American Economic Review 2023, 113(9): 2355–2381
https://doi.org/10.1257/aer.20171014

A Road to Efficiency through Communication and Commitment†

By Ala Avoyan and João Ramos*

We experimentally examine the efficacy of a novel pre-play institution in a well-known coordination game—the minimum-effort game—in which coordination failures are robust and persistent phenomena. This new institution allows agents to communicate while incrementally committing to their words, leading to a distinct theoretical prediction: the efficient outcome is uniquely selected in the extended coordination game. We find that commitment-enhanced communication significantly increases subjects' payoffs and achieves higher efficiency levels than various nonbinding forms of communication. We further identify the key ingredients of the institution that are central to achieving such gains. (*JEL* C73, C92, D83)

Economic situations often require agents to coordinate their actions, and coordination failures leading to underperformance are pervasive in society. Players face strategic trade-offs in coordination environments; in particular, to achieve better outcomes, they must choose an action that they are typically unwilling to take unless other players do the same. Although players in these environments desire the same outcome, strategic uncertainty leads to coordination failures. The literature seeking to identify institutions to overcome coordination failures uses controlled experimental environments to compare interactions under different institutions.

Given that players' interests are aligned in coordination games and that failure is rooted in uncertainty, institutions formalizing communication have been a natural starting point for attempted solutions to coordination failure. However, experimental evidence on the effects of communication is mixed, and communication alone

*Avoyan: Department of Economics, Indiana University (email: AAvoyan@iu.edu); Ramos: Department of Finance and Business Economics, Marshall School of Business, University of Southern California (email: Joao.Ramos@marshall.usc.edu). Pietro Ortoleva was the coeditor for this article. We thank Andrew Schotter for his invaluable support and input to this project. We would like to thank Guillaume Fréchette, Alessandro Lizzeri, Andreas Blume, David Cesarini, Yan Chen, Cary Frydman, Boyan Jovanovic, Chad Kendall, Volodymyr Lugovskyy, Keith O'Hara, Daniela Puzzello, Debraj Ray, Séverine Toussaert, James M. Walker, and Roberto Weber for their helpful comments. This paper also benefited from comments by seminar and conference participants of NRET at New York University, the 27th International Conference on Game Theory, the 2016 ESA North American Meetings, University of Southern California, USC Marshall, University College London, Indiana University, Oxford University, Royal Holloway, and Kellogg School of Management at Northwestern University. We gratefully acknowledge financial support from the Center for Experimental Social Science at New York University, from the Institute for Outlier Research in Business (iORB) at the USC Marshall School of Business, and from Indiana University. The data collection was approved by the IRB of New York University (IRB-10-8117) and by IRB of Indiana University (IRB-2003653598). We thank four anonymous reviewers and the editors for their careful reading of the paper and their insightful comments and suggestions.
†Go to https://doi.org/10.1257/aer.20171014 to visit the article page for additional materials and author disclosure statements.

may not be enough to ensure success, even in a controlled experimental setting.[1] Additionally, many of the studied pre-play interactions lack theoretical implications; therefore, even if a given intervention empirically improves coordination in the laboratory, it is unclear how to isolate the features underlying its success.

In this paper, we experimentally examine a novel institution studied theoretically by Calcagno et al. (2014), asynchronous revision pre-play, which predicts that the Pareto-efficient profile is the unique outcome of an extended coordination game. In addition to the unique outcome prediction, the theory is used to obtain testable predictions of agents' dynamic behavior throughout the pre-play phase. The institution in Calcagno et al. (2014) formalizes the intuition that agents must prepare the actions that they intend to take at a predetermined deadline, and these preparations are public. As the deadline approaches, each player receives opportunities to update her prepared action at asynchronous and stochastically determined times. Once the players reach the deadline, their most recently prepared actions are implemented, and players' payoffs are determined by these actions only.

In the laboratory setting, we embed the mechanism into a minimum-effort game, in which a player's payoff depends on her own effort choice and the minimum effort chosen by the members of her group. We focus on this game because it is a prominent example of a coordination game with multiple Pareto-ranked equilibria. Moreover, a vast experimental literature observes coordination failures in this environment. To bring this institution to the laboratory, we extend its theoretical results. We introduce a discrete- instead of a continuous-time pre-play phase. Furthermore, we solve the game for the specific payoff structure and parameters used in the experiment, which leads to the prediction: subjects' initial choices should be the efficient effort. In the main treatment, *revision mechanism* (RM), the pre-play phase starts with all group members choosing an initial effort. If an opportunity arises, they can update this effort during a preparation phase of 60 seconds. Revision opportunities are awarded randomly to each group member, and the probability of two group members revising in the same instant is zero. In a single dynamic graph, each player can see real-time information on all group members' posted effort choices, including the history of posted efforts, revision opportunities, and updates. Players can change their effort on the screen at any time (i.e., change their intention). Still, these revisions will not be publicly posted on the graph unless the player is awarded a revision opportunity. Throughout the 60-second preparation phase, each player is expected to receive 8 revision opportunities. At the end of the pre-play phase, the players' most recently revised efforts are implemented.

In this paper, we test whether asynchronous revision pre-play can reduce coordination failures and allow the individuals involved to reach higher payoffs. Focusing on the efficiency gains, our first main result is, thus, that the mechanism increases efficiency by 18 percentage points (pp) over one round of public cheap-talk messages, which, in turn, increases efficiency by 16 pp over an environment without any interaction. The efficiency gains achieved with the revision mechanism

[1] Some communication protocols, such as two-way communication or public announcements, have been documented to increase coordination (see, e.g., Cooper et al. 1992; Chaudhuri, Schotter, and Sopher 2009; Charness 2000; Blume and Ortmann 2007; and Burton and Sefton 2004). However, other protocols, such as one-way communication or private advice, have led to coordination failures (see Cooper et al. 1992 and Chaudhuri, Schotter, and Sopher 2009).

(efficiency is 82 percent) are significant not only in comparison with the Baseline (48 percent) but also with one round of public cheap-talk messages (64 percent).[2] Furthermore, we show that this efficiency gain follows from a combination of subjects choosing higher effort and subjects better coordinating on any effort profile (Result 1).

Going beyond the efficiency gains, we test two exact predictions of the theory. First, we examine whether this novel mechanism entirely eliminates coordination failures, leading to 100 percent efficiency. Second, we examine whether 100 percent of subjects' initial choices are the efficient effort from the beginning. We show that efficiency in the revision mechanism treatment (82 percent) and the frequency of players choosing the efficient effort from the start of the pre-play (86 percent) are both high but significantly different from the prediction of 100 percent (Result 2). The unique outcome prediction, while stark, is also rigid, not accounting for factors that could influence behavior; hence, some distance between the point predictions and subjects' behavior is to be expected.[3]

To further understand how the revision mechanism affects behavior and, in particular, to determine whether that effect can be attributed to the forces behind the theoretical results, we study six additional insights from the theory. We begin by investigating the robustness of the efficiency gain provided by the revision mechanism. The theory predicts that the same outcome should be observed independent of the initial effort profile and for various payoff parameters. To test these predictions, we consider two treatments called *random-revision mechanism* (R-RM) and *revision mechanism VHBB* (RM-VHBB). R-RM is similar to RM, except that subjects' initial choices are randomly assigned to them. RM-VHBB is identical to RM, except that we use the main payoff parameters from Van Huyck, Battalio, and Beil (1990). The results indicate that subjects' performance in the revision mechanism is invariant to having exogenous initial choices and to a different set of payoff parameters (Result 3).

We then examine two assumptions that are key to the theoretical results. The unique outcome prediction relies on revision opportunities being frequent and asynchronous. To test these predictions, we consider two treatments called *infrequent revision mechanism* (I-RM) and *synchronous revision mechanism treatment* (S-RM). I-RM is similar to RM, except that the frequency of revisions is reduced to one-eighth of the RM frequency. S-RM is identical to RM, except when a group receives a revision, every member of the group receives a revision opportunity. We find that I-RM and S-RM lead to significantly lower efficiency levels than RM; hence, we conclude that the frequency and asynchronicity of revisions contribute significantly to achieving 82 percent efficiency in RM (Result 4).

The fifth insight is that commitment matters. Since a player may not have a chance to revise her prepared action, players should not treat their own preparations or others' prepared actions as cheap talk. To test this, we introduce the *revision cheap talk* (R-CT) treatment. R-CT follows RM's protocol with one exception;

[2] The effects of one round of public cheap-talk messages found in this paper are similar to the effects documented in the literature. In Blume and Ortmann (2007), similar communication treatment leads to 69 percent efficiency, an improvement over the 34 percent efficiency of their baseline treatment.

[3] We further evaluate exact theoretical predictions of other treatments and papers in the literature, with similar failures, in online Appendix C.

when the 60-second pre-play is over, the subjects in R-CT can choose any effort they wish, and they are in no way committed to what they stated during the pre-play. While R-CT reaches 67 percent efficiency, RM reaches 82 percent. Therefore, the commitment in RM is a significant factor in achieving higher efficiency rates in RM compared to R-CT (Result 5). The sixth insight is that a player's best response depends not only on the effort profile chosen by others but also on the time left before the deadline. If the deadline is close, then a player should revise her effort to match the minimum of the group—maximizing her own payoff. In contrast, if the deadline is far enough in the future, it is optimal for a player to revise her effort in a forward-thinking way—that is, to revise her effort upward. The results show that early revisions are vastly forward thinking in R-RM, while late revisions are payoff improving (Result 6).

Let us take an overall look at these results and highlight some observations. First, removing any of the key elements behind the revision mechanism—commitment, asynchronicity, and frequency of revision opportunities—leads to significantly lower efficiency than in RM. Therefore, all three components are necessary to achieve the levels of coordination observed in RM. Second, removing either of the key elements reduces efficiency to the levels of standard cheap talk (S-CT), which is significantly higher than the level of the Baseline. We take this evidence to indicate that while each of the key elements of RM individually generates an improvement over Baseline, only the combination of all leads to a significant improvement over S-CT.

