

Should I Stay or Should I Go? Bandwagons in the Lab*

Tom-Reiel Heggedal[†] Leif Helland[‡] Knut-Eric Neset Joslin[§]

February 16, 2018

Abstract

We experimentally investigate the impact of strategic uncertainty and complementarity on leader and follower behavior using the model of Farrell and Saloner (1985). At the core of the model are endogenous timing, irreversible actions and private valuations. We find that strategic complementarity strongly determines follower behavior. Once a subject decides to abandon the status quo the probability that others jump on the bandwagon increases sharply. However, there is a reluctance to lead when leading is a conditional best response. We explain this deviation from the neo-classical equilibrium by injecting some noise in the equilibrium concept. We also find that cheap talk improves efficiency.

Keywords: strategic complementarity; type uncertainty; endogenous timing; laboratory experiment

JEL Codes: D82, L14, L15

*We are grateful for helpful comments from Urs Fischbacher, Jean-Robert Tyran, Henrik Orzen, Mark Bernard, Stefan Palan, and from participants at the BI Workshop on Experimental Economics, Oslo, May 2014, the ESA European Meeting, Prague, September 2014, the 9th NCBEE meeting, Aarhus, September 2014, the seminar of the Thurgau Institute of Economics, Kreuzlingen, April 2015, and the 2nd IMEBESS, Toulouse, April 2015. This research was financed by the Research Council of Norway, grant 212996/F10.

[†]BI Norwegian Business School and Centre for Experimental Studies and Research (CESAR)

[‡]BI Norwegian Business School and CESAR, e-mail: leif.helland@bi.no (corresponding author)

[§]Uppsala University and CESAR

1 Introduction

Many economic environments are characterized by the presence of asymmetric information and strategic complementarity. Examples include bank runs (Garratt and Keister, 2009; Goldstein and Pauzner, 2005); speculative currency attacks (Morris and Shin, 1998); setting of industry standards (Farrell and Saloner, 1985; Farrell and Klemperer, 2007); technology adoption (Katz and Shapiro, 1985, 1986); political revolts (Edmond, 2013; Egorov and Sonin, 2011); and foreign direct investment (Rodrik, 1991; Goldberg and Kolstad, 1995). In such environments, there is a potential for joint welfare improvements through coordination of actions. When players' moves are endogenous, the timing of moves may in itself serve as an important coordinating device. For players with conditional best responses, strategic uncertainty enters the picture and may impact on the ability to coordinate actions.¹

We investigate the seminal model of Farrell and Saloner (1985) (FS) in a controlled laboratory experiment.² In the model, players have incomplete information about types and endogenously time their actions in the presence of strategic complementarity. In stage one, players simultaneously decide whether to Stay with the status quo or Go to the alternative, where Go is an irreversible action.³ In stage two, players that are not committed to Go again choose between Stay or Go. If no player committed in stage one, second stage decisions are again simultaneous. All payoffs are obtained after the second stage. The key decision in the model is whether to Lead or Follow. A leader is defined as a player that chooses to Go in the first stage. A follower is defined as a player that Stays in the first stage and matches the first stage decision of her co-player in the second stage. Due to strategic complementarity, when a player leads, this may create incentives for the co-player to “jump on the bandwagon.” The strength of the incentive depends on the private valuations of the co-player with respect to the status quo and its alternative. Thus, a player may regret the decision to Lead if the co-player fails to Follow.

In FS, the combination of a specific information structure and the endogeneity of moves produces a unique equilibrium.⁴ This provides an unequivocal benchmark for our analysis and facilitates separate assessment of the role of strategic uncertainty and complementarity. Our two main treatments explore variations in strategic uncertainty with respect to leadership decisions. This treatment variation turns out to be consequential also for follower decisions, through strategic complementarities.

We present two main results. Foremost, we find that subjects often fail to Lead when the optimality of this actions depends on beliefs about their match. This effect of strategic uncertainty is unaccounted for by the model. Leading carries the risk of failure; the leader might end up alone. We find that it is the variation in the cost of failed leadership, rather than the sharp cut-off between dominant and non-dominant equilibrium strategies, that appears to cause the reluctance to Lead. We clarify this argument by introducing some noise in the decision making process. Such noise makes beliefs relevant everywhere, eroding the sharp divide between dominant and non-dominant equilibrium strategies. In particular, we show that an agent quantal response equilibrium (AQRE) organizes our data well.

Second, we find that the effect of strategic complementarity is strong. If a subject takes the Lead, all strategic uncertainty is resolved, and types who should Follow in equilibrium do so with high probability. This contrasts with recent findings in a similar environment where subjects have incomplete information about fundamentals rather than types. We comment further on this below.

In addition we investigate an extension of the model which permits cheap talk. We find that cheap

¹We follow Morris and Shin (2002) in defining strategic uncertainty as “uncertainty concerning the actions and beliefs (and beliefs about the beliefs) of others.” In neo-classical theory, while a player with a dominant best reply may be strategically uncertain, this has no bearing on her choice of action.

²For textbook treatments see Shy (2001) and Belleflamme and Peitz (2015).

³E.g. the action Go could—depending on the application—be “switch to the new technology platform”; “rise against the ruler”; or “make an investment”. The action Stay would have the prefix “do not” attached.

⁴Coordination problems are defined by the presence of multiple, Pareto-ranked equilibria. Coordination failure results if players beliefs lead them to play a payoff dominated equilibrium. Thus, in a strict sense, there are no coordination problems in the game we use.

talk improves subjects' ability to coordinate on mutually beneficial actions and increases efficiency.

To the best of our knowledge ours is the first experiment to address the FS-model. The paper closest to ours is Brindisi et al. (2014).⁵ While they use the same sequence of moves as we do, type uncertainty is replaced by uncertainty about fundamentals. Agents get a private signal about the true state of fundamentals, resembling the global games set-up. In contrast to us, they find that strategic complementarity does not strongly determine outcomes, as it should do in equilibrium. This indicates that the information structure is crucial in determining the strength of bandwagon behavior in the presence of complementarities and irreversible choices. While strategic complementarity is a strong force in environments with private information about types, it appears not to be so under private information about fundamentals.

More generally, most, if not all, economic situations of interest embody a mix of type uncertainty and uncertainty about fundamentals. Usually, it is not evident what the crucial source of uncertainty is in a particular situation. Accordingly, the choice of information structure should be determined with a view to the context.⁶ For these reasons, we believe that models such as the one analyzed in this paper have the potential to shed further light on situations in which the current practice is to rely on a global games approach.

There is an experimental literature on leadership effects in weak-link games. In contrast to our setting multiple Pareto ranked equilibria coexist in these games. Like in our setting strategic complementarities are strong in these games. Several instruments of leadership have been found to increase efficiency in weak-link games. This holds for leadership by example (Cartwright et al., 2013); leadership by communication (Brandts et al., 2015; Brandts and Cooper, 2007; Chaudhuri and Paichayontvijit, 2010); and leaders committing to help (low ability) followers (Brandts et al., 2016). There is also an experimental literature on leadership in public goods provision in which there are no strategic complementarities (see Helland et al. (2017) for a review).

The remainder of the paper is organized as follows. In the next section, we describe the model. For concreteness, we present the model using the parameters of the experiment. Thereafter, in the third section, we review our design and the experimental procedures. In section four, we present the experimental results. The fifth section considers how noisy behavior impacts the equilibrium. The final section concludes.

2 Model

There are two players, $i \in \{1, 2\}$, and two stages, $t \in \{1, 2\}$.⁷

Prior to the first stage, nature draws a payoff relevant type θ_i for each player. θ_i is private information observed by player i only. Type draws are *i.i.d* from a uniform distribution on $[0, 10]$ such that $\theta_i \sim U[0, 10]$ for all i . We discuss the role of θ_i for payoffs in detail below.

In each stage, each player has a possible action Stay or Go. These actions are indicated by S_i^t and G_i^t , respectively. If a player chooses Stay in the first stage (S_i^1), then the player again chooses between Stay and Go in the second stage (S_i^2 or G_i^2). However, if a player chooses Go in the first stage (G_i^1), this commits the player to Go in the second stage as well (G_i^2). The decision to Go in the first stage is thus irreversible. We refer to the choice of G_i^1 as a decision to Lead.

At the end of each stage, players observe all previous actions. Second stage decisions can therefore be conditioned on first stage actions. We refer to the choice of G_i^2 as a decision to Follow whenever the first

⁵Brindisi et al. (2009) provides a thorough exposition of the theory.

⁶This is also the view taken in the seminal work on global games (see the discussion in Carlsson and Van Damme (1993) pp.251-2).

⁷The model can be generalized to the case with n players and n stages. The model can also accommodate more general payoff functions than the ones we use.

stage actions are S_i^1 and G_{-i}^1 . Hence, a decision to Follow is a decision to Go conditional on the other player choosing to Lead.

