have

$$F(1/2)\boldsymbol{E}_{\{v\leq 1/2\}}\phi_{u,v}(L,L) + (1 - F(1/2))\,\boldsymbol{E}_{\{v>1/2\}}\phi_{u,v}(L,L) \geq F(1/2) + \alpha(1 - F(1/2))$$

for any $u \leq 1/2$, by the fact that $\boldsymbol{E}_{\{v\leq 1/2\}}\phi_{u,v}(L,L) \leq 1$, we must then have that

$$\boldsymbol{E}_{\{v>1/2\}}\phi_{u,v}(L,L) \geq \alpha$$

for any $u \leq 1/2$ and, by an analogous argument that

$$\boldsymbol{E}_{\{v<1/2\}}\phi_{u,v}(R,R) \geq 1-\alpha,$$

for any $u > 1/2$, again with at least one of the inequalities holding strictly for a positive probability set of types.

This implies that

$$\boldsymbol{E}_{\{u<1/2\}}\boldsymbol{E}_{\{v>1/2\}}\phi_{u,v}(L,L) \geq \alpha$$

as well as

$$\boldsymbol{E}_{\{u>1/2\}}\boldsymbol{E}_{\{v<1/2\}}\phi_{u,v}(R,R) \geq 1-\alpha,$$

with a at least one of the two inequalities holding strictly. By the symmetry of ϕ we have $\phi_{u,v}(R,R) = \phi_{v,u}(R,R)$ and thus,

$$\boldsymbol{E}_{\{u<1/2\}}\boldsymbol{E}_{\{v>1/2\}}\phi_{u,v}(L,L) + \boldsymbol{E}_{\{u<1/2\}}\boldsymbol{E}_{\{v>1/2\}}\phi_{u,v}(R,R) > 1,$$

a contradiction to ϕ being a probability distribution, and thus, to our supposition. □

The proof of Proposition 2 uses the following lemma (of some independent interest).

Lemma 8. *Let $\sigma \in \mathcal{E}$ be a coordinated equilibrium strategy. Then there is a renegotiation proof strategy σ' such that either σ and σ' are interim (pre-communication) payoff equivalent or σ' interim (pre-communication) Pareto dominates σ.*

Proof. Let $\sigma = (\mu, \xi) \in \mathcal{E}$ be a coordinated equilibrium strategy. For each message $m \in M$, let $p_m \in [0,1]$ be the probability that the players coordinate on L, conditional on the agent sending message m:

$$p_m = \sum_{m' \in M} \mu(m') \mathbf{1}_{\xi(m,m')=L}.$$

As σ is coordinated $1 - p_m$ is then the probability that the players coordinate on R, conditional on the agent sending message m.

Let $\bar{p} = \max_{m \in M} p_m$ be the maximal probability, and let $\underline{p} = \min_{m \in M} p_m$ be the minimal probability. By definition $\underline{p} \leq \bar{p}$. As σ is an equilibrium strategy, $\underline{p} < \bar{p}$ would imply that all types $u < 1/2$ send a message inducing probability $\bar{p}$ and all type $u > 1/2$ send a message inducing probability $\underline{p}$. We, therefore, must have that the expected payoff of a type $u \leq 1/2$ is given by

$$\pi_u(\sigma, \sigma) = \bar{p}(1-u) + (1-\bar{p})u,$$

and the expected payoff of any type $u > 1/2$ is equal to

$$\pi_u(\sigma, \sigma) = \underline{p}(1-u) + \left(1 - \underline{p}\right)u.$$

Note that this is also true if $\underline{p} = \bar{p}$. Note furthermore that for every type $u < 1/2$ this type's expected payoff strictly increases in $\bar{p}$ and for every type $u > 1/2$ this type's expected payoff strictly decreases in $\underline{p}$.

We consider three cases. Suppose first that $\underline{p} \leq \bar{p} \leq F(1/2)$. Then let $\sigma' = \sigma_R$. This strategy is also coordinated and its induced payoffs can be written in the same form as those for strategy σ with $\underline{p}' = 0$ and $\bar{p}' = F(1/2)$. Thus, we get that $\pi_u(\sigma', \sigma') \geq \pi_u(\sigma, \sigma)$ for every $u \in [0,1]$. This implies that σ is either interim (pre communication) payoff equivalent to or Pareto dominated by $\sigma' = \sigma_R$.

The second case where $F(1/2) \leq \underline{p} \leq \bar{p}$ is analogous to the first one, with now $\sigma' = \sigma_L$ the equivalent or dominating strategy.

This leaves the final case where $\underline{p} < F(1/2) < \bar{p}$. Let $\alpha \in [0,1]$ be such that $F(1/2) + (1 - F(1/2))\alpha = \bar{p}$ and let σ' be a renegotiation proof strategy with left tendency α. We then must have that $\underline{p} \geq \alpha F(1/2)$ and by construction σ' is either interim (pre communication) payoff equivalent to or Pareto dominates σ. □

Proof of Proposition 2. By Lemma 8 we have that every coordinated equilibrium strategy σ is interim (pre communication) Pareto dominated by some renegotiation proof strategy with some left tendency $\alpha \in [0,1]$ denoted by σ_α. We thus have that

$$\pi(\sigma, \sigma) \leq \pi(\sigma_\alpha, \sigma_\alpha).$$

The ex-ante expected payoff of to a u type under strategy σ_α is given by

$$\pi_u(\sigma_\alpha, \sigma_\alpha) = (1-u)\left[F(1/2) + \alpha(1-F(1/2))\right] + u(1-\alpha)(1-F(1/2)),$$

for $u \leq 1/2$ and

$$\pi_u(\sigma_\alpha, \sigma_\alpha) = (1-u)\alpha F(1/2) + u\left[1 - F(1/2) + (1-\alpha)F(1/2)\right],$$

for $u > 1/2$. It is straightforward to verify that for every u,

$$\pi_u(\sigma_\alpha, \sigma_\alpha) = \alpha \pi_u(\sigma_1, \sigma_1) + (1-\alpha)\pi_u(\sigma_0, \sigma_0).$$