The institution experimentally studied in this paper is predicated on the combination of three ingredients: a setting in which players would like to coordinate their actions; a pre-play phase during which players publicly display their prepared actions; and incremental commitment, as preparations cannot be changed instantaneously. The revision mechanism can be interpreted in two distinct ways. First, it can represent an intentionally designed institution for implementing the efficient outcome. In this sense, our results have practical implications for the designer, as we illuminate the assumptions and forces relevant to the success of this complex theoretical mechanism. Second, the mechanism can be understood as a feature of real-world coordination environments. Although not always formal, scenarios in which preparation, communication, and incremental commitment go hand in hand form an integral part of our social lives.

I. Literature Review

A large experimental literature, spurred by Van Huyck, Battalio, and Beil (1990), has established that coordination failures—ubiquitous in the real world[4]—are also common in experimental settings. The main contribution of our paper is the examination of how commitment-enhanced pre-play communication can help reduce coordination failures and improve subjects' payoffs. In addition, given the mechanism studied in this paper, our results relate to the literature that focuses on the effects of commitment and real-time interaction on coordination.

[4] For instance, see Rosenstein-Rodan (1943); Murphy, Shleifer, and Vishny (1989); Matsuyama (1991); Rodrik (1996); and Li (2012).

In coordination environments, players face a very particular trade-off because their preferences are more aligned with those of others than in most other strategic situations. The main hurdle for coordination and efficiency is the presence of strategic uncertainty. Some researchers argue that costless pre-play communication could eliminate this hurdle. Blume and Ortmann (2007) and Deck and Nikiforakis (2012) implement a cheap-talk communication phase before the actual play in a minimum-effort game. The pre-play communication in Blume and Ortmann (2007) is done with one round of simultaneous public messages, whereas the protocol in Deck and Nikiforakis (2012) allows for richer interaction, with the subjects having one minute to choose an effort level and the ability to revise their chosen effort at any time.[5] Cheap-talk communication improves coordination and boosts efficiency to 69 percent and 71 percent (from 34 percent and 44 percent)[6] in Blume and Ortmann (2007) and Deck and Nikiforakis (2012), respectively. Despite the ability to update the messages at any second in Deck and Nikiforakis (2012), the efficiency levels in these papers are similar, suggesting that multiple rounds of cheap-talk communication do little to improve efficiency over a single round. Moreover, the gains from the baseline are higher in Blume and Ortmann (2007) than in Deck and Nikiforakis (2012). Similarly, in this paper, we find that cheap-talk treatments produce similar efficiency levels, regardless of whether pre-play communication consists of simultaneous one-shot public messages, multiround rich communication, or a richer message space.

The communication studied in this paper contains no explicit cost of sending messages. However, there is an implicit cost of communication—the inability to revise the intended effort choices instantly. Van Huyck, Battalio, and Beil (1993) and Devetag (2005) consider a costly form of pre-play communication (a pre-play auction in each round) and conclude that such an extension enables the players to achieve better coordination on the payoff-dominant profile in coordination games. Kriss, Blume, and Weber (2016) study the effects of costly and voluntary communication with full and partial subsidies on coordination in a minimum-effort game. The authors find that even a small cost of message sending deters subjects from communication and leads to high coordination failures. Fehr (2017) endogenizes the presence of pre-play communication in an environment in which two groups with prior coordination history are merged. The author finds that most subjects are unwilling to pay a small cost of establishing and maintaining pre-play communication.

To our knowledge, our paper is the first to experimentally study the effects of incremental commitment in a coordination game. The impact of incremental commitment on cooperation has been studied in the context of public goods games. Building on insights by Schelling (1960), Dorsey (1992) is the first to introduce revisions and real-time monitoring in a voluntary contribution mechanism. Looking at those results from a different perspective, Duffy, Ochs, and Vesterlund (2007)

[5] See also Brandts, Cooper, and Weber (2015) and Bornstein, Gneezy, and Nagel (2002) for alternative mechanisms to increase coordination, through leadership and intergroup competition.

[6] We use normalized efficiency throughout this paper, as it summarizes the strength of the treatments. Also, normalized efficiency allows us to compare results from frameworks with different payoffs or group sizes. In particular, our paper, Blume and Ortmann (2007), and Deck and Nikiforakis (2012) each use a different payoff specification. Group size in the current paper is the same as in Deck and Nikiforakis (2012) but is different than in Blume and Ortmann (2007). We provide more details of the measure in the results section.

test theoretical predictions about the dynamic voluntary-contribution game in Marx and Matthews (2000) and show that, whereas a dynamic setting increases the rate of contributions compared with a static setting, the results do not seem to be driven by the theoretically identified forces. Fundamental differences exist between the forces that impede coordination on the Pareto-efficient equilibrium in coordination games and the forces that drive the lack of cooperation in the public-goods provision and social dilemmas. In the latter, the trade-off is between efficiency and individual rationality. In contrast, in coordination games, the miscoordination is a result of the multiplicity of equilibria along with a lack of selection criteria, leading to strategic uncertainty. Contrasting our results with those of Duffy, Ochs, and Vesterlund (2007) highlights how different the two settings are. In this paper, we show that pre-play revisions significantly improve efficiency in coordination games, and we are also able to highlight the critical assumptions for its success.

Roy (2023) experimentally studies market competitiveness in a Cournot duopoly in which firms can simultaneously revise their targeted quantities before the final production. Building on the revision-games theoretical results in Kamada and Kandori (2017), Roy (2023) tests the prediction that a synchronous revision Cournot duopoly may result in higher collusion than in the case without stochastic interaction. Although the theories that predict more collusion in Roy (2023) and higher coordination in our paper have some overlap, the forces behind them are fundamentally different. First, when a revision phase is introduced to a Cournot duopoly, the set of possible equilibrium outcomes increases, and collusion becomes theoretically sustainable. By contrast, the introduction of a revision phase to coordination games shrinks the set of equilibrium-supported outcomes to the unique Pareto-efficient profile. Inefficiencies in coordination settings and the Cournot duopoly arise from fundamentally different forces, calling for distinct mitigation mechanisms.

II. General Framework

A. *Component Game*

Consider a normal-form game $\left(I,(E)_{i\in I},(\pi_i)_{i\in I}\right)$, where I is a finite set of players, $I = \{1,2,\ldots,n\}$; E is a finite set of effort levels available to each player i; and $\pi_i(\mathbf{e})$ is the payoff for player i given the strategy profile $\mathbf{e} \in \mathbf{E}$, where $\mathbf{e} = (e_i)_{i\in I}$ and $\mathbf{E} = \prod_{i\in I} E$. Let $\bar{e}$ ($\underline{e}$) be the highest (lowest) element of E, and let $\bar{\mathbf{e}}$ ($\underline{\mathbf{e}}$) be the profile for which all players choose $\bar{e}$ ($\underline{e}$).

The results presented in Calcagno et al. (2014) and the results discussed here hold for a wide class of games with common interest.[7] In our experiments, we focus on a particular payoff structure, the minimum-effort game, with payoffs given by

$$\pi_i(\mathbf{e}) = \gamma + \alpha \cdot \min_{j\in I} e_j - \beta \cdot e_i, \tag{1}$$

where $\alpha > \beta > 0$. A player's payoff decreases with a higher choice of effort and increases with the minimum effort among all the players.

[7] A game is a common interest game if it has a strictly Pareto-dominant action profile.

Equilibrium Analysis of the Minimum-Effort Game.—In the minimum-effort game described above, every profile in which all players choose the same pure strategy is a strict Nash equilibrium. These equilibria can be Pareto ranked by the effort choice: the higher the effort, the more efficient is the equilibrium. In particular, $\pi_i(\bar{\mathbf{e}})$ is the highest equilibrium payoff, whereas $\pi_i(\underline{\mathbf{e}})$ is the lowest equilibrium payoff. Note that the minimum-effort game is a game of common interest, as the strategy profile $\bar{\mathbf{e}}$ strictly Pareto dominates any other profile.

We now introduce a definition capturing the level of payoff similarity in a wide set of common interest games, following Calcagno et al. (2014).[8] We then adapt it to our context, applying it to the minimum-effort game.

DEFINITION 1: *Calcagno et al.* (2014): *A component game with common interest is a K-coordination game if, for any pair of players $i,j \in I$ and strategy profile* $\mathbf{e} \in \mathbf{E}$,

$$\frac{\pi_i(\bar{\mathbf{e}}) - \pi_i(\mathbf{e})}{\pi_i(\bar{\mathbf{e}}) - \pi_i(\underline{\mathbf{e}})} \leq K\frac{\pi_j(\bar{\mathbf{e}}) - \pi_j(\mathbf{e})}{\pi_j(\bar{\mathbf{e}}) - \pi_j(\underline{\mathbf{e}})}. \tag{2}$$

A game is a K-coordination game if each player can decrease other players' payoffs by at most K times their own cost of punishment. The constant K captures how similar the players' payoffs are between different action profiles. Applying this to the minimum-effort game, if a player choosing the minimum effort reduces their effort choice by one unit, then their own payoff decreases by $\alpha - \beta$, while other players' payoffs decrease by α. In a general game, the smaller K is, the more similar players' preferences are. In particular, if $K = 1$, the game is a pure coordination game, and players have identical payoffs for any outcome. Any finite game with common interest is a K-coordination game for some finite constant $K \geq 1$. Finally, given the payoff structure of the minimum-effort game, the definition can be further simplified: a component game is a K-coordination game if $\alpha/(\alpha - \beta) \leq K$.

B. *Asynchronous Revision Game*

Consider an environment in which players must prepare their actions before they execute them. We follow Calcagno et al. (2014) in modeling this as an asynchronous revision game: there is a pre-play phase, during which a player can revise their prepared action only if a revision opportunity is awarded to them. At the end of the pre-play phase, the most recently prepared action profile is played, and players collect the payoff associated with that action profile. While Calcagno et al. (2014) model an asynchronous revision game with a continuous-time pre-play phase and revisions governed by independent Poisson processes, we extend their results to an environment with discrete time and multinomial revisions.