The payoff to player i depends on which of four possible outcomes is realized at the end of the second stage. The payoffs are presented in Table 1. All payoffs are either constants or linear functions of θ_i .

	S_2^2	G_2^2
S_1^2	7, 7	5, $\alpha\theta_2$
G_1^2	$\alpha\theta_1$, 5	$\theta_1 + 2$, $\theta_2 + 2$

Table 1: Payoff Matrix

In particular, θ_i determines the payoffs from action Go. Higher realizations of θ_i translate into higher payoffs from both a joint choice of Go and a unilateral choice of Go. The payoff from a unilateral choice of Go is also affected by the parameter α . We vary α between our two main treatments: $\alpha = 1$ in treatment D and $\alpha = 1/2$ in treatment N . Because of the difference in α , Lead is a dominant strategy for high types in the D treatment but not in the N treatment.⁸ In the N treatment, a decision to Go is always conditional on the belief that joint play of Go is sufficiently likely. As we discuss in the next section, this enables us to compare groups of players for whom the decision to Lead is optimal but for whom the role of beliefs is different. Observe in addition that the game exhibits strategic complementarity: The difference in payoffs between Go and Stay increases if the other player also chooses to Go.⁹

An optimal strategy in this setting depends on both a player's type and the strength of the complementarity effect.¹⁰ In particular, a player who favors joint play of Go faces a key decision in the first stage. On the one hand, a decision to Lead means that joint play of Go is more likely. The decision to Lead resolves strategic uncertainty (because the player is committed to Go) and increases the payoff from a joint choice of Go (due to complementarity in actions). This makes it more attractive for the other player to also choose Go. On the other hand, when a player chooses to Lead, it forgoes an opportunity to observe the first stage actions. When complementarity is relatively important for payoffs, a player may therefore want to Stay rather than Lead.

Based on this intuition, it is reasonable to expect that players with sufficiently high types Lead while players with intermediate preferences delay their decision to the second stage. Below, we show that there is a "bandwagon" equilibrium of this type. In the bandwagon equilibrium, players use monotone threshold strategies in which types above a threshold θ^* Lead, players with types above a threshold $\underline{\theta}$ Follow, and players with types below $\underline{\theta}$ Stay in both stages.¹¹ This divides players into three strategic ranges according

⁸This explains our treatment names. We use D to denote the treatment in which some players have a dominant strategy to Lead. We use N to denote the treatment in which no types have a dominant strategy to Lead.

⁹This first difference is the discrete analog of an increasing cross partial derivative of payoffs. See the definition of strategic complementarity given by Bulow et al. (1985) (for a discussion, see chapter 2 in Cooper (1999), the section on supermodular games). If the other player chooses Stay, then the difference in payoffs between Go and Stay equals $\alpha\theta_i - 7$. However, if the other player chooses Go, then the difference in payoffs is $\theta_i - 3$. Because $\theta_i - 3 > \alpha\theta_i - 7$, the payoffs exhibit increasing first differences (recall $\alpha \leq 1$). For a recent study that also examines the role of strategy complementarities in a discrete game with private information see Brindisi et al. (2014).

¹⁰Formally, a strategy in this game is a mapping from types and history of actions to a first and second stage action. In practice, however, we use "strategy" to refer to optimal play for a particular type (or range of types).

¹¹Our linear specification of payoffs satisfies the four assumptions given by Farrell and Saloner (1985) that guarantee the existence of a bandwagon equilibrium. The only exception is our N treatment, which violates part of assumption 3: There are no players in the N treatment for whom a unilateral decision to Go is dominant. However, this part of the assumption 3 is superfluous. The model has "interesting" equilibria as long as the bandwagon thresholds are such that $0 < \underline{\theta} < \theta^* < 10$. This is the case in our N treatment. In contrast, if *both* parts of assumption 3 were violated, then the interesting aspects

to their type: (1) a Lead range $(\theta^*, 10]$ in which players Go in the first stage (G_i^1, G_i^2) , (2) a Follow range $(\underline{\theta}, \theta^*]$ in which players first Stay and then Go if the other player Leads $(S_i^1, (G_i^2|G_{-i}^1, S_i^2|S_{-i}^1))$, and (3) a Stay range in which players choose to Stay in both stages (S_i^1, S_i^2) .¹²

To demonstrate formally that there exists a unique symmetric (Bayesian perfect) bandwagon equilibrium, we compare the payoffs associated with the three strategies, Lead, Follow, and Stay.¹³ Throughout, we denote payoffs by a function π with arguments indicating the second stage actions of both players.¹⁴ Since payoffs depend on type, π is conditional on θ_i . For example, $\pi(G_i, S_{-i}; \theta_i)$ denotes the payoff from a unilateral choice of Go by a player with type θ_i (the bottom left cell in Table 1).

We begin by checking which players have dominant strategies. The decision to Lead is dominant if a player prefers a unilateral choice of Go compared to a joint choice of Stay. Based on Table 1, it is straightforward to identify the threshold $\bar{\theta}$ above which players strictly prefer to Lead. Specifically, $\bar{\theta}$ follows from the indifference condition

$$\begin{aligned}\pi(G_i, S_{-i}; \bar{\theta}) &= \pi(S_i, S_{-i}; \bar{\theta}) \\ \alpha \bar{\theta} &= 7.\end{aligned}$$

In the D treatment ($\alpha = 1$), Lead is a dominant strategy for players with types $\theta > \bar{\theta} = 7$. No comparable dominance region exists in the N treatment ($\alpha = \frac{1}{2}$) because the dominance threshold exceeds the highest possible type.

We use a similar argument to establish the region in which Stay is dominant. Let $\underline{\theta}$ denote the type of a player who is indifferent between a unilateral choice of Stay and a joint choice of Go, such that all $\theta_i < \underline{\theta}$ have a dominant strategy to Stay:

$$\begin{aligned}\pi(S_i, G_{-i}; \underline{\theta}) &= \pi(G_i, G_{-i}; \underline{\theta}) \\ 5 &= \underline{\theta} + 2\end{aligned}$$

In both treatments, $\underline{\theta} = 3$ and the strategy (S_i^1, S_i^2) is dominant if a player has a type $\theta_i \leq \underline{\theta} = 3$. This also implies that any players with types greater than $\underline{\theta}$ should Follow in the second stage whenever the other player chooses to Lead.

Next, we consider players who do not have dominant strategies and for whom strategic uncertainty is the key challenge. These are players with types above $\underline{\theta}$ but below $\bar{\theta}$. In this region, the complementarity effect is critical. The payoffs associated with coordinated actions (G_i^2, G_{-i}^2) and (S_i^2, S_{-i}^2) exceed those associated with unilateral choices (S_i^2, G_{-i}^2) or (G_i^2, S_{-i}^2) . In addition, some types in this range prefer a joint choice of Stay while other types prefer a joint choice of Go. Specifically, types below $\theta^\circ = 5$ prefer Stay while types above $\theta^\circ = 5$ prefer Go.¹⁵

Two strategies are relevant for types who lack dominant strategies. The first is to Follow. The other is to Lead. Follow is attractive because it guarantees coordination whenever players are using the bandwagon strategy: If the other player chooses to Lead, then the outcome is (G_i^2, G_{-i}^2) and if the other

of the model would disappear. If no player has a dominant strategy to Stay, then any player who favors a joint choice of Go would choose to Lead while any other player has a best response to Follow. As a result, the outcome would be either a joint choice of Go (whenever at least one player favors Go) or a joint choice of Stay (when neither player favors Go).

¹²We use the notation of the form $(G_i^2|G_{-i}^1, S_i^2|S_{-i}^1)$ to indicate that the second stage action is conditioned on the realization of the first stage action of the other player.

¹³These are the only strategies that need to be considered. Lead is the only strategy in which Go is a first stage action. There are four possible strategies in which Stay is the first stage action. These include the strategy Stay, $(S_i^1, (S_i^2|G_{-i}^1, S_i^2|S_{-i}^1))$, and the strategy Follow, $(S_i^1, (G_i^2|G_{-i}^1, S_i^2|S_{-i}^1))$. A third possible strategy is Go in the second stage regardless of the first stage actions of the other player $(S_i^1, (G_i^2|G_{-i}^1, G_i^2|S_{-i}^1))$, and a fourth possible strategy is to “anti-Follow,” $(S_i^1, (G_i^2|S_{-i}^1, S_i^2|G_{-i}^1))$. However, these last two strategies both (weakly) dominated whenever the other players use the bandwagon strategy. Hence, only the strategies Lead, Stay, and Follow need to be compared.

¹⁴Since payoffs depend only on the second stage outcome only, we drop the t superscripts.

