Online appendix to *Why don't we talk about it? Communication and coordination in teams*

Jens Dietrichson
VIVE - The Danish Center for Social Science Research

Jens Gudmundsson
University of Copenhagen

Torsten Jochem
University of Amsterdam

Contents

1 Theorem 1 and 2 with different π's in the communication and action phase

We first define separate π_q and π_p for the communication and action phase, respectively. Let $q(a, a) = \pi_q$ when $\boldsymbol{m}^{-1} = (\underline{a}^{-1}, \ldots, \underline{a}^{-1})$ or $\boldsymbol{m}^{-1} = (\emptyset, \ldots, \emptyset)$, and $p(a) = \pi_p$ when $\boldsymbol{m} = (a, \emptyset, \ldots, \emptyset)$.

First, note that the proof of Theorem 1 in the main text only includes π in the communication phase. Thus, the theorem and the proof can be re-written by simply changing every π to π_q. We have therefore omitted this re-formulation here.

Having separate π's does not change the fact that moving from a higher to a lower-ranked absorbing state always require one mistake. Furthermore, the condition for when one mistake is enough to move from a lower to a higher absorbing state in the proof of Theorem 2 can be re-written by changing π to π_p. That is, whenever $\lambda\pi_p - 1 \geq 0$.

However, as Theorem 2 depends on Theorem 1, we get a more complicated structure of the theorem. With non-equal π_q and π_p the theorem reads:

Theorem 2b. The set of stochastically stable states is described as follows.

1. If $\lambda\pi_p - 1 < 0$ and $(K-1)\left(\lambda\pi_q - 1\right) < \gamma$, then ω_1 is the unique stochastically stable state: $\Omega^S = \{\omega_1\}$.
2. If $\lambda\pi_p - 1 < 0$ and $\lambda\pi_q - 1 < \gamma \leq (K-1)\left(\lambda\pi_q - 1\right)$, then ω_l and ω_K are the stochastically stable states: $\Omega^S = \{\omega_l, \omega_K\}$.
3. If $\lambda\pi_p - 1 < 0$ and $\lambda\pi_q - 1 \geq \gamma$, then ω_K is the unique stochastically stable state: $\Omega^S = \{\omega_K\}$.
4. If $\lambda\pi_p - 1 \geq 0$, then all absorbing states are stochastically stable: $\Omega^S = \Omega^A$.

If the condition in 1. holds, then only ω_1 among the absorbing states can be reached by one mistake, i.e., moving up always requires more than one mistake.

If the condition in 2. holds, then ω_l and ω_K can be reached by one mistake in the action phase from all other absorbing states whereas two mistakes are needed to transition to all other absorbing states from at least one other absorbing state. As ω_l is the lowest ranked absorbing state, it can be reached by one mistake from all other absorbing states because moving down requires only one agent playing action l. Moving to ω_K requires only one mistake because $l > l-1 \geq 1$, and if an agent mistakenly plays any action ranked lower than l, then the best response for all agents in the next period is to send K. Hence, play ends up in ω_{KK}, after which play transition to ω_K (see proof of Proposition 1 for why sending K is optimal from ω_{l-1} when 2. holds, and why play then transitions to ω_K). To reach absorbing states $\omega_{l+1}, ..., \omega_{K-1}$ from ω_l requires first a mistake so that the minimum action becomes a lower-ranked action than l. Then play transitions to ω_K by the same process as described in the previous paragraph, after which another mistake is needed to transition from ω_K to, e.g., ω_{K-1}.

If the condition in 3. holds, then ω_K is the only absorbing state, and therefore the only stochastically stable state by Theorem 1.

If the condition in 4. holds, then all absorbing states are stochastically stable by the same reasoning as in the proof of Theorem 2.

2 Alternative definition of stochastic stability

Foster and Young (1990), Kandori et al. (1993), and Young (1993) develop a more general concept of stochastic stability than the one we use, based on arguments in Freidlin and Wentzell (1984). Here, we argue that our simpler solution concept yield the same result as theirs for Theorem 2.