Formally, we model this as an asynchronous revision game with discrete time, $t \in \{-T, \ldots, -1, 0\}$. The game proceeds as follows. First, at time $-T$, an initial effort profile is in place. It can be exogenously given to the players, or each player

[8] See Takahashi (2005) for the definition and discussion of the concept.

can simultaneously and independently choose an effort level before the pre-play starts. Second, during the pre-play phase, $t < 0$, each player obtains revision opportunities according to a random process with a symmetric arrival rate. At each instant, a revision opportunity is awarded to the group with probability $p \in (0,1]$. If a revision opportunity is awarded to the group, then it is allocated to one of the players with equal probability. Third, at the end of the countdown, $t = 0$, the posted effort profile is implemented, and each player receives the payoff as specified in the component game.

This is a sequential game with multiple rounds of asynchronous play and perfect information, as players observe all the past events in the revision game. The natural solution concept is subgame perfect equilibrium. We refer to a subgame perfect equilibrium of a revision game as *revision equilibrium*.

We now present the main theoretical result, which is an extension of the result in Calcagno et al. (2014) to the framework presented above. Proposition 1 formalizes the intuition that a player can, if far from the deadline, revise her effort choice upward with an eye on leading others to follow her. If all follow and the group reaches the efficient profile, then players do not revise their choices until the deadline; if players do not follow, then she can backpedal her effort choice. Doing so has a low cost, as the deadline is far and, hence, the chance of no further revisions is very small. The proof is presented in online Appendix A.

PROPOSITION 1: *In a discrete-time asynchronous revision game with a symmetric arrival rate of revision opportunities, if the component game is a K-coordination game with the strict Pareto-dominant action profile, $\bar{\mathbf{e}}$, and the game satisfies $(n-2)K < (n-1)$, then for any $\varepsilon > 0$, there exists $T' > 0$ such that for all $T > T'$, all revision equilibria have $\mathbf{e}(0) = \bar{\mathbf{e}}$ with probability higher than $1 - \varepsilon$.*

Let us briefly discuss the argument of the proof (the proof closely follows the steps from Calcagno et al. 2014 and is presented in online Appendix A). The proof can be divided into two parts. The first part shows that the Pareto-dominant equilibrium profile, $\bar{\mathbf{e}}$, is absorbing. The second step of the proof constructs a payoff lower bound for a player who chooses the highest effort well before the deadline. This is done by induction on the number of players choosing the highest effort, $\bar{e}$. To construct a payoff lower bound for player i, one needs to consider the case that another player obtains a revision opportunity before i does. In that case, we rely on the similarity of different players' payoffs, guaranteed by the condition in Definition 1. Finally, the condition stated in the proposition is sufficient to guarantee that the pre-play length needed for the induction step is finite. Formally, the proposition gives us that, when far before the deadline, independent of the current effort choices, by selecting the highest effort, a player is guaranteed a payoff close to the efficient payoff with a probability close to one.

According to Proposition 1, in any revision equilibrium of a long enough revision game, all the players choose the efficient effort in the payoff-relevant moment with probability close to one. This result holds independent of the effort configuration at the beginning of the revision phase. If the time horizon is long enough, then at $t = 0$, all players will be choosing the efficient effort (with probability $1 - \varepsilon$). That is, even if all players start with the minimum or randomly determined effort, or if players choose simultaneously at $-T$, all players will be choosing the efficient-effort at $t = 0$ with probability at least $1 - \varepsilon$. It is essential to highlight that, if the conditions for the proposition fail—for instance, if the pre-play phase is too short—then Proposition 1 does not indicate anything about equilibrium selection. In particular, the efficient effort profile, $\bar{\mathbf{e}}$, would still be one possible effort profile played at the end of pre-play in an equilibrium of the game, but it would not be the unique outcome of a revision equilibrium. A different equilibrium could have all players preparing the lowest effort at the end of pre-play.

In online Appendix B, we go beyond the proposition and, given the parameters used in the experiment, we numerically solve the game by backward induction. Consequently, we gain two additional insights into the revision equilibrium strategy. First, if players can choose their effort before the pre-play phase, then the revision equilibrium prescribes that they all choose the efficient effort from the start. Second, if a player has a revision opportunity far enough from the deadline, then it is optimal to revise to the efficient effort, irrespective of what other players are preparing. Solving the game by backward induction also allows us to go beyond the proposition. We show that, focusing on the minimum-effort game, the condition stated in Proposition 1 is sufficient but not necessary.

C. *Pre-play Communication*

Theoretical work regarding cheap-talk pre-play communication in coordination games has focused on evaluating the credibility of a message profile. The idea is that pre-play communication will promote Pareto-efficient Nash equilibrium play if players' messages are credible when they communicate their intentions to take a certain action. The literature has proposed several requirements for a message to be considered credible.

The early literature considers one-way communication and analyzes the credibility of a message in isolation. For instance, Farrell (1988) postulates that a message is credible if it is self-committing: if the message is to be believed, a sender's best response to the actions induced by her message is to follow the intention stated in the message. If we consider the minimum-effort game with one player sending a public message to all other players, then sending an efficient-effort message is self-committing. If all receivers believe the message and choose the efficient effort, the sender's optimal choice is to follow the message and choose the efficient effort. Aumann (1990) challenges the above reasoning, focusing on whether the sender has a strict preference over the other players' strategy choices. The author argues that, when the sender wants the receiver to believe the message, whether or not they intend to act in accordance with it, the message has no credibility. For Aumann (1990), a message leads to effective communication only if it is self-signaling: the sender wants their message to be believed if, and only if, they plan to follow the intention conveyed in the message. Note that, in the minimum-effort game, a player weakly prefers that other players choose the efficient effort level, independent of the player's choice; hence, a message signaling the intent to choose the efficient effort is not self-signaling.

Self-committing and self-signaling are both concepts that relate to individual messages. Although the definitions could be generalized to the case of multilateral communication, message credibility needs to be defined for profiles of messages, not for individual messages. A player can simultaneously be a sender and a receiver of a message, and a player might send a message linked to one equilibrium and receive a message linked to another. In particular, Blume (1998) argues that communication makes an equilibrium profile more attractive for a player only if all players communicate homogeneously, agreeing on the equilibrium in question. In a minimum-effort game with one-shot multilateral communication, if the chosen message profile is homogeneous, then the associated equilibrium profile could be considered more salient. However, no consensus exists on how to interpret heterogeneous message profiles, and the standard prediction is that communication will be ignored.[9]

D. *Exact Predictions and Qualitative Insights*

Before we describe the experimental design of the paper, we distinguish between *the exact theoretical predictions* and *the qualitative insights* laid down by the forces behind such results. Two predictions arise from the theoretical setup. First, players' payoff should be close enough to the efficient payoff. Second, focusing on players' choices, when enough time remains before the deadline, it is optimal to revise the effort choice to the efficient effort at the first revision opportunity awarded (see online Appendix B). As a consequence, if players can choose their initial effort, all players should choose the efficient effort from the start.

Going beyond these exact predictions, we consider six qualitative insights. First, the theory suggests that one should expect the same outcome, regardless of whether the players choose the initial choices or they are picked for them randomly before the pre-play starts. Second, the theory proposes that the same outcome should be expected for different payoff specifications (within some parametric limits).

Third, as explicitly stated in Proposition 1 and verified numerically in online Appendix B, the uniqueness of the revision equilibrium is conditional on the presence of frequent-enough revision opportunities. If revisions are infrequent, then all equilibria of the component game are an outcome of a revision equilibrium of the extended game.

Fourth, revisions being asynchronous is key to the backward induction argument behind the Proof of Proposition 1. If a player cannot condition her behavior on other players' effort choices when a revision opportunity arises—for instance, if other players also have a revision opportunity at that time—then the repetition of any static Nash equilibrium is a revision equilibrium of the extended game.

Fifth, commitment matters: even if revisions are frequent, it is key that players do not treat their own preparations, or others' prepared actions, as cheap talk. A revision is not cheap talk, as there is a strictly positive chance of not having any

[9]Empirical evidence on the topic is mixed and is context and game dependent. In coordination games in which there is some conflict of interest, He, Offerman, and van de Ven (2019) find support for the "feigned-ignorance principle"—i.e., players ignore messages unless they reach an agreement in which both players are weakly better off. In a coordination game, Cooper et al. (1992) find that actions following heterogeneous messages significantly differ from actions without communication. Focusing on one round of cheap-talk communication, using the data in Blume and Ortmann (2007), and in our paper, we find that in cases with heterogeneous message profiles, there is a significant correlation between minimum message and minimum effort chosen. Hence, even when subjects face a heterogeneous message profile, they seem to extract information instead of entirely ignoring communication.

revisions before the deadline. As time passes, this chance increases, and players are further committed to their prepared actions.

Sixth, a player's best response to a revision opportunity (given the effort profile being prepared) depends on the time left before the deadline. For instance, consider that, at a time $-t$, all players are choosing the minimum effort, and player i has a revision opportunity. If t is large enough, the proposition dictates that it is optimal for player i to choose in a forward-thinking way, to revise her effort upward, and to initiate a chain reaction that will end with all players choosing the highest effort. In contrast, if t is small, the probability of further revision opportunities for all players is negligible. Hence, it is optimal for player i to choose the minimum effort.[10]

III. Experimental Design

In this section, we first present three main treatments that establish the effect of the revision mechanism. Then, we describe additional treatments to test the exact prediction and theoretical insights discussed in Section IID. The instructions used in our experiment can be found in the online Appendix.

The experimental sessions were conducted at the Center for Experimental Social Science (CESS) laboratory at New York University (NYU) and at the Interdisciplinary Experimental Laboratory (IELAB) at Indiana University (IU), using the software z-Tree (Fischbacher 2007).[11] All participants were NYU or IU students. The experiment lasted about 45 minutes, and subjects earned, on average, \$18, which included the \$8 show-up fee. In each session, written instructions were distributed to the subjects and also read aloud.