¹⁵ θ° follows from the a comparison of $\pi(G_i, G_{-i}; \theta^\circ) = \theta^\circ + 2$ and $\pi(S_i, S_{-i}; \theta^\circ) = 7$.

player chooses to Stay, then the outcome is S_i^2, S_{-i}^2 . The drawback of strategy Follow is, however, that players who prefer the outcome G_i^2, G_{-i}^2 (i.e., players with types greater than θ°) may end up coordinated on S_i^2, S_{-i}^2 . Players who prefer Go strongly enough will therefore prefer to Lead even though they do not have a dominant strategy and will regret the decision to Lead if the second stage outcome is G_i^2, S_{-i}^2 . For such players, Lead maximizes expected payoffs because it sufficiently increases the likelihood of the preferred outcome G_i^2, G_{-i}^2 .

To identify the threshold θ^* above which Go is optimal—though not necessarily dominant—we compare the expected payoffs from Follow and Lead:

$$\underbrace{\mathbb{P}(\theta_{-i} > \underline{\theta})\pi(G_i, G_{-i}; \theta^*) + (1 - \mathbb{P}(\theta_{-i} > \underline{\theta}))\pi(G_i, S_{-i}; \theta^*)}_{\text{Expected Payoff Lead}} = \underbrace{\mathbb{P}(\theta_{-i} > \theta^*)\pi(G_i, G_{-i}; \theta^*) + (1 - \mathbb{P}(\theta_{-i} > \theta^*))\pi(S_i, S_{-i}; \theta^*)}_{\text{Expected Payoff Follow}}. \quad (1)$$

The left hand side of equation 1 is the expected payoff from Lead. This is comprised of two terms. The first is the probability of meeting a player who is either in the Lead or Follow ranges and the final outcome is joint play of Go. The second term on the left-hand side is the probability of meeting a player in the Stay range; in this case, the final outcome is a unilateral choice of Go. Analogously, the right hand side of equation 1 gives the expected payoff from the strategy Follow. This is also composed of two terms. The first is the probability of meeting a player who Leads and the second stage outcome is G_i, G_{-i} . The second term on the right-hand side is the probability of meeting a player in the Follow or Stay ranges. In this case, the final outcome is that both players Stay.

Plugging in for the payoff functions and the probabilities (which follow from the assumption that $\theta_i \sim U[0, 10]$ for all i) yields

$$\frac{(10 - \underline{\theta})}{10}(\theta^* + 2) + \frac{\underline{\theta}}{10}\alpha\theta^* = \frac{(10 - \theta^*)}{10}(\theta^* + 2) + \frac{\theta^*}{10}7.$$

In the case of the D treatment ($\alpha = 1$), this reduces to

$$\theta^{*2} - 5\theta^* - 2\underline{\theta} = 0.$$

Given that $\underline{\theta} = 3$, the only positive root of this equation is $\theta^* = 6$. For the N treatment, the identical computation yields $\theta^* = 7.3$. This demonstrates that the bandwagon strategy with thresholds $\underline{\theta}$ and θ^* is a best response to itself. In addition, observe that the left-hand side of equation 1 is monotonically increasing in θ^* while the right hand side is monotonically decreasing. This means that the bandwagon strategy is not just a best response to itself, but that it is unique in the class of symmetric monotone threshold strategies.

Finally, to establish that any symmetric equilibrium has the bandwagon form, notice that regardless of a player's beliefs, the benefits of leading are non-decreasing in the player's type θ . If it is optimal for a player of type θ' to Go in the first stage, then it is optimal for types $\theta > \theta'$ to also Go in the first stage. Any symmetric equilibria must therefore have the threshold form. This establishes that the bandwagon strategy is the unique symmetric equilibrium because we already showed that the bandwagon strategy is unique among symmetric threshold strategies.

To summarize the results from this section, there are two strategically relevant bandwagon thresholds $\underline{\theta}$ and θ^* . The lower bandwagon threshold $\underline{\theta}$ separates players who should Stay from those who should Follow. The upper bandwagon threshold θ^* separates players who should Follow from those who should Lead. Because it is important in the remainder, we emphasize that $\theta^* < \bar{\theta}$. This means that there exist players for whom Leading is optimal but who will regret the decision to Lead if they alone choose Go. For players with types in the range $(\theta^*, \bar{\theta}]$, the decision to Lead thus relies on beliefs. Moreover, because

$\bar{\theta}$ is not part of the bandwagon strategy, it is not strategically relevant. As a consequence, the model does not anticipate behavior to be different above and below this threshold within the Lead range.

Signaling We also investigate a version of the game with communication. The game with communication is identical to the game presented above but with the addition of a cheap talk stage just after agents have observed their type θ_i but prior to the first stage. In the cheap talk stage, players must send either the message Stay or the message Go. This allows players to announce their preference for one of the outcomes. Importantly, the message is non-binding and this is common knowledge.

Our focus is on the truth-telling equilibrium of the signaling game (a formal analysis of the signaling equilibrium is provided in the supplementary materials section S.3). However, in addition to the truth telling equilibrium, there can exist equilibria in which the messages are uninformative. In such babbling equilibria, the thresholds are the same as in the game without cheap talk described above. Our reason for focusing on the truth telling equilibrium is that it payoff dominates babbling equilibria.

In the truth-telling equilibrium, players send a message that corresponds to their favored platform: Players with types $\theta > \theta^\circ$ send the message Go while the remaining types send the message Stay. If players send the same message, they coordinate in the first stage. This eliminates Pareto inefficiency. If the players send conflicting messages, however, then the game resembles the game without communication except that players can partially update their beliefs about the type of the other player. Thus, conflicting communication does not improve coordination. Specifically, a player who sends the message Stay must have a type in the range $[0, \theta^\circ]$. This has consequences for the optimal θ^* . Because the probability that the other player will Follow, conditional on giving signal Stay, is lower than the unconditional probability in the absence of communication, the threshold θ^* is higher in the game with communication.¹⁶

3 Design and procedures

Design We conduct three treatments, D , N , and S . Table 2 facilitates a comparison of the key features and predictions associated with each treatment.

As is evident from table 2a, the only payoff difference between the treatments is that $\pi(G_i, S_{-i}; \theta_i)$ is reduced by half in the N treatment. Otherwise, the main difference between the treatments is that pre-play communication is allowed in the S (Signal) treatment but is absent from the D and N treatments.

The model predicts treatment differences in the first stage due to differences in the threshold θ^* . The values of θ^* are summarized in table 2b. For instance, $\theta^* = 6$ in the D treatment but $\theta^* = 7.3$ in the N treatment. The Lead range should therefore be largest in the D treatment and smallest in the N treatment. No comparable differences are expected in the second stage because $\underline{\theta} = 3$ in all treatments: Players with types above $\underline{\theta}$ should Follow and players with types below $\underline{\theta}$ should Stay.

Our design allows for a clean test of two main relationship. First, what is the impact of strategic uncertainty on first stage behavior? Second, to what extent do complementarities in actions shape second stage behavior? In addition our design permits an assessment of the effect of communication on efficiency.

Strategic uncertainty: Despite the fact that the model makes identical predictions for players in the Lead range, the relevance of beliefs is distinctly different in the D and N treatments. These differences are summarized in the fourth and fifth columns of Table 1b. In the D treatment, the decision to Lead is dominant for players with types greater than $\bar{\theta} = 7.0$. Such players face no strategic uncertainty. In contrast, the decision to Lead is always predicated on beliefs in the N treatment. Hence, strategic uncertainty enters the picture. Comparison of the first stage behavior of subjects in the D and N treatments thus facilitates a test of the behavioral impact of beliefs.

¹⁶The computations are analogous to those presented for the model without communication, but take into account the partial updating that results from observing the message of the match.

<i>Payoffs</i>				
	$\pi(S_i, S_{-i}; \theta_i)$	$\pi(S_i, G_{-i}; \theta_i)$	$\pi(G_i, S_{-i}; \theta_i)$	$\pi(G_i, G_{-i}; \theta_i)$
Dominant (<i>D</i>)	7	5	θ_i	$\theta_i + 2$
Non-dominant (<i>N</i>)	7	5	$\frac{1}{2}\theta_i$	$\theta_i + 2$
Signal (<i>S</i>)	7	5	θ_i	$\theta_i + 2$

(a)

<i>Predictions</i>					
	Thresholds			Best Response = G_i^1	
	$\underline{\theta}$	θ°	θ^*	Conditional	Dominant
Dominant (<i>D</i>)	3.0	5.0	6.0	$\theta_i \in [6.0, 7.0]$	$\theta_i \in [7.0, 10.0]$
Non-dominant (<i>N</i>)	3.0	5.0	7.3	$\theta_i \in [7.3, 10.0]$	$\theta_i \in \emptyset$
Signal (<i>S</i>)	3.0	5.0	6.2	$\theta_i \in [6.2, 7.0]$	$\theta_i \in [7.0, 10.0]$

(b)

Table 2: Payoffs (2a) and predictions (2b) for treatments *D*, *N*, and *S*.