Following Young (1993), we can limit the attention to absorbing states, i.e., $\omega_a \in \Omega^A$, as the candidates for stochastically stable states. Define the latter in the following way: let $r_{aa'}$ be the minimum number of mistakes or experiments (mistakes from here on for brevity) required to make a transition between ω_a and $\omega_{a'}$. That is, $r_{aa'}$ is either a positive integer or, if no mistakes are needed, zero. The latter can only be the case when the transition is between one non-absorbing and one absorbing state.

To find the stochastically stable states: first, construct a complete directed graph with one node for each absorbing state. The directed edge from ω_a and $\omega_{a'}$ is called aa' and the weight on the edge is equal to $r_{aa'}$. A rooted tree T is a set of directed edges such that from every node different from ω_a, there is a unique directed path in the tree to ω_a. The total resistance of T is the sum of the $r_{aa'}$ on the edges that compose it. The stochastic potential $\gamma(\omega_a)$ of the absorbing state ω_a is defined as the minimum resistance over all trees rooted at a. That is, if we denote the set of all trees rooted at ω_a with $T(a)$, then the stochastic potential is

$$\gamma(\omega_a) = \min_{T \in T(a)} \sum_{k,k' \in T} r_{kk'}$$

Stochastically stable states are the absorbing states that have the minimum stochastic potential, i.e., $\min_{\omega_a \in \Omega^A} \gamma(\omega_a)$ (Young, 1993, Theorem 2).

To find the stochastically stable states in our setting, we can use the following: First, as stated by Riedl et al. (2012), because the distribution of mistakes is uniform, the path of least resistance between ω_a and $\omega_{a'}$ is always one where mistakes are of action a'. Second, transitions to both higher and lower ranked absorbing states are possible for all ω_a, except for the lowest and highest ranked absorbing states, ω_ℓ and ω_K. Thus, if it always requires less mistakes to move to a lower ranked state, the lowest ranked absorbing state (ω_ℓ) will be the stochastically stable state.

In case 1 of our Theorem 2, all coordinated states – no player communicates and all players choose the same action – are absorbing. As it takes one mistake to transition to ω_1 from all other absorbing states and N mistakes to transition from ω_1 to higher ranked states, ω_1 is stochastically stable.

In case 2 of Theorem 2, either there is only one absorbing and stochastically stable state (i.e., ω_K), or all transitions between absorbing states can happen through one mistake (see the proof of Theorem 2 in the appendix of the main text for an explanation). Hence, the stochastic potential is the same for all absorbing states, so all are stochastically stable.

Calculating the stochastic potential may be cumbersome, and require the comparison of a large number of very unlikely mutations (Nöldeke and Samuelson, 1993). Nöldeke and Samuelson (1993) argued that results obtained for absorbing sets that are "locally stable", i.e., robust against one mistake or experiment, are likely to be more robust. In our case, either there is only one absorbing and stochastically stable state (i.e., ω_K), or all absorbing states are stochastically stable, or the unique stochastically stable state is precisely the the absorbing state that cannot be left with just one mistake (i.e., ω_1).

3 Alternative model of belief formation

A more general model of belief formation than the one we used in the main text could allow for small probabilities that other actions than the previous round's minimum action or an action indicated by a message were possible minimum actions. For example, it seems reasonable that the agents may realize that mistakes and experiments are possible, and therefore would put a small probability on other actions becoming the minimum.

To incorporate this idea in a simple way, we generalize the model by changing condition 2 that governs belief formation in communication phase in the following way. Let $q(a,a) \equiv \pi >> q(b,a) \equiv \delta > 0$ for all $a, b \in A \setminus \{\underline{a}^{-1}\}$, which implies that $q(\underline{a}^{-1}, a) = 1 - (K-2)\delta - \pi$ so that probabilities still sum to one. This formulation incorporates mistakes and experiments in a natural way. When contemplating sending message a, the agent also plans to play action a but recognizes that she as well as the other agents may end up playing differently and that the minimum action therefore may become $b \neq a$.

