Games and Economic Behavior 137 (2023) 23–49

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Coordination via delay: Theory and experiment ☆

Ye Jin [a], Zhen Zhou [b,*], Adam Brandenburger [c]

[a] NYU Shanghai, Shanghai, 200122, China
[b] PBC School of Finance, Tsinghua University, Beijing, 100083, China
[c] Stern School of Business, Tandon School of Engineering, NYU Shanghai, New York University, New York, NY 10012, USA

ARTICLE INFO

Article history:
Received 5 June 2021
Available online 14 November 2022

JEL classification:
C73
C92
D83

Keywords:
Coordination
Forward induction
Iterated weak dominance

ABSTRACT

This paper studies the effect of introducing an option of delay in coordination games—that is, of allowing players to wait and then choose between the risk-dominant and payoff-dominant actions. The delay option enables forward-induction reasoning to operate, whereby a player's waiting and not choosing the risk-dominant action right away signals an intention to choose the payoff-dominant action later. If players have ϵ-social preferences—they help others if they can do so at no cost to themselves—then iterated weak dominance yields a unique outcome in which everyone waits and then chooses the payoff-dominant action if everyone else waited. Thus, efficient coordination results. Experimental evidence from a binary-action minimum-effort game confirms that adding a delay option can significantly increase the occurrence of efficient outcomes. Moreover, consistent with our theory, the clear majority of subjects in our experiment take the unique iteratedly undominated strategy and not other strategies that are implied by equilibrium analysis.

1. Introduction

Coordination games are models of the challenge of coordination among economic (or other) players. The coordination challenge consists of two parts. First, there is the challenge of achieving a Nash equilibrium of the game. Consider a simple coordination game as shown in Table 1. In this 2×2 game with $x \in (2, 4)$, there are two pure-strategy Nash equilibria—one in which all players choose the safe but inefficient action B and one in which all players choose the risky but efficient action A. However, miscoordination may occur if one player chooses A, while the other chooses B.

The (pure-strategy) Nash equilibria can be Pareto-ranked, so that the second challenge is to attain an efficient equilibrium. In this simple example, the payoff-dominant equilibrium (A, A) differs from the risk-dominant equilibrium (B, B) in the sense of Harsanyi and Selten (1988). Carlsson and van Damme (1993) show that the risk-dominant equilibrium (B, B) is uniquely selected if we relax common knowledge of the payoffs (for example, common knowledge of the parameter x). The selection of the risk-dominant equilibrium is also supported by ample experimental evidence (Van Huyck et al., 1990).

☆ We thank the editor Vince Crawford, an advisory editor, and two referees for extremely helpful comments. Ala Avoyan, Deepal Basak, Giacomo Bonanno, Jordi Brandts, Stefan Bucher, Emiliano Catonini, Archishman Chakraborty, Vered Kurtz, Barry Nalebuff, Xiangdong Qin, Joao Ramos, Satoru Takahashi, John Wooders (discussant), and seminar and conference participants at NUS, NYU Shanghai, UC Davis, ESWC2020 and WESSI2020 provided important input and suggestions. We thank Shuhuai Zhang and Dandan Zhao for outstanding research assistance, and the National Science Foundation of China (project no. 71703101), the NYU Stern School of Business, NYU Shanghai, and J.P. Valles for financial support. We are grateful to the Smith Lab at Shanghai Jiao Tong University for hosting the lab experiment.
* Corresponding author.
E-mail addresses: wendyjin@nyu.edu (Y. Jin), zhouzh@pbcsf.tsinghua.edu.cn (Z. Zhou), adam.brandenburger@nyu.edu (A. Brandenburger).

https://doi.org/10.1016/j.geb.2022.11.001

Table 1
A simple coordination game.

		Player 2 A	Player 2 B
Player 1	A	4, 4	0, x
	B	x, 0	x, x

How can the players overcome the challenge of coordination? In reality, coordination games are often played dynamically, and the option of "wait and see" is often available. For instance, in the bank-run game, which is a classic coordination game, each depositor might be able to wait and then make their final withdrawal decision conditional on the information they observe.[1] This paper explores a class of dynamic games that allow each player to exercise the option to delay their choice of the efficient action. We find that the addition of the delay option can help players overcome miscoordination and also achieve the payoff-dominant outcome.

The delay option, if exercised, enables a player to observe other players' past history of play. However, more than observability, this paper highlights the idea that exercising the delay option and not taking the inefficient action early enables a player to signal their intention to take the efficient choice in future play—signaling this to other players who also exercise that option. We show, both theoretically and experimentally, that signaling through exercising the delay option—i.e., adopting the strategy of waiting and then taking the efficient action if all others wait, can work effectively to overcome the challenge of coordination.

Theoretical analysis The main result is proved for a multiple-player coordination game in which the efficient outcome is achieved if all players choose action A. For now, we will rely on the simple 2×2 game above to illustrate the intuition. The game unfolds in two periods, $t = 0, 1$. At $t = 0$, each player can choose between the irreversible choice B and "wait." A player who chooses to wait observes whether or not the other player chooses B at $t = 0$, and then makes their final choice between A and B at $t = 1$.[2] There is no cost associated with the delay option. A player who does not choose B at $t = 0$ should, in some sense, be signaling that they intend to choose A at $t = 1$. That is, there is a forward-induction flavor to choosing "wait." Intuitively, if a player intends to play B and secure the safe payoff x, then they can do so right away, rather than waiting and doing so later, which does not result in any extra benefit.[3] By contrast, waiting and then playing B, regardless of the history can be costly if the player is concerned about the other player's payoff. That is because using this strategy hurts the other player if they choose to wait and then play A if the first player chose B earlier. Next, we describe our analysis in more detail.

First, observe that if a player waits and then receives the "B" message (i.e., the other player chooses B at $t = 0$), they will optimally choose B at $t = 1$. Formally, any strategy that involves choosing A after the "B" message is weakly dominated. Next, consider the situation in which a player receives the "no-B" message (i.e., the other player does not choose B at $t = 0$). The game then enters a simultaneous-move subgame in which each player make a final choice between A and B. A player might decide to choose B in this subgame if they believe that the other player will do so. But then this player chooses B after either message. Compare this with the strategy of choosing B at $t = 0$. The two strategies are equivalent in that, they yield the same payoff to our player for each strategy profile of the other player. But choosing B after the message "no-B" hurts the other player if they choose A after observing "no-B." Formally, we will assume that each player has a utility function given by their original payoff function plus an infinitesimal weight on the other player's payoff function. We will refer to such preferences as ϵ-social preferences. Under our assumption, a player will not gratuitously hurt another player; in this context, "gratuitously" means that one makes a choice that hurts other players without helping oneself.

Formally, we have just argued that, with ϵ-social preferences, the strategy of waiting and then playing B, regardless of the message received, is weakly dominated—in the second round of elimination—by playing B immediately (at $t = 0$). Once this strategy is eliminated, the strategy of playing B immediately is weakly dominated by waiting and then playing B after the "B" message and A after the "no-B" message. In effect, by choosing "wait," a player signals their intention to follow the message and, in particular, to choose the efficient action A after the "no-B" message. This is the sole strategy that survives iterated weak dominance, and the result is that each player chooses "wait" at $t = 0$, collectively generating the "no-B" message, and then takes action A. In that way, the payoff-dominant outcome is achieved, even when it is risk-dominated.

There are two key components to this analysis. The first is forward induction. Introduced by Kohlberg and Mertens (1986) as a property of stable sets of equilibria, forward induction has been developed in two different directions. One is as an equilibrium refinement (Van Damme, 1989; Govindan and Wilson, 2009). The other is as iterated elimination of "bad"

[1] For other examples of applications of coordination games, see, inter alia, new technology adoption (Farrell and Saloner, 1985; Katz and Shapiro, 1986), team production (Bryant, 1983), search (Diamond, 1982), currency attacks (Morris and Shin, 1998), and debt crises (Corsetti et al., 2006).

[2] To illustrate this reversibility structure, consider, again, the case of a bank run. It is reasonable to believe that if a depositor chooses "run," withdrawing money from the bank, they will not return the money to the bank during financial distress. Likewise, one can interpret this reversibility condition in a coordination game of FDI investors, in which investors decide whether or not to relocate their investments back to their own country (Mathevet and Steiner, 2013).

[3] The simple idea that any player, if choosing to wait, tends to take different actions based on different observed histories is also in Chamley and Gale (1994) and Gul and Lundholm (1995), who study delay options in a (non-strategic) social learning setup.

strategies. One sub-approach here is extensive-form rationalizability (Pearce, 1984; Battigalli, 1997). A second sub-approach, the one employed in this paper, is iterated elimination of weakly dominated strategies (Ben-Porath and Dekel, 1992).[4]

The second key component of our analysis is the inclusion of social preferences. Other-regarding preferences have been identified in various experimental studies (see Fehr and Schmidt (2006) for a survey). We adopt a very weak form of social preference in which there is no trade-off between a player's own payoff and those of other players. In our model, other players' payoff functions are decisive only when the player is choosing between two equivalent strategies. This particular concept should be contrasted with the usual models of altruism, which, in many games, will modify a player's original preferences in more ways than our condition does.

In addition to iterated weak dominance, we analyze the game via symmetric pure-strategy subgame-perfect Nash equilibrium. Unlike iterated weak dominance, this solution concept does not lead to a unique prediction of efficient coordination, even with ϵ-social preferences. To see this, consider our 2×2 example again, and focus on the subgame starting from the "no-B" message. If the other player chooses B at this information set, playing A hurts the first player without benefiting the other player. Thus, subgame perfection, together with ϵ-social preferences, does not rule out the Nash equilibrium in which both players wait and choose the inefficient action B, regardless of the $t = 0$ outcome.

Experimental evidence Motivated by this sharp difference in predictions across different solution concepts, we designed an experiment by adding a delay option to a binary-action minimum-effort (a.k.a. weakest-link) game, and examining its effectiveness in overcoming the challenge of coordination. The minimum-effort game, first examined experimentally in Van Huyck et al. (1990), is a well-studied example of a coordination game with multiple Pareto-ranked equilibria. The subsequent literature has replicated the difficulty of promoting the efficient outcome and has proposed various mechanisms to bring about efficient coordination.[5]

In the experiments, we compare the coordination efficiency between a four-player static game and the dynamic version with the delay option.[6] The experiments were implemented with fixed groups, following the minimum-effort game literature, as well as randomly formed groups. Our experimental evidence confirmed that, while the efficient outcome is hard to achieve in the static version of the game, with the presence of the delay option, waiting can be a meaningful signal of future play of the risky but efficient action A. In addition, efficient coordination was significantly improved in both fixed- and random-matching treatments. We adopted the strategy method to elicit the subjects' full plans of play. Consistent with our theory, the vast majority of the subjects adopted the unique strategy surviving iterated weak dominance—namely, waiting and then choosing A after the "no-B" message and B after the "B" message.

We conducted further treatments to check the robustness of our main findings by considering a different representation of the "No-B" message and providing finer information (rather than binary information) regarding the first-period outcome. Our results found that the efficiency-enhancing effect of the delay option was robust to these variations.

An extra block was added to the main experiment to elicit subjects' social preferences in the random-matching treatments. We found that over 90 percent of the subjects held ϵ-social preferences, pointing to the validity of this assumption. Consistent with our theory, subjects' ϵ-social preferences and their beliefs that others held ϵ-social preferences were positively associated with their adoption of the unique strategy surviving iterated weak dominance in the main experiment.

Related literature Efficient equilibrium in dynamic coordination games has been well studied in the literature. However, the dynamic setup that we consider, involving synchronicity of moves and the irreversibility and observability of actions, appears to be novel. For example, a backward-induction argument predicts efficient coordination when players can move sequentially in an exogenous order (Farrell and Saloner, 1985). A similar argument can be made in the case in which players can choose the timing of taking the efficient action A, while the inferior choice B is reversible.[7] Different from Calcagno et al. (2014), who investigate a preparation stage in which actions are partially reversible,[8] our approach relies on the fact that the inefficient action B is the only irreversible choice.

Our mechanism in which the players signal their intention to play the efficient action by waiting is distinct from achieving efficient coordination through pre-play communication. The latter is equivalent to a dynamic setting in which both actions are reversible, and taking a particular action in the first period can be viewed as a non-binding message expressing an intent concerning future play.[9] There is ample experimental evidence confirming the effectiveness of costless communi-

[4] Brandenburger et al. (2008) provide an epistemic foundation for iterated weak dominance.

[5] See, for example, Van Huyck et al. (1993) and Cachon and Camerer (1996) for pre-game auctions, Blume and Ortmann (2007) for costless communication, Weber (2006) for gradually increasing group size, Brandts and Cooper (2006a) for changes in incentives, Chen and Chen (2011) for social identify, and Avoyan and Ramos (2019) for asynchronous pre-play revision. See Devetag and Ortmann (2007) for a survey.

[6] The efficiency rate of a round in a certain treatment is defined as the percentage of subject groups in which every group member's realized choice is the efficient action A.

[7] In this case, the players, knowing that their early actions can be observed and expecting others to follow, may take the lead by selecting the efficient action earlier, thereby arriving at the efficient outcome. We also investigate the case in which the efficient action is the only irreversible one. See Section 2.2 for the theoretical analysis and see Section 4.4 for the corresponding experimental results.

[8] In a recent study, Avoyan and Ramos (2019) append a stochastic revision mechanism to the minimum-effort game and show, theoretically and experimentally, that this mechanism of partial commitment can help to promote efficient coordination on high effort.

[9] Regarding the credibility of non-binding messages in one-way communication, see Farrell (1988) for the notion of self-commitment and Aumann (1990) for the notion of self-signaling. For follow-up papers on coordination games, see Baliga and Morris (2002), Sobel (2017), and Lo (2020). It is worth noting

cation in coordination games (see, among others, Cooper et al. (1992b); Charness (2000); and Blume and Ortmann (2007)). To distinguish our mechanism from pre-play communication, we conduct both theoretical and experimental analyses (see Sections 2.2 and 4.4, respectively).

The experimental literature on coordination games with pre-play moves proceeds mainly via an informal use of forward induction, asserting that the payoff-dominant outcome is achieved if the players adhere to this logic. An early paper along these lines is Cooper et al. (1992b), who find that granting an outside option with an appropriate payoff to one player significantly improves coordination efficiency. Subsequent papers add other forward-induction features such as pre-play auctions (Cachon and Camerer, 1996; Van Huyck et al., 1993) and costly messages (Blume et al., 2017). The last paper is closest to ours in treating forward-induction reasoning formally, though via the Govindan and Wilson (2009) route rather than via iterated weak dominance. Finally, Crawford and Broseta (1998) develop a model of stochastic, history-dependent learning dynamics to econometrically separate the forward-induction effect of the pre-play auction from the selection effect ("optimistic" subjects) in the data of Van Huyck et al. (1993).

By contrast, our paper formalizes the idea of forward induction by iterated weak dominance. Additionally, we find experimental evidence that, in our dynamic coordination setting, iterated weak dominance predicts the subjects' choices better than equilibrium concepts, supporting forward-induction reasoning. Forward-induction effects have been investigated in other games. Ben-Porath and Dekel (1992), Cooper et al. (1993), Brandts and Holt (1995), Huck and Müller (2005), Brandts et al. (2007), and Krol and Krol (2020) all study forward induction in Battle-of-the-Sexes and entry games. The first and the last three papers cited above follow the route of treating forward induction via iterated weak dominance, as in our paper. Ben-Porath and Dekel (1992), Brandts et al. (2007), and Krol and Krol (2020) all build in money burning as the specific mechanism to signal intentions. In money burning, only one player has the option to signal, while in our model, every player has the option to delay. It follows that the second-period information is also different. In money burning, the first-period choice of just one player is revealed. As we will see, in our game with a delay option, players get to see a message that depends on all the first-period choices made by the players.[10]

Outline The remainder of the paper is organized as follows. Section 2 presents the static benchmark model and the theory of the dynamic protocol. The experimental design and procedure are discussed in Section 3. Section 4 reports the experimental results and, Section 5 concludes the paper. Some of the proofs and statistical analyses are relegated to the Appendix.