Complementarities in action: Due to irreversibility, when a player's match chooses to Lead, this resolves all strategic uncertainty in the second stage of the game. What remains is the pure effect of strategic complementarities. In the second stage, we therefore investigate the behavior of subjects conditioned on the their match Leading. Hence, a direct measure of the strength of complementarities is the frequency of Follow decisions by subjects in the Follow range. This complementarity effect should not vary over treatments.

Communication: In the *S* treatment, players send a cost-free signal simultaneously, prior to taking their first stage action. According to theory, access to a cost-free signal should eliminate Pareto inefficiency. We implement this treatment with the same parameters as the *D* treatment. This allows a direct assessment of differences in efficiency, including Pareto inefficiency, due to communication. We compute overall measures of efficiency based on gross payoffs.

Experimental procedures All sessions were conducted in the research lab of BI Norwegian Business school using participants recruited from the general student population at the BI Norwegian Business School and the University of Oslo, both located in Oslo, Norway. Recruitment and session management were handled via the ORSEE system Greiner (2015). In each of our three treatments, we ran five sessions per treatment with between 16 and 20 subjects per session. No subject participated in more than one session. z-Tree was used to program and conduct the experiment (Fischbacher, 2007). Anonymity of subjects was preserved throughout.

On arrival to the lab, subjects were randomly allocated to cubicles in order to break up social ties. After being seated, instructions were distributed and read aloud in order to achieve public knowledge of the rules. All instructions were phrased in neutral language. Rather than choose between Stay and Go, subjects were asked to choose either shape *Circle* (that is, Stay) or shape *Square* (that is, Go). Sample instructions and screen shots are provided in the supplementary materials.

Each session of the experiment began with two non-paying test games for subjects to get acquainted with the software. This was immediately followed by $n - 1$ games in which the subjects earned payoffs,

where n is the total number of participants in the session.¹⁷ In each game, subjects were matched with one other subject according to a highway protocol. Every subject thus met every other subject once and only once.¹⁸ In total, our data consists of 2673 unique games (excluding test games). Each game consisted of a single repetition of the two-player, two-stage game with the rules and payoff functions outlined above. Subjects earned experimental currency units (ECUs). After the final game, accumulated earnings in ECU were converted to NOK using a fixed and publicly announced exchange rate. Subjects were paid in cash privately as they left the lab. On average subjects earned 250 Norwegian Kroner (about 36 USD at the time). A session lasted on average 50 minutes.

Gameplay was formulated in the following fashion: At the beginning of each each new game, each subject received a private number drawn from a uniform distribution on the interval $(0, 10)$ with two decimal points of precision. This number corresponded to the subject's type θ_i . A dedicated screen was used to display this information. Thereafter, subjects observed a 2×2 matrix of their own payoffs and a button to choose a first-stage action. The first stage concluded when both subjects in the match had made their decisions. If both subjects decided to Go in the first stage, they bypassed the second stage and continued directly to the feedback.

The second stage began with a screen that revealed the first stage actions of both subjects in the pair. Next, subjects who decided to Stay in the first stage again chose between Stay and Go. If a subject's match decided to Go in the first stage, then the subject observed a truncated 2×1 matrix in which the payoffs conditioned on the match choosing Stay were removed. This reflected the fact that the subject's match had committed to Go. Otherwise, if both subjects decided to Stay, then subjects observed the same 2×2 matrix as in the first stage.

After all second stage decisions were resolved, the subjects moved to a feedback screen. The feedback consisted of a history of decisions and profits for each of the games played. It also displayed total accumulated profits.

The signal treatment S included an additional stage between the type draw stage and the first stage action choice. In this stage, subjects simultaneously selected either the message "I choose circle" or the message "I choose square". The chosen message was revealed to their match on a dedicated screen. Apart from this additional stage, the screens and information are identical to those used in the two other treatments.

4 Results

First stage behavior The first stage behavior of the subjects is consistent with the use of bandwagon strategies and the essential predictions of the model. Table 3 presents the proportion of test subjects in each of the three strategic ranges that chose to Go in the first stage. As anticipated by the model, test subjects in the Stay and Follow ranges chose to Go with a low probability while test subjects in the Lead range chose to Go with a high probability.¹⁹ Moreover, the probability of Go increases rapidly in the vicinity of θ^* —as one would expect if subjects use threshold strategies. This pattern is clear in Figure S1 in the supplemental materials. This figure presents the information in Table 3 on a finer grid. Note that throughout we use the prefix "S" to denote material found in the supplemental materials.

To formally assess the predictions of the model, we compare behavior across treatments using Wilcoxon Rank Sum (WRS) tests. Using session level data, between treatment comparisons find no significant differences in behavior between the D and N treatments in either the Stay range or the Follow range

¹⁷Hence, in a treatment with 20 participants, each participants played 19 repetitions of the game with payoffs.

¹⁸This protocol eliminates certain dynamic problems, such as strategic teaching and reciprocity (see Fréchet (2012) for a discussion).

¹⁹Note that the observed differences between the Stay and Follow range arise not because of a general difference throughout the ranges but due to a relatively high rate of Go among subjects in the Follow range just below θ^* .

	D	N
Stay	0.04	0.04
Follow	0.27	0.21
Lead	0.89	0.71

Table 3: First Stage, Proportion choosing Go

($p = 0.92$ and $p = 0.35$, respectively).²⁰ Likewise, in the part of the Lead range in which decisions are conditioned on beliefs, i.e. in the region $(\theta^*, \bar{\theta}]$, we are unable to identify a significant treatment difference ($p = 0.17$).²¹ When beliefs are relevant, we find no behavioral differences over treatments.

In contrast, when we compare behavior over the entire Lead range, $(\theta^*, 10]$, we reject equality of the treatments. In this region, the probability that a subject chooses to Lead is a full 18 percentage points higher in the D treatment than in the N treatment. This difference is strongly significant ($p = 0.01$).²² Moreover, given a 5 percent significance level and the observed variances of the D and N treatments, this test has a power of 99 percent.²³

We attribute the difference in behavior to the fact that Go is a dominant action in the D treatment but is predicated on beliefs in the N treatment.²⁴ When beliefs are relevant for actions, subjects tend to be more tentative and adopt a “wait-and-see” approach.²⁵

Result 1 (Leader behavior) *Subjects are relatively more reluctant to Lead when leading is a conditional best response.*

Relative to the model, test subjects with types above θ^* do not Lead often enough. In doing so, these subjects forgo an opportunity to induce their favored outcome whenever their match is in the Follow range.²⁶ This is costly. In the D treatment, subjects with $\theta_i > \theta^*$ who Stay on average earn 3.6 ECU less than what they would earned if they had chosen to Lead. The comparable number for participants in the N treatment is 1.6 ECU.²⁷ Also, a high fraction of subjects in the Follow range choose the out of equilibrium action Go in the first stage of the game. Below, we use the AQRE to rationalize observed deviations from the Nash equilibrium of the model.

Second Stage Behavior The results from the second stage are characterized by bandwagon behavior. The second stage results are summarized in table 4.²⁸ When a subject in the Follow range has a match that Stays in the first stage, the subject also Stays with high probability: In 92 percent of cases in the D treatment and in 93 percent of cases in the N treatment (that is, they Go in 8 percent and 7 percent

²⁰See tables S1 and S2. Inspecting Figure S2 it may appear that there are treatment differences in the error rate also to the left of θ^* . But these differences are not statistically significant at conventional levels, and we refrain from further discussion of them.

²¹See tables S3 and S4. We do not identify a difference regardless of whether we compare the range $\theta_i^D \in (6, 7]$ in the D treatment with the entire Lead range in the N treatment $\theta_i^N \in (7.3, 10]$ or if we base the test on a balanced set of data and use the restricted ten base point region just beyond the Go threshold, $\theta_i^N \in (7.3, 8.3]$.

²²See table S5.

²³The power computation was carried out using the simulation routine of Bellemare et al. (2016).

²⁴When we limit the comparison to individuals with high types—in the D treatment, only those subjects with a dominant strategy—the results remain unchanged relative to the comparison over the entire range ($p = 0.01$). See table S6.

²⁵Duffy and Ochs (2012) observe a similar “wait-and-see” dynamic in a study of binary entry games.

²⁶Because types are distributed uniformly, subjects are expected to be in the Follow range 30 percent of the time in the D treatment and 43 percent of the time in the N treatment. The actual rates realized in the treatments were 35 percent and 43 percent.

²⁷These are subjects in the region $\theta \in [\theta^* = 6, \bar{\theta} = 7]$ in the D treatment and $\theta \in [\theta^* = 7.3, 10]$ in the N treatment.