If agents' believe that all actions are possible minimum actions in the communication phase, we should also reformulate the agents' expected utility of sending message m to reflect these beliefs. That is, as other actions are believed to be possible minimum actions, the agent would not only consider messaging and subsequently choosing $\min a, b$. An expected utility function capturing this notion is

$$\mathbb{E}u(a,m) = \lambda \sum_{b \in A} q(b,m)b - a - c(m)$$

Below, we sketch a proof of Theorem 3 with this model of belief formation, in which the tie-breaking assumption does not play a substantial role. The only need for a tie-breaking assumption arises if we want to separate between ω_{KK} and ω_K as absorbing states with fully compensated communication. However, both of these states are efficient when $\gamma = 0$. Theorem 1 and 2 would have similar forms and qualitative implications with this expected utility function as with the one used in the manuscript, but would include the extra parameter δ. For example, in Theorem 2, ω_1 is the unique stochastically stable state if $\lambda(\pi - \delta) - 1 < 0$. As the theorems can be proved using slight modifications of the proofs in the manuscript, we have omitted them here.

Proof of Theorem 3 with modified utility function: As in Proposition 1, starting from an arbitrary state ω, we transition into some ω_{aa}.

Then

$$\sum_{b \in A} q(b,m)b = \begin{cases} \pi m + (1 - \pi - (K-2)\delta)a + \sum_{b \in A \setminus m,a} \delta b & \text{if } m \neq a \\ (1 - (K-1)\delta)a + \sum_{b \in A \setminus a} \delta b & \text{if } m = a \end{cases}.$$

This is non-decreasing in m. As $\lambda > 0$ and c is either constant in m (and have to be incurred under mandatory communication) or zero, $\mathbb{E}u(a,m)$ is also non-decreasing in m:

$$\mathbb{E}u(a,m) = \lambda \sum_{b \in A} q(b,m)b - a - c(m).$$

If $a = K$, so all agents communicate and play K, they will continue to do so the next round; compare Proposition 1. That is, ω_{KK} is absorbing.

Next, consider $a < K$. As $\mathbb{E}u(a, K) - \mathbb{E}u(a, a) =$

$$\lambda \left[\pi K + (1 - \pi - (K-2))a + \sum_{b \in A \setminus a, K} \delta b \right] - a - c(m) -$$
$$\lambda \left[(1 - (K-1))a + \sum_{b \in A \setminus a} \delta b \right] + a + c(m)$$
$$= \lambda \left[(\pi - \delta)(K - a) \right] > 0,$$

sending $m = K$ is strictly preferred to a in ω_{aa}, and we always transition to ω_{KK}. As ω_{KK} is the unique absorbing state, it is also stochastically stable.

4 Additional simulation results

Figure 1 shows the results of the simulation specification mentioned in footnote 26 on page 24 in the main text. We use configurations with 8 rounds, 8 agents, 10 actions, $\alpha = 20$ and $\beta \in \{8, 9, 10, 11, 12\}$,[1] and without message costs, mistakes and experiments, and no mandatory communication. The mean/median minimum action in the final round is 8.2/8.5 and slightly less than 40 percent of the teams coordinate on the maximum action of 10.

Further, Table 1 shows additional results on the variability of outcomes after 50 rounds of play, as mentioned in footnote 29 on page 30 in the main text.

[1] That is, $\lambda \in \{20/8, 20/9, 20/10, 20/11, 20/12\}$.