In all treatments, participants are randomly divided into groups of six, and they participate in a sequence of ten rounds as a part of that group. In each round, the group plays the minimum-effort game with effort choices from 1 to 7 $\left(E = \{1,2,3,4,5,6,7\}\right)$. For all but one treatment, the subjects have the same payoff function: $\pi_i(e) = 0.18 - 0.04 \cdot e_i + 0.2 \times \min_{j \in I} e_j$. For treatment revision mechanism VHBB (RM-VHBB), the payoff function is $\pi_i(e) = 0.60 - 0.1 \cdot e_i + 0.2 \times \min_{j \in I} e_j$, as the primary parameters in Van Huyck, Battalio, and Beil (1990). The payoffs are described to subjects in matrix form, and the subjects take a comprehension test to ensure that they understand the payoff structure. After ten rounds, subjects answer a short survey and are paid their final payoff, which is the sum of the payoffs from all ten rounds plus the show-up fee.

[10] To gain some intuition, consider a player facing a scenario in which every other player is choosing the lowest effort, 1. On the one hand, if there are 59 seconds left before the deadline, the player should revise her choice to 7, if given the opportunity. Figure B.1, panel A in the online Appendix shows us that when the deadline is far away, revising the effort choice to the highest effort is dominant, irrespective of other players' choices. On the other hand, if there is only one second left in the pre-play, then the probability of reaching the optimal strategy is zero; therefore, the player should revise the effort to the minimum effort choice by the group, i.e., 1, if given the opportunity. The best response during the pre-play period depends on both the time left and the prepared profile. In particular, given our parameters, the number of iterations required for 7 to be the dominant effort choice is 13 (see Figure B.1, panel A in the online Appendix). If there is less time left in the pre-play interaction, choosing 7 is not a dominant choice, and the best response depends on the current profile.

[11] The sessions at NYU were conducted in December 2015, February 2016, and April through July 2018. At NYU, a session included either 12 or 18 subjects. The sessions at IU were conducted in October–November 2020 and March–April 2021. Due to COVID-19 social distancing norms, each session included only one group with six subjects.

Baseline.—The baseline treatment replicates the standard control treatment in the literature. Subjects play the normal-form one-shot game. After each round of playing a standard simultaneous minimum-effort game, participants receive feedback on the minimum number chosen in their group in that round. This information is the only history available to them in the baseline treatment.

Revision Mechanism (*RM*).—We design a treatment that closely replicates the conditions of our theoretical setup. However, implementing this institution in the laboratory presents several challenges.

One challenge is that the game involves frequent interaction among players, and they need to have all the information at every round. Thus, each player's revision opportunity and posted effort, as well as the history of posted efforts and revisions, should be available to all players at all times. We compile this information in a graph that summarizes all the key points and represents the players' efforts in different colors. Every time a player receives a revision opportunity, a dot appears on that player's action line. The graph summarizes all the key information and makes it easily accessible to the subjects. Figure 1 presents an example of the graph after 30 seconds have passed.[12]

Another challenge is the implementation of revision opportunities. Theory dictates that revisions should happen frequently. If we stop ("freeze") the phase every time a subject receives a revision opportunity, the phase could last a long time. To control how long a phase lasts, we let subjects change a number any time they want; all they need to do is place the cursor over the button on the screen. However, the subject's new selection is updated on the graph only after the subject receives a revision opportunity. A by-product of this method is that we gather two streams of data: payoff-relevant decisions and what subjects want to do (we use these data streams to test whether our choices for the frequency of revisions and the length of the pre-play interaction restrict players' behavior; see online Appendix for more details).

In the RM treatment, each round begins with all group members simultaneously choosing a number from 1 to 7.[13] Once all group members make their initial choices, a graph appears, and a one-minute countdown begins.[14] In Figure 1, we present an example of the graph after 30 seconds of the countdown. The time in seconds is on the horizontal axis, and the number chosen by each group member is on the vertical axis. The initially chosen numbers are along the vertical line above the zero-second mark. Each player is represented on the graph by a different color. As the countdown progresses, at any time, any member of the group can change their chosen number by placing the cursor on the desired number on the left side of the screen. When a subject selects a number, the respective button turns green on the subject's own screen (see the number 4 in Figure 1). The subjects do not see their group members' "planned choices." The number posted on the graph updates only when the subject receives a revision opportunity, and the entire group can see this update on the graph.

[12] We thank Bigoni et al. (2015) and Friedman and Oprea (2012) for sharing their code with us.

[13] Data for RM were collected both at NYU and IU. Data for 8 groups were gathered at NYU and for 8 groups at IU, resulting in 16 groups for this treatment. We do not find any considerable differences between the two locations, and we, therefore, combine the data for the analysis throughout the paper.

[14] We explained the graph in great detail in the instructions, and all subjects took a comprehension test regarding the graph and payoff table

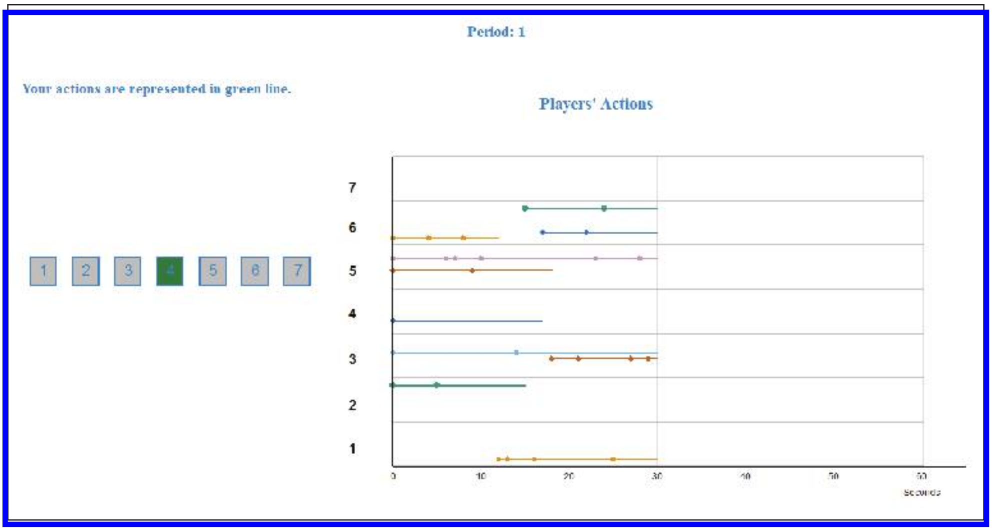

FIGURE 1. SAMPLE SCREEN AFTER 30 SECONDS IN RM

On average, a subject receives eight revision opportunities in one round. Formally, at every second, the group has an 80 percent chance of receiving a revision opportunity; if that occurs, then the six group members have an equal probability of $1/6$ of receiving the revision opportunity. Only the numbers posted at the end of the countdown matter for the payoff. The initially chosen efforts and all the revision effort choices are irrelevant for the payoff calculation.

Standard Cheap Talk (*S-CT*).—The standard cheap-talk treatment offers subjects multilateral one-shot communication, similar to the main communication treatment in Blume and Ortmann (2007). In this treatment, before subjects make their payoff-relevant effort choices, they simultaneously send a public message (a number from 1 to 7). This is followed by 60 seconds during which subjects see all the messages sent by their group members (including their own message).[15] After the subjects see the messages, buttons appear, and they make their payoff-relevant effort choices (all the group messages are visible on the screen when subjects make their payoff-relevant decision). At the end of the round, the subjects see a feedback page with their choice and the minimum number chosen by the group.

Random Revision Mechanism (*R-RM*).—In the random revision mechanism, initial choices are randomly chosen for the subjects using a uniform distribution over all possible efforts, with all groups facing the same initial effort. However, the rest of the round is executed identically to the RM protocol. We know, theoretically, that the outcomes of this treatment should be similar to those of RM; however, we get a much richer best-response behavior due to the initial heterogeneity of effort choices.

[15] To ensure that subjects spend the same amount of time in the lab and have an experience similar to that of the RM treatment, we display the messages sent by subjects on a graph and give the subjects 60 seconds to make the payoff-relevant effort choice.

Revision Mechanism VHBB (*RM-VHBB*).—The revision mechanism VHBB treatment is identical to the RM treatment, except for the payoff parameters. In RM-VHBB, we use the main payoff parameters used in Van Huyck, Battalio, and Beil (1990) ($\alpha = 0.2$, $\beta = 0.1$, and $\gamma = 0.6$).

Infrequent Revision Mechanism (*I-RM*).—In the infrequent revision mechanism, we reduce a group's probability of having a revision opportunity from 0.8 of RM to 0.1. In RM, the chance of having no more revisions 60 seconds before the deadline is approximately 0.0 percent, while in I-RM, it is 36 percent. In online Appendix B, we numerically solve the game with 0.1 as the probability of a revision opportunity, and we find that at least 164 seconds per round would be needed for the theoretical results to hold.

Synchronous Revision Mechanism (*S-RM*).—The synchronous revision mechanism is identical to RM, except that the revisions for all subjects in a group coincide. Recall that in RM, a revision is awarded to a group with an 80 percent chance every second. When a revision opportunity is awarded to a group, it is given to one of the group members with equal probability. In S-RM, we have essentially combined each group member's one revision occurring asynchronously (six revisions in total) into one simultaneous revision when all six group members can revise simultaneously. The realizations of revision opportunities used in S-RM are taken from the realizations of revisions used in RM for one of the group members. That is, in S-RM, when a group receives a revision opportunity, all group members receive it at once with a $1/6$ chance, or no group member receives it with a $5/6$ chance. This ensures that the expected number of revisions per subject in S-RM is eight, the same as in RM. The only difference between RM and S-RM is that S-RM revisions are synchronous, while revisions in RM are asynchronous.