²⁸Figure S3 provides a complementary illustration of second stage behavior.

	Match Go		Match Stay	
	D	N	D	N
Stay	0.04	0.06	0.04	0.02
Follow	0.90	0.87	0.08	0.07
Lead	0.90	0.95	0.64	0.36

Table 4: Second Stage, Proportion choosing Go

of cases, respectively). Furthermore, when a subject in the Follow range has a match that Leads, the subject also chooses to Go with a high probability: In 90 percent of cases in the D treatment and in 87 percent of cases in the N treatment. Subjects in the Follow range thus mirror the first stage action of their match. As anticipated, WRS tests do not identify differences in the behavior across treatments.²⁹ We conclude that there is strong evidence of strategic complementarity and that second stage behavior is consistent with the use of bandwagon strategies.³⁰

Result 2 (Follower behavior) *In the presence of a Leader complementarity strongly determines second stage behavior of subjects in the Follow range.*

One feature of Table 4 that is not accounted for is the observation of individuals in the Lead range who chose to Stay in the first stage. This is a mistake relative to the bandwagon equilibrium. Nevertheless, the optimal second stage behavior is straightforward to characterize: When a subject’s match chooses to Lead, all subjects with types above $\underline{\theta} = 3$ —including those above θ^* —should Follow. This is due to strategic complementarity. As expected, we observe high rates of Go among this group (Table 4, bottom row, the first two columns, 0.90 and 0.95). Otherwise, if a subject’s match chooses to Stay in the first stage, only individuals in the D treatment for whom Go is a dominant strategy should correct their mistake by choosing to Go in the second stage. We also see evidence of this in the data. The rate of Go among subjects in the Lead range whose match chose to Stay is substantially higher in the D treatment than the N treatment (Table 4, bottom row, the last two columns, 0.64 and 0.36).³¹ In addition, it is worth noting that the relatively high level of errors among this group—individuals who also fail to Go in the second stage despite the other player Leading—does not reflect an unconditional level of errors but the rate of errors among only the group that has already made a first stage error. The unconditional probability of this type of “double error” is about five percent in the D treatment and ten percent in the N treatment.

Signal Recall that players in the S treatment with types $\theta_i > \theta^\circ = 5$ should signal Go while players with types $\theta_i < \theta^\circ = 5$ should signal Stay. This means that there are four possible outcomes from the communication stage: Two outcomes in which the subjects give the same signal, either Go or Stay, and the two outcomes in which the subjects give opposite signals.

²⁹Details are provided in tables S8, S9, S10, and S11. The only comparison that approaches significance is for players in the Follow range for whom the match chooses to Go ($p = 0.07$). However, the average difference in the rate of Go is only 0.03 (0.90 compared with 0.87, Table S11). For the other three tests, the p -values are all in excess of 0.24. Note, however, that for the observed differences, our tests are not sufficiently powerful. To achieve sufficient power (90% or more) in our tests observations would have to be increased by over an order of magnitude. This was not feasible. The lack of power ultimately reflects a lack of difference between the treatments relative to variability within the treatments. This means that even if there is an unidentified difference between the treatments, the economic magnitude of the difference is relatively small.

³⁰These effects are confirmed in the supplementary materials, using a logistic regression. Specifically, see figure S4).

³¹See the left panel figure S3 in which individuals with dominant strategies correct their mistake in the second stage even if their match chose to Stay.

Table 5 shows the proportion of test subjects that chose to Go for each combination of messages. Included in parentheses is the number of cases.

	Match Message Go		Match Message Stay	
	Own Message Go	Own Message Stay	Own Message Go	Own Message Stay
$\theta_i \leq 5$	0.44 (18)	0.19 (367)	0.31 (13)	0.02 (440)
$\theta_i > 5$	0.96 (384)	0.47 (62)	0.74 (416)	0.28 (60)

Table 5: First Stage, Proportion choosing Go

The first observation is that subjects almost always send the correct message: Of the participants with types $\theta_i < \theta^\circ$, 807 send the message Stay while only 31 send the message Go. Of the participants with types $\theta_i > \theta^\circ$, 800 send the message Go while 112 send the message Stay.

The next observation is that test subjects almost always coordinate actions in the first stage if they send the same message (columns 1 and 4). This is most evident in the bottom left and the upper right cells. This is consistent with theory. The opportunity for pre-play communication enables participants to update their beliefs about their match's type (section S.3 shows the computation of the bandwagon thresholds in this case). If both subjects send the same message, then both subjects should choose that action in the first stage.

Our results relate to previous findings on the effect of cheap talk. In particular, two way communication has been found to increase coordination when preferences are aligned, but not so when preferences conflict (Crawford (1998); Blume and Ortman (2007)).

Result 3 (Communication) *Cheap talk promotes subjects ability to coordinate actions on mutually beneficial outcomes.*

Our last observation is that there is a similar pattern of first stage errors in the S treatment as in the other treatments. When subjects send conflicting signals, subjects with types above 6.2 should Go in the first stage. We find that in 85 percent of such cases, subjects with $\theta_i > 6.2$ do in fact Lead. However, as in the other treatments, we also find over eagerness to Go among individuals who are below this threshold in the Follow range. In 35 percent of such cases, subjects choose to Lead when they should have chosen to Stay.

Efficiency The different treatments affect the incentives and ability of subjects to achieve efficient outcomes. Table 6 presents the theoretical and empirical efficiency of each treatment, computed as the

	S	D	N
Model	98.0	97.8	93.7
Realized	95.2	94.1	88.4

Table 6: Efficiency (as % of maximum payoff)

fraction of the maximum total earnings.³² The theoretical efficiency associated with equilibrium play of

³²The *maximum* total earnings is computed as the payoff that would be realized if a social planner chose the subjects' actions to maximize total payoff. The *theoretical* payoff is computed as the payoff that would be realized if all subjects played the equilibrium bandwagon strategy. The *empirical* payoff is computed based on the actual payoffs of test subjects. Note that the maximum total surplus is *nearly* identical across treatments.

the model is marginally higher in the S treatment than the D treatment, and somewhat lower in the N treatment. In terms of realized efficiency, however, the differences are more pronounced. 95.3 percent of the maximum payoffs are realized in the S treatment, compared to 94.1 percent in the D treatment and 88.4 percent in the N treatment.³³ Furthermore these differences are statistically significant at conventional levels using one-sided Wilcoxon Rank Sum tests (between S and D $p = 0.087$, between S and N $p = 0.005$, and between D and N $p = 0.005$).³⁴ Efficiency thus seems to improve with access to cheap talk and with the redundancy of beliefs. These findings should, however, be treated with some caution given that we have not power tested the relationships.

Result 4 (Efficiency) *Cheap talk and unconditional best responses to Lead improve efficiency.*

5 Agent Quantal Response Equilibrium

Although the predictions of the model tend to be supported by the data, behavior is not uniformly consistent with the equilibrium of the model. In particular, we observe an asymmetric pattern of errors in the vicinity of θ^* for both the D and N treatment, although this pattern is most pronounced in the D treatment.³⁵ Although it is natural for subjects to make mistakes in the computation of θ^* , we would expect a symmetric pattern of error if mistakes were idiosyncratic. Asymmetry suggests instead a systematic deviation from the equilibrium. The overall level of errors is also higher in the N treatment than the D treatment.

A key observation is that the frequency of errors is inversely related to their costs. Subjects with types close to θ^* , who are nearly indifferent between Lead and Stay, often make mistakes while subjects with extreme types, who strongly prefer either abandonment or preservation of the status quo, rarely do. This is consistent with the core intuition for a quantal response equilibrium. We therefore estimate the AQRE of the model (McKelvey and Palfrey, 1998). This framework enables us to assess whether the observed pattern of behavior is consistent with an equilibrium in which decisions are noisy. Furthermore, the AQRE perspective emphasizes that beliefs are consequential everywhere. Since the model we study has a unique equilibrium, this allows us to gauge the impact of beliefs on behavior in a smooth way.

Employing the notation in Turocy (2010), let a, a' denote actions and $I(a)$ denote the information set that includes action a . In a game of perfect recall, like the bandwagon game, any node appears at most once along any path of play. Let ρ denote a behavior strategy profile. Such a profile denotes, for each action a , the probability ρ_a that action a is played if information set $I(a)$ is reached. Finally, let $\pi_a(\rho)$ denote the expected payoff to the player of taking action a on reaching information set $I(a)$, contingent on the behavior profile ρ being played at all other information sets. We say that the strategy profile is a logit AQRE if, for all players, for some $\lambda \geq 0$, and for all actions a and every information set:

$$\rho_a = \frac{e^{\lambda \pi_a(\rho)}}{\sum_{a' \in I(a)} e^{\lambda \pi_{a'}(\rho)}}$$

In an AQRE $\rho_a > 0$ for all actions a . Thus, beliefs are relevant everywhere. Equilibrium requires that beliefs are correct at each information set. The set of logit AQRE maps $\lambda \in [0, \infty]$ into the set of

³³Note that the percentages are affected by how we have defined our efficiency measure. In particular, realize that all pairs will earn a certain minimum amount in equilibrium. Because of this, our efficiency measure does not range from 0 to 100. To see why, notice that the efficiency measure could be made arbitrarily close to 1 by simply adding a sufficiently large number to each of the payoffs. Although we considered other efficiency measures, we rejected them in favor of the simplest computation. Regardless, the important question is whether there is a statistical difference between the treatments.