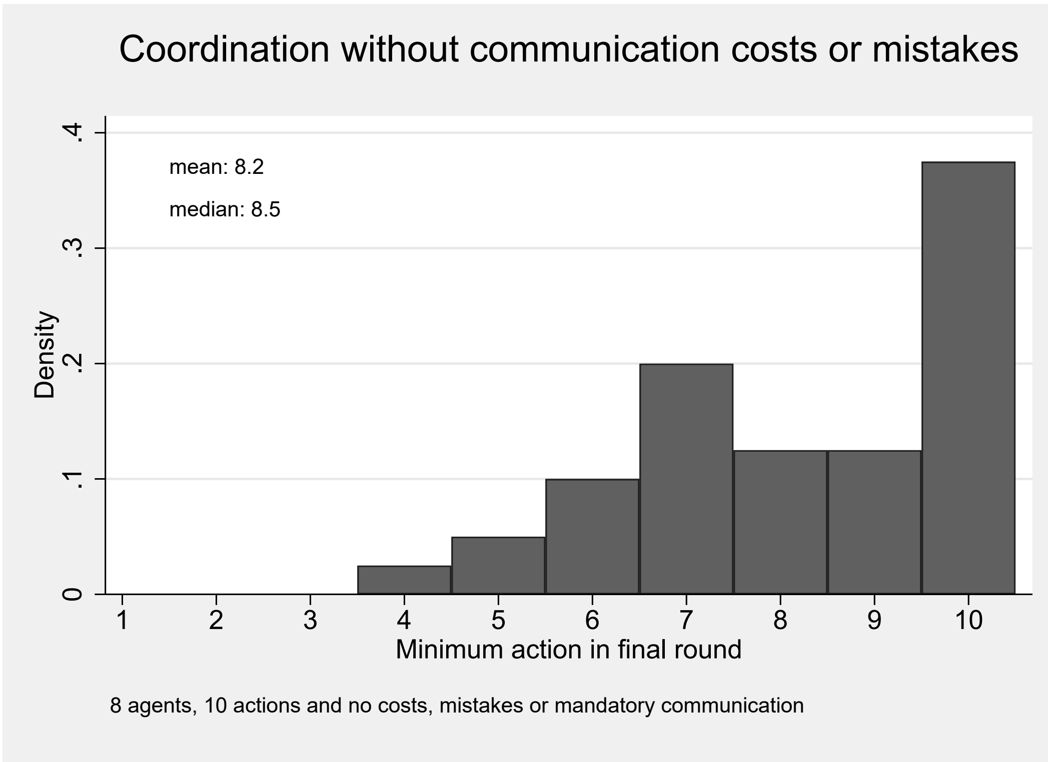

Figure 1: Results without message costs or mistakes and experiments

This figure shows the simulated minimum team actions after round 8 when following a specification with $\alpha = 20$ and $\beta \in \{8, 9, 10, 11, 12\}$ (i.e., $\lambda \in \{20/8, 20/9, 20/10, 20/11, 20/12\}$), 8 agents, 10 actions, and no message costs, no mistakes or experiments, and no mandatory communication.

Table 1: Distribution of long-run outcomes

The panels below show the distribution of the average minimum action after 50 rounds of play. P5, P25, P50, P75 and P95 represent respectively the 5th, 25th, median, 75th and 95th percentile. In Panel A, the first (second) mistake probability shown vertically is the unconditional probability that a mistake occurs in the communication stage (action stage). Since a larger action space allows for higher outcomes, the below results keep the action space (unless in Panel C) constant at 10 actions.

Panel A: Team size and mistake probabilities

Mistake	Agents=4					Agents=8				
probabilities	P5	P25	P50	P75	P95	P5	P25	P50	P75	P95
25%, 25%	1.16	3.54	5.22	7.24	8.88	1.00	1.20	2.60	4.14	6.56
25%, 50%	1.36	2.74	3.94	6.40	7.74	1.00	1.16	1.80	3.20	5.24
50%, 25%	1.52	4.32	5.74	7.66	9.04	1.00	2.00	3.80	6.64	8.24
50%, 50%	1.58	3.32	4.24	7.14	8.34	1.00	1.62	2.54	5.80	7.34

Panel B: Team size and incentives

	Agents=4					Agents=8				
Lambda	P5	P25	P50	P75	P95	P5	P25	P50	P75	P95
20/8	2.68	4.14	5.56	7.54	9.04	1.00	1.60	3.24	5.40	7.92
20/9	2.80	4.12	5.62	7.62	8.88	1.00	1.64	3.00	5.00	8.02
20/10	3.04	4.14	5.50	7.10	8.54	1.00	1.40	2.70	4.78	7.78
20/11	1.28	2.16	3.56	6.08	8.28	1.00	1.28	2.16	3.60	7.68
20/12	1.12	1.94	3.46	5.90	8.30	1.00	1.18	2.12	3.72	7.24