As $\sigma_1 = \sigma_L$ and $\sigma_0 = \sigma_R$ and as for all $u \in [0,1]$ $\pi_u(\sigma_\alpha, \sigma_\alpha)$ is the same convex combination of $\pi_u(\sigma_L, \sigma_L)$ and $\pi_u(\sigma_R, \sigma_R)$ we get that

$$\pi(\sigma_\alpha, \sigma_\alpha) = \alpha \pi(\sigma_1, \sigma_1) + (1-\alpha)\pi(\sigma_0, \sigma_0),$$

which implies that

$$\pi(\sigma, \sigma) \leq \pi(\sigma_\alpha, \sigma_\alpha) \leq \max\left\{\pi(\sigma_L, \sigma_L), \pi(\sigma_R, \sigma_R)\right\}.$$

□

Proof of Proposition 3. The payoff of each type u induced by equilibrium strategy x can be bounded as follows:

$$\pi_u(x, x) = \mathbf{1}_{u \leq x} \cdot F(x) \cdot (1-u) + \mathbf{1}_{u > x} \cdot (1 - F(x)) \cdot u \leq F(x) \cdot (1-u) + (1-F(x)) \cdot u$$

$$< F(x) \cdot \pi_u(\sigma_L, \sigma_L) + (1-F(x)) \cdot \pi_u(\sigma_R, \sigma_R) \leq \max\left\{\pi(\sigma_R, \sigma_R), \pi(\sigma_R, \sigma_R)\right\}.$$

□

B.4 Proof of Proposition 6

In what follows we prove that σ_L is neutrally stable, and that it is evolutionarily stable if $|M| = 2$ (the analogous proof for σ_R is omitted for brevity). Let $\sigma' = (\mu', \xi')$ be a best reply strategy against $\sigma_L = (\mu_L, \xi_L)$. i.e., assume that $\pi(\sigma', \sigma_L) = \pi(\sigma_L, \sigma_L)$. The proof includes the following steps.

1. We begin by showing that the second-stage behavior of almost all types is the same according to σ_L and according to σ'; that is we show that for almost all types $u \in [0,1]$ (I) $u \leq \xi'(m, m_L)$ for any message $m \in M$ satisfying $\mu_u(m) > 0$, (II) $u \leq \xi'(m, m_R)$ for any message $m \neq m_R \in M$ satisfying $\mu_u(m) > 0$, and (III) if $\mu_u(m_R) > 0$, then $u \geq \xi'(m_R, m_R)$. Assume to the contrary

that there is a positive mass of types of agents following strategy σ' who have a second-stage behavior different than σ_L by either (I) playing R after observing the partner sending m_L, (II) playing R after the agent sending message $m \neq m_R$,or (III) playing L after both the agent and the partner sending message m_R. Note, that these types obtain a payoff of zero whenever their second-stage behavior differs from σ_L (which is strictly less than the positive payoff obtained by agents with these types who follow strategy σ_L), as they always mis-coordinate the partner who follows strategy σ_L. This implies that there is a positive mass of agents satisfying $\pi_u(\sigma', \sigma_L) < \pi_u(\sigma_L, \sigma_L)$, which implies, in turn, that $\pi(\sigma', \sigma_L) < \pi(\sigma_L, \sigma_L)$ and we get a contradiction.

2. Next we show that the first-stage behavior of almost all types is essentially the same according to σ_L and according to σ'; that is we show that $(\mu_L)_u(m_R) = \mu'_u(m_R)$ for almost all types. Assume to the contrary that there is a positive mass of types satisfying $(\mu_L)_u(m_R) \neq \mu'_u(m_R)$. This implies that there is either a positive mass of types below $1/2$ sending m_R or there is a positive mass of types above $1/2$ sending messages different than m_R. Each such type u obtains a payoff of at most $\min(u, 1-u) < 1/2$ when facing a partner who follows strategy σ_L (because the partner always plays their less preferred action), which is strictly less than the payoff that is obtained by agents following strategy σ (who obtain of $\max(u, 1-u)$ for any $u < 1/2$ and obtain an expected payoff of $1/2$ for any $u > 1/2$). This implies that g $\pi_u(\sigma', \sigma_L) < \pi_u(\sigma_L, \sigma_L)$ for a positive mass of types, which implies, in turn, that $\pi(\sigma', \sigma_L) < \pi(\sigma_L, \sigma_L)$ and we get a contradiction. This implies that there exists a positive mass f or $\int_{u=1/2}^{1} 1 - \mu'_u(m_R) < 1$. In what follows we show that, essentially, σ' differs from σ_L only by having agents with type $\leq 1/2$ different distribution of choosing messages in M_L (and, in particular, $|M| = 2$, and $M = \{m_R, m_L\}$ only equivalent strategies $\sigma' \approx \sigma^*$ are best replies against $\approx \tilde{\sigma}$.

3. The previous two steps imply that almost all types behave in essentially the same way when following strategy σ_L and when following σ'. This implies that $\pi(\sigma_L, \sigma') = \pi(\sigma', \sigma')$, which, in turn, implies that σ_L is neutrally stable.

4. Assume now that $|M| = 2$. In this case Step (2) implies that for almost all type $\mu'_u(m) = (\mu_L)_u(m)$, which implies (together with Step (1)) that $\sigma_L \approx \sigma'$. Thus any strategy that is a best reply against σ_L must be equivalent to σ_L, which implies that σ_L is evolutionary stable.