Revision Cheap Talk (*R-CT*).—The revision cheap-talk treatment follows the RM pre-play phase protocol. First, all members of the group simultaneously choose an integer from 1 to 7; then, once everyone makes a choice, the one-minute countdown begins. As in the revision mechanism, all members of the group see the same real-time graph, and the chosen effort is updated only when a revision opportunity is awarded.[16] In contrast to RM, the choice at the end of the countdown is not payoff relevant in R-CT. Once the 60-second countdown is over, a new screen appears, and subjects choose an integer from 1 to 7 that determines their payoffs.[17]

Richer Revision Cheap Talk (*R-R-CT*).—In the richer revision cheap-talk treatment, subjects can inform others what they intend to play as well as what they think everyone should play. R-R-CT has a protocol similar to that of R-CT, except that in

[16] The revision realizations used in R-CT are the same as in RM.

[17] This treatment is conducted in two ways. First, for eight groups, while subjects are making payoff-relevant choices, the graph from the round is not present; we refer to this treatment as R-CT-O. To avoid concerns about memory issues, we rerun this treatment with an additional eight groups; we refer to this treatment as R-CT-M. For these subjects, when payoff-relevant buttons appear, the graph with the history of 60 seconds is present on the same screen. Because we do not find any considerable differences between R-CT-O and R-CT-M treatments, we combine the data, 16 groups, and refer to this as R-CT treatment.

TABLE 1—EXPERIMENTAL DESIGN

Treatment	Communication	Commitment	Subjects	Groups
Baseline	None	NA	48	8
Standard cheap-talk (S-CT)	One-shot	None	48	8
Revision mechanism (RM)	Revisions	Gradual	96	16
Random RM (R-RM)	Revisions	Gradual	48	8
RM VHBB (RM-VHBB)	Revisions	Gradual	48	8
Infrequent RM (I-RM)	Revisions	Abrupt	48	8
Synchronous RM (S-RM)	Revisions	Gradual	48	8
Revision cheap-talk (R-CT)	Revisions	None	96	16
Richer R-CT (R-R-CT)	Revisions	None	48	8

R-R-CT, subjects observe two sets of buttons and two graphs similar to Figure 1. After 60 seconds, the graphs stay on the screen and buttons appear, which subjects use to make payoff-relevant choices. After every group member has made their choice, subjects observe a feedback page presenting their choice and the minimum of their group.

Table 1 summarizes our experimental design and highlights the differences between our treatments based on the two main dimensions: communication and commitment.[18]

IV. Results

In this section, we first establish the main result of the paper: introducing a revision mechanism to a minimum-effort game significantly improves efficiency compared with both the one-shot game and the one-shot game preceded by a round of one-shot cheap-talk communication (RM versus Baseline and RM versus S-CT). We then test the exact predictions of the theory highlighted in Section IIB. We proceed by using additional treatments described in Section III to shed light on the six insights provided by the theory, discussed in Section IID. Further, we take a deeper look into how subjects communicate and how commitment affects their communication.

A. *Overall Effect of Revision Mechanism*

We first focus on the overall efficiency of treatments. We follow the literature and calculate normalized efficiency as

$$(3) \qquad Efficiency = \frac{Actual - Min}{Max - Min},$$

where *Actual* is the average amount earned in a treatment and *Min* (*Max*) is the average minimum (maximum) possible amount that a subject can earn. Results are displayed in Figure 2. The RM efficiency is 82.1 percent, whereas, in Baseline and S-CT treatments, the efficiency is 47.8 percent and 64.1 percent, respectively (Mann-Whitney U [MWU] tests lead to p-values of less than 1 percent for both RM

[18] One session of the baseline treatment was voided because one of the subjects publicly announced an intended action and asked others to play the same. We ran an extra session to replace the voided session.

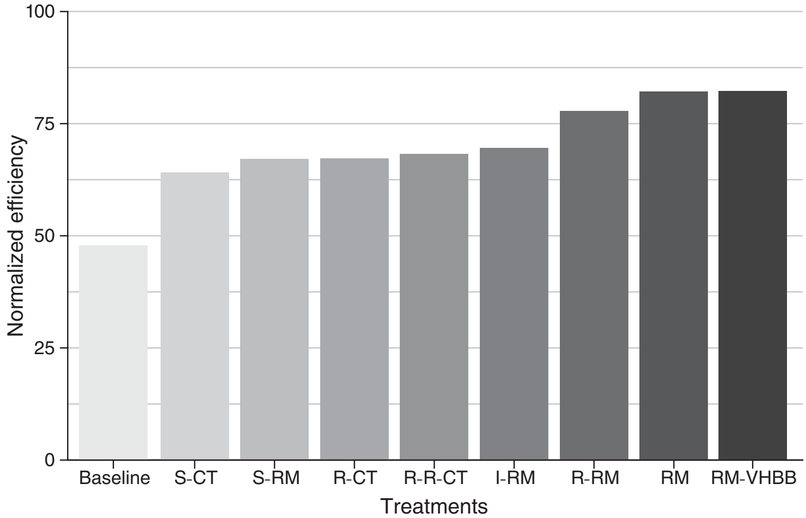

Figure 2. Treatment Normalized Efficiency

versus S-CT and RM versus Baseline). The introduction of the revision mechanism restores more than half (65.7 percent) of the efficiency loss in the Baseline. While these results are based on all ten rounds combined, the differences among earnings in the Baseline, S-CT, and RM treatments get stronger over the ten rounds. In the first round, the normalized efficiency is 43.1 percent, 51.1 percent, and 69.7 percent in the Baseline, S-CT, and RM, respectively. In contrast, in the last round, the normalized efficiency is 43.8 percent, 68.0 percent, and 92.0 percent, respectively.[19]

Given the payoff function for the minimum-effort game (see equation (1)), deviations from the efficient equilibrium reduce payoffs in two ways. First, subjects choose an inefficient minimum effort, and, second, subjects miscoordinate and select different efforts. We now focus on outcome variables to capture the effect of RM over these two forces and compare it with other treatments. To capture whether subjects try to coordinate on the efficient profile, we analyze the minimum effort of the group as well as the frequency of efficient effort choices. To assess the coordination on any effort profile achieved by the group, we rely on the fraction of fully coordinated groups and a novel measure, *equilibrium deviation*, which captures how far a group is from full coordination. Equilibrium deviation calculates, for each group, the average distance between the effort choices and the myopic best response—the minimum effort chosen in the group. Note that equilibrium deviation calculates how far the choice is from the minimum of the group. This measure does not capture the distance between the subject's current choice and the revision equilibrium.

A comparison of the treatments clearly reveals that S-CT falls short of RM on all four of these measures. In particular, the average minimum effort is lower (4.61 versus 5.83), as are the average frequency of efficient effort (0.44 versus 0.78) and

[19] For most of this section, we focus on aggregate statistics and the overall effects of the treatments. In online Appendix D.D.5, we provide more details on variables of interest over the 10 rounds of play and also within the 60-second pre-play phase.

the average fraction of fully coordinated groups (0.28 versus 0.66).[20] In contrast, the average group equilibrium deviation is higher in S-CT (0.86 versus 0.49). We reject the hypothesis of equal distributions for all measures using Mann–Whitney U tests, with $p < 0.001$, using the group average in a round as a unit of observation.[21]

We run a regression analysis with payoffs and the four aforementioned measures as endogenous variables. We cluster standard errors at the group level so that one group is treated as one independent observation. The results of the regression analysis are presented in Table 2. The control group is S-CT treatment. Baseline and S-CT perform similarly on payoffs, minimum effort, and frequency of efficient effort. But compared to Baseline, S-CT leads to more fully coordinated groups and lower equilibrium deviation. The regression highlights that RM performs better than the two treatments on an aggregate level and all four considered measures. Note that there are eight groups in the S-CT treatment, and, therefore, some results are only marginally significant, with errors clustered at the group level.

Finally, we run two alternative regression analyses of the five outcome variables and display them in Tables D.3 and D.4 in online Appendix D.D.3. We compare the treatments that incorporate the revision mechanism with those that introduce cheap talk. We begin by comparing RM with standard cheap talk and revision cheap talk jointly. In the last regression, we compare all of the revision mechanism treatments (RM, R-RM, and RM-VHBB) with the control group of all treatments that incorporate cheap talk (S-CT, R-CT, and R-R-CT). In all regressions, errors are clustered at the group level. All regression analyses indicate the same outcome: revision mechanisms significantly outperform cheap-talk treatments. This result can be observed in players' payoffs and in the four measures we focus on.

RESULT 1: *The revision mechanism significantly increases efficiency over Baseline and S-CT treatments, by 34 pp and 18 pp, respectively. The increased efficiency is achieved through increases in the minimum effort chosen by the group as well as by the overall coordination on any effort profile.*

B. *Evaluating Exact Predictions of the Revision Mechanism*

We proceed by testing the exact predictions listed in Section IID. We restate them and test each one separately. The first exact prediction is that subjects' payoffs should be close enough to the efficient payoff. Or, alternatively, efficiency of the revision mechanism should be 100 percent. Recall that efficiency of RM is 82.1 percent, which is significantly lower than 100 percent with $p < 0.01$. The second exact prediction is that subjects' initial effort choice should be the efficient effort. We test whether the fraction of subjects initially choosing 7 is 100 percent. In the last round, 93.8 percent of players' initial choices are the efficient effort. On average, 85.7 percent of the initial choice is the efficient effort, which is significantly lower than 100 percent with $p < 0.01$.

[20] See Tables D.1 and D.2 in the online Appendix for details on average payoffs and other relevant variables as well as the statistical test results between the revision mechanism and other treatments.

[21] Note that there are 16 groups in RM and 8 groups in S-CT, and each group plays the game for 10 rounds. Hence, we consider 160 versus 80 observations for each of the measures.