Panel C: Team size and action space

	Agents=4					Agents=8				
Actions	P5	P25	P50	P75	P95	P5	P25	P50	P75	P95
2	1.00	1.00	1.04	1.16	1.44	1.00	1.00	1.00	1.04	1.20
4	1.00	1.40	1.88	2.46	3.08	1.00	1.00	1.20	1.60	2.36
6	1.08	2.08	2.84	3.84	4.92	1.00	1.08	1.60	2.48	3.96
8	1.28	2.72	3.84	5.40	6.80	1.00	1.16	2.08	3.42	5.70
10	1.40	3.36	4.80	6.98	8.76	1.00	1.36	2.60	4.44	7.74

Panel D: Team size and messaging costs

Messaging costs	Agents=4					Agents=8				
	P5	P25	P50	P75	P95	P5	P25	P50	P75	P95
1	1.50	3.38	4.76	7.10	8.72	1.00	1.36	2.56	4.40	7.58
3	1.36	3.48	4.68	6.96	8.54	1.00	1.36	2.62	4.56	7.72
5	1.40	3.54	4.80	7.02	8.80	1.00	1.40	2.56	4.50	7.72
7	1.54	3.36	4.86	6.96	8.84	1.00	1.36	2.72	4.20	7.70
9	1.32	3.32	4.92	6.86	8.48	1.00	1.28	2.48	4.68	7.86

Panel E: Team size and mistake distribution types

Game stage	Agents=4					Agents=8				
	P5	P25	P50	P75	P95	P5	P25	P50	P75	P95
By mistake distribution in communication stage:										
Uniform	2.46	3.36	4.20	6.24	7.58	1.34	1.92	2.68	3.82	5.20
Dbl-dist-half-as-likely	1.12	1.52	2.78	4.36	5.54	1.00	1.00	1.00	1.08	1.48
Exponential	3.26	4.48	6.36	7.88	8.86	1.58	2.58	4.16	6.46	7.90
Highest message	3.08	4.44	6.52	8.16	9.12	1.72	2.48	3.86	6.40	8.14
By mistake distribution in action stage:										
Uniform	1.40	2.86	3.88	4.92	6.38	1.00	1.30	2.08	2.84	4.14
Dbl-dist-half-as-likely	1.42	4.58	6.88	8.00	8.92	1.00	1.60	4.30	6.48	8.02

5 Description of the simulation code

The simulation code, written in Python, is available as a supplementary file (file name "DGJ-simulation.py"). Here, we describe and explain the code in words and equations with emphasis on how agents form their beliefs, and decides on which messages to send and which actions to take.

The first sections of the code defines concepts, error distributions, and functions, which are used later in the code. Note that the simulation code was written to be as similar as possible to the set up in the experiment we compare our results to and therefore use a slightly different notation than our main text. For example, the payoffs are determined by two parameters, α and β rather than the single parameter λ that we use in the analytical model. There is also a τ parameter, which is added in experiments to ensure that no subject earns negative profits. We follow this notation below but omit the τ as it does not affect the agents' choices of messages and actions (i.e., our agents are not worried about negative profits).

In the code section "Main simulation program", we first initialize arrays used to store actions, profits, beliefs in the communication and action phase, messages, and actions. Thereafter, we create beliefs in the initial round by using a uniform randomization to create a vector of non-empty messages that agents use to form beliefs about which messages they think other agents will send in period 1. Agents then send best response messages conditional on these beliefs. The procedure to find the best response message follow the procedure described below.

The subsection PART (A): "DETERMINE OPTIMAL MESSAGES TO SEND" describes how best response messages are chosen conditional on beliefs. Agents reason about how other agents would react to their messages: each agent i first forms a belief about what other agents' would believe under different messages sent by i. Second, agent i then computes the expected payoff for each of the other agents, thereby learning which message by i would trigger what payoff-maximizing action by any other agent.