2. Theory

We start with a static binary-action coordination game. There are $N \geq 2$ players, indexed by $i \in \mathcal{N} = \{1, 2, \ldots, N\}$. Each player i simultaneously makes a decision d_i from $D_i = \{A, B\}$. For any action/strategy profile $d \in \prod D_i$, the monetary payoff of player i is denoted by $\pi_i(d)$. For each player i, the payoff $\pi_i(\cdot)$ satisfies the following conditions: (1) if a player chooses $d_i = B$, then $\pi_i(d_i = B; d_{-i}) = b$, regardless of the other players' choices d_{-i}; (2) if a player chooses $d_i = A$, then $\pi_i(d_i = A; d_{-i}) = c < b$ for any d_{-i} that involves $d_j = B$ for some player $j \neq i$; and (3) $\pi_i(d_i = A; (d_j = A)_{j \neq i}) = a > b$.

In words, B is a safe choice, which yields a payoff of b, regardless of other players' choices, while A is a risky choice that yields a high payoff of a only if coordination is successful—that is, if all players choose A. Otherwise, if any other player chooses B, choosing A yields a low payoff of c. We further assume that each player i has the following utility function:

$$u_i(d) = \pi_i(d) + \epsilon \sum_{j \neq i} \omega_{ij} \pi_j(d), \tag{1}$$

where ϵ is positive and infinitesimal, and the w_{ij} are strictly positive numbers. We say that player i has *ϵ-social preferences*. Under this assumption, player i's preference is lexicographic, in the sense that there is no tradeoff between i's own monetary payoff and that of any other player j. Sometimes, for the sake of comparison, we will make a different assumption—namely, that player i will make a choice that hurts other players without helping i. We will refer to this as the case of "spiteful" (ϵ-social) preferences.

We assume that each player holds ϵ-social preferences and believes that all other players likewise hold ϵ-social preferences.

Proposition 1. *In the static binary-action coordination game, the pure-strategy Nash equilibria are $d_i = A$ for all i and $d_i = B$ for all i.*

Clearly, when players are selfish ($\epsilon = 0$), there are two equilibria in this simple static coordination game. Proposition 1 confirms that the set of equilibria is the same with ϵ-social preferences.

that, in the case of $N = 2$ (with one sender and one receiver), the message "A" is self-committing in one-way communication (i.e., the sender will take A at $t = 1$ if they expect the receiver to trust the message); and it is self-signaling with the assumed social preferences (i.e., the sender wants the receiver to trust the message if and only if the announcement is credible and they would choose A at $t = 1$). However, in our game, there are multiple players, and each of them can "send messages" to all the other players. In this case, since we need to consider a profile of messages, the notions of self-commitment and self-signaling are not so clearly defined (see the discussion in Blume and Ortmann (2007)).

[10] Our theory involves three rounds of iterated weak dominance. The standard money-burning mechanism necessitates more rounds, which likely accounts for a portion of the failure to obtain experimental observations consistent with theoretical predictions.

Table 2
2-player payoff matrix.

		Player 2				
		B	WBB	WBA	WAB	WAA
Player 1	B	b, b	b, b	b, b	b, c	b, c
	WBB	b, b	b, b	b, c	b, b	b, c
	WBA	b, b	c, b	a, a	c, b	a, a
	WAB	c, b	b, b	b, c	b, b	b, c
	WAA	c, b	c, b	a, a	c, b	a, a

Coordinating on the risky choice A yields the highest payoff for all players, regardless of whether or not players have the ϵ-social preferences as defined in (1). We say *efficient coordination* is achieved if and only if $d_i = A$ for all $i \in \mathcal{N}$. As shown in Proposition 1, efficient coordination is not guaranteed since coordinating on the safe action B constitutes another equilibrium.

We say *miscoordination* occurs if players fail to coordinate on a certain equilibrium—that is, if there are some players who choose A, while some other players choose B. Miscoordination incurs a loss to the players who choose A.

Next, we add a dynamic structure to the static coordination game. This enables each player to exercise a delay option so that they can choose between A and B at a later date. We will investigate how this delay option changes the outcome.

2.1. Dynamic structure with irreversible choice of B

There are two periods, $t = 0, 1$. At $t = 0$, each player chooses between B and "wait." The choice of B is irreversible. That is, if a player chooses B at $t = 0$, they cannot make any further changes. However, the players who wait at $t = 0$ get to choose between A and B at $t = 1$. There is no cost associated with waiting. By waiting, players can observe a binary message m that takes the value $m = 0$ if all players choose to wait at $t = 0$ and the value $m = 1$ otherwise.

We denote the set of pure strategies as

$$\mathcal{S} = \{B, WBB, WBA, WAB, WAA\}.$$

The strategy of not waiting and taking B at $t = 0$ is denoted by B. Any strategy involving waiting at $t = 0$ is a plan contingent on the message m. We denote such a strategy by "W, $d_i(m = 1)$, $d_i(m = 0)$," respectively, where $d_i(m)$ is defined as the action chosen conditional on m at $t = 1$. For example, if a player chooses strategy WAB, they will wait at $t = 0$ and then choose A after observing $m = 1$; otherwise, they will choose B.

For any strategy profile $s_{-i} = (s_j)_{j \in \mathcal{N} \setminus \{i\}}$ of other players, if player i chooses to wait, then the total number of B choices at $t = 0$ can be written as

$$n(s_{-i}) = \left|\{j \in \mathcal{N} \setminus \{i\} | s_j = B\}\right|,$$

and, accordingly, the binary message that player i will receive after waiting is

$$m = \mathbb{1}\{n(s_{-i}) \geq 1\}.$$

The cases $m = 0$ and 1 correspond to the "no-B" and "B" messages, respectively.

Proposition 2. *For any $s_i \in \mathcal{S}$, the strategy profile $(s_i)_{i=1}^N$ constitutes a pure-strategy Nash equilibrium. The subgame-perfect equilibria are $(s_i = B)_{i=1}^N$, $(s_i = WBA)_{i=1}^N$, and $(s_i = WBB)_{i=1}^N$.*

It is easy to see that choosing A in the subgame following the message $m = 1$ (the "B" message) cannot be part of an equilibrium in this subgame. Still, as Proposition 2 demonstrates, subgame perfection does not yield a unique outcome or imply efficient coordination. In the subgame-perfect equilibria in which all players choose $s_i = B$, or in which they all choose $s_i = WBB$, each player ends up choosing B, and, therefore, efficiency does not result.

In the following theorem, we formalize forward induction as iterated simultaneous maximal deletion of weakly dominated strategies, which we henceforth simply call iterated weak dominance. The theorem shows that this procedure yields a unique strategy profile, which achieves efficient coordination.

Theorem 1. *The unique strategy profile that survives iterated weak dominance is $(s_i = WBA)_{i=1}^N$. Under this strategy profile, efficient coordination is achieved.*

The argument involves three rounds of elimination. Here, for the purpose of illustration, we use the payoff matrix of a two-player example (see Table 2) to illustrate the elimination process. We give the main argument for each step of elimination and relegate the complete proof to the Appendix.

First round (eliminate WAB and WAA) Strategy WAB is weakly dominated by WBB. To see this, note that after the message $m = 0$, these two strategies yield equivalent outcomes. When $m = 1$, WAB involves choosing A and yields a monetary payoff $\pi_i = c$, while WBB yields a monetary payoff $\pi_i = b > c$. The same argument can be used to show that WAA is weakly dominated by WBA.[11,12]

Second round (eliminate WBB) After the first round of elimination, the remaining pure strategies are B, WBB, and WBA. Regardless of what other players choose, the realized choice under both strategies B and WBB is B. Thus, these two strategies yield the same payoff $\pi_i = b$ to any player i.

Both strategies induce the same monetary payoff to player $j \neq i$ in all but one case, in which all players $j \neq i$ choose to wait at $t = 0$, and at least some players $j \neq i$ choose strategy WBA. In this case, if player i chooses B, a player j who chooses WBA gets payoff $\pi_j = b$ from playing B after seeing $m = 1$. However, player j's payoff is reduced to c if player i chooses WBB because, in this case, player j's realized choice is A, following $m = 0$. Therefore, under the assumption of ϵ-social preferences, strategy B weakly dominates WBB.

Third round (eliminate B) The two strategies remain after the second round are B and WBA. Based on the same logic as before, each player understands that other players will play either B or WBA, provided that they believe other players hold ϵ-social preferences. If at least one player $j \neq i$ chooses B at $t = 0$, both B and WBA yield the same payoff to player i and to all other players. However, if all $j \neq i$ choose WBA, then WBA yields a strictly higher payoff to i. Thus, B is weakly dominated by WBA.

Coordination outcome Since all players choose strategy WBA, the realized message is $m = 0$, and, thus, the realized choice is $d_i = A$ for all $i \in \mathcal{N}$. Therefore, efficient coordination is achieved.

2.2. Discussion

We consider a simple binary-action coordination game with $N \geq 2$ players. By incorporating a delay option into the static game, we create a dynamic variant in which the safe but inefficient choice B is the only irreversible action. The players who exercise the delay option can observe a binary message about whether or not all players have taken the delay option. Somewhat surprisingly, there is a unique strategy WBA that survives iterated weak dominance in the resulting dynamic game. Under this strategy, a player, by giving up the safe but inefficient choice and exercising the delay option at $t = 0$, signals their intention to play the risky but efficient choice A (conditional on observing that all other players chose to wait). Under this unique strategy profile, efficient coordination is achieved. This result is built on forward-induction reasoning, which has bite only when players have ϵ-social preferences.

Next, we discuss how our result depends on the extensive form that governs the play—i.e., on the observability of the history of play and the (ir)reversibility of the actions.

Observability of past actions

In our benchmark model, players who choose to wait can observe only a binary message regarding the history of play. This is a deliberate assumption meant to capture the difficulty of observing the precise history of play in a multiple-player setting. But a delay option, per se, does not rule out cases in which players can observe more information about the past history.

Here, we consider an environment in which any player i who exercises the delay option can observe the exact number of irreversible choices that occurred at $t = 0$. We denote this number by $n(s_{-i})$ and say that this scenario exhibits *finer information*.[13] With finer information, the strategy of waiting at $t = 0$ and then choosing A at $t = 1$ if and only if $n = 0$ remains the unique strategy profile that survives iterated weak dominance. To reduce the notational burden, in the finer-information setting, we continue to write "WBA" for this strategy.

Proposition 3. *With finer information, the unique strategy profile that survives iterated weak dominance is $(s_i = WBA)_{i=1}^N$. Under this strategy profile, efficient coordination is achieved.*

Note that efficient coordination cannot be achieved as long as someone chooses B at $t = 0$; that is, $n(s_{-i}) \geq 1$, irrespective of the exact number of B choices. After any information set $n(s_{-i}) \geq 1$, the best response for any player i who waited is to play B at $t = 1$. Therefore, as Proposition 3 states, providing finer information by partitioning the information set of

[11] This round of elimination holds for both $\epsilon = 0$ and $\epsilon > 0$. In fact, both WAB (resp. WAA) and WBB (resp. WBA) yield the same payoff to each of the other players, regardless of any possible strategies they use.

[12] Strategy WBB is not weakly dominated by B in the first round of elimination. To see this, consider the case in which all other players choose WAB. In this case, compared with B, strategy WBB generates strictly higher payoffs to the other players. As such, ϵ-social preferences cannot eliminate WBB before the second round of elimination.

[13] Note that "finer information" here is different from perfect information, since the identities of the players who choose B and who choose to wait at $t = 0$ remain unknown.

$m = 1$ (“B” message) in the binary-message setting changes neither the unique strategy that players choose or change the coordination outcome. Our mechanism is robust to finer information because the intention to coordinate efficiently is signaled via an information set that is a singleton ($n = 0$), which is exactly the same as the information set $m = 0$ (“no B” message) in the binary message setting.

(Ir)reversibility structure

We have argued that a delay option can resolve the coordination problem if the inefficient choice B is the only binding choice at $t = 0$. What if both actions A and B are reversible, or if the efficient choice A, instead of B, is the only irreversible choice at $t = 0$? In this subsection, we discuss the essentiality of the reversibility structure to our result.

Neither choice is irreversible We first consider the case in which neither A nor B is irreversible at $t = 0$. More precisely, players choose between A and B at $t = 0$, but the first-period choice is not binding. At $t = 1$, they first observe the number of A and B choices at $t = 0$ and then make a final choice between A and B. Since a player’s payoff depends only on their action at $t = 1$, their choice at $t = 0$ is payoff-irrelevant.

In fact, we can interpret the play at $t = 0$ as costless pre-play communication and the play at $t = 1$ as the actual coordination game. It is easy to check that neither subgame-perfect equilibrium nor iterated weak dominance generates a unique prediction, even with the assumed ϵ-social preferences. For instance, there is a subgame-perfect equilibrium in which all players choose A at $t = 0$ and then all switch to B at $t = 1$, regardless of the information they observe. In another equilibrium, everyone chooses B at $t = 0$ and makes no change at $t = 1$, regardless of the observed information. These strategy profiles also survive iterated weak dominance and lead to inefficient outcomes.

Action A is irreversible Now consider the case in which the efficient choice A is the only irreversible choice at $t = 0$. More precisely, players first choose between “wait” and A at $t = 0$. After observing the number of wait and A choices at $t = 0$, players who waited choose between A and B at $t = 1$. This setting has an irreversibility structure that is the opposite of our benchmark setup. The following proposition shows that iterated weak dominance and subgame-perfect equilibrium predict the efficient outcome only in the case of $N = 2$. For $N \geq 3$ players, these theories fail to yield (uniquely) the coordination outcome.

Proposition 4. *In a coordination game in which A is the only irreversible choice at $t = 0$:*

1. *if $N = 2$, both iterated weak dominance and subgame-perfect equilibrium yield efficient strategy profiles;*
2. *if $N \geq 3$, both iterated weak dominance and subgame-perfect equilibrium allow inefficient strategy profiles.*

As shown in the proof of Proposition 4 in the Appendix, efficient coordination can be achieved under iterated weak dominance as well as under subgame-perfect equilibrium in the case of $N = 2$ players. The underlying mechanism, however, is fundamentally different from the one highlighted in this paper. Rather than signaling intention through waiting, the mechanism in this case is based on the fact that each player has a weakly dominant choice of A earlier. This is because: (1) if one player takes A early and the other player waits, it is optimal for other player to follow by taking A at $t = 1$; and (2) if one player takes A early and the other player also takes A at $t = 0$, efficient coordination has already been achieved at $t = 0$.

However, for $N \geq 3$ players, neither iterated weak dominance nor subgame-perfect equilibrium guarantees efficiency. For example, the strategy profile in which all players choose to wait at $t = 0$ and choose A at $t = 1$ only if all other players have chosen A at $t = 0$ survives iterated weak dominance, and it constitutes a subgame-perfect equilibrium.[14] Under this strategy profile, each player’s final choice is B, and, thus, efficient coordination is not achieved.

