Games and Economic Behavior 137 (2023) 23–49

C D

Games and Economic Behavior

: . . / /

100083, Ch a

[c] S e Sch f B e , Ta d Sch f E g ee g, NYU Sha gha , Ne Y U e , Ne Y , NY 10012, USA

ARTICLE INFO

A ce h :
Received 5 June 2021
Available online 14 November 2022

JEL c a fica :
C73
C92
D83

Ke d :
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 $\in (2, 4)$, 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

(B, B)

is uniquely selected if we relax common knowledge of the payoffs (for example, common knowledge of the parameter). 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- a add e e : 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
0899-8256/

Table 1
A simple coordination game.

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

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.

The e ca a a 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, $= 0, 1$. At $= 0$, 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 $= 0$, and then makes their final choice between A and B at $= 1$.[2] There is no cost associated with the delay option. A player who does not choose B at $= 0$ should, in some sense, be signaling that they intend to choose A at $= 1$. That is, there is a forward-induction flavor to choosing "wait." Intuitively, if a player intends to play B and secure the safe payoff , 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 $= 0$), they will optimally choose B at $= 1$. 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 $= 0$). 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

[redacted]
strategies. One [redacted] the one employed in this paper, [redacted].

The second [redacted] identified in various experimental studies (see Fehr and Schmidt (2006) for a survey). We adopt a very weak form of social preference in [redacted] 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,

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

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 *effic e c d a* is achieved if and only if $d = A$ for all $\in \mathcal{N}$. As shown in Proposition 1, efficient coordination is not guaranteed since coordinating on the safe action B constitutes another equilibrium.

We say *c d a* 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. D a c c e h e e b e ch ce f B

There are two periods, $= 0, 1$. At $= 0$, each player chooses between B and "wait." The choice of B is irreversible. That is, if a player chooses B at $= 0$, they cannot make any further changes. However, the players who wait at $= 0$ get to choose between A and B at $= 1$. There is no cost associated with waiting. By waiting, players can observe a binary message that takes the value $= 0$ if all players choose to wait at $= 0$ and the value $= 1$ 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 $= 0$ is denoted by B. Any strategy involving waiting at $= 0$ is a plan contingent on the message . We denote such a strategy by "$W, d(= 1), d(= 0)$," respectively, where $d( )$ is defined as the action chosen conditional on at $= 1$. For example, if a player chooses strategy WAB, they will wait at $= 0$ and then choose A after observing $= 1$; otherwise, they will choose B.

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

$$( ) = \left| \{ \in \mathcal{N} \setminus \{ \} | = B \} \right|,$$

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

$$= \mathbb{1}\{ ( ) \geq 1\}.$$

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

Proposition 2. *F a $\in \mathcal{S}$, he a eg fi e $( )_{=1}^{N}$ c e a e- a eg Na h e b . The bga e- e fec e b a a e $(= B)_{=1}^{N}$, $(= WBA)_{=1}^{N}$, a d $(= WBB)_{=1}^{N}$.*

It is easy to see that choosing A in the subgame following the message $= 1$ (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 $= B$, or in which they all choose $= WBB$, 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 e a eg fi e ha e e a ed ea d a ce $(= WBA)_{=1}^{N}$. U de h a eg fi e, effic e c d a ach e ed.*

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.

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

Sec d d (e a e 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 $\pi = b$ to any player .

Both strategies induce the same monetary payoff to player $\neq$ in all but one case, in which all players $\neq$ choose to wait at $= 0$, and at least some players $\neq$ choose strategy WBA. In this case, if player chooses B, a player who chooses WBA gets payoff $\pi = b$ from playing B after seeing $= 1$. However, player 's payoff is reduced to c if player chooses WBB because, in this case, player 's realized choice is A, following $= 0$. Therefore, under the assumption of ϵ-social preferences, strategy B weakly dominates WBB.

Th d d (e a e 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 $\neq$ chooses B at $= 0$, both B and WBA yield the same payoff to player and to all other players. However, if all $\neq$ choose WBA, then WBA yields a strictly higher payoff to . Thus, B is weakly dominated by WBA.

C d a c e Since all players choose strategy WBA, the realized message is $= 0$, and, thus, the realized choice is $d = A$ for all $\in \mathcal{N}$. Therefore, efficient coordination is achieved.

2.2. D c

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 $= 0$, 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.