To do this, i uses i's own prospective message in period t, and the empirical distribution of messages and the minimum action in $t-1$ to form beliefs of each other agent j's subjective probabilities in t (see Step A.1. in the code). As each period follows the same rules, we describe an arbitrary period and omit t everywhere (play in the previous period is marked with superscript -1). In equation form, let

$$q_{ij}(a,m) = \frac{1}{N}\left(\mathbf{1}\left(m=a\right) + \sum_{j\in N\setminus\{i\}} \mathbf{1}\left(m_j^{-1}=a\right) + |\varnothing| \times \mathbf{1}(\underline{a}^{-1}=a)\right) \tag{1}$$

be *i*'s expectation over *j*'s subjective probability of action a becoming the minimum in the case i should send message m. $\mathbf{1}(\cdot)$ are indicator functions equal to 1 whenever the conditions in parentheses hold. Unless agent i contemplates to make a change from communication to non-communication or the other way around, the term $|\varnothing|$ is just the number of empty messages sent in the previous period. If i contemplates a change from sending a substantive to the empty message (or from the empty to a substantive message), $|\varnothing|$ decreases (increases) by one. That is, if, for some $a \in A_i$

$$m = a \text{ and } m_i^{-1} = \varnothing \Rightarrow |\varnothing| = \sum_{j=1}^{N} \mathbf{1}\left(m_j^{-1}=\varnothing\right) - 1 \tag{2}$$

$$m = \varnothing \text{ and } m_i^{-1} = a \Rightarrow |\varnothing| = \sum_{j=1}^{N} \mathbf{1}\left(m_j^{-1}=\varnothing\right) + 1 \tag{3}$$

where $\mathbf{1}(\cdot)$ is an indicator function equal to 1 when an agent sent the empty message in the previous period. This formulation constrains $\sum_{a=1}^{K} q_{ij}(a,m) = 1$ for each $m \in M$ and all periods, except for the initial period. (Beliefs in the initial period are generated as described above.) It implies that agents believe that other agents will send the same messages as in the previous period.

While π does not feature as an explicit parameter in this formulation, $q_{ij}(a,m)$ follows a similar condition as condition 2 that governs belief formation in the analytical model. That is, if $\boldsymbol{m}^{-1} = (\underline{a}^{-1},\ldots,\underline{a}^{-1})$ or $\boldsymbol{m}^{-1} = (\emptyset,\ldots,\emptyset)$, then $q_{ij}(\underline{a}^{-1},\emptyset) = q_{ij}(\underline{a}^{-1},\underline{a}^{-1}) = 1$. Further, $q_{ij}(a,a) = 1/N \; \forall a \in A \setminus \{\underline{a}^{-1}\}$ and $q(\underline{a}^{-1},a) = 1 - 1/N$.

In Step A.2, agent i calculates each agent j's expected payoff for $a > 1$ as

$$\mathbb{E}u_{ij}(a,m) = \sum_{h=a}^{K} q_{ij}(h,m)a(\alpha-\beta) + \sum_{l=1}^{a-1} q_{ij}(l,m)\left(\alpha l - \beta a\right). \tag{4}$$

As each agent's payoff of playing $a = 1$ is always safe (payoff-determining), and equal to $\alpha - \beta$. Agent i does not have to take into account any message costs for agent j, as these represent sunk costs in the action stage for j and are not considered when choosing a best response action.

In Step A.3, the agents use the expected payoff $\mathbb{E}u_{ij}(a,m)$ to evaluate the expected minimum action by checking each agent j's best response to each of i's messages. Formally, let

$$\Pi_j(m) = \{a \in A_i : \mathbb{E}u_{ij}(a,m) \geq \mathbb{E}u_{ij}(a',m) \forall\, a' \in A_i\} \tag{5}$$

be the set of actions such that they are an expected best response to message m for agent j (from the point of view of agent i). If $\mathbb{E}u_{ij}(a,m) = \mathbb{E}u_{ij}(a',m)$ for some actions a and a', agents randomize uniformly among them (so $\Pi_j(m)$ becomes a singleton). Let $\Pi_{-i}(m) = \Pi_1(m) \cup ... \cup \Pi_{i-1}(m) \cup \Pi_{i+1}(m) \cup ... \cup \Pi_n(m)$ be the union of all agents' $j \neq i$ expected best response sets. There is thus one $\Pi_{-i}(m)$ for each $m \in M_i$ and $K+1$ in total for every agent i.