One might think that, in this setting, efficient coordination is easier to achieve because, now, any player can take the lead by committing to the binding choice A early, at $t = 0$. Intuitively, these early A choices would encourage other players to follow. However, Proposition 4 says that this intuition is false. The only message that ensures that a player will follow and choose A at $t = 1$ is that all other players have chosen A at $t = 0$. When the game is played by $N \geq 3$ players, it is possible that all players who waited would choose B at $t = 1$ if more than one player waited at $t = 0$, as they anticipate that other players will do the same. Efficient coordination cannot be guaranteed.

2.3. Summary

Our main results show that an option to delay facilitates efficient coordination by allowing players to signal their intentions. The above discussion demonstrates, at a theoretical level, that signaling intentions by waiting is different from both signaling intentions by non-binding communication (when both actions are reversible) and signaling intentions by taking the efficient action early (when the efficient action A is irreversible). Therefore, although the mechanism we emphasize is robust to the availability of finer information, it relies crucially on the irreversibility of the inefficient choice.

[14] This strategy is, in fact, the most frequently used one in our experiment.

For completeness, we extend the model further to show that the delay mechanism can work in a more general coordination game, in which successful coordination does not require all players to choose the efficient choice A. We also discuss the case in which both actions are irreversible choices and the case in which delay is costly.[15] Since these extensions are not essential to our theoretical analysis and experimental tests, we relegate them to the Online Appendix.

3. Experimental design

Our theory demonstrates that the dynamic structure with an irreversible B choice admits a unique prediction of efficient coordination via iterated dominance. However, the inferior outcome still qualifies as a subgame-perfect equilibrium, even with the assumption of ϵ-social preferences. Therefore, we experimentally test the efficacy of this delay structure and check whether participants' choices are consistent with the theoretical prediction based on iterated dominance.

Since the theory speaks to games with multiple players, we do not restrict ourselves to a two-player group and, instead, explore four-player coordination games in the experiment. A well-developed literature that studies multi-player coordination conduct experiments on the minimum-effort, or the weakest-link, games. The coordination game we consider can be interpreted as a binary-action minimum-effort game played by $N \geq 2$ players, with high-effort level A and low-effort level B, as the group coordination is determined by the lowest choice in the group. To make our experimental findings comparable to those of the existing studies, in our main treatments, we follow the design of the minimum-effort games literature, which started with Van Huyck et al. (1990).

Following the standard protocol in the literature on minimum-effort games, our subjects played a game for 15 rounds in fixed four-person groups in the main treatments. Subjects' strategies were elicited using the strategy method in the dynamic games. The parameters chosen were $a = 55$, $b = 45$, $c = 5$, and $N = 4$. A more detailed description of the experimental implementation will be given in Section 3.3. In addition, since fixed matching might invoke learning from the previous rounds of play and other dynamic concerns for future play, we also implemented the main treatments with randomly matched groups.

3.1. Main treatments

The main treatments compare the coordination efficiency in static games and the dynamic games with the irreversible B choice.

Static game ("St-b" and "St-b-rand")

The two static treatments were the static version of the binary-action coordination game with "binary feedback" ("b" for short): at the end of each round, subjects were told whether the efficient outcome had been achieved. Feedback only about the coordination outcome was the standard protocol in the minimum-effort literature; that is, subjects observed only the minimum effort chosen in the previous rounds. Random-matching treatments are denoted as "rand" throughout the paper. For example, "St-b" and "St-b-rand" stand for the static treatments with fixed- and random-matching groups, respectively.

Dynamic game with irreversible B action ("BI-b" and "BI-b-rand")

The main treatments followed the dynamic structure proposed in Section 2.1 (see Proposition 2 and Theorem 1) with fixed-matching ("BI-b") and random-matching ("BI-b-rand") groups. In these dynamic games, B was the only irreversible action ("BI" for short), and, at the end of each period, subjects receive binary information ("b" for short) about whether or not B had been chosen thus far. Each round of the game consisted of two periods. At $t = 0$, each subject chose between B and the "Wait" option. If a subject chose to wait, they would receive a binary message about whether someone in the group had chosen B in $t = 0$, and they would then choose between A and B in $t = 1$. At the end of $t = 1$, all subjects were informed of the binary outcome of whether efficient coordination had been achieved. Following the protocol in the minimum-effort literature, they did not observe the exact numbers of the irreversible choices.

3.2. Additional treatments

Additional treatments were implemented to further test the robustness of the efficacy of delay option in promoting efficient coordination and the underlying mechanism, including a treatment with a reversible choice of A in the first period to mitigate the potential framing effect, treatments with finer information, and alternative reversibility structures.

Binary information with three options ("BI-b-3c" and "BI-b-3c-rand")

One potential concern about the dynamic treatment "BI-b" was that, some subjects might interpret "B not chosen by anyone in the first period" (the "no-B" message), literally (simply following the face value of this message) as "no one will

[15] In a more general class of coordination games, or when both A and B are irreversible in our benchmark setup, our result is sensitive to the information available to the players who exercise the delay option. Specifically, our result holds in the binary-information setting but does not hold in the finer information-setting. Moreover, although the mechanism of signaling intentions can still work with a costly delay option, this mechanism cannot ensure efficient coordination.

choose B," rather than "everyone decided to wait and see." Thus, it might bias the results in a direction that favors our theory's prediction. To mitigate this framing effect, we added the treatments "BI-b-3c" and "BI-b-3c-rand," which included A in the first period as a reversible choice. The additional reversible A choice does not affect our theoretical results, but with this additional reversible option, the "Wait" choice at $t = 0$ should not be interpreted literally as a choice that is biased toward the choice of A at $t = 1$. Additionally, the binary message after the first period in all "3c" treatments was framed as "nobody (someone) chose B in the first period, and (not) everyone chose 'Wait' or A in the first period." Altogether, these settings stressed the fact that the face value of the "no-B" message is merely "B was not chosen in the first period," which helped mitigate the framing effect.

Finer-information treatments ("St-f", "BI-f" and "BI-f-rand")

In addition to the main treatments with binary feedback, we also tested the finer-information versions of these two treatments: "St-f" (static, finer information), "BI-f" (B-irreversible, finer information), and "BI-f-rand" (B-irreversible, finer information, random-matching). In contrast to the binary information setting, all finer-information treatments ("f" for short) enabled the subjects to observe the number of B choices at the end of first period (in dynamic treatments) and the number of B choices as final choices at the end of the second period. More precisely, in the "BI-f" treatment, if a subject decided to wait at $t = 0$, they would face four possible situations: everybody waited; or 1, 2, or 3 group members chose B. Therefore, the subject's strategy would be whether to wait at $t = 0$, and, if they waited, a full plan on these four contingencies.

There is mixed evidence in the minimum-effort literature about whether providing finer information alters subjects' behavior. Van Huyck et al. (1990) found that the finer-information setting did not affect coordination efficiency, while in the "full feedback" treatment of Brandts and Cooper (2006b), efficiency was significantly improved. The finer-information treatments serve as a further test of our theoretical results. Based on Proposition 3, the same efficient outcome could be generated with the delay option in the finer-information treatment, "BI-f."[16] Moreover, since the alternative irreversibility structures were theoretically studied and experimentally tested based on the finer-information setting,[17] examining the finer-information ("BI-f") treatments allowed for a fair comparison across different irreversibility structures.

Alternative irreversibility treatments ("NI-f" and "AI-f")

We also tested the two alternative irreversibility structures discussed in Section 2 to distinguish our delay option with other potentially efficiency-enhancing dynamic mechanisms. In the "NI-f" (neither action being irreversible, finer information) treatment, both the choices of A and B at $t = 0$ were reversible. At $t = 0$, subjects chose between A and B. There was no wait option in the first period. Then, at $t = 1$, upon observing the distribution of the choices from $t = 0$, they could freely switch to the other choice at no cost. Under this dynamic setting, a player could still express their intention to play A or B, but in a non-binding way.

In the "AI-f" (A-irreversible, finer information) treatment, only the A choice was binding at $t = 0$. In $t = 0$, subjects chose between A and the wait option. Then, in $t = 1$, those who chose the wait option could decide between A and B after observing the number of A choices at $t = 0$. This delay structure allowed a player to credibly signal their intention to choose the efficient action A. However, as discussed in Section 2.2, there is no unique prediction of efficient coordination by SPNE or weak dominance for a four-person group ($N = 4$). The "AI-f" treatment further helped us understand whether the irreversibility structure is essential for the delay mechanism.

3.3. Experimental procedure

Fixed-matching sessions

The fixed-matching sessions were implemented by a web-based program and by Otree (Chen et al., 2016) in the Smith Lab at Shanghai Jiao Tong University in 2019 and 2021. A total of 396 undergraduate and graduate students participated in 20 sessions. At the beginning of each session, each subject arriving at the lab was randomly assigned a seat number. Subjects were then randomly put into groups of four that remained fixed throughout the sessions.

We adopted a between-subject design. In each session, subjects played the game from one treatment for 15 rounds with their group mates. The choices were labeled "1" and "2" instead of "A" and "B." There was no time limit for making the choices.

In the static treatment, subjects simply submitted their choices of "1" or "2" in each round. In the dynamic treatments, subjects' complete strategies were elicited using the strategy method. For example, on the choice page of our main treatment ("BI-b"), subjects were first asked to choose between "1" and "Wait." If their choice was "Wait," then two additional choices would appear, asking them to choose an action for each of the two possible realizations of the message, $m = 0, 1$. Subjects were made aware that only one of the choices would be realized, based on the outcome in the first period. Instead, if

[16] In addition, there is a minor concern that relates to the framing effect. In the "BI-b" treatment, subjects may have felt tempted to choose differently for the $m = 0$ and $m = 1$ messages, thereby inducing more choices of WBA and WAB than of WBB and WAA. The finer-information treatment helped avoid this.

[17] For alternative irreversibility structures, it is reasonable to focus our analysis on the finer-information setting. For example, when both actions are reversible, it is natural to allow subjects to observe the number of A and B choices at $t = 0$, as in Blume and Ortmann (2007).

Table 3
Experimental design.

Fixed-matching	# Sessions	# 4-player Groups
"St-b" (static, binary info)	5	21
"BI-b" (B irreversible, binary info)	5	21
"St-f" (static, finer info)	2	11
"BI-f" (B irreversible, finer info)	2	12
"BI-b-3c" (B irreversible, binary info, three choices available in the first period)	2	10
"NI-f" (neither irreversible, finer info)	2	12
"AI-f" (A irreversible, finer info)	2	12
Random-matching	**# Sessions**	**# 8-player Matching cohorts**
"St-b-rand"	3	8
"BI-b-rand"	3	8
"BI-f-rand"	2	6
"BI-b-3c-rand"	2	6

any subject's first-period choice was "1," then there would be a notice telling them that they did not need to make any choice for the second period. However, the subject still needed to click a "confirm" button for each possible realization of the binary message to finish this round. With these two "confirm" buttons, the total number of clicks would be the same whether a subject chose to wait or not to wait at $t = 0$. Thus, subjects would not be able to infer others' choices from the number of clicks.

In the finer-information treatments, after choosing "Wait" (or either of the two actions in the "NI-f" treatment), the four possible outcomes from the first period would appear, and the subject needed to choose an action for each of the four contingencies.[18] If a subject chose not to wait, then they needed to click on the four "confirm" buttons.

At the beginning of the experiment, the instructions were first read aloud in the lab. Then, the subjects completed a short comprehension test before the 15-round play of the experiment. After all participants finished the experiment, we gave them unincentivized and anonymous questionnaires about their decision rules. Participants had not been informed about the questionnaires beforehand. At the end of each session, subjects were paid based on their cumulative payoffs from all rounds (1 point was converted into 0.07 RMB). Each session took about 45 minutes, and the average earnings were 55 RMB (or 8.5 USD), including a participation fee of 5 RMB. The numbers of subjects in each session and treatment are summarized in Table 3.

Random-matching sessions

The random-matching sessions, "BI-b-rand," "St-b-rand," "BI-f-rand," and "BI-b-3c-rand," were computerized using Otree (Chen et al., 2016) and took place in 2021 in Shanghai Jiao Tong University. A total of 224 undergraduate and graduate students participated in the ten new sessions.

In the random-matching treatments (Table 3), each matching cohort consisted of eight subjects, randomly matched to play the four-player game within the cohort. We chose this cohort size to generate the largest number of independent observations. To limit the motive of learning and exploration, subjects played the game for only ten rounds. Though the matching cohort was relatively small, it still served as a comparison with the fixed-group sessions, and provided us with insights from different matching protocols.

In addition, in all random-matching sessions with the delay option, the "no-B" message was modified to alleviate framing effects. In their fixed-matching counterparts, the "no-B" ($m = 0$) message read "nobody chose B in the first period," while the revised "no-B" message read as "nobody chose B in the first period, and everyone chose 'Wait'." (See the discussion of "BI-b-3c" and "BI-b-3c-rand" in Section 3.2 for more details.)

After the main experiments in the random-matching sessions, an extra block was added to elicit subjects' social preference and their belief in others' social preference. The details will be given in Section 4.3.

At the end of each session, subjects were paid based on their payoffs from two randomly selected rounds and their earnings in the social preference block (1 point was converted into 0.5 RMB). Each session took about 50–60 minutes, and the average earnings were 70 RMB (or 11 USD), including a participation fee of 10 RMB.

4. Experimental results

We first compare the efficiency rates, defined as the percentages of groups that achieved efficient coordination, between the static and the dynamic "BI" treatments, followed by a decomposition of the strategies adopted in the "BI" treatments.

[18] In all treatments with finer information, there were $N = 4$ contingencies that could arrive at the end of $t = 0$—specifically, all four possible numbers $(0, 1, 2, 3)$ of the irreversible choices made by the other three group members in the "BI-f" and "AI-f" treatments. The same held true for the number of B choices in the "NI-f" treatment.

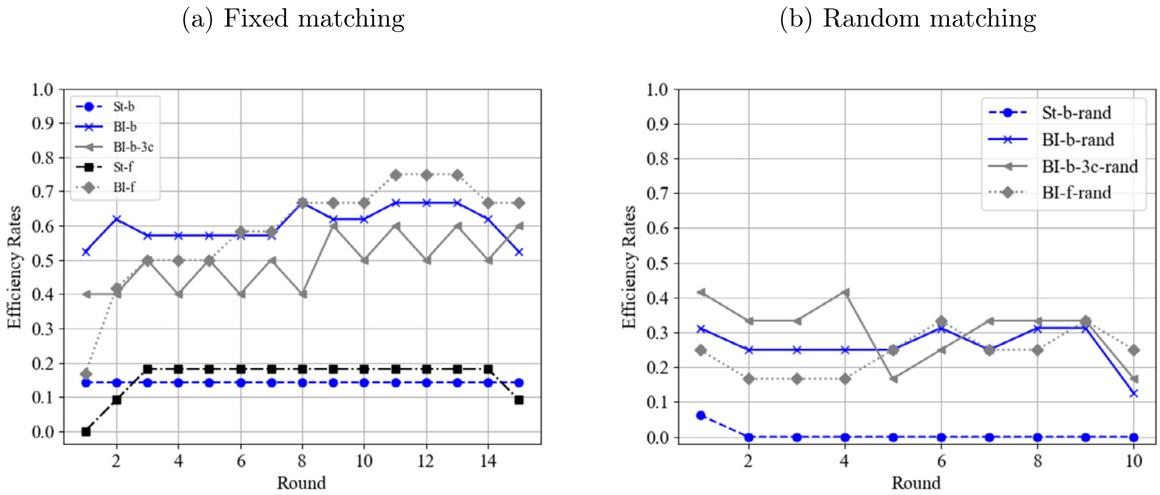

Fig. 1. Group efficiency rates.