Ob e ab f a ac

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 who exercises the delay option can observe the exact number of irreversible choices that occurred at $= 0$. We denote this number by ( ) and say that this scenario exhibits *fi e f a* .[13] With finer information, the strategy of waiting at $= 0$ and then choosing A at $= 1$ if and only if $= 0$ 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. *W h fi e f a , he e a eg fi e ha e e a ed ea d a ce ($= WBA)^N_{=1}$. U de h a eg fi e, effic e c d a ach e ed.*

Note that efficient coordination cannot be achieved as long as someone chooses B at $= 0$; that is,

$= 1$ (“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 ($= 0$), which is exactly the same as the information set $= 0$ (“no B” message) in the binary message setting.

(I) e e b c e

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

Ne he ch ce e e b e We first consider the case in which neither A nor B is irreversible at $= 0$.

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 games (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 games 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 a 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 $= 0$ should not be interpreted literally as a choice that is biased toward the choice of A at $= 1$. 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.

F e - f a ea e (S -f , BI-f a d BI-f- a d)

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 $= 0$, 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 $= 0$, 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.

A e a e e e b ea e (NI-f a d 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 $= 0$ were reversible. At $= 0$, subjects chose between A and B. There was no wait option in the first period. Then, at $= 1$, upon observing the distribution of the choices from $= 0$, 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 $= 0$. In $= 0$, subjects chose between A and the wait option. Then, in $= 1$, those who chose the wait option could decide between A and B after observing the number of A choices at $= 0$. 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. E e e a ced e

F ed- a ch g e

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, $= 0, 1$. 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 $= 0$ and $= 1$ 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 $= 0$, 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 $= 0$. 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.

Ra d - a ch g e

The random-matching sessions, “BI-b-rand,” “St-b-rand,” “BI-f-rand,” and

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; * < 0.10, ** < 0.05, *** < 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; * < 0.10, ** < 0.05, *** < 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 *(G - e e effic e c : b , e). The efficac f he de a , a f d be b he ea e ha d ced a add a e e b e ch ce A (BI-b-3c), a ed b ec b e efi e f a (BI-f).*

F a g effec 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

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.

F e f a 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. Ma e : ad f a eg e

Result 3 *(Ad ed a eg e). I a d a c ea e h a e e b e B ch ce (.e., he BI ea e), he a f he b ec he e e a ed d a ed a eg WBA b h fi ed- a ch g a d a d - a ch g e .*

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.

F a g effec The “BI-b-3c” and “BI-b-3c-rand” treatments allowed subjects to choose the reversible option A at $= 0$ 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