³⁴See supplementary materials section S.4 for WSR statistics for pairwise comparisons of the treatments and the associated tabulations of realized efficiency by session.

³⁵We document this asymmetry further in figure S2 which plots the distribution of errors along with a smoothed trend.

totally mixed behavior profiles. Letting $\lambda \rightarrow \infty$ identifies a subset of the set of sequential equilibria as limiting points (McKelvey and Palfrey, 1998). Thus, when noise vanishes one is back in the neo-classical equilibrium theory. On the other hand, and for a given game, moderate noise can get amplified in an AQRE, resulting in substantial deviations from neo-classical equilibrium theory.

We estimate the logit AQRE on 20 equally sized bins (i.e. the empirically observed Go frequency in that range) for the three decision nodes: The first stage action and two second stage actions that depend on whether the match chose Stay or Go in the first stage. Our estimation performs a fixed point iteration in which we loop through the QREs for each stage, taking behavior in the other stages as given. We fit λ by minimizing the distance between the binned empirical data and the estimates. Figure 1 presents the best fit for each treatment individually.³⁶ We choose to present the individually estimated logit AQREs because the treatments are quite different, both in terms of the costs of unilaterally choosing Go (which are higher for the N treatment) and in terms of the complexity of the environment (in the D treatment the majority of the subjects have a dominant strategy whereas the majority of subjects in the N treatment have only a conditional best response).³⁷ Based on a χ^2 test of the first stage frequency of Go, we are able to conclude that the AQRE predictions are statistically different from what would be expected under the hypothesis of random or Nash behavior. Details of the tests are included in the supplemental materials.³⁸

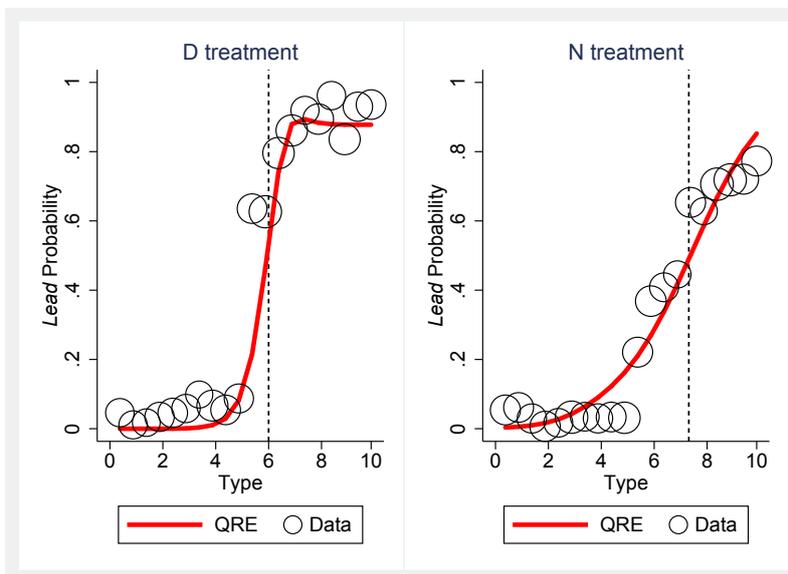


Figure 1: AQRE and Data by Treatment

The AQRE reproduces key features of the data. Crucially, it captures the behavior around θ^* : The AQRE correctly predicts that types just below θ^* deviate from the model to a greater extent than types just above this cut-off. The AQRE also identifies key features such as the stable level of leading for high types in the D treatment.³⁹ Between treatments, the AQRE correctly predicts that there should be a

³⁶Section S.5 in the supplementary materials provides a detailed discussion of the estimation procedure.

³⁷Jointly estimated logit AQRE are presented in the supplementary materials figure S6. With joint fitting of the data, it is primarily the fit for high types in the N treatment that suffers. Qualitatively, however, the jointly estimated logit AQREs are consistent with the ones presented in the main text. Haile et al. (2008) demonstrate the lack of falsifiability of QRE when any error distribution is permitted. However, even a treatment by treatment estimation of the logit AQRE is disciplined by the extreme value distributional assumption necessary to arrive at the logit form of choice probabilities.

³⁸See section S.5.

³⁹The flat (and even declining for high noise) Lead probability for players in the D treatment with high types is the

rapid change in behavior around the cut point $\theta^* = 6$ in the D treatment whereas behavior in the N treatment should change more gradually. The AQRE thus predicts the difference in the level of errors we observe in the data. Relative to our earlier discussion of the role of beliefs with regard to conditional and unconditional best responses, the AQRE provides a more nuanced perspective: It suggests that beliefs vary continuously and that this is an important feature for modeling actual behavior.

Result 5 (Noisy Leadership) *The AQRE rationalizes observed behavior. In particular, it explains the reluctance to Lead when leading is a conditional best response in the neo-classical equilibrium.*

6 Conclusion.

We have investigated the model of Farrell and Saloner (1985) in a controlled laboratory experiment. We find that subjects by and large respond to the incentives of the model as predicted. However, there is a reluctance to Lead not accounted for by the model. This reluctance is primarily present when leadership failure is costly. For our parameters leadership failure is more costly when leading is a conditional best response. We use a quantal response equilibrium to account for this phenomenon. In the quantal response equilibrium beliefs are relevant everywhere. We find that the observed deviations from neo-classical equilibrium is explained well by injecting some noise in the equilibrium concept.

Once a subject decides to Go he or she produces a strong incentive for moderate types to jump on the bandwagon. This is because the leader resolves all uncertainty on behalf of potential followers. We find that this complementarity in actions strongly determines follower behavior. Hence, the main driver of deviations from neo-classical equilibrium is weak leadership. As a consequence, efficiency losses are greater when potential leaders have non-dominant best responses. However, we find that cheap talk improves subjects' ability to coordinate on mutually beneficial actions and increases efficiency.

outcome of the subgame structure. For subjects with high types, if they fail to Lead in the first stage, there is still a high probability of Leading in the second (since they prefer y alone). The payoff consequence is therefore about the same for all subjects in this range: It is the size of the payoff externality from not inducing the preferred outcome. This predicts similar behavior for these subjects. In addition, when behavior is noisy, lower types are less likely to correct their mistakes in the second stage than higher types. This can explain why for lower levels of noise it is actually types in the vicinity of $\bar{\theta}$ for whom a error to not Lead is most costly.

References

- Belleflamme, P. and M. Peitz (2015). *Industrial organization: markets and strategies*. Cambridge, UK: Cambridge University Press.
- Bellemare, C., L. Bissonnette, and S. Kröger (2016). Simulating power of economic experiments: the powerBBK package. *Journal of the Economic Science Association* 2(2), 157–168.
- Blume, A. and A. Ortman (2007). The effects of costless pre-play communication: Experimental evidence from games with pareto-ranked equilibria. *Journal of Economic Theory* 132, 274–290.
- Brandts, J. and D. J. Cooper (2007). It’s what you say, not what you pay: an experimental study of manager-employee relationships in overcoming coordination failure. *Journal of the European Economic Association* 5(6), 1223–1268.
- Brandts, J., D. J. Cooper, E. Fatas, and S. Qi (2016). Stand by me—experiments on help and commitment in coordination games. *Management Science* 62(10), 2916–2936.
- Brandts, J., D. J. Cooper, and R. A. Weber (2015). Legitimacy, communication, and leadership in the turnaround game. *Management Science* 61(11), 2627–2645.
- Brindisi, F., B. Çelen, and K. Hyndman (2009). On the role of information and strategic delay in coordination games: Theory and experiment. Unpublished manuscript, <https://www.mcgill.ca/economics/files/economics/hyndman.pdf>.
- Brindisi, F., B. Çelen, and K. Hyndman (2014). The effect of endogenous timing on coordination under asymmetric information: An experimental study. *Games and Economic Behavior* 86(July), 264–281.
- Bulow, J. I., J. D. Geanakoplos, and P. D. Klemperer (1985). Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political economy* 93(3), 488–511.
- Carlsson, H. and E. Van Damme (1993). Equilibrium selection in stag hunt games. In K. Binmore, A. Kirman, and P. Tani (Eds.), *Frontiers of game theory*, Chapter 12, pp. 237–253. Cambridge, MA: MIT Press.
- Cartwright, E., J. Gillet, and M. Van Vugt (2013). Leadership by example in the weak-link game. *Economic Inquiry* 51(4), 2028–2043.
- Chaudhuri, A. and T. Paichayontvijit (2010). Recommended play and performance bonuses in the minimum effort coordination game. *Experimental Economics* 13(3), 346–363.
- Cooper, R. (1999). *Coordination games*. Cambridge University Press.
- Crawford, V. (1998). A survey of experiments on communication via cheap talk. *Journal of Economic Theory* 78, 286–298.
- Duffy, J. and J. Ochs (2012). Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior* 76(1), 97–116.
- Edmond, C. (2013). Information manipulation, coordination, and regime change. *Review of Economic Studies* 80(4), 1422–1458.
- Egorov, G. and K. Sonin (2011). Dictators and their viziers: Endogenizing the loyalty–competence trade-off. *Journal of the European Economic Association* 9(5), 903–930.