4.1. Main results: efficiency of the B-irreversible structure

Result 1 *(Group-level efficiency). The efficiency rates were significantly higher in the dynamic treatment ("BI-b" and "BI-b-rand") than in the static treatment ("St-b" and "St-b-rand").*

Fixed-matching: main treatment Fig. 1 plots the frequency of an efficient outcome in the static and B-irreversible treatments. Efficient coordination was difficult to achieve in the static games "St-b," which replicates the findings in the minimum-effort game literature (see, for example, Van Huyck et al. (1990); Brandts and Cooper (2006a)). Only 14 percent of the groups managed to coordinate on the efficient A choice in "St-b." In sharp contrast, the average efficiency rate over the 15 rounds was significantly higher—above 60 percent—in the dynamic game with the B-Irreversible structure, "BI-b." This high rate of efficient coordination was sustained overall. The regression results presented in Column 2 of Table 4 suggest that the difference in the efficiency rates was statistically significant.

In Table 4, we examine other measures of coordination outcomes: frequencies of A as final choices; group average payoff; and rates of coordination.[19] Significantly more subjects made A as their final choices in "BI-b" than in "St-b" (Column 1), and the groups coordinated much better on either of the choices (Column 4). The higher incidence of efficient coordination in the dynamic setting led to significantly higher average payoffs (Column 3 of Table 4).[20]

Random matching: main treatment Overall, in both the static and dynamic games, the efficiency rates were relatively lower in all of the random-matching treatments (Fig. 1(b)) than in the fixed-matching treatments, and the differences were present in the first round.[21] This result indicates that the dynamic concerns and the exploration motive might play a role in facilitating coordination in the fixed-matching sessions. More details about these motives will be discussed in Section 4.2, where we examine the decomposition of the adopted strategies.

Despite this, we still observe a significant difference between "St-b-rand" and "BI-b-rand" in the efficiency rates and the frequency of A as the final choice (Columns 1 and 2 of Table 5). In particular, the higher efficiency rates in "BI-b-rand" could be sustained over the rounds, which is consistent with the pattern found in "BI-b."[22]

However, the differences in payoff and coordination rates were not significant (Columns 3 and 4 of Table 5), despite being in the right direction. To understand this, we performed an in-depth data analysis and found that, although "BI-b-rand" yielded significantly higher efficiency rates and slightly lower miscoordination, many more subjects ended up having A as their final choice in the miscoordinated groups in "BI-b-rand" than in "St-b-rand." These subjects received the lowest

[19] The rate of coordination was calculated as the percentage of groups in which all members made the same choice. Miscoordination occurs when there is at least one subject whose realized choice is A, while some others choose B.

[20] The average payoff in the static games was 42.4 per round, while it was 49.6 per round in the dynamic games. Given the fact that the highest possible payoff (when achieving efficient coordination) was 55, and the B choice secured a payoff of 45, the dynamic structure with the irreversible B choice significantly recovered the efficiency loss in the static games.

[21] Statistical analysis comparing the fixed- and random-matching sessions can be found in Appendix C.1.

[22] To understand the smaller efficiency rate gap observed in the random-matching treatments, it is important to note that the chosen payoff parameters ($a = 55, b = 45, c = 5$) made efficient coordination extremely difficult in the static games. With these parameters, the efficiency rate in the fixed-matching treatment, "St-b," was only 14 percent, and there was not much room for the decline, and in the random-matching treatment, "St-b-rand," it was extremely low (0.6 percent). This may also contribute to the insignificant payoff gap observed in the random-matching treatments.

Table 4
Group-level regression analysis (fixed-matching).

	(1) A_rate	(2) effi_rate	(3) payoff	(4) coor_rate
St-b	-0.283***	-0.460***	-7.238***	-0.200***
	(0.0266)	(0.1246)	(0.7278)	(0.0303)
BI-f	0.0115	-0.0143	-0.941	-0.0333
	(0.0460)	(0.1508)	(1.0871)	(0.0234)
St-f	-0.268***	-0.446***	-7.391***	-0.194***
	(0.0225)	(0.1414)	(0.7392)	(0.0347)
BI-b-3c	-0.0835	-0.110	-0.670	0.0267
	(0.0670)	(0.1674)	(1.2799)	(0.0222)
Constant			49.60***	
			(0.4434)	
R^2			0.121	
Pseudo R^2	0.123	0.152		0.0893
N	1125	1125	1125	1125

Notes: Standard errors clustered at the group level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Reference category is "BI-b." Each observation is a group-average level in a round. Dependent variables (and the regression models used) are (1) percentages of A as final choices (tobit); (2) efficient outcome dummy (probit); (3) group average payoff (OLS); and (4) the dummy for coordination on either action (probit). Marginal effects are reported for tobit and probit regressions.

Table 5
Group-level regression analysis (random-matching).

	(1) A_rate	(2) effi_rate	(3) payoff	(4) coor_rate
BI-b-3c-rand	0.0103	0.0378	1.146	0.0218
	(0.1456)	(0.1295)	(1.7964)	(0.0438)
BI-f-rand	-0.0217	-0.0166	-0.0208	-0.0137
	(0.1400)	(0.1303)	(2.3104)	(0.0462)
St-b-rand	-0.216**	-0.171**	-1.125	-0.00906
	(0.1026)	(0.0821)	(1.7084)	(0.0650)
Constant			43.69***	
			(1.0857)	
R^2			0.0126	
Pseudo R^2	0.239	0.278		0.0197
N	280	280	280	280

Notes: Standard errors clustered at matching cohort level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Reference category is "BI-b-rand." Each observation is a matching-cohort-average level in a round. Dependent variables (and the regression models used) are (1) percentages of A as final choices (tobit); (2) efficient outcome dummy (probit); (3) group average payoff (OLS); and (4) the dummy for coordination on either action (probit). Marginal effects are reported for tobit and probit regressions.

possible payoff 5. As a result, we do not observe a significant improvement in the average payoffs in "BI-b-rand." For a detailed discussion, see Section C.2 in the Appendix.

Result 2 *(Group-level efficiency: robustness). The efficacy of the delay option was found to be robust to the treatments that introduced an additional reversible choice "A" ("BI-b-3c"), or allowed subjects to observe finer information ("BI-f").*

Framing effect Regression analyses (Table 4 and Table 5) demonstrate no significant difference in any of the four measures between "BI-b" and "BI-b-3c" and between "BI-b-rand" and "BI-b-3c-rand." In addition, a significant improvement in coordination efficiency with respect to "St-b" ("St-b-rand") was still observed in "BI-b-3c" ("BI-b-3c-rand") (see Table 18 in the Appendix). Thus, although framing might contribute to the efficacy of the delay option observed in the main treatments "BI-b" and "BI-b-rand," there is a lack of support for the notion that the significant improvement in the efficiency rate was caused mainly by the framing effect.

Furthermore, since all the random-matching sessions adopted the new presentation of the "no-B" message, the results here also suggest that additional framing effects potentially associated with the presentation of "no-B" message did not account for the observed improvement in coordination efficiency.

Finer information The results from the "BI-f" treatment confirm the theoretical prediction (Proposition 3) that a delay option implements the efficient coordination. As shown in Fig. 1(a), a significantly large gap in the group-level efficiencies can be observed between "BI-f" and "St-f." The regression results reported in Table 9 show that the difference is significant. In addition, the regression results in Tables 4 and 5 show that the efficiency rates of the finer-information treatment "BI-f" (or "BI-f-rand") were not significantly different from those of the binary-information treatment "BI-b" (or "BI-b-rand").

4.2. Main results: adoption of strategies

Result 3 *(Adopted strategies). In all dynamic treatments with an irreversible B choice (i.e., the "BI" treatments), the majority of the subjects took the unique iteratedly undominated strategy WBA in both fixed-matching and random-matching sessions.*

The strategy method allowed us to decompose the strategies adopted in the dynamic treatments. Fig. 2 plots the distribution of strategies B, WBB, WBA and the dominated strategies (WAB and WAA) adopted by subjects in the "BI" treatments. Consistent with the theoretical prediction, the vast majority of the subjects adopted the unique iteratedly undominated strategy WBA. In "BI-b," the proportion of WBA choices was above 70 percent across all rounds.[23] By contrast, the other two strategies consistent with the SPNE predictions, B and WBB, were adopted much less frequently, with 15 percent of the subjects choosing B and ten percent choosing WBB, on average, over time.[24]

Strategy WBA was chosen by the majority of subjects in the random-matching treatment "BI-b-rand," as well as in the additional treatments with the B-irreversible structure.[25] Across different treatments, although the frequency varied, the proportions of subjects who took the strategies that could be categorized as WBA were, overall, greater than 60 percent (see Fig. 2).

However, we do observe a difference between the two matching protocols. Compared with "BI-b," a higher proportion of subjects chose B in "BI-b-rand," and the difference was present in the first round, though it was only marginally significant (Table 19 in the Appendix). This suggests that the exploration motive and other dynamic concerns might have been contributing marginally to the high frequency of waiting observed in the fixed matching treatments. More statistical analysis is reported in Table 13 in the Appendix.

Framing effect The "BI-b-3c" and "BI-b-3c-rand" treatments allowed subjects to choose the reversible option A at $t = 0$ and, thus, adopt the strategies ABA or ABB. Theoretically, these two strategies are identical to WBA and WBB, respectively. They can, however, be different if the name of the reversible action carries some meaning. To test for a framing effect, we compared the frequency of B choices in the first period and that of A choices after the "no-B" message in the second period between "BI-b-3c" and "BI-b" and between their random-matching counterparts. The results reported in Table 6 indicate no significant differences in the frequency between the treatments in comparison. Therefore, the high adoption rate of the WBA strategy was unlikely to have been driven by the framing effect of the binary information.

We performed an additional analysis to examine whether the labeling of the reversible choice at $t = 0$ had different implications for the play at $t = 1$. Recall that, the choice of A and of "Wait" (in "BI-b-3c" and "BI-b-3c-rand") both lead to the option to make a choice in the second period, and they were indistinguishable to the teammates under binary information. Therefore, in theory, the labeling of the first-period reversible choice should not have had an impact on the choice at the second period. We decomposed the second-period observations from "BI-b-3c" and "BI-b-3c-rand" into two groups based on the first-period choices and report the results in Column 4 of each panels in Table 6. Those who chose A in the first period tended to choose A after the "no-B" message more often than those who chose "Wait" (see Column 3 of each panel in Table 6). However, the difference was not significant in the fixed-matching sessions and was only marginally significant in the random-matching sessions, suggesting minor effects of the face value of the strategy labeling.

[23] Across rounds, the proportion of WBA in the "BI-b" treatment dropped slightly over time and ended up being 71 percent in the last round. Note that subjects could identify whether a WBB choice in the group was the reason for the inferior coordination outcome via the feedback following $t = 0$. They could also learn the presence of WBB choices if they ended up getting a payoff of 5 for that round. This learning could cause them to switch from WBA to WBB or B. An analysis of individual choices (see Appendix B for details) suggests that some participants who started with WBA switched to B or WBB because they were hurt by their group mates who chose WBB.

[24] In particular, there was an increase in the proportion of the subjects who chose B, from seven percent in the first round to over 15 percent in the later rounds, possibly due to the inferior outcome in the previous rounds. It is not surprising that a small fraction of subjects chose strategies B and WBB. This can be understood with the help of our model. Choosing B over WBB might suggest that the subject had a social preference not to hurt others but did not believe that group mates had such preferences. For the ten percent of subjects choosing WBB, one plausible explanation would be the presence of selfish or spiteful social preferences.

[25] Recall that we refer to the following strategy as WBA in the "BI-f" and "BI-f-rand" treatments: "wait in the first period; choose A if $n = 0$; and choose B if $n \geq 1$," where n is the number of B choices in $t = 0$. Any strategy involving choosing A after observing $n \geq 1$ is called a dominated strategy.

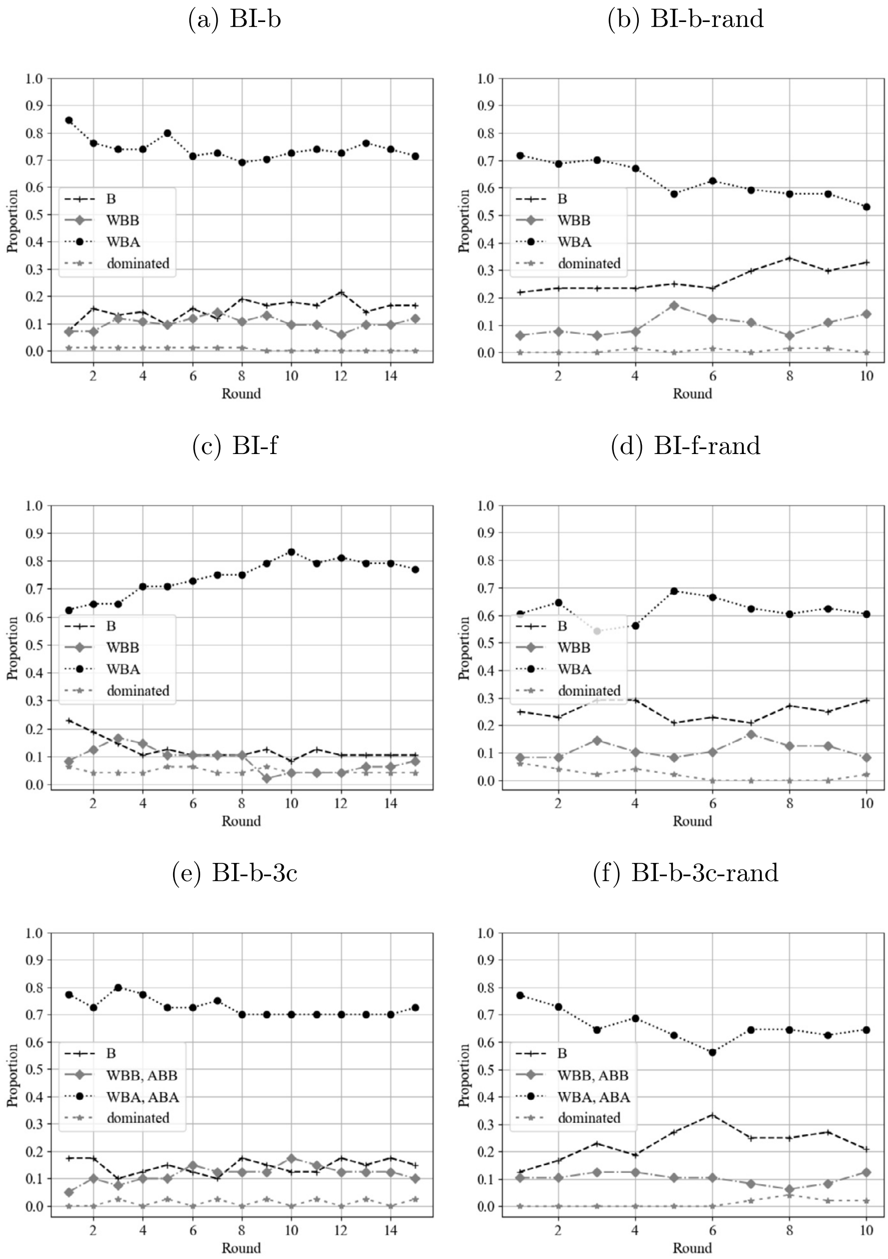

Fig. 2. Decomposition of strategies in the "BI" treatments.