Table 2—Regression Analysis

	Dependent variable				
	Payoffs	*Minimum effort*	*Freq efficient effort*	*Full coordination*	*Eqbm deviation*
Baseline	−0.168	−1.03	−0.173	−0.25	0.589
	(0.124)	(0.75)	(0.125)	(0.116)	(0.268)
Revision mechanism	0.21	1.21	0.339	0.388	−0.395
	(0.012)	(0.712)	(0.136)	(0.141)	(0.221)
Quiz	−0.0337				0.089
	(0.0266)				(0.0944)
Constant	1.02	4.61	0.442	0.275	0.834
	(0.19)	(0.618)	(0.122)	(0.115)	(0.712)
Observations	1,920	320	320	320	1,920
Demographics	Yes	NA	NA	NA	Yes

Notes: Standard errors clustered at the group level are in parentheses. Reference category is standard cheap Talk treatment. *Payoffs* variable is a subject payoff in a round. *Minimum effort*, *Freq efficient effort*, and *Full coordination* are group-level measures, and subject demographic information is not applicable. *Eqbm deviation* is a subject-level variable.

RESULT 2: *The revision mechanism achieves 82.1 percent efficiency, which is significantly lower than 100 percent. In addition, the frequency of the initial choice of efficient effort, 85.7 percent, is significantly lower than 100 percent.*

The evidence is not sufficient to support the exact theoretical predictions of the asynchronous revision mechanism; we see that subjects' behavior is significantly different from the predictions. Given the specificity of the predictions, some distance between subjects' behavior and the point predictions is expected, in some sense. For instance, as an alternative theoretical prediction, pure-strategy Nash equilibrium dictates that, even in the Baseline treatment, all subjects should choose the same effort. In the Baseline, however, the groups fully coordinate on an effort choice (any effort profile, not necessarily efficient) in only 2 out of 80 cases. Similarly low full coordination results are found in the literature; for example, in Van Huyck, Battalio, and Beil (1990), out of 70 cases, there are 0 cases with full coordination. In the baseline treatment for the minimum-effort game of Blume and Ortmann (2007), in 3 out of 32 cases do the groups fully coordinate on an effort level.

C. *Evaluating the General Insights from the Theoretical Framework*

With the goal of understanding how the revision mechanism affects subjects' behavior in the laboratory and, in particular, to determine whether such an effect can be attributed to the forces behind the theoretical results, we now turn our attention to the six general insights presented in Section IID. We start with the two insights related to the robustness of the theory.

Exogenous Initial Choices and Different Payoff Parameters.—The theory postulates that, if the initial choices for the players are picked at random, then—with a long enough pre-play phase—the outcome should be the same as with the endogenous initial choice, with a probability close to 1. We test this prediction by comparing the RM with the R-RM treatment. We find that the RM and R-RM treatments lead to similar behavior in all dimensions. We find no statistically significant difference in efficiency, with R-RM achieving 77.8 percent efficiency, compared with 82.1 percent in RM. Not only is the average payoff similar (10.43 versus 10.93), but the average minimum effort (5.54 versus 5.83), the average frequency of efficient effort (0.71 versus 0.78), the average fraction of fully coordinated groups (0.60 versus 0.66), and the average equilibrium deviation (0.57 versus 0.49) are both statistically and economically indistinguishable. We cannot reject MWU tests of equal distributions for any of the five measures, with $p = 0.173$, $p = 0.169$, $p = 0.152$, $p = 0.343$, and $p = 0.319$, respectively. (For further details, see Tables D.1 and D.2 in online Appendix D.D.2.)

We now focus on the robustness of the prediction to different payoff specifications: the efficient profile is the unique revision equilibrium for the payoff parameters from the original experiment on the minimum-effort game, Van Huyck, Battalio, and Beil (1990). Our data indicate that the RM and RM-VHBB treatments lead to similar behavior in most dimensions. We find no statistically significant difference in efficiency, with R-RM achieving 82.2 percent efficiency, compared with 82.1 percent for RM. Not only is the efficiency similar, but the fraction of fully coordinated groups (0.61 versus 0.66) and the average equilibrium deviation (0.40 versus 0.49) are both statistically and economically indistinguishable. We cannot reject MWU tests of equal distributions for these two measures, with $p = 0.447$ and $p = 0.772$. The minimum effort is marginally lower in RM-VHBB (5.44 versus 5.83), with $p = 0.061$, and the frequency of efficient effort is significantly lower (0.59 versus 0.78), with $p < 0.01$.

RESULT 3: *Subjects' performance in the revision mechanism is invariant to having exogenous initial choices and to alternative payoff parameters that do not satisfy the condition of Proposition 1.*

Frequency and Asynchronicity of Revision Opportunities.—The frequency and the asynchronicity of the revision opportunities are key for the backward-induction argument used to select a unique revision equilibrium outcome in Proposition 1. We test the importance of these two elements by comparing RM to I-RM and to S-RM.

We observe that I-RM leads to significantly lower efficiency, 69.6 percent, compared with 82.1 percent in RM. Moreover, subjects' performance in I-RM falls short of the performance in RM on every one of the other five measures we consider. In particular, the average payoff is lower (9.47 versus 10.93), and so are the average minimum effort (4.99 versus 5.83), the average frequency of efficient effort (0.49 versus 0.78), and the average fraction of fully coordinated groups (0.39 versus 0.66). In contrast, the average equilibrium deviation is higher in I-RM (0.78 versus 0.49). We reject the hypothesis of equal distributions for each of the five measures using MWU tests, with $p < 0.001$.

We find that subjects' behavior in S-RM leads to significantly lower efficiency than in RM: 67.1 percent compared with 82.1 percent. Furthermore, subjects' performance in S-RM is worse than the performance in RM for every one of the other five measures we consider. In particular, the average payoff is lower (9.47 versus 10.93), and so are the average minimum effort (4.99 versus 5.83), the average frequency of efficient effort (0.49 versus 0.78), and the average fraction of fully coordinated groups (0.39 versus 0.66). However, the average equilibrium deviation is higher in S-RM compared with RM (0.69 versus 0.49). We reject the hypothesis of equal distributions for payoffs, minimum effort, and frequency of efficient effort measures using the MWU test, with $p < 0.001$, and for fully coordinated groups and equilibrium deviation measure using the MWU test, with $p = 0.003$ and $p = 0.006$, respectively.

RESULT 4: *The frequency and asynchronicity of revisions are significant contributing components to achieving 82 percent efficiency in RM. When the frequency of revisions is reduced from 8 to 1 or when revisions are synchronous instead of asynchronous*, *the efficiency is reduced by 13 and 15 percentage points*, *respectively.*

The Importance of Commitment.—A key factor in the Proof of Proposition 1 is that players cannot and, therefore, do not treat their own preparations—or others' prepared actions—as cheap talk. Comparing the performance of the R-CT treatment with that of RM empirically highlights the importance of commitment.

We observe that R-CT leads to significantly lower efficiency, 67.2 percent, compared to RM at 82.1 percent. Accordingly, subjects' performance in R-CT is worse than the performance in RM on every one of the other five measures we consider. We find that R-CT leads to lower payoffs (9.19 versus 10.93) as well as to lower minimum effort (4.86 versus 5.83), lower frequency of efficient effort (0.55 versus 0.78), and lower fraction of fully coordinated groups (0.32 versus 0.66). Furthermore, R-CT leads to a significantly higher equilibrium deviation compared with RM (0.97 versus 0.49). In keeping with this, we reject the hypothesis of equal distributions for all five measures using MWU tests, with $p < 0.001$.

RESULT 5: *The commitment in the revision mechanism is a significant contributing component to achieving 82 percent efficiency in RM. Removing commitment from the mechanism reduces efficiency by 15 percentage points.*

The Best Response Dynamics.—This section discusses the dynamic behavior implied by the Proof of Proposition 1 and the numerical exercise in online Appendix B. A player's best response depends not only on the effort profile of other agents but also on the time left before the deadline.

We look at subjects' dynamic behavior in R-RM and classify all the revisions taken into three categories—forward thinking, myopic down, and others—detailed below.[22] We display in Figure 3 the types of moves taken as a function of the amount of time that has passed in the pre-play phase.

[22] The R-RM treatment creates variance in the initial choice, which is lacking in RM since the initial effort was determined at random. We explore this variance in initial choices to analyze the dynamic best-response behavior of our subjects. The distribution of initial efforts over rounds: 7 percent were 1; 13 percent, 2; 25 percent, 3; 13 percent, 4; 18 percent, 5; 13 percent, 6; and 10 percent, 7. This pattern is different from the initial choices in RM. In RM, the distribution of the initial choices was 1 percent were 1; 1 percent, 2; 1 percent, 3; 1 percent, 4; 4 percent, 5; 6 percent, 6; and 86 percent, 7.

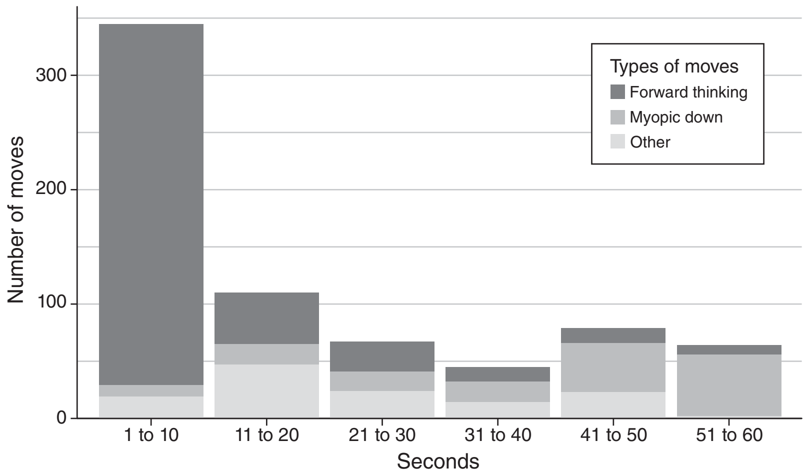

Figure 3. Classification of Moves over Time

Forward Thinking: Moves that would decrease the player's payoff if they occur in the last instant; however, these moves will increase the payoff if other players follow. This category combines two types of moves: (i) the subject's current choice is above the group's minimum, but they still increase their chosen effort; and (ii) the subject's choice is the group's minimum, and they increase their effort above the second minimum effort.