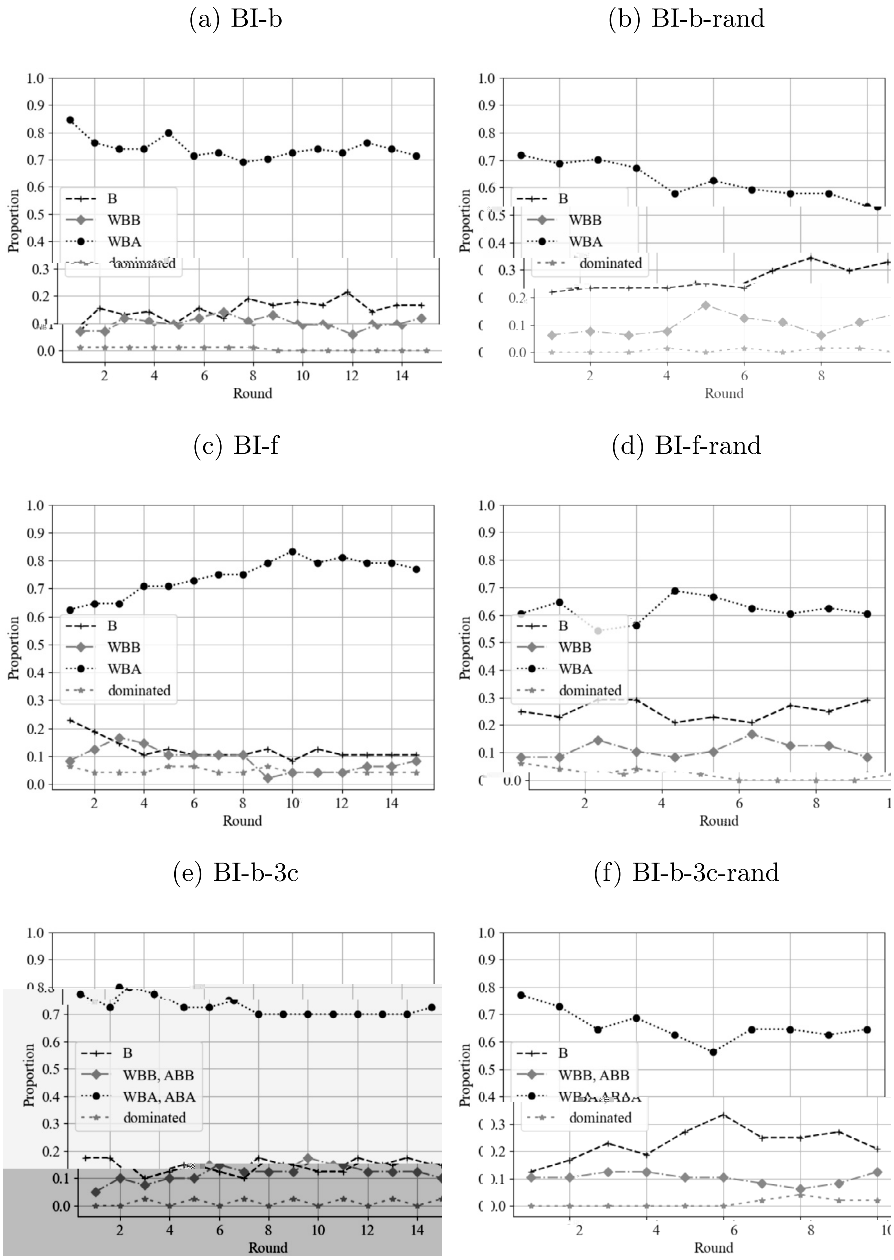

Fig. 2. Decomposition of strategies in the "BI" treatments.

F e f a 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 $=0$ and the A choice after the "no-B" message at $=1$ between "BI-f" ("BI-f-rand") and the main treatment "BI-b"

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; * < 0.10, ** < 0.05, *** < 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 0 (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 $= 0$, then under binary information, they would not observe whether other subjects also chose B at $= 0$. Therefore, the subjects who wished to learn about this information might have chosen "Wait" at $= 0$, implemented by

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. S c a efe e ce a d ad f a eg e

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 *(S c a efe e ce). The ac f ϵ- c a efe e ce e a c a ed h he ch ce f WBB. The ac f be ef ha he a e h d ϵ- c a efe e ce e a c a ed h he ch ce f B a d 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

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; * < 0.10, ** < 0.05, *** < 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. A e a e e e b c e

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 *(A e a e e e b). B h NI-f a d AI-f ed he effic e c d a e e e . H e e , ha e- e a e ha de BI-f a d e ed a d ffe e echa .*

NI-f ea e 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 $= 1$ (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 $= 0$) following the "no-B" message in "BI-f" with that of subjects who chose A at $= 0$ 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 $= 0$ 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; * < 0.10, ** < 0.05, *** < 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; * < 0.10, ** < 0.05, *** < 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 0 (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 ea e When *A he Whe*

When

Whe W

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

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

Proof of Proposition 2. First, consider the case in which all other players take $= B$. In this case, $= 1$, and, $\pi = b$ for all $\in \mathcal{N} \setminus \{\ \}$ independent of . For player , the private payoff from choosing $= B$ 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, $(= B)^{N}_{=1}$ 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 $\pi = a$ and $\pi = a$ for all $\in \mathcal{N} \setminus \{\ \}$; (2) deviating to B yields $\pi = b < a$ (and $\pi = c < a$); (3) deviating to WBB or WAB yields $\pi = b < a$ (and $\pi = c < a$); and (4) deviating to WBA yields the same π and π as choosing WAA. Hence, $(= WAA)^{N}_{=1}$ is a Nash equilibrium. Following similar arguments, we can show that $(= WBA)^{N}_{=1}$ 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 $\pi = b$ and $\pi = b$ for all $\in \mathcal{N} \setminus \{\ \}$; (2) deviating to B yields $\pi = b$ and $\pi = c < b$; (3) deviating to WAA or WBA yields $\pi = c < b$ (and $\pi = b$); and (4) deviating to WBB yields the same π and π as choosing WAB. Hence, $(= WAB)^{N}_{=1}$ is a Nash equilibrium. Following similar arguments, we can show that $(= WBB)^{N}_{=1}$ is a Nash equilibrium.

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

Proof of Theorem 1. *F d f e a* Consider any player and any strategy profile $_{-} = (\ )_{\in \mathcal{N}\setminus\{\ \}}$. 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 $_{-} = (\ )_{\in \mathcal{N}\setminus\{\ \}}$, which satisfies $\left|\{ \in \mathcal{N} \setminus \{\ \} | = B\}\right| \geq 1$; that is, some other players choose B at $= 0$ (or $= 1$, regardless of). Given $_{-}$, $\pi(= WBB, _{-}) = b > \pi(= WAB, _{-}) = c$, whereas $\pi(= WBB, _{-}) = \pi(= WAB, _{-})$ for any and any $\in \mathcal{N} \setminus \{\ \}$. In the other case, $_{-} = (\ )_{\in \mathcal{N}\setminus\{\ \}}$ satisfies $\left|\{ \in \mathcal{N} \setminus \{\ \} | = B\}\right| = 0$, meaning that $= 0$, regardless of . Given this $_{-}$, $\pi(= WBB, _{-}) = \pi(= WAB, _{-}) = b$ and $\pi(= WBB, _{-}) = \pi(= WAB, _{-})$ for any and any $\in \mathcal{N} \setminus \{\ \}$.