Finer information Proposition 3 predicts that subjects would take the strategy equivalent to WBA with finer information about the strategies chosen by their group mates. The majority of subjects' choices in the experiment are consistent with this theoretical prediction. To better understand how finer information changes the subjects' adopted strategies, we further compared the frequency of B choices at $t = 0$ and the A choice after the "no-B" message at $t = 1$ between "BI-f" ("BI-f-rand") and the main treatment "BI-b" ("BI-b-rand"). The results, reported in Table 6, indicate no significant difference in frequency.

Learning motive There was a distinguishable difference in the frequency of WBA between "BI-b" and "BI-f" in the first two rounds. More subjects chose B and the dominated strategies in "BI-f," which resulted in lower efficiency rates in these early rounds, though the rate rapidly caught up in the later rounds.[26] Such differences in the early rounds might have been due

[26] The pattern of increasing efficiency in the first couple of rounds is also present in Avoyan and Ramos (2019) and Brandts and Cooper (2006b), in which finer information is available. In the treatments with finer information, subjects could observe the exact number of A and B choices in the previous rounds,

Table 6
Individual-level regression.

Panel A: fixed-matching				
	Reference = BI-b			Reference = BI-b-3c: A
	(1)	(2)	(3)	(4)
	B in t0	A after no-B	A after no-B	A after no-B
BI-f	-0.0270	0.0247		
	(0.0576)	(0.0644)		
BI-b-3c	-0.00572	-0.0105		
	(0.0643)	(0.0979)		
BI-b-3c: A			0.0865	
			(0.0562)	
BI-b-3c: Wait			-0.212	-0.300
			(0.1974)	(0.1934)
Pseudo R^2	0.00171	0.00381	0.0691	0.207
N	2580	2214	1583	513
Panel B: random-matching				
	Reference = BI-b-rand			Reference = BI-b-3c-rand: A
	(1)	(2)	(3)	(4)
	B in t0	A after no-B	A after no-B	A after no-B
BI-f-rand	-0.0147	-0.0198		
	(0.0904)	(0.0729)		
BI-b-3c-rand	-0.0381	0.00414		
	(0.1013)	(0.0547)		
BI-b-3c-rand: A			0.0553	
			(0.0561)	
BI-b-3c-rand: wait			-0.0298	-0.0872*
			(0.0605)	(0.0495)
Pseudo R^2	0.00357	0.00156	0.00962	0.0215
N	1600	1196	836	366

Notes: Standard errors clustered at the group level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Probit regressions. Reference categories are "BI-b" for Panel A and "BI-b-rand" for Panel B. Each observation is an individual subject in a round. Dependent variables are choice of B in $t0$ (dummy) and choice of A after the no-B message (dummy). Other control variables include Rounds 1–5 (dummy), Rounds 6–10 (dummy), and Rounds 11–15 (dummy). Marginal effects are reported.

to subjects' incentive to learn about other players' strategies so as to facilitate future play.[27] This incentive is referred to as the learning motive.

In "BI-b," if a subject chose B at $t = 0$, then under binary information, they would not observe whether other subjects also chose B at $t = 0$. Therefore, the subjects who wished to learn about this information might have chosen "Wait" at $t = 0$, implemented by either WBA or WBB. By contrast, with finer information, subjects would have learned about the number of B choices at $t = 0$ no matter whether they chose B or "Wait." This learning motive can be stronger in the earlier rounds, and it may increase the choices of WBA in "BI-b," while it should not influence the adopted strategies when finer information is available. Therefore, such a learning motive may help explain the observation that B (WBA) was chosen more (less) frequently in the early rounds of "BI-f," as compared with "BI-b."

This learning motive should have been weakened in the "BI-b-rand" treatment since groups were randomly shuffled each round. Consistent with this line of reasoning, the WBA choices accounted for around 70 percent in the first round of "BI-b-rand," significantly lower than that in the "BI-b" treatment.[28] In addition, as Fig. 2 shows, the difference in the early rounds between random-matching treatments "BI-f-rand" and "BI-b-rand" was less pronounced than that between "BI-f" and "BI-b."

which might have facilitated coordination. For example, if a subject realized that they are the only one playing B in one round, they might have wanted to switch from B to A in the next round.

[27] Another possibility is the complexity associated with the finer-information treatment, due to which some subjects were confused by the four contingencies and, therefore, chose the most conservative strategies at the beginning of the experiment. In addition, the challenge from implementing the strategy with four contingencies might have lead to more choices of dominated strategies.

[28] A possibly better metric to examine this learning motive was the frequency of "Wait" at $t = 0$ vs. that of WBA choices. The frequency of waiting choices was 77 percent in the first round of the "BI-f" treatment, which is much lower than the 92 percent in "BI-b" treatment. Interestingly, the frequency was almost identical to that in the "BI-b-rand" treatment (78 percent), under which the learning motive is largely absent. We thank an anonymous referee for suggesting this analysis.

Table 7
Social preferences and use of strategies.

Choice	Belief	N	pct	B	WBB	WBA
Y	Y	144	90.0%	0.23	0.072	0.68
				(0.028)	(0.014)	(0.03)
Y	N	6	3.8%	0.72	0.23	0.05
				(0.17)	(0.17)	(0.034)
N	Y	2	1.2%	0.45	0.1	0.45
				(0.35)	(0.1)	(0.45)
N	N	8	5.0%	0.17	0.59	0.23
				(0.073)	(0.1)	(0.08)

Notes: the data are from the "BI-b-rand," "BI-f-rand," and "BI-b-3c-rand" sessions. The "Choice and "Belief" columns represent whether a subject made a choice consistent with the ϵ-social preferences, and whether the subject believed in the ϵ-social preferences of other participants. The final three columns of Table 7 report the frequency of adopted strategies across the four categories of social preferences and beliefs about the social preferences of others. Standard errors are in the parentheses.

While the learning motive may have influenced the waiting decisions, it is unlikely to be the primary reason for the observation that the majority of subjects chose the strategy WBA. This is because, even in the treatments "BI-f" and "BI-f-rand," in which learning motive should have been significantly weakened, the vast majority of subjects still chose WBA.

4.3. Social preferences and adoption of strategies

According to our theory, the strategy of WBA becomes the unique iteratedly undominated strategy when the players have ϵ-social preferences and believe that other players hold these preferences. To explore the underlying mechanism of the delay option, we conducted additional experiments to examine whether subjects' social preferences and their beliefs about others' preferences were correlated with the adopted strategies.

Result 4 *(Social preferences). The lack of ϵ-social preferences is positively associated with the choice of WBB. The lack of belief that other players hold ϵ-social preferences is positively associated with the choices of B and WBB.*

In the sessions with randomly-matched groups, an additional block was added after the main experiments. The block consisted of two choice problems. The first one asked the subjects to choose from two allocations of experimental points between themselves and a randomly selected participant. The two options were (15, 15) and (15, 5) in experimental points for oneself and the other participant. Since the choice affected only the payoff of the other participant, a player with ϵ-social preferences would have selected the first option, while a spiteful subject would have selected the latter one. We consider the choice of (15,15) to be an indicator of ϵ-social preferences.

The second question elicited subjects' beliefs about a randomly selected subject's response to the first question. A correct prediction would yield a payoff of 5 experimental points. Subjects who predicted that a randomly selected participant would choose (15, 15) are assumed to have believed that other subjects in the game had ϵ-social preferences.

The social preference block's findings are summarized in Table 7. The subjects are categorized into four groups according to their elicited social preferences and beliefs about the social preferences of others. The last three columns of Table 7 report the frequencies of adopted strategies within the four categories. Over 90 percent of the 160 subjects who participated in the three "BI" treatments of the random-matching sessions made the choices consistent with ϵ-social preferences and also believed that other players had ϵ-social preferences. These findings support our assumption that the majority of the subjects held the ϵ-social preferences and beliefs in this type of preference. Consistent with our theory, this group exhibited a greater propensity than others to choose the unique iteratively undominated strategy WBA than other groups.

According to our theory, if a player has ϵ-social preferences but does not believe that other players hold ϵ-social preferences, both WBA and B survive iterated weak dominance. Without ϵ-social preferences, regardless of players' beliefs about other players' preferences, all three strategies, WBA, WBB, and B, survive iterated weak dominance. The regression results in Table 8 substantiate our theory. The absence of ϵ-social preferences was significantly associated with the choice of WBB, but negatively predicts the choice of WBA. Additionally, not believing that other players held ϵ-social preferences resulted in an increase in both B and WBB, and a decrease in WBA.

The regressions were complemented by the non-parametric Mann–Whitney U tests, which show similar results (Table 16 in the Appendix). Furthermore, we performed additional tests to show that the results from the social-preference block were not affected by subjects' experience in the previous rounds of experiments and, thus, were likely to reflect their innate preferences (see Table 17). These additional analyses further confirm that the association between the social preference (and belief about social preference) and the chosen strategy is robust to these considerations. Please refer to Section C.3 in the Appendix for details.

Table 8
Social preferences and individual choices.

	(1)	(2)
Independent var.	no ϵ-SP	no belief in ϵ-SP
B_predict	-0.0310	0.171**
	(0.0765)	(0.0809)
WBB_predict	0.437***	0.366***
	(0.1126)	(0.0696)
WBA_predict	-0.406***	-0.537***
	(0.1023)	(0.0744)
Pseudo R^2	0.0441	0.0700
N	1581	1581

Notes: Standard errors clustered at the matching cohort level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Multinomial logit regressions. Each observation is an individual subject in a round from the "BI" treatments with random-matching (excluding the observations in which dominated strategies were chosen). Dependent variable is the adopted strategy. Independent variables are the measured social preferences and the belief in social preferences. ϵ-SP stands for the ϵ-social preferences. Control variables are the dummies for the three "BI" treatments. Marginal effects are reported.

4.4. Alternative irreversibility structures

Our theory relies on the specific irreversibility structure, under which the efficient choice A is the only reversible choice at an earlier date. However, other reversibility structures may also improve efficient coordination, thanks to mechanisms different from signaling intention by waiting. For example, experimental evidence has shown that when both choices are reversible, multi-sided costless pre-play communication in common-interest coordination games (Cooper et al., 1992a; Charness, 2000; Blume and Ortmann, 2007) can help to facilitate efficient coordination. We experimentally examined the effect of alternative irreversibility structures, and compare these with the mechanism of our delay option with the irreversible B choice.

Result 5 *(Alternative irreversibility). Both "NI-f" and "AI-f" improved the efficient coordination to some extent. However, that improvement was lower than under "BI-f" and relied on a different mechanism.*

"NI-f" treatment We find evidence that "BI-f" is more effective than "NI-f" at facilitating efficient coordination. First, according to the estimation results (Column 6 of Table 9), similar to "BI-f," "NI-f" increased efficiency rates in comparison to "St-f;" but, different from "BI-f," the difference was not statistically significant. Moreover, comparing "NI-f" with "BI-f," the efficiency rates in "NI-f" were lower than those in "BI-f," although the differences were not statistically significant (Column 2 of Table 9). Moreover, the frequency of the realized A choices, the group average payoffs, and the coordination rates were found to be significantly lower in "NI-f" than in "BI-f."

Next, we make a detailed comparison between the two mechanisms, signaling intention by waiting and expressing intention via costless pre-play communication, which has also been found to improve coordination efficiency to some extent. Based on our theory, the choice of "Wait" signals the intention to play A at $t = 1$ (if all others choose to wait). From this perspective, the "Wait" choice in "BI-f" (or "BI-b") is comparable to the non-binding A choice in the "NI-f" treatment.

To distinguish these two mechanisms, we first compare the proportion of A choices (among subjects who chose to wait at $t = 0$) following the "no-B" message in "BI-f" with that of subjects who chose A at $t = 0$ after the "all-A" message (i.e., everyone else in the group chose A in the first period) in "NI-f." As shown in Column 3 of Table 10, these two proportions do not differ significantly.

Recall that the "no-B" message in "BI-f" simply means that no one has made a binding choice. In fact, all messages in the "NI-f" setting have that meaning since all actions are reversible. However, the "all-A" message in "NI-f," on its face value, says that "all subjects intend to choose A." If the face value of messages can, indeed, affect players' beliefs and their subsequent moves (not in a strategic sense but in a linguistic sense), it is striking that the "no-B" message in "BI-f" can be as effective as the "all-A" message in "NI-f." Indeed, our theory implies that, under the unique iteratedly undominated strategy profile, all subjects would take A following the "no-B" message in "BI-f."

Furthermore, among the subjects who chose B in "NI-f", the proportion of A choices following the "no-B" message was significantly lower than that among those who chose to wait at "BI-b" (Column 3 of Table 10). Therefore, combining both reversible choices, the "all-A" message in "NI-f" was no longer as effective as in "BI-f" (Column 2 of Table 10). Additionally, as can be seen from Column 1 of Table 10, for the $t = 0$ choices, the frequency of the irreversible B choices in "NI-f" was much higher than that of the reversible B choices in "BI-f," showing that the irreversibility structure made a difference. The "no-B" message was generated more frequently in "BI-f," as compared with the "all-A" message in "NI-f," thereby inducing a higher frequency of A as final choices and a higher rate of efficient coordination in "BI-b."

Table 9
Group-level regressions (fixed-matching).

	Reference = BI-f				Reference = St-f			
	(1) A_rate	(2) effi_rate	(3) payoff	(4) coor_rate	(5) A_rate	(6) effi_rate	(7) payoff	(8) coor_rate
St-f	-0.279***	-0.431***	-6.450***	-0.161***				
	(0.0449)	(0.1539)	(1.1605)	(0.0368)				
NI-f	-0.135***	-0.189	-2.274*	-0.0722**	0.146***	0.242	4.177***	0.0884**
	(0.0465)	(0.1682)	(1.1598)	(0.0297)	(0.0209)	(0.1604)	(0.8421)	(0.0394)
AI-f	-0.199***	-0.333**	-6.051***	-0.189***	0.0812***	0.0980	0.399	-0.0283
	(0.0530)	(0.1605)	(1.7731)	(0.0583)	(0.0307)	(0.1522)	(1.5877)	(0.0639)
Constant			48.66***				42.21***	
			(0.9969)				(0.5962)	
R^2			0.0653				0.0336	
Pseudo R^2	0.0772	0.0899		0.0360	0.0310	0.0425		0.0130
N	705	705	705	705	525	525	525	525

Notes: Standard errors clustered at the group level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Reference category is "BI-f" (1–4) or "St-f" (5–8). Each observation is a group-average level in a round. Dependent variables (and the regression models used) are (1 & 5) percentages of A as final choices (tobit), (2 & 6) efficient outcome dummy (probit), (3 & 7) group average payoff (OLS), and (4 & 8) the dummy for coordination on either action (probit). Marginal effects are reported for tobit and probit regressions.

Table 10
Individual-level regression (fixed-matching).

	(1) B in t0	(2) A after no-B or all-A	(3) A after no-B or all-A
NI-f	0.118*	-0.147*	
	(0.0713)	(0.0790)	
NI-f-A			0.00838
			(0.0595)
NI-f-B			-0.635***
			(0.1019)
Pseudo R^2	0.0309	0.0411	0.251
N	1440	1351	1351

Notes: Standard errors clustered at the group level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Probit regressions. Reference category is "BI-f." Each observation is an individual subject in a round. Dependent variables are choice of B in $t0$ (dummy) and choice of A after the no-B message (dummy). Other control variables include Rounds 1–5 (dummy), Rounds 6–10 (dummy), and Rounds 11–15 (dummy). Marginal effects are reported.