Myopic Down: Moves decreasing the player's effort that would increase their payoff if the moves occur in the last instant. This category contains the moves that get the subject's effort closer to the group minimum.[23]

Other: Moves not included in the above categories. This includes moves increasing the subject's effort choice, which would increase their payoff, and moves decreasing the subject's effort choice, which would decrease their payoff (e.g., because the subject is already at the group minimum or moves below the current minimum).

In the first 10 seconds of a round, more than 300 revision opportunities are used for forward-thinking moves, representing around 92 percent of all the revision opportunities taken. In contrast, in the last 10 seconds, 84 percent of the revision opportunities taken are used for myopic down moves. Note that the subjects have about the same number of revision opportunities available in the first ten and in the last ten seconds of the pre-play interaction. Despite this, in the first 10 seconds of the round, subjects use a similar number of revision opportunities as in the last 50 seconds combined.

Also, note that subjects do not use all early opportunities to revise the prepared effort to the highest effort immediately, as predicted by the theory. We observe that

[23] Notice that if there is not enough time left for others to adjust, moving down to the minimum of the group is a best-response behavior.

78.2 percent of revisions taken in the first 10 seconds are used to revise the effort choice to 7, which is sizable but significantly lower than 100 percent. Given the evidence, the subjects' behavior is largely aligned with the general insight of the theory, even if significantly different from the point prediction.

RESULT 6: *In the revision mechanism, early revisions are forward thinking, while late revisions are myopically payoff improving. In the first 10 seconds, 91.6 percent of moves are forward thinking, and 2.9 percent of moves are myopic down. In contrast, in the last 10 seconds, 12.5 percent of moves are forward thinking, and 84.4 percent of moves are myopic down.*

D. *Communication and Commitment*

We now shift from the theoretical insights to focus on how subjects communicate in different treatments. We examine three aspects: how commitment affects communication, the differences between communication and actions in the cheap-talk treatments; and the impacts of a richer message space.

Effects of Commitment on Communication.—Commitment affects communication in two interconnected ways. First, commitment makes communication more credible, with subjects acting differently in the presence of commitment after a particular profile is revealed. Second, commitment changes the optimal communication in that subjects communicate differently in the presence of commitment. Below, we try to parse these two distinct but interconnected forces.

We first compare the credibility of messages in RM and in R-CT. To do so, we follow the theoretical discussion in Blume (1998), and we examine whether the whole group converging on a particular message profile makes that effort profile (an equilibrium in the component game) more salient, thus leading to higher coordination on that profile. We look at all groups that converged to stating their intention to choose the same effort level (not necessarily the efficient effort profile). This happens with a frequency of 75.6 percent and 56.2 percent in RM and R-CT, respectively. Out of all the times that a group converges to a common effort profile during the 60-second pre-play, that profile is implemented at the payoff-relevant moment in 87.6 percent of these rounds in RM, compared with 51.1 percent in R-CT.

We now focus on how commitment affects what is communicated. In Figure 4, panel A, we display the average equilibrium deviation of the prepared—and, hence, publicly posted on the graph—efforts over the 60-second interval.[24] The graph shows a decline in the equilibrium deviation over the pre-play in RM, as players coordinate more. We also document a difference between RM and R-CT, especially in the latter part of the pre-play phase. Facing a group member who is choosing a smaller effort close to the deadline, a subject reduces their chosen effort (thus reducing the equilibrium deviation) in RM but not in R-CT. In the absence of commitment, lowering the choice of effort is unnecessary because revising it in the final instant is possible.

[24] In Figure 4, we use the last five of the total ten rounds. In online Appendix D, we present the graph for all ten rounds (Figure D.6, panel A) and additional treatment, R-R-CT.

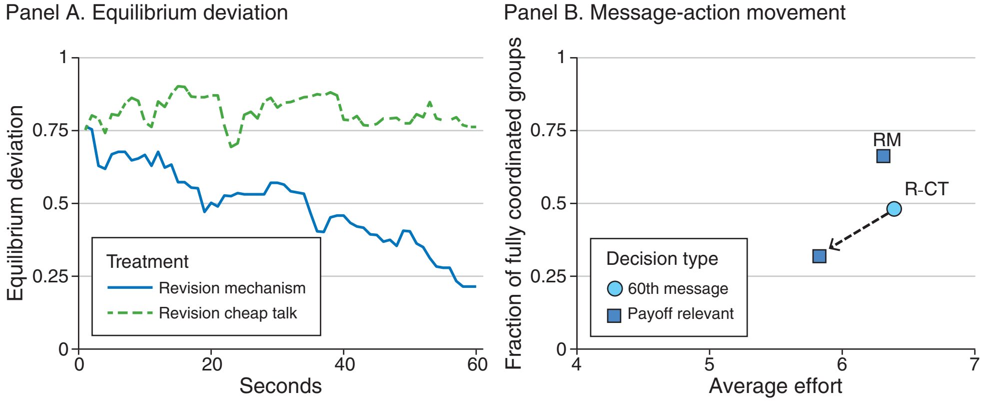

Figure 4. Communication and Action Differences

RESULT 7: *Convergence to a common message leads to more credibility of the communication in the presence of commitment. When a group converges to a common effort profile during the 60-second pre-play, that profile is implemented at the payoff-relevant moment 87.6 percent and 51.1 percent of the time in RM and R-CT, respectively.*

Differences between Communication and Action.—We now turn to the differences between the profile communicated in the last instant of the pre-play phase and the payoff-relevant effort choices in cheap-talk treatment. As depicted in Figure 4, panel B, both the average effort and the fraction of fully coordinated groups of the payoff-relevant profile are significantly lower than those of the sixtieth second profile.

Comparing the payoff-relevant efforts in R-CT and the sixtieth second message in R-CT, we have that the average payoff is lower (9.19 versus 10.18), as are the average minimum effort (4.9 versus 5.5), the average frequency of efficient effort (0.56 versus 0.77), and the average fraction of fully coordinated groups (0.32 versus 0.48). However, the average equilibrium deviation is similar (0.97 versus 0.92). We reject the hypothesis of equal distributions for payoffs, minimum effort, and frequency of efficient effort measures using the MWU test, with $p < 0.001$, and for fully coordinated groups using the MWU test, with $p = 0.002$. We cannot reject the equal distribution hypothesis of equilibrium deviation, as the MWU leads to $p = 0.103$. The difference between the messages and actions leads to a 10.7 percent loss in payoffs.

RESULT 8: *Substantial differences exist between the preparations at the last second and the effort implemented, in the absence of commitment in R-CT.*

Richness of Communication.—The results above point to the lack of credibility of the messages as a possible culprit behind the differences in subjects' behavior between the RM and R-CT. It might be that players do not believe in others' communication since they cannot distinguish positive communication about intentions ("I intend to play … ") from normative communication about the group's effort profile ("We should all coordinate on … "). We now consider whether enriching the message space to include a second message—"I think we should all choose"—improves coordination.

We find that R-CT and R-R-CT treatments lead to similar behavior in all dimensions. Not only is the average payoff similar (9.19 versus 9.31), but the average minimum effort (4.86 versus 4.94), the average frequency of efficient effort (0.56 versus 0.60), the average fraction of fully coordinated groups (0.32 versus 0.36), and the average equilibrium deviation (0.97 versus 0.96) are statistically indistinguishable. We cannot reject MWU tests of equal distributions for any of the five measures, with $p = 0.685$, $p = 0.672$, $p = 0.385$, $p = 0.5$, and $p = 0.687$, respectively. Accordingly, R-R-CT performs significantly worse on all five measures compared with RM.

RESULT 9: *Enriching the cheap-talk messages to include explicitly labeled normative messages about what the group's choice should be does not significantly improve subjects' payoffs. Richer R-CT treatment leads to 68.2 percent efficiency, which is statistically indistinguishable from the efficiency of 67.2 percent in R-CT.*

V. Concluding Discussion

Coordination environments are prevalent in real-world situations, and coordination failures leading to inefficiencies are widespread. In this paper, we provide experimental evidence that commitment-enhanced communication can significantly reduce coordination failures in a particular coordination game—the minimum-effort game. A helpful way to summarize our results and to compare the treatments is to look at the efficiency:

$$RM > S\text{-}CT \approx R\text{-}CT \approx I\text{-}RM \approx S\text{-}RM > Baseline. \tag{4}$$

RM delivers significantly higher efficiency than S-CT treatment, and this boost is achieved through increases in the minimum effort and overall coordination on any effort profile. While the revision mechanism undoubtedly provides gains in efficiency, these gains are based on the presence of all key ingredients at the same time: commitment (RM versus R-CT), asynchronicity (RM versus S-RM), and frequency of revisions (RM versus I-RM).

Let us take a closer look at the $S\text{-}CT > Baseline$ part of expression (4). The theory we examine in this paper is not equipped to explain this result. Even if we consider alternative theories, as discussed in Section IIC, introducing cheap-talk communication does not necessarily reduce the set of equilibria in a minimum-effort game. Previous literature has already documented the beneficial effects of communication even without theoretical support. In a Stag-Hunt game—in which a message indicating intent to cooperate is self-committing but not self-signaling—Charness (2000) finds that one-sided communication increases the coordination on the efficient outcome. Communication has beneficial effects even beyond coordination games. It leads to more egalitarian allocations and increases efficiency in a divide-a-dollar bargaining game with a unanimity voting rule (see Agranov and Tergiman 2019). Venturing into social-dilemma games, communication leads to higher rates of contributions in a public goods game (see Isaac and Walker 1988; Ostrom, Walker, and Gardner 1992; Oprea, Charness, and Friedman 2014; and Palfrey, Rosenthal, and Roy 2017). This growing empirical evidence could further inspire theoretical work to account for the effects of cheap-talk communication in various environments.