- Farrell, J. and P. Klemperer (2007). Coordination and lock-in: Competition with switching costs and network effects. In M. Armstrong and R. Porter (Eds.), *Handbook of Industrial Organization, Volume 3* (1 ed.), Chapter 31, pp. 1967–2072. Amsterdam: North Holland.
- Farrell, J. and G. Saloner (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics* 16(1), 70–83.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2), 171–178.
- Fréchette, G. R. (2012). Session-effects in the laboratory. *Experimental Economics* 15(3), 485–498.
- Garratt, R. and T. Keister (2009). Bank runs as coordination failures: An experimental study. *Journal of Economic Behavior & Organization* 71(2), 300–317.
- Goldberg, L. and C. Kolstad (1995). Foreign direct investment, exchange rate variability and demand uncertainty. *International Economic Review* 36(4), 855–873.
- Goldstein, I. and A. Pauzner (2005). Demand–deposit contracts and the probability of bank runs. *Journal of Finance* 60(3), 1293–1327.
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association* 1(1), 114–125.
- Haile, P. A., A. Hortacsu, and G. Kosenok (2008). On the empirical content of quantal response equilibrium. *American Economic Review* 91(1), 180–200.
- Helland, L., J. Hovi, and H. Sælen (2017). Climate leadership by conditional commitments. *Oxford Economic Papers*. Forthcoming.
- Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *American Economic Review* 75(3), 424–440.
- Katz, M. L. and C. Shapiro (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy* 94(4), 822–841.
- McKelvey, R. D. and T. R. Palfrey (1998). Quantal response equilibria for extensive form games. *Experimental Economics* 1(1), 9–41.
- Morris, S. and H. S. Shin (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88(3), 587–597.
- Morris, S. and H. S. Shin (2002). Measuring strategic uncertainty. Unpublished manuscript, <http://www.nuff.ox.ac.uk/users/shin/pdf/barcelona.pdf>.
- Rodrik, D. (1991). Policy uncertainty and private investment in developing countries. *Journal of Development Economics* 36(2), 229–242.
- Shy, O. (2001). *The Economics of Network Industries*. Cambridge, UK: Cambridge University Press.
- Turocy, T. (2010). Computing sequential equilibria using agent quantal response equilibria. *Economic Theory* 42(1), 255–69.

S SUPPLEMENTARY MATERIALS

S.1 First-Stage Behavior

S.1.1 Figures

Figure S1 shows the first stage behavior of subjects in the D and N treatments. On the horizontal axis is a set of twenty bins, each corresponding to 0.5 intervals over subject types: The first bin includes subjects with types $\theta \in [0, 0.5)$, the second bin includes subjects with types $\theta \in [0.5, 1)$, etc. On the vertical axis is the proportion of subjects in each bin who chose to Go in the first stage. We interpret this proportion as a probability. The bubbles are scaled by the number of observations within a bin, relative to the total number of observations within a treatment.

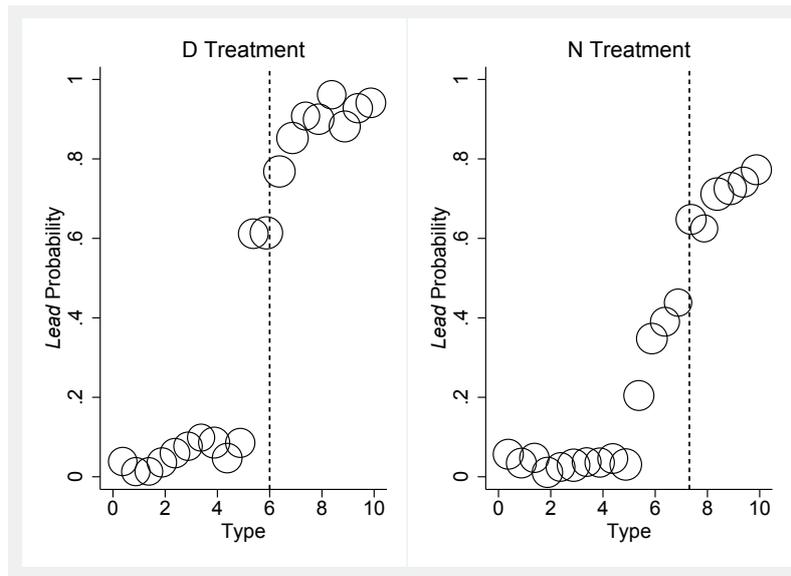


Figure S1: First Stage Behavior

The theoretical prediction for the first stage behavior is a step function at θ^* : In the equilibrium of the model, subjects with types below θ^* Stay in the first stage while those above θ^* Lead. For each treatment, this threshold is indicated by a dashed line. The plots in figure S1 illustrate that subjects with low types tend to Stay in the first stage while subjects with high types tend to Go. Moreover, the frequency of leading increases steeply in the vicinity of θ^* in both treatments. This is consistent with the use of bandwagon strategies.

S.1.2 Between Treatment Comparisons

To assess whether behavior in the strategic ranges is the same in the D and N treatments, we use Wilcoxon rank-sum (WSR) tests to compare the frequency with which participants choose to Go. These tests are based on between treatment comparisons of session-level data. For each test we present the relevant session data and the associated means and standard deviations. We denote the WSR test statistic by W . The p -value indicates how likely it is that the given observations come from the same distribution.

Comparison of behavior in the Stay Range (Table S1). A two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior in the Stay range, $\theta \in [0, \theta]$, in the D and N treatments: $W = 0.1, p = 0.92$.

Session	D	N
1	0.00	0.01
2	0.02	0.03
3	0.04	0.04
4	0.08	0.05
5	0.08	0.07
Mean	0.04	0.04
Std	0.04	0.02

Table S1: $\theta_D \in [0, 3)$ vs. $\theta_N \in [0, 3)$

Comparison of behavior in the Follow Range (Table S2). A two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior in the Follow range, $\theta \in [\theta, \theta^*]$, in the D and N treatments: $W = 0.94, p = 0.35$.

Session	D	N
1	0.22	0.12
2	0.22	0.18
3	0.24	0.23
4	0.33	0.24
5	0.36	0.29
Mean	0.27	0.21
Std	0.07	0.06

Table S2: $\theta_D \in [3, 6)$ vs. $\theta_N \in [3, 7.3)$

Comparison of behavior when Lead is a conditional best response (Tables S3 and S4). A two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior when Lead is a conditional best response regardless of whether we compare the region $\theta \in [\theta^*, \bar{\theta}]$ in the D treatment with the entire Lead region in the N treatment or just the restricted region $\theta \in [\theta^*, \theta^* + 1]$: Both tests deliver identical results, $W = 1.36, p = 0.17$.

Session	D	N
1	0.69	0.63
2	0.73	0.70
3	0.76	0.71
4	0.85	0.73
5	0.92	0.77
Mean	0.79	0.71
Std	0.09	0.05

Session	D	N
1	0.69	0.44
2	0.73	0.57
3	0.76	0.71
4	0.85	0.74
5	0.92	0.79
Mean	0.79	0.65
Std	0.09	0.14

Table S3: $\theta_D \in [6, 7)$ vs. $\theta_N \in [7.3, 10]$

Table S4: $\theta_D \in [6, 7)$ vs. $\theta_N \in [7.3, 8.3)$

Comparison of behavior in the Lead range (Table S5). A two-sample Wilcoxon rank-sum (Mann-Whitney) test rejects equality of behavior for the Lead range, $\theta \in [\theta^*, 10]$: $W = 2.61, p = 0.01$.