"AI-f" treatment When A is the only binding choice at $t = 0$ as discussed in Section 2.2, neither the subgame-perfect Nash equilibrium nor iterated weak dominance can provide a clear prediction about efficient coordination for any multiple-player group ($N \geq 3$). Although choosing A early may serve as a signal to induce the players who have waited to follow, it is, indeed, a risky choice because all the subjects who waited still faced a coordination problem at $t = 1$. The experimental evidence shows that the "AI-f" protocol raised the frequency of A as final choices, but, overall, it did not significantly promote efficient coordination or improve the average payoffs and coordination rates, compared with the static benchmark (Table 9).

Overall, the results provide some evidence that our delay mechanism is more effective than the alternative ones, and that the reversibility structures do make a difference in promoting efficient coordination.

5. Conclusion

This paper highlights a distinctive function of a delay option in strategic interactions that can be modeled as coordination games. The option enables forward-induction reasoning to operate, and, in this way, each player, by delaying their choice, can signal their intention to take the risky and efficient action. We show that this mechanism of signaling intentions via delay can work to achieve the efficient outcome. This idea is formalized via iterated weak dominance when players have ϵ-social preferences.

We also provide experimental evidence to support our theoretical analysis regarding use of the strategy that survives iterated weak dominance and the resulting coordination outcome. The results are robust to the presentation of first-period outcomes and to randomly matched groups. Additionally, iterated weak dominance relies on an assumption that the players hold ϵ-social preferences, and we found in our experiment that subjects' ϵ-social preferences, and their beliefs that other players held such preferences, were positively related to the choices of the unique surviving iterative undominated strategy.

The unique strategy surviving iterated weak dominance—waiting and then taking the efficient action if and only if none of the other players took the inefficient action earlier—can be interpreted as "no first use (of the inefficient action)." Obviously, if everyone commits to such a strategy, which can lead to the efficient outcome, then this way of achieving efficiency becomes possible when each player is granted the option to delay. We believe that this simple idea should be applicable to more complex coordination settings.[29] We leave this direction to future work.

In addition, since our experiments tested only a limited set of parameters, we acknowledge that the mechanism of coordination via delay could be better understood if the experimental tests are extended to a larger set of parameters. For example, it would be helpful to see how the efficacy of the delay mechanism is affected by changing the payoff gap between successful coordination and miscoordination. By changing group size and cohort size in random matching, it would be interesting to see how our mechanism can work with large-sized groups and how it varies with the matching protocol. We leave this to future research, as well.

Declaration of competing interest

The authors do not have any conflict of interest.

Appendix A. Proofs

Proof of Proposition 1. Consider player $i \in \mathcal{N}$. When all other players take $d_j = B$, their monetary payoffs will be $\pi_j = b$ (for all $j \in \mathcal{N} \setminus \{i\}$), which is independent of player i's choice. For player i, choosing B yields $\pi_i = b$, while choosing A yields $\pi_i = c$. Since $b > c$, $u_i(d_i = B, (d_j = B)_{j\in\mathcal{N}\setminus\{i\}}) > u_i(d_i = A, (d_j = B)_{j\in\mathcal{N}\setminus\{i\}})$, and, thus, $\{d_i = B\}_{i=1}^N$ is a Nash equilibrium.

Similarly, when all others are taking $d_j = A$, choosing $d_i = A$ yields the same monetary payoff a for player i and all other players, while choosing $d_i = B$ yields $\pi_i = b$ for player i and $\pi_j = c$ for all $j \in \mathcal{N} \setminus \{i\}$. Hence, $u_i(d_i = A, (d_j = A)_{j\in\mathcal{N}\setminus\{i\}}) > u_i(d_i = B, (e_j = A)_{j\in\mathcal{N}\setminus\{i\}})$, and, thus, $\{d_i = A\}_{i=1}^N$ is a Nash equilibrium. □

Proof of Proposition 2. First, consider the case in which all other players take $s_j = B$. In this case, $m = 1$, and, $\pi_j = b$ for all $j \in \mathcal{N} \setminus \{i\}$ independent of s_i. For player i, the private payoff from choosing $s_i = B$ is b, while deviating to other strategies never strictly increases this private payoff but possibly strictly decreases it to c (for example, deviating to WAB). Hence, $(s_i = B)_{i=1}^N$ is a Nash equilibrium.

Next, consider the case in which all other players choose WAA. In this case, WAA is a best response because (1) choosing WAA yields $\pi_i = a$ and $\pi_j = a$ for all $j \in \mathcal{N} \setminus \{i\}$; (2) deviating to B yields $\pi_i = b < a$ (and $\pi_j = c < a$); (3) deviating to WBB or WAB yields $\pi_i = b < a$ (and $\pi_j = c < a$); and (4) deviating to WBA yields the same π_i and π_j as choosing WAA. Hence, $(s_i = WAA)_{i=1}^N$ is a Nash equilibrium. Following similar arguments, we can show that $(s_i = WBA)_{i=1}^N$ is a Nash equilibrium.

Let us further consider the case in which all other players choose WAB. Given that, WAB is a best response because (1) choosing WAB yields $\pi_i = b$ and $\pi_j = b$ for all $j \in \mathcal{N} \setminus \{i\}$; (2) deviating to B yields $\pi_i = b$ and $\pi_j = c < b$; (3) deviating to WAA or WBA yields $\pi_i = c < b$ (and $\pi_j = b$); and (4) deviating to WBB yields the same π_i and π_j as choosing WAB. Hence, $(s_i = WAB)_{i=1}^N$ is a Nash equilibrium. Following similar arguments, we can show that $(s_i = WBB)_{i=1}^N$ is a Nash equilibrium.

Lastly, choosing A on the information set $m = 1$ is not subgame-perfect. That is because, in the subgame starting at $t = 1$ following $m = 1$—i.e., after someone has already taken B at $t = 0$—deviating from A to B increases one's payoff from c to b (without changing others' payoffs). □

Proof of Theorem 1. *First round of elimination* Consider any player i and any strategy profile $s_{-i} = (s_j)_{j\in\mathcal{N}\setminus\{i\}}$. We want to show that WAB (WAA) is weakly dominated by WBB (WBA). Consider two mutually exclusive and collectively exhaustive cases. In the first case, the other players adopt the strategy profile $s_{-i} = (s_j)_{j\in\mathcal{N}\setminus\{i\}}$, which satisfies $\left|\{j \in \mathcal{N} \setminus \{i\} | s_j = B\}\right| \geq 1$; that is, some other players choose B at $t = 0$ (or $m = 1$, regardless of s_i). Given s_{-i}, $\pi_i(s_i = WBB, s_{-i}) = b > \pi_i(s_i = WAB, s_{-i}) = c$, whereas $\pi_j(s_i = WBB, s_{-i}) = \pi_j(s_i = WAB, s_{-i})$ for any s_j and any $j \in \mathcal{N} \setminus \{i\}$. In the other case, $s_{-i} = (s_j)_{j\in\mathcal{N}\setminus\{i\}}$ satisfies $\left|\{j \in \mathcal{N} \setminus \{i\} | s_j = B\}\right| = 0$, meaning that $m = 0$, regardless of s_i. Given this s_{-i}, $\pi_i(s_i = WBB, s_{-i}) = \pi_i(s_i = WAB, s_{-i}) = b$ and $\pi_j(s_i = WBB, s_{-i}) = \pi_j(s_i = WAB, s_{-i})$ for any s_j and any $j \in \mathcal{N} \setminus \{i\}$. Hence, based on the utility function u_i defined in (1), WAB is weakly dominated by WBB. The same argument can be applied to show that WAA is weakly dominated by WBA.

To see why WBA, WBB, and B cannot be eliminated in this round, consider any mixed strategy that might dominate any of these three strategies. If such a mixed strategy exists, and if it assigns positive probabilities to WAA or WAB (or both), then we can reassign those probabilities to WBA or WBB (or both), respectively, to generate another strategy that still

[29] For example, Basak and Zhou (2021) apply a similar idea to design information disclosure policy in a dynamic regime change game with irreversible attacks in an incomplete-information environment.

satisfies the weak dominance relationship (since WAA and WAB are weakly dominated by WBA and WBB, respectively). So, we need only consider the mixtures of WBA, WBB, and B.

First, note that any mixed strategy consisting of B and WBB cannot dominate WBA because WBA is the best response to $s_{-i} = (s_j = WBA)_{j\in\mathcal{N}\setminus\{i\}}$.

Now, suppose that B can be weakly dominated by $s^0 = p_0 \cdot WBB \oplus (1-p_0)\cdot WBA$ for some $p_0 \in [0,1]$. Consider the case in which the other players' strategy profile is $s_{-i} = (s_j = WBB)_{j\in\mathcal{N}\setminus\{i\}}$. Then, $\pi_i(s_i = B, s_{-i}) = b$, $\pi_i(s_i = s^0, s_{-i}) = p_0 b + (1-p_0)c$, while $\pi_j(s_i = B, s_{-i}) = \pi_j(s_i = s^0, s_{-i}) = b$. So, weak dominance requires $p_0 = 1$, which means that B can be dominated only by the pure strategy WBB. Next, fix $p_0 = 1$ in s^0 and consider the other case in which $s'_{-i} = (s_j = WBA)_{j\in\mathcal{N}\setminus\{i\}}$. Then, $\pi_i(s_i = B, s'_{-i}) = \pi_i(s_i = s^0, s'_{-i}) = b$, while $\pi_j(s_i = B, s'_{-i}) = b > \pi_j(s_i = s^0, s'_{-i}) = c$, which means that B is preferred to WBB in this case. Therefore, no such $p_0 \in [0,1]$ exists, and B cannot be weakly dominated by any mixed strategy.

To see that WBB cannot be weakly dominated either, suppose that a mixed strategy $s^1 = p_1 \cdot B \oplus (1-p_1)\cdot WBA$ for some $p_1 \in [0,1]$ weakly dominates WBB. Consider the case in which the other players' strategy profile is $s_{-i} = (s_j = WBB)_{j\in\mathcal{N}\setminus\{i\}}$. Then, $\pi_i(s_i = WBB, s_{-i}) = b$, $\pi_i(s_i = s^1, s_{-i}) = p_1 b + (1-p_1)c$, while $\pi_j(s_i = WBB, s_{-i}) = \pi_j(s_i = s^1, s_{-i}) = b$. So, weak dominance requires $p_1 = 1$, which means that WBB could be dominated only by the pure strategy B. Next, consider the other case, in which $s'_{-i} = (s_j = WAB)_{j\in\mathcal{N}\setminus\{i\}}$. Then, $\pi_i(s_i = WBB, s'_{-i}) = \pi_i(s_i = B, s'_{-i}) = b$, while $\pi_j(s_i = WBB, s'_{-i}) = b > \pi_j(s_i = B, s'_{-i}) = c$, which means that WBB is preferred to B in this case. Therefore, no such $p_1 \in [0,1]$ exists, and WBB cannot be weakly dominated by any mixed strategy.

Second round of elimination The remaining strategies are B, WBB and WBA. For player i, given any s_{-i}, $\pi_i(s_i = WBB, s_{-i}) = \pi_i(s_i = B, s_{-i}) = b$. If $n(s_{-i}) \geq 1$, then $m = 1$ regardless of s_i, and, therefore, $\pi_j(s_i = B, s_{-i}) = \pi_j(s_i = WBB, s_{-i})$ for all $j \in \mathcal{N}\setminus\{i\}$. This means that player i is indifferent between B and WBB. The indifference also holds when $s_j = WBB$ for all j. However, if, among other players, no one chooses B and some players choose WBA—i.e., $n(s_{-i}) = 0$ and $\left|\{j \in \mathcal{N}\setminus\{i\} | s_j = WBA\}\right| \geq 1$—then $\pi_{j'}(s_i = B, s_{-i}) = b > \pi_{j'}(s_i = WBB, s_{-i}) = c$ for all $j' \in \{j \in \mathcal{N}\setminus\{i\} | s_j = WBA\}$, and $\pi_j(s_i = B, s_{-i}) = \pi_j(s_i = WBB, s_{-i}) = b$ for all $j \in \{j \in \mathcal{N}\setminus\{i\} | s_j = WBB\}$. Hence, under the ϵ-social preferences assumption, WBB is weakly dominated by B.

No other strategies can be eliminated in this round. WBA is the unique best response if all others take WBA. When all others choose WBB, compared with strategy WBA, choosing B yields a strictly higher payoff to player i but the same payoffs to other players. Therefore, B cannot be dominated by WBA. Since we have already shown that B weakly dominates WBB, B cannot be eliminated in this round.

Third round of elimination The remaining strategies are B and WBA. Again, consider two mutually exclusive and collectively exhaustive cases regarding s_{-i}. First, suppose that s_{-i} satisfies $\left|\{j \in \mathcal{N}\setminus\{i\} | s_j = B\}\right| \geq 1$, which means that $m = 1$, regardless of s_i. Then, player i is indifferent between B and WBA. Second, suppose that s_{-i} satisfies that $\left|\{j \in \mathcal{N}\setminus\{i\} | s_j = B\}\right| = 0$—i.e., all other players choose WBA; then, $\pi_i(s_i = WBA, s_{-i}) = a > \pi_i(s_i = B, s_{-i}) = b$, and $\pi_j(s_i = WBA, s_{-i}) = a > \pi_j(s_i = B, s_{-i}) = b$ for all $j \in \mathcal{N}\setminus\{i\}$. Hence, B is weakly dominated by WBA. □

Proof of Proposition 3. As in the proof of Theorem 1, in the first round of elimination, we can eliminate any strategies that involve waiting and taking A following any message that indicates that $n(s_{-i}) \geq 1$; that is, someone else has already chosen B at $t = 0$. Then, the proofs of second-round and third-round elimination follow immediately from that of Theorem 1. □

Proof of Proposition 4. For the $N = 2$ case, waiting and then taking B after observing that the other player chose A at $t = 0$ is dominated by waiting and then taking A based on this history. Given that, choosing A at $t = 0$ weakly dominates waiting and then choosing A after observing that the other player chose A, and choosing B (or A) after observing that the other player chose to wait. The symmetric subgame-perfect equilibria are (1) A at $t = 0$ and (2) wait and always choose A. It is worth mentioning that the strategy "waiting and choosing A if the other player chooses A; otherwise, choosing B" cannot constitute a symmetric equilibrium, as each player would profit from deviating to choosing A at $t = 0$.

In this proof for player sets with $N \geq 3$, we consider a simple case with $N = 3$, and we find all symmetric strategy profiles that are consistent with iterated weak dominance. The result can easily be generalized to cases with $N > 3$.

In the three-player case, we can write the strategies as A, $WBBB$, $WBBA$, $WBAB$, $WABB$, $WAAB$, $WABA$, $WBAA$, and $WAAA$. The strategy of choosing A at $t = 0$ is denoted by A. For any strategy profile s_{-i} of the other players, let $n^A(s_{-i}) := \left|j \in \mathcal{N}\setminus\{i\} | s_j = A\right|$ denote the number of the irreversible A choices at $t = 0$. Then, we denote any player i's strategy associated with waiting at $t = 0$ as follows. "W" stands for waiting at $t = 0$. The first letter after "W" is for the choice of action when no one chose A at $t = 0$ ($n^A = 0$), and the second (third) letter is for the choice of action when $n^A = 1$ ($n^A = 2$).