In Step A.4, each agent then compares the payoffs of the lowest ranked action in each $\Pi_{-i}(m)$ – the minimum, denoted $\underline{a}$ – and then chooses the message corresponding to the set with the minimum yielding the highest payoff. Denote this collected set of minimum actions by Π_i^{min}. Best response messages are found in

$$BR_i^m = \{m \in M_i : u_i\left(\underline{a}\right) \geq u_i\left(\underline{a'}\right)\ \forall\, \underline{a}, \underline{a'} \in \Pi_i^{min}\} \tag{6}$$

where

$$u_i\left(\underline{a}\right) = \underline{a}\left(\alpha - \beta\right) - c\left(m\right). \tag{7}$$

If there is more than one message in this best response correspondence, we again assume that agents randomize uniformly between them. The implication of the above procedure is that the only probabilistic judgement is made when assessing the impact of a certain message on other agents' choice of action.

In Step A.5, the agents send the best response message, or with a probability equal to 0%, 10%, or 20%, chooses a message according to one of the distributions we described in section 4.2 in the main text. This concludes the communication phase.

PART (B), Step B.1, contains the code for the action phase. In the action phase, we assume that agents best respond to expectations given by the frequencies of received messages and the minimum action in the previous period. When an agent receives messages from some but not all other agents, agents assume that the non-communicating agents will play the minimum action in the previous period. All agents sees all messages and observe the same minimum action. Let $p(a)$ denote the probability assigned by all agents to a being the minimum action. Step B.1 computes the expected payoff of an action a as

$$\mathbb{E}(u(a) = \left(\sum_{h=a}^{K} p(h)\right) a(\alpha-\beta) + \sum_{l=1}^{a-1} p(l)(\alpha l - \beta a). \tag{8}$$

where $p(h) = \frac{1}{N}\sum_{j\in N} p_j(h)$ and $p(l) = \frac{1}{n}\sum_{j\in N} p_{ij}(l)$, and

$$p_j(h) = \begin{cases} 1 & \text{if } m_j = h \\ 1 & \text{if } m_j = \varnothing \text{ and } \underline{a}^{-1} = h \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

and

$$p_j(l) = \begin{cases} 1 & \text{if } m_j = l \\ 1 & \text{if } m_j = \varnothing \text{ and } \underline{a}^{-1} = l \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

That is, the subjective probabilities of each action being the minimum are equal to the frequencies with which the action has been indicated by messages, or indicated by the combination of the empty message and being the last period's minimum action. The procedure implies that $\sum_{a=1}^{K} p_i(a) = 1$. In turn, this formulation follows condition 3 that governs belief formation in the analytical model and implies that if $a \neq \underline{a}^{-1}$ and $\boldsymbol{m} = (a, \emptyset, \ldots, \emptyset)$, then $p(a) = 1/N$ and $p(\underline{a}^{-1}) = 1 - 1/N$.

Unless a mistake or experiment occurs, agents choose an action in the best response correspondence for actions

$$BR_i^a = \{a \in A_i : \mathbb{E}u_i(a, m) \geq \mathbb{E}u_i(a', m) \,\forall\, a' \in A_i\} \quad (11)$$

If there is more than one action in BR_i^a, the agent choose one by uniform randomization. Depending on the simulation set up, the agent may also with probability equal to 0%, 10%, or 20% choose an action by a uniform randomization over all actions in A_i or by the *DoubleDist* distribution described in the main text.

PART (C) computes the resulting profit for each agent for the round in question (in Step C.1), and then keeps a log of what happened in the round (Step C.2).

The last parts of the code store the outcomes of the simulation in a .csv-file, determines the parameters for the simulation one wants to run, and which combinations of parameters that should be included.

References

Foster, D. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology*, 38(2):219–232.

Freidlin, M. and Wentzell, A. (1984). *Random Perturbations of Dynamical Systems*. Springer Verlag, New York.

Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56.

Nöldeke, G. and Samuelson, L. (1993). An evolutionary analysis of backward and forward induction. *Games and Economic Behavior*, 5(3):425–454.

Riedl, A., Rohde, I. M. T., and Strobel, M. (2012). Efficient coordination in weakest-link games. CESifo Working Paper 3685.

Young, H. P. (1993). The evolution of conventions. *Econometrica*, 61(1):57–84.