Continuing our discussion of expression (4), let us explore a rationale for R-CT, S-RM, and I-RM faring significantly better than Baseline. The gains of the R-CT treatment over Baseline are unsurprising. One round of restricted message exchange in S-CT delivers a significant improvement over the one-shot game without any interactions; therefore, it is expected that the multiple exchanges of messages allowed in R-CT should do at least as well. More surprisingly, the richness of communication in R-CT does not improve upon S-CT treatment. Our examination of communication credibility suggests that, if a message is not seen as credible, then it does not matter whether it is stated once (S-CT) or multiple times (R-CT). The efficiency gains from S-RM and I-RM over Baseline might be expected, as both contain many features similar to cheap talk. For instance, in I-RM, although revisions are infrequent, the initially submitted choices can still be thought of as messages—even if it is not cheap talk.[25]

Finally, the additional efficiency gains from RM over S-RM suggest an alternative interpretation of our results. The success of RM may stem from transforming a simultaneous move game into a sequential game, without creating asymmetries by pinning down an order of play. Asynchronicity has long been recognized as an important element in coordination games (see Dutta 2012 and Ambrus and Ishii 2015 for theoretical work on the relevance of asynchronicity in coordination environments and Schotter, Weigelt, and Wilson 1994; Weber, Camerer, and Knez 2004; and Li 2007 for experimental work). Although the experimental evidence suggests a significant improvement in making the game sequential, a substantial amount of coordination failure remains. Our results highlight that, beyond asynchronicity, there are other key driving forces behind coordination (i.e., commitment and frequency of interactions). Nevertheless, a study with sequential moves—focusing on the horizon and the order of play—would shed further light on the importance of the random, repeated, and sequential nature of moves in RM in achieving higher efficiency. We leave this for future work.

REFERENCES

Agranov, Marina, and Chloe Tergiman. 2019. "Communication in Bargaining Games with Unanimity." *Experimental Economics* 22 (2): 350–68.

Ambrus, Attila, and Yuhta Ishii. 2015. "On Asynchronicity of Moves and Coordination." Unpublished.

Aumann, Robert. 1990. "Nash Equilibria are not Self-Enforcing." In *Economic Decision Making: Games, Econometrics and Optimisation*, edited by J. J. Gabszewicz, J. F. Richard, and L. A. Wolsey, 201–06. Amsterdam: North-Holland.

[25] Additionally, in S-CT, the subjects receive feedback only on the group minimum, while in I-RM, all the payoff-relevant choices are observable. Berninghaus and Ehrhart (2001) show that providing information about the distribution of group members' efforts enhances coordination.

Avoyan, Ala, and João Ramos. 2023. “Replication data for: A Road to Efficiency through Communication and Commitment.” American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E185662V1.

Berninghaus, Siegfried K., and Karl-Martin Ehrhart. 2001. “Coordination and Information: Recent Experimental Evidence.” *Economics Letters* 73 (3): 345–51.

Bigoni, Maria, Marco Casari, Andrzej Skrzypacz, and Giancarlo Spagnolo. 2015. “Time Horizon and Cooperation in Continuous Time.” *Econometrica* 83 (2): 587–616.

Blume, Andreas. 1998. “Communication, Risk, and Efficiency in Games.” *Games and Economic Behavior* 22 (2): 171–202.

Blume, Andreas, and Andreas Ortmann. 2007. “The Effects of Costless Pre-Play Communication: Experimental Evidence from Games with Pareto-Ranked Equilibria.” *Journal of Economic Theory* 132 (1): 274–90.

Bornstein, Gary, Uri Gneezy, and Rosmarie Nagel. 2002. “The Effect of Intergroup Competition on Group Coordination: An Experimental Study.” *Games and Economic Behavior* 41 (1): 1–25.

Brandts, Jordi, David J. Cooper, and Roberto A. Weber. 2015. “Legitimacy, Communication, and Leadership in the Turnaround Game.” *Management Science* 61 (11): 2627–45.

Burton, Anthony, and Martin Sefton. 2004. “Risk, Pre-Play Communication and Equilibrium.” *Games and Economic Behavior* 46 (1): 23–40.

Calcagno, Riccardo, Yuichiro Kamada, Stefano Lovo, and Takuo Sugaya. 2014. “Asynchronicity and Coordination in Common and Opposing Interest Games.” *Theoretical Economics* 9 (2): 409–34.

Charness, Gary. 2000. “Self-Serving Cheap Talk: A Test of Aumann’s Conjecture.” *Games and Economic Behavior* 33 (2): 177–94.

Chaudhuri, Ananish, Andrew Schotter, and Barry Sopher. 2009. “Talking Ourselves to Efficiency: Coordination in Inter-Generational Minimum Effort Games with Private, Almost Common and Common Knowledge of Advice.” *Economic Journal* 119 (534): 91–122.

Cooper, Russell, Douglas V. DeJong, Robert Forsythe, and Thomas W. Ross. 1992. “Communication in Coordination Games.” *Quarterly Journal of Economics* 107 (2): 739–71.

Deck, Cary, and Nikos Nikiforakis. 2012. “Perfect and Imperfect Real-Time Monitoring in a Minimum-Effort Game.” *Experimental Economics* 15 (1): 71–88.

Devetag, Giovanna. 2005. “Precedent Transfer in Coordination Games: An Experiment.” *Economics Letters* 89 (2): 227–32.

Dorsey, Robert E. 1992. “The Voluntary Contributions Mechanism with Real Time Revisions.” *Public Choice* 73 (3): 261–82.

Duffy, John, Jack Ochs, and Lise Vesterlund. 2007. “Giving Little by Little: Dynamic Voluntary Contribution Games.” *Journal of Public Economics* 91 (9): 1708–30.

Dutta, Prajit K. 2012. “Coordination Need not be a Problem.” *Games and Economic Behavior* 76 (2): 519–34.

Farrell, Joseph. 1988. “Communication, Coordination and Nash Equilibrium.” *Economics Letters* 27 (3): 209–14.

Fehr, Dietmar. 2017. “Costly Communication and Learning from Failure in Organizational Coordination.” *European Economic Review* 93: 106–22.

Fischbacher, Urs. 2007. “z-Tree: Zurich Toolbox for Ready-Made Economic Experiments.” *Experimental Economics* 10 (2): 171–78.

Friedman, Daniel, and Ryan Oprea. 2012. “A Continuous Dilemma.” *American Economic Review* 102 (1): 337–63.

He, Simin, Theo Offerman, and Jeroen van de Ven. 2019. “The Power and Limits of Sequential Communication in Coordination Games.” *Journal of Economic Theory* 181: 238–73.

Isaac, R. Mark, and James M. Walker. 1988. “Communication and Free-Riding Behavior: The Voluntary Contribution Mechanism.” *Economic Inquiry* 26 (4): 585–608.

Kamada, Yuichiro, and Michihiro Kandori. 2017. “Revision Games Part I: Theory.” Unpublished.

Kriss, Peter H., Andreas Blume, and Roberto A. Weber. 2016. “Coordination with Decentralized Costly Communication.” *Journal of Economic Behavior and Organization* 130: 225–41.

Li, Mei. 2012. “Investment Complementarities, Coordination Failure, and Systemic Bankruptcy.” *Oxford Economic Papers* 65 (4): 767–88.

Li, Tao. 2007. “Are There Timing Effects in Coordination Game Experiments?” *Economics Bulletin* 3 (13): 1–9.

Marx, Leslie M., and Steven A. Matthews. 2000. “Dynamic Voluntary Contribution to a Public Project.” *Review of Economic Studies* 67 (2): 327–58.

Matsuyama, Kiminori. 1991. “Increasing Returns, Industrialization, and Indeterminacy of Equilibrium.” *Quarterly Journal of Economics* 106 (2): 617–50.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny. 1989. "Industrialization and the Big Push." *Journal of Political Economy* 97 (5): 1003–26.
Oprea, Ryan, Gary Charness, and Daniel Friedman. 2014. "Continuous Time and Communication in a Public-Goods Experiment." *Journal of Economic Behavior and Organization* 108: 212–23.
Ostrom, Elinor, James Walker, and Roy Gardner. 1992. "Covenants with and without a Sword: Self-Governance is Possible." *American Political Science Review* 86 (2): 404–17.
Palfrey, Thomas, Howard Rosenthal, and Nilanjan Roy. 2017. "How Cheap Talk Enhances Efficiency in Threshold Public Goods Games." *Games and Economic Behavior* 101: 234–59.
Rodrik, Dani. 1996. "Coordination Failures and Government Policy: A Model with Applications to East Asia and Eastern Europe." *Journal of International Economics* 40 (1-2): 1–22.
Rosenstein-Rodan, Paul N. 1943. "Problems of Industrialisation of Eastern and South-Eastern Europe." *Economic Journal* 53 (210-211): 202–11.
Roy, Nilanjan. 2023. "Fostering Collusion Through Action Revision in Duopolies." *Journal of Economic Theory* 208: 105611.
Schelling, Thomas C. 1960. *The Strategy of Conflict.* Cambridge, MA: Harvard University Press.
Schotter, Andrew, Keith Weigelt, and Charles Wilson. 1994. "A Laboratory Investigation of Multiperson Rationality and Presentation Effects." *Games and Economic Behavior* 6 (3): 445–68.
Takahashi, Satoru. 2005. "Infinite Horizon Common Interest Games with Perfect Information." *Games and Economic Behavior* 53 (2): 231–47.
Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. 1990. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review* 80 (1): 234–48.
Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil. 1993. "Asset Markets as an Equilibrium Selection Mechanism: Coordination Failure, Game Form Auctions, and Tacit Communication." *Games and Economic Behavior* 5 (3): 485–504.
Weber, Roberto A., Colin F. Camerer, and Marc Knez. 2004. "Timing and Virtual Observability in Ultimatum Bargaining and 'Weak Link' Coordination Games." *Experimental Economics* 7 (1): 25–48.