At $t = 1$, it is strictly better to choose A after observing $n^A = 2$. Therefore, we can eliminate $WBBB$ (by $WBBA$), $WBAB$ (by $WBAA$), $WAAB$ (by $WAAA$), and $WABB$ (by $WABA$). The remaining strategies are A, $WBBA$, $WABA$, $WBAA$ and $WAAA$.

We will show that none of the other strategies can be eliminated in this round. Consider the case in which the second player chooses $WBBA$ and the third player chooses a mixed strategy $p \cdot A \oplus (1-p)\cdot WBBA$ with $p \in (0, \frac{b-c}{a-c})$. As can be seen from the table below, $WBBA$ and $WBBB$ are the only two strategies that serve as best responses. They both generate the highest (expected) payoff π_i. Also, they both generate the same payoff to other players (since n^A obtains a value of 0 or 1, but these two strategies differ only when $n^A = 2$).

Strategy	Payoff π_i
A	$pa + (1-p)c$
$WBBA$ (or $WBBB$)	b
$WBAA$ (or $WBAB$)	$pc + (1-p)b$
$WABA$ (or $WABB$)	$pb + (1-p)c$
$WAAA$ (or $WAAB$)	c

Therefore, $WBBA$ can be weakly dominated only by a mixture of $WBBA$ and $WBBB$. This is not possible since $WBBA$ weakly dominates $WBBB$. Thus, we have shown that $WBBA$ cannot be weakly dominated.

Similarly, $WABA$ and $WABB$ are the only best responses when the second player chooses $WABA$, and the third player chooses $p \cdot A \oplus (1-p) \cdot WABA$ with $p \in (0, \frac{a-c}{2a-b-c})$. Moreover, $WBAA$ and $WBAB$ are the only best responses when the second player chooses the mixed strategy $p \cdot A \oplus (1-p) \cdot WBBA$ with $p \in (0,1)$, and the third player chooses $WBAA$. Lastly, $WAAA$ and $WAAB$ are the only best responses when the second player chooses $WAAA$, and the third player chooses a mixed strategy $p \cdot A \oplus (1-p) \cdot WABA$ with $p \in (0,1)$. Following this logic, we can show that $WABA$, $WBAA$, and $WAAA$ cannot be weakly dominated. In addition, A is the unique best response when all other players choose $WBAA$.

Therefore, in the first round of elimination, we can remove any strategy that involves choosing B after seeing all other players choose A ($n^A = N-1$) at $t = 0$. However, the strategy of not waiting (i.e., A) and strategies that involves waiting and then choosing either B or A after any $n^A < N-1$ (i.e., $WBBA$, $WABA$, $WBAA$ and $WAAA$) cannot be eliminated.

After eliminating $WBBB$, $WBAB$, $WABB$, and $WAAB$, by repeating the same arguments for why other strategies cannot be eliminated in the first round, we can show that each strategy that survives the first round of elimination is, in fact, a unique best response to some strategies chosen by the other players. Thus, none of them can be eliminated later.

To summarize, the strategy profiles consistent with iterated weak dominance are: (1) all players choose A at $t = 0$; and (2) all players wait and choose A or B when $n^A < N-1$ but choose A when $n^A = N-1$.

The subgame-perfect equilibria take the following forms. All players choose A at $t = 0$. In all other cases, all players choose "wait" at $t = 0$, choose A when $n^A = N-1$, and choose A or B if $n^A = 2, \ldots, N-2$. There are multiple possibilities for $m^A = 0, 1$. In one case, all players also choose A following $n^A = 0, 1$. In another case, all players choose B following $n^A = 0, 1$. In the third case, all players choose A following $n^A = 0$ and choose B following $n^A = 1$.

It is easy to check that any of the strategies A, $WAAA$, $WBBA$, or $WABA$ can constitute a subgame-perfect equilibrium. To see why each player choosing strategy $WBAA$ is not such an equilibrium, consider the case in which the other two players choose $WBAA$. Then, a player would choose A and receive a (monetary) payoff a rather than choosing strategy $WBAA$ and receiving a (monetary) payoff of b. □

Appendix B. Choice dynamics analysis

The experiments in this study consisted of the fixed-matching sessions that follow the design in the minimum-effort literature and of the random-matching sessions, which, hypothetically, would be less influenced by the learning and exploration motives, as well as other dynamic concerns. Though our theory provides no basis for understanding how various types of dynamic concerns affect subjects' choices and group coordination over time, in this section, we empirically analyze how the coordination outcome in previous rounds—in particular, the most recent round—affected subjects' choices, controlling for their initial choices.

We first categorize subjects, based on their choices of undominated strategies B, WBA, and WBB in the most recent round, into three categories. We then investigate how the following three types of observable coordination outcomes in the most recent round and in all the past rounds influenced their next-round choice of strategy. In each round, the outcomes could be classified as:

- Outcome h_1: Efficient coordination was achieved.
- Outcome h_2: Efficient outcome was not achieved, and it was observed that someone chose B in the first period.
- Outcome h_3: Efficient outcome was not achieved, but no one chose B in the first period. This suggests that at least one player chose WBB (or WAB).

Table 11 presents the multinomial regressions of subjects' choices on the histories of the three types of outcomes. History enters the regressions in two ways. First, there are two dummy variables on whether h_2 or h_3 was observed in the latest round. Second, we include two variables of the percentages of h_2 and h_3 in the past rounds (excluding the latest round). That is, we assume that the outcome from the latest round has a higher weight in the history.

The WBA choosers (Columns 1 and 4 in Table 11) were affected mainly by the occurrence of h_3 in the latest round, which greatly reduced the probability of continuing with the WBA choice in the next round. Among these subjects, the occurrence of h_3 significantly increased the frequency of B choices in the next round to avoid further harm. A smaller fraction of them switched to WBB after being hurt by others' use of WBB, which might be explained by the motive of

Table 11
Choices and learning ("BI" treatments).

	(1) fix_WBA	(2) fix_B	(3) fix_WBB	(4) rand_WBA	(5) rand_B	(6) rand_WBB
Outcome h_2						
B_predict	0.0952***			0.0294*		
	(0.0298)			(0.0168)		
WBB_predict	-0.00546			0.00333		
	(0.0139)			(0.0223)		
WBA_predict	-0.0897***			-0.0328		
	(0.0319)			(0.0321)		
Outcome h_3						
B_predict	0.334***		-0.00244	0.173***		0.0406
	(0.0944)		(0.0857)	(0.0422)		(0.0786)
WBB_predict	0.0981***		-0.208*	0.0524*		-0.205*
	(0.0374)		(0.1117)	(0.0268)		(0.1121)
WBA_predict	-0.432***		0.210**	-0.225***		0.164*
	(0.1014)		(0.0999)	(0.0502)		(0.0882)
%_Outcome h_2						
B_predict	-0.00163	0.190**	0.0246	0.0286	0.250***	0.246
	(0.0123)	(0.0940)	(0.1088)	(0.0218)	(0.0746)	(0.2214)
WBB_predict	0.00497	-0.0598	-0.0408	0.0578*	-0.0124	-0.180
	(0.0176)	(0.0762)	(0.1275)	(0.0339)	(0.0717)	(0.1688)
WBA_predict	-0.00334	-0.130	0.0162	-0.0863**	-0.237***	-0.0654
	(0.0184)	(0.0810)	(0.0649)	(0.0419)	(0.0582)	(0.1265)
%_Outcome h_3						
B_predict	0.0341	-0.0555	0.127	-0.0256	0.443**	0.188
	(0.0368)	(0.1508)	(0.1160)	(0.0555)	(0.1832)	(0.3255)
WBB_predict	0.0353	0.247**	-0.111	0.155**	-0.0280	0.127
	(0.0216)	(0.1008)	(0.1370)	(0.0667)	(0.1165)	(0.3145)
WBA_predict	-0.0694	-0.191*	-0.0163	-0.129*	-0.415***	-0.314
	(0.0451)	(0.1089)	(0.0562)	(0.0710)	(0.1303)	(0.2631)
Pseudo R^2	0.341	0.104	0.114	0.125	0.131	0.133
N	1775	332	242	915	352	147

Notes: Standard errors clustered at the group or matching cohort level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Multinomial logit regressions. Each observation is an individual subject in a round who adopted the strategies WBA, B, or WBB in the previous round. Dependent variable is the adopted strategy. Explanatory variables include the percentage of h_2 and h_3 in the previous rounds, and the dummy variables of h_2 or h_3 occurring in the last round. Other control variables include Rounds 2–5 (dummy), Rounds 6–10 (dummy), Rounds 11–15 (dummy), choice in the first round, and treatments. Marginal effects are reported.

retaliation,[30] which suggests negative reciprocity and/or spitefulness. The impacts of the observation h_3 are much smaller in the random matching treatments (Column 4 in Table 11), possibly due to a weakened dynamic concern when groups are randomly matched.

On the other hand, our data show that, among the groups in which everyone chose WBA, 99.29% and 96.25% of the members chose WBA in the next round in "BI-b" and "BI-b-rand," respectively. Therefore, everyone choosing WBA appeared to be stable based on this observation.

For WBB choosers (Columns 3 and 6), only h_2 or h_3 are possible outcomes. Interestingly, if h_3 was observed, which occurred as a result of these subjects' and possibly other players' simultaneous choice of WBB, these subjects tended to avoid using the same strategy and were more likely to adopt WBA in the next round, suggesting that some choices of WBB might be due to confusion or other reasons associated with dynamic learning or retaliation,[31] rather than spitefulness.

[30] We collected subjects' self-reported motives for choosing a certain strategy from the survey after the experiment. The survey was unincentivized and anonymous and served only as anecdotal evidence. The results from the subjects who self-reported having chosen WBB showed that most of them, indeed, had spiteful social preferences, but that some of them were also confused and failed to realize that adopting WBB could have hurt others (compared with strategy B). Some subjects explicitly mentioned that they chose WBB to hurt other subjects, and a small fraction mentioned retaliation against the teammate who chose WBB.

[31] One learning motive that can possibly incentivize subjects to choose WBB was discussed in Section 4.2 and also in Footnote 28.

Table 12
Group-level regressions (fixed- vs. random-matching).

	Reference = BI-b				Reference = St-b			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	A_rate	effi_rate	payoff	coor_rate	A_Rate	effi_rate	payoff	coor_rate
BI-b-rand	-0.195**	-0.239**	-5.916***	-0.0602**				
	(0.0880)	(0.1039)	(1.1698)	(0.0305)				
St-b-rand					-0.118***	-0.0808*	0.197	0.0651
					(0.0292)	(0.0491)	(1.4364)	(0.0542)
Constant			49.60***				42.37***	
			(0.4480)				(0.5831)	
R^2			0.0717				0.0000855	
Pseudo R^2	0.0422	0.0599		0.280	0.0666	0.0551		0.00778
N	395	395	395	395	395	395	395	395

Notes: Standard errors clustered at the group or matching cohort level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Reference category is "BI-b" (1–4) or "St-b" (5–8). Each observation is a group- or matching-cohort-average level in a round. Dependent variables (and the regression models used) are (1 & 5) percentages of A as final choices (tobit), (2 & 6) efficient outcome dummy (probit), (3 & 7) group average payoff (OLS), and (4 & 8) the dummy for coordination on either action (probit). Marginal effects are reported for tobit and probit regressions.

Table 13
Individual-level regression (fixed- vs. random-matching).

	Reference = BI-b			Reference = all BI		
	(1)	(2)	(3)	(4)	(5)	(6)
	B in t0	A after no-B	WBA	B in t0	A after no-B	WBA
BI-b-rand	0.116	-0.0121	-0.115			
	(0.0818)	(0.0623)	(0.1096)			
all BI-rand				0.119**	-0.0255	-0.116
				(0.0531)	(0.0475)	(0.0738)
Pseudo R^2	0.0195	0.000354	0.0113	0.0207	0.00147	0.0109
N	1900	1540	1900	3700	3044	3700

Notes: Standard errors clustered at group or matching cohort level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Probit regressions. Reference category is "BI-b" or all of the BI treatments with fixed-matching. Each observation is an individual subject in a round. Dependent variables are choice of B in $t0$ (dummy) and choice of A after the no-B message (dummy). Marginal effects are reported.

If a subject chose B in the most recent round (Columns 2 and 5), the only outcome they could observe would be h_2. Therefore, the variables of interest are only the two percentages of outcomes. Higher occurrences of h_2 or h_3 in the history, which indicates that efficient coordination failed more frequently in the previous rounds, reduced the likelihood that players switched to WBA in the next period, while it increased the likelihood that they continued to player B or switched to WBB.

Appendix C. Additional analysis

C.1. Comparing fixed- and random-matching treatments

As discussed in Sections 4.1 and 4.2, in both fixed- and random-matching treatments, significantly higher efficiency and the high adoption of strategy WBA were found in the dynamic games versus the static games. In this section, instead of the comparison between the dynamic and the static games, we focus on differences in the matching protocols.

In Tables 12 and 13, we compare the coordination efficiency and individual choices between the two matching protocols. The efficiency rates and the frequency of realized A choices were both significantly lower in the static and dynamic games with random-matching (Table 12), though there was still a significant improvement in "BI-b-rand" with respect to "St-b-rand" (Section 4.2).

To further understand the lower efficiency in the random-matching sessions, we looked into the comparison of the adopted strategies between the two matching protocols. Table 13 reports the regressions on the adopted strategies in "BI-b" and "BI-b-rand" (first three columns) and in all dynamic "BI" treatments (last three columns). The choices of B in $t = 0$ were found to be more frequent in the random-matching sessions when combining the data from all dynamic treatments (Column 4), and the difference was present in the first round.[32] This might be due to the exploration motives or dynamic

[32] Reported in Table 19, using first-round data, the difference was significant for the regressions on "BI-b" and "BI-b-rand," but was not significant for the comparison between all "BI" treatments in fixed- and random-matching sessions.

Table 14
Miscoordination analysis.

	No miscoordination			Miscoordination			obs.
	on A	on B	ave. pay	%	% A choice	ave. pay	
BI-b-rand	26.2%	58.1%	48.1	15.6%	63.0%	19.8	160
St-b-rand	0.6%	82.5%	45.1	16.9%	37.0%	30.2	160

Notes: The first two columns report the percentages of the groups coordinated on A or B. The 3rd-5th columns report the percentage of miscoordinated groups, the frequencies of A as final choices in the miscoordinated groups, and the average payoffs of the miscoordinated groups.

Table 15
Number of A as final choices in the miscoordinated groups.

	1A	2A	3A	obs.
BI-b-rand	4 (16.0%)	4 (16.0%)	17 (68.0%)	25
St-b-rand	16 (59.3%)	9 (33.3%)	2 (7.4%)	27

concerns[33] with fixed-matching. However, there was no significant difference in the frequencies of A choices after the “no-B” message or in the frequencies of the WBA choices.

Additionally, the overall differences in the waiting frequencies might also be due to learning from previous experiences in the later rounds. According to the dynamic analysis in Appendix B, subjects were more likely to switch to strategy B after being hurt by someone in their group choosing WBB. Given that around 10% of the subjects chose WBB in both fixed- and random-matching sessions, when groups were randomly formed, the chance to meet such a groupmate at least once was greatly increased. It might explain the overall higher frequency of taking B in the random-matching treatments as well as the gap in the adoption rates of WBA between fixed and random matching in the later rounds.

C.2. Miscoordination in the random-matching treatments

This subsection presents the data analysis on the differences in the patterns of miscoordination in “BI-b-rand” and “St-b-rand” and how they led to the insignificant difference in average payoffs. Table 14 reports the percentage of groups that successfully coordinated on either action A or B or miscoordinated in these two treatments.