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 $\sigma_i^0 = p_0 \cdot WBB \oplus (1-p_0)\cdot WBA$ 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, $\pi_i(s_i = B, s_{-i}) = b$, $\pi_i(s_i = \sigma_i^0, s_{-i}) = p_0 b + (1-p_0)c$, while $\pi_j(s_i = B, s_{-i}) = \pi_j(s_i = \sigma_i^0, s_{-i}) = b$. 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 σ_i^0 and consider the other case in which $s'_{-i} = (s_j = WBA)_{j\in\mathcal{N}\setminus\{i\}}$. Then, $\pi_i(s_i = B, s'_{-i}) = \pi_i(s_i = \sigma_i^0, s'_{-i}) = b$, while $\pi_j(s_i = B, s'_{-i}) = b > \pi_j(s_i = \sigma_i^0, s'_{-i}) = c$, 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 $\sigma_i^1 = p_1 \cdot B \oplus (1-p_1)\cdot WBA$ 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, $\pi_i(s_i = WBB, s_{-i}) = b$, $\pi_i(s_i = \sigma_i^1, s_{-i}) = p_1 b + (1-p_1)c$, while $\pi_j(s_i = WBB, s_{-i}) = \pi_j(s_i = \sigma_i^1, s_{-i}) = b$. 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 $n_B(s_{-i}) \geq 1$, then $x = 1$ 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., $n_B(s_{-i}) = 0$ 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 $x = 1$, 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. $\square$

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 $n_B(s_{-i}) \geq 1$; 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. $\square$

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 . Also, they both generate the same payoff to other players (since ^{A} obtains a value of 0 or 1, but these two strategies differ only when $^{A} = 2$).

Strategy	Payoff π
A	$a + (1 - \ )c$
$WBBA$ (or $WBBB$)	b
$WBAA$ (or $WBAB$)	$c + (1 - \ )b$
$WABA$ (or $WABB$)	$b + (1 - \ )c$
$WAAA$ (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 $\cdot A \oplus (1 - \ ) \cdot WABA$ with $\in (0, \frac{a-c}{2a-b-c})$. Moreover, $WBAA$ and $WBAB$ are the only best responses when the second player chooses the mixed strategy $\cdot A \oplus (1 - \ ) \cdot WBBA$ with $\in (0, 1)$, 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 $\cdot A \oplus (1 - \ ) \cdot WABA$ with $\in (0, 1)$. 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 ($^{A} = N - 1$) at $= 0$. However, the strategy of not waiting (i.e., A) and strategies that involves waiting and then choosing either B or A after any $^{A} < N - 1$ (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 $= 0$; and (2) all players wait and choose A or B when $^{A} < N - 1$ but choose A when $^{A} = N - 1$.

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

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, whichnext greatly reduced the probability of continuing with the WBA choice the in 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; * < 0.10, ** < 0.05, *** < 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

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

	Reference = BI-b				Reference = St-b			
	(1) A_rate	(2) effi_rate	(3) payoff	(4) coor_rate	(5) A_Rate	(6) effi_rate	(7) payoff	(8) coor_rate
BI-b-rand	-0.195** (0.0880)	-0.239** (0.1039)	-5.916*** (1.1698)	-0.0602** (0.0305)				
St-b-rand					-0.118*** (0.0292)	-0.0808* (0.0491)	0.197 (1.4364)	0.0651 (0.0542)
Constant			49.60*** (0.4480)				42.37*** (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; * < 0.10, ** < 0.05, *** < 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 5

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. M c d a he a d - a ch g ea e

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.

C d a e he ac A When coordination

B

neTc-(dis)9.1(d) TJ /F3 1 Tf 9.9626 0 0 9.213.2924.97(g)58..898 Tm 0 Tc ( )Tj /F1 1 Tf 7.9701 0 0 7.216.7484727

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 preference

C.3.1. Mann–Whitney U test of social preference 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

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; * < 0.10, ** < 0.05, *** < 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; * < 0.10, ** < 0.05, *** < 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 0 (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.

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.