We found that the overall rates of miscoordination did not differ much (15.6% v.s. 16.9%) between “BI-b-rand” and “St-b-rand.” However, conditional on miscoordination (or coordination on either action), the distributions of choices were quite different. We discuss the findings in detail below.

Coordination on either action When coordination occurred, subjects were more likely to coordinate on the efficient choice A in the dynamic treatment. In “St-b-rand,” when there was no miscoordination, almost all groups coordinated on the inferior choice B. Only 0.72% of the groups coordinated on the efficient choice A. In contrast, in the dynamic treatment, “BI-b-rand,” the ratio of coordination on A was 31.08%. Clearly, the delay option increased (decreased) the incidences in which coordination was achieved on the efficient action A (inferior choice B).

Based on the payoff parameters, the payoff difference between coordination on A ($a = 55$) and coordination on B ($b = 45$) was only 10. Consequently, despite the fact that the delay option significantly increased the frequency of coordination on A, the average payoff difference between the dynamic and static treatments (conditional on coordination on either action) was not very large (48.1 v.s. 45.1).

Miscoordination Overall, we found that miscoordination in “BI-b-rand” was more detrimental to subjects’ payoffs than in “St-b-rand.” Recall that miscoordination occurs when both A and B are selected as the final choices in a group. In a miscoordinated group, the A choosers receive a very low payoff ($c = 5$), whereas the B choosers receive $b = 45$. 63 percent of subjects in the miscoordinated groups in “BI-b-rand” chose A, which was a much higher proportion than in “St-b-rand” (37%). It led to a much lower average payoff (19.8 v.s. 30.2) for the miscoordinated groups in “BI-b-rand.”

We further examined the numbers of A choices in the miscoordinated groups (see Table 15) and found that in “BI-b-rand,” 3 subjects out of 4 ended up choosing A in 68% of the cases; however, in the static treatment, this frequency was only 7.4%, and, in 59.3% of the cases, only one subject chose A. Therefore, many more subjects received the low payoff of 5 in the miscoordinated group in “BI-b-rand.”

Overall, since the incidence of coordination on either action was of four to five times of that of miscoordination, the dynamic treatment still produced a higher payoff. However, the payoff gap (between “BI-b-rand” and “St-b-rand”) was insignificant, because miscoordination in the dynamic treatment yielded a much lower payoff (despite not occurring more frequently).

[33] For example, subjects might cooperate more during the initial rounds of repeated play in order to influence their group mates’ subsequent moves. This might lead to a high frequency of WBA choices in the initial rounds in the fixed-matching games. Treatments with random matching significantly weaken this pattern.

Table 16
One-sided Mann–Whitney U test of social preference and adoption of strategies.

	ϵ-SP, Y vs. N	p-value	belief, Y vs. N	p-value
B	715.5000	0.3993	770.0000	0.0529
WBB	178.5000	0.0000	429.0000	0.0000
B+WBB	316.0000	0.0009	251.5000	0.0000

The first (last) two columns provides the U statistics and the p-value of the test between subjects consistent or not consistent with ϵ-SP (belief in others' ϵ-SP).

Table 17
Social preferences and experience in the coordination games.

	(1) no ϵ-SP	(2) no belief in ϵ-SP
pay5	-0.174	-0.555
	(0.1790)	(0.3478)
Pseudo R^2	0.0456	0.0620
N	160	160

Notes: Standard errors clustered at the matching cohort level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Probit regressions. Each observation is an individual subject. Control variable: treatments. Marginal effects are reported.

C.3. Additional analysis of social preferences

C.3.1. Mann–Whitney U test of social preferences and adoption of strategies

In Section 4.3, we performed the regression analysis to compare the adoption of strategies between the subjects whose behaviors were either consistent or inconsistent with the ϵ-SP or belief in ϵ-SP in the social preference block. Since the behaviors of only a small group of subjects whose behaviors were not consistent with the ϵ-SP or belief in ϵ-SP, these two comparison groups are not balanced. To address this issue, we conducted additional statistical analyses with the one-sided version of the non-parametric Mann–Whitney U test (Mann and Whitney, 1947; Wilcoxon, 1945). We tested whether the adoption of B (WBB and the sum of B and WBB) among the group of subjects whose choices were not consistent with ϵ-social preference or with the beliefs about other teammates' ϵ-social preference was more frequent than that among the subjects whose choices were consistent with ϵ-social preference and beliefs about others' ϵ-social preference.

Table 16 reports the Mann–Whitney U statistics and the p-values. We find that these strategies were adopted more frequently in groups without ϵ-SP (or belief in ϵ-SP), with the exception of B, for which the absence of ϵ-SP did not significantly reduce the frequency of its selection. These new findings are largely consistent with the regression results in Table 8.

C.3.2. Social preference and experience

A caveat to our measure of ϵ-social preferences is that, since the social preference block was added after the main experiment, it is possible that the experience in the experiment affected subjects' choices in the later rounds and their choices in the social-preference block. To test this possibility, we looked at whether players' choices in this block depended on their negative experiences in the coordination games. The negative experience is measured by the percentage of rounds in which they received the payoff of 5, which could happen only if the subject decided to cooperate by playing A after the "no-B" message but someone else in the group played B following the "no-B" message. As Table 17 shows, this negative experience could not predict their social-preference choices. This suggests that the social-preference measures were derived primarily from the subjects' innate preferences rather than from their experience during the experiment.

C.4. Additional regression results

In Table 18, we present the additional regression results on the comparison between "St-b" and "BI-b-3c," as well as between "St-b-rand" and "BI-b-3c-rand." Table 19 presents the regression results on the comparison of the first-round results in the dynamic games between the two matching protocols.

Appendix D. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.geb.2022.11.001.

Table 18
Additional group-level regression analysis.

Panel A: Fixed-matching	(1) A_rate	(2) effi_rate	(3) payoff	(4) coor_rate
BI-b-3c	0.246	0.350**	6.568***	0.227***
	(0.1594)	(0.1574)	(2.1130)	(0.0531)
Constant			42.37***	
			(1.2640)	
R^2			0.109	
Pseudo R^2	0.0337	0.118		0.0902
N	465	465	465	465

Panel B: Random-matching	(1) A_rate	(2) effi_rate	(3) payoff	(4) coor_rate
BI-b-3c-rand	0.277*	0.487*	2.271	0.0833
	(0.1476)	(0.2800)	(3.2021)	(0.0959)
Constant			42.56***	
			(2.3136)	
R^2			0.0304	
Pseudo R^2	0.155	0.366		0.0692
N	140	140	140	140

Notes: Standard errors clustered at the group level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Reference category is "St-b" for Panel A and "St-b-rand" for Panel B. Each observation is a group-average level in a round. Dependent variables (and the regression models used) are (1) percentages of A as final choices (tobit); (2) efficient outcome dummy (probit); (3) group average payoff (OLS); and (4) the dummy for coordination on either action (probit). Marginal effects are reported for tobit and probit regressions.

Table 19
First-round differences (fixed- vs. random-matching).

	Reference = BI-b			Reference = all BI		
	(1) B in t0	(2) A after no-B	(3) WBA	(4) B in t0	(5) A after no-B	(6) WBA
BI-b-rand	0.147*	0.00974	-0.126			
	(0.0803)	(0.0399)	(0.0827)			
all BI-rand				0.0926	-0.0149	-0.0978
				(0.0666)	(0.0356)	(0.0679)
Pseudo R^2	0.0576	0.000495	0.0229	0.0149	0.000791	0.00978
N	148	128	148	284	234	284

Notes: Standard errors clustered at the group or matching cohort level are in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
Probit regressions. Reference category is "BI-b" or all of the BI treatments with fixed matching. Each observation is an individual subject in a round. Dependent variables are choice of B in $t0$ (dummy) and choice of A after the no-B message (dummy). Marginal effects are reported.

References

Aumann, Robert, 1990. Nash equilibria are not self-enforcing. In: Economic Decision Making: Games, Econometrics and Optimisation, pp. 201–206.
Avoyan, Ala, Ramos, Joao, 2019. A road to efficiency through communication and commitment. Available at SSRN 2777644.
Baliga, Sandeep, Morris, Stephen, 2002. Co-ordination, spillovers, and cheap talk. J. Econ. Theory 105 (2), 450–468.
Basak, Deepal, Zhou, Zhen, 2021. Panics and early warnings. Working Paper.
Battigalli, Pierpaolo, 1997. On rationalizability in extensive games. J. Econ. Theory 74 (1), 40–61.
Ben-Porath, Elchanan, Dekel, Eddie, 1992. Signaling future actions and the potential for sacrifice. J. Econ. Theory 57 (1), 36–51.
Blume, Andreas, Ortmann, Andreas, 2007. The effects of costless pre-play communication: experimental evidence from games with Pareto-ranked equilibria. J. Econ. Theory 132 (1), 274–290.
Blume, Andreas, Kriss, Peter H., Weber, Roberto A., 2017. Pre-play communication with forgone costly messages: experimental evidence on forward induction. Exp. Econ. 20 (2), 368–395.
Brandenburger, Adam, Friedenberg, Amanda, Keisler, H. Jerome, 2008. Admissibility in games. Econometrica 76 (2), 307–352.
Brandts, Jordi, Holt, Charles A., 1995. Limitations of dominance and forward induction: experimental evidence. Econ. Lett. 49 (4), 391–395.
Brandts, Jordi, Cooper, David J., 2006a. A change would do you good.... An experimental study on how to overcome coordination failure in organizations. Am. Econ. Rev. 96 (3), 669–693.
Brandts, Jordi, Cooper, David J., 2006b. Observability and overcoming coordination failure in organizations: an experimental study. Exp. Econ. 9 (4), 407–423.
Brandts, Jordi, Cabrales, Antonio, Charness, Gary, 2007. Forward induction and entry deterrence: an experiment. Econ. Theory 33 (1), 183–209.

Bryant, John, 1983. A simple rational expectations Keynes-type model. Q. J. Econ. 98 (3), 525–528.
Cachon, Gerard P., Camerer, Colin F., 1996. Loss-avoidance and forward induction in experimental coordination games. Q. J. Econ. 111 (1), 165–194.
Calcagno, Riccardo, Kamada, Yuichiro, Lovo, Stefano, Sugaya, Takuo, 2014. Asynchronicity and coordination in common and opposing interest games. Theor. Econ. 9 (2), 409–434.
Carlsson, Hans, van Damme, Eric, 1993. Global games and equilibrium selection. Econometrica 61 (5), 989–1018.
Chamley, Christophe, Gale, Douglas, 1994. Information revelation and strategic delay in a model of investment. Econometrica 62 (5), 1065–1085.
Charness, Gary, 2000. Self-serving cheap talk: a test of Aumann's conjecture. Games Econ. Behav. 33 (2), 177–194.
Chen, Daniel L., Schonger, Martin, Wickens, Chris, 2016. oTree—an open-source platform for laboratory, online, and field experiments. J. Behav. Exp. Finance 9, 88–97.
Chen, Roy, Chen, Yan, 2011. The potential of social identity for equilibrium selection. Am. Econ. Rev. 101 (6), 2562–2589.
Cooper, Russell, De Jong, Douglas V., Forsythe, Robert, Ross, Thomas W., 1992a. Communication in coordination games. Q. J. Econ. 107, 739–771.
Cooper, Russell, De Jong, Douglas V., Forsythe, Robert, Ross, Thomas W., 1992b. Forward induction in coordination games. Econ. Lett. 40 (2), 167–172.
Cooper, Russell, DeJong, Douglas V., Forsythe, Robert, Ross, Thomas W., 1993. Forward induction in the battle-of-the-sexes games. Am. Econ. Rev., 1303–1316.
Corsetti, Giancarlo, Guimaraes, Bernardo, Roubini, Nouriel, 2006. International lending of last resort and moral hazard: a model of IMF's catalytic finance. J. Monet. Econ. 53 (3), 441–471.
Crawford, Vincent, Broseta, Bruno, 1998. What price coordination? The efficiency-enhancing effect of auctioning the right to play. Am. Econ. Rev., 198–225.
Devetag, Giovanna, Ortmann, Andreas, 2007. When and why? A critical survey on coordination failure in the laboratory. Exp. Econ. 10 (3), 331–344.
Diamond, Peter A., 1982. Aggregate demand management in search equilibrium. J. Polit. Econ. 90 (5), 881–894.
Farrell, Joseph, 1988. Communication, coordination and Nash equilibrium. Econ. Lett. 27, 209–214.
Farrell, Joseph, Saloner, Garth, 1985. Standardization, compatibility, and innovation. Rand J. Econ., 70–83.
Fehr, Ernst, Schmidt, Klaus M., 2006. The economics of fairness, reciprocity and altruism—experimental evidence and new theories. In: Handbook of the Economics of Giving, Altruism and Reciprocity, vol. 1, pp. 615–691.
Govindan, Srihari, Wilson, Robert, 2009. On forward induction. Econometrica 77 (1), 1–28.
Gul, Faruk, Lundholm, Russell, 1995. Endogenous timing and the clustering of agents' decisions. J. Polit. Econ. 103 (5), 1039–1066.
Harsanyi, John C., Selten, Reinhard, 1988. A general theory of equilibrium selection in games. In: MIT Press Books, vol. 1.
Huck, Steffen, Müller, Wieland, 2005. Burning money and (pseudo) first-mover advantages: an experimental study on forward induction. Games Econ. Behav. 51 (1), 109–127.
Katz, Michael L., Shapiro, Carl, 1986. Technology adoption in the presence of network externalities. J. Polit. Econ. 94 (4), 822–841.
Kohlberg, Elon, Mertens, Jean-Francois, 1986. On the strategic stability of equilibria. Econometrica 54 (5), 1003–1037.
Krol, Michal, Krol, Magdalena Ewa, 2020. On the strategic value of 'shooting yourself in the foot': an experimental study of burning money. Int. J. Game Theory 49 (1), 23–45.
Lo, Melody, 2020. Language and coordination games. Econ. Theory, 1–44.
Mann, Henry B., Whitney, Donald R., 1947. On a test of whether one of two random variables is stochastically larger than the other. Ann. Math. Stat., 50–60.
Mathevet, Laurent, Steiner, Jakub, 2013. Tractable dynamic global games and applications. J. Econ. Theory 148 (6), 2583–2619.
Morris, Stephen, Shin, Hyun Song, 1998. Unique equilibrium in a model of self-fulfilling currency attacks. Am. Econ. Rev., 587–597.
Pearce, David G., 1984. Rationalizable strategic behavior and the problem of perfection. Econometrica, 1029–1050.
Sobel, Joel, 2017. A note on pre-play communication. Games Econ. Behav. 102, 477–486.
Van Damme, Eric, 1989. Stable equilibria and forward induction. J. Econ. Theory 48 (2), 476–496.
Van Huyck, John B., Battalio, Raymond C., Beil, Richard O., 1990. Tacit coordination games, strategic uncertainty, and coordination failure. Am. Econ. Rev. 80 (1), 234–248.
Van Huyck, John B., Battalio, Raymond C., Beil, Richard O., 1993. Asset markets as an equilibrium selection mechanism: coordination failure, game form auctions, and tacit communication. Games Econ. Behav. 5 (3), 485–504.
Weber, Roberto A., 2006. Managing growth to achieve efficient coordination in large groups. Am. Econ. Rev. 96 (1), 114–126.
Wilcoxon, Frank, 1945. Individual comparisons by ranking methods. In: Biometrics Bulletin, pp. 80–83.