Session	D	N
1	0.83	0.63
2	0.86	0.70
3	0.90	0.71
4	0.91	0.73
5	0.94	0.77
Mean	0.89	0.71
Std	0.04	0.05

Table S5: $\theta_D \in [6, 10]$ vs. $\theta_N \in [7.3, 10]$,

Comparison of behavior when Lead is an unconditional best response in D (Table S6). A two-sample Wilcoxon rank-sum (Mann-Whitney) test rejects equality of behavior for subjects with high type draws; when we compare test subjects in the D treatment in the region $\theta_D \in [\theta^* + 1, 10]$ with test subjects in the N treatment in the region $\theta_N \in [\theta^* + 1, 10]$, we strongly reject equality of behavior: $W = 2.61, p = 0.01$.

Session	D	N
1	0.87	0.68
2	0.89	0.69
3	0.92	0.74
4	0.94	0.78
5	0.97	0.79
Mean	0.92	0.74
Std	0.04	0.05

Table S6: $\theta_D \in [7, 10]$ vs. $\theta_N \in [8.3, 10]$

S.1.3 The Cost of Failing to Lead

In the bandwagon game, the strategic decision to Stay or Lead is most difficult for players in the vicinity of the first stage leading threshold θ^* who have a conditional best response to Lead. These are players who prefer a joint choice of G but would stick with S if they knew that their match will choose S with certainty. In the D treatment this range is relatively small while in the N treatment it is relatively large. To get a measure of how costly it is for players in this region to forgo leading, we tabulate the frequency with which subjects in the relevant ranges encounter a subject in the Follow range. Next, we tabulate the frequency with which subjects take the correct strategic timing decision and the matched subject does in fact Follow.⁴⁰ Finally, as a crude measure of the importance of correctly taking the strategic timing decision, we list the average payoff from the (correct) decision to Lead relative to the average payoff from the choice to Stay:

Data	Correct Go	Correct Go & Followed	Follow Error	Δ Payoff
Joint	0.79	0.72 (0.91)	6.51	3.28
D	0.79	0.79 (1.00)	9.17	1.59
N	0.78	0.71 (0.91)	5.24	3.61

Table S7: Strategic Timing

⁴⁰In about 9% of cases subjects in the Follow range fail to Follow.

S.1.4 Errors

Figure S2 shows the empirical error frequency in 0.1 unit bins along with a thick black line which is a 0.5 unit moving average of the error frequencies. This figure has three main features: First, the pattern of errors is asymmetric around the cut point in both treatments. Second, there is sharper pattern of errors in the *D* treatment relative to the *N* treatment. This indicates that the range in which the decision to Lead is uncertain is more narrow in the *D* treatment relative to the *N* treatment. Third, there is a generally higher level of errors in the *N* treatment relative to the *D* treatment for subjects with high types. This testifies to an overall higher level of uncertainty in the *N* treatment.

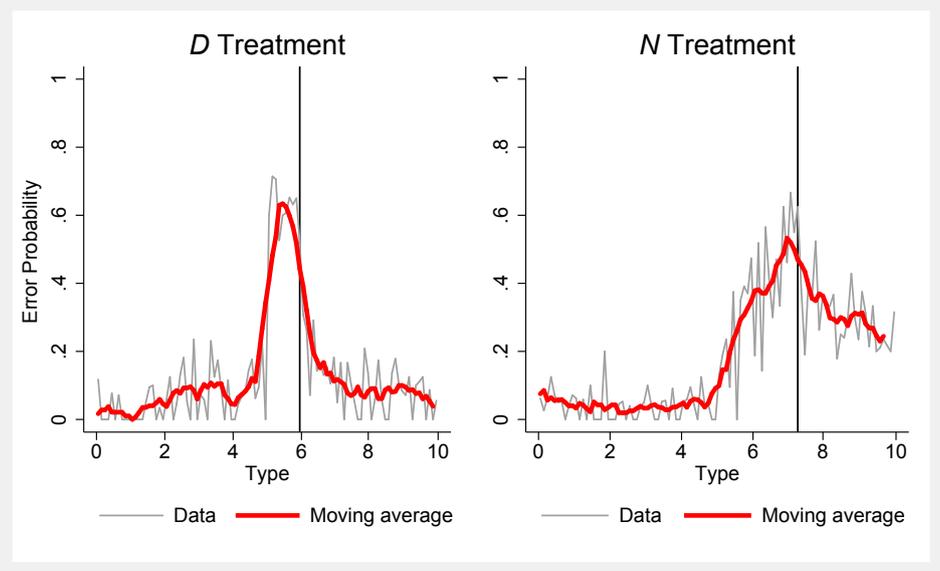


Figure S2: First Stage Errors

S.2 Second-Stage Behavior

S.2.1 Figures

Figure S3 presents the second-stage behavior by treatment, conditional on the match's action. The figure includes only those subjects who take a second stage decision. On the horizontal axis is subject type, grouped in half unit bins, and on the vertical axis is the proportion of subjects that Go in the second stage. The left panel presents the second stage behavior for subjects whose match chose to Stay in the first stage while the right panel presents the second stage behavior for subjects whose match chose to Lead in the first stage. Data from the N treatment are presented as hollow bubbles and data from the D treatment are presented as shaded bubbles. The size of the bubbles reflects the proportion of observations in a bin relative to the total number of observations within a treatment. The thresholds identified by the equilibrium of the model are marked by vertical lines: We denote $\underline{\theta}$ by a short dashed black line (this threshold is identical for both treatments) while we denote θ^* by a short dashed black line for the D treatment ($\theta^* = 6$) and a long dashed gray line for the N treatment ($\theta^* = 7.3$).

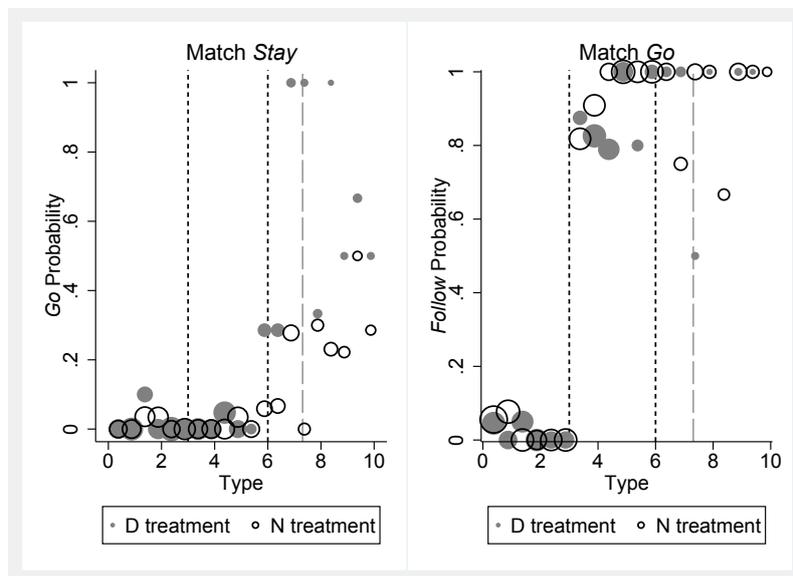


Figure S3: Second Stage Behavior by Match Action

As is evident from figure S3, complementarity has a strong effect on the outcomes.

S.2.2 Second Stage Between Treatment Comparisons

Comparison of behavior in the Stay Range (Table S9 and S8). Two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior when test subjects are in the Stay range. This holds when the match chooses to Stay ($W = 1.16, p = 0.24$) and when the match chooses to Go ($W = -0.63, p = 0.53$).

Match Stay		
Session	D	N
1	0.00	0.00
2	0.03	0.00
3	0.04	0.02
4	0.05	0.04
5	0.08	0.04
Mean	0.04	0.02
Std	0.03	0.02

Table S8: $\theta_D \in [0, 3)$ vs. $\theta_N \in [0, 3)$

Match Go		
Session	D	N
1	0.00	0.00
2	0.03	0.03
3	0.03	0.07
4	0.05	0.08
5	0.10	0.13
Mean	0.04	0.06
Std	0.04	0.05

Table S9: $\theta_D \in [0, 3)$ vs. $\theta_N \in [0, 3)$

Comparison of behavior in the Follow Range (Table S10 and S11). Two-sample Wilcoxon rank-sum (Mann-Whitney) test can not reject equality of behavior when test subjects are in the Follow range. This holds when the match chooses to Stay ($W = 0.52, p = 0.60$) and when the match chooses to Go ($W = 1.79, p = 0.07$).

Match Stay		
Session	D	N
1	0.02	0.05
2	0.03	0.06
3	0.10	0.07
4	0.12	0.08
5	0.13	0.08
Mean	0.08	0.07
Std	0.05	0.01

Table S10: $\theta_D \in [3, \theta_D^* = 6)$ vs. $\theta_N \in [3, \theta_N^* = 7.3)$

Match Go		
Session	D	N
1	0.88	0.84
2	0.89	0.85
3	0.89	0.86
4	0.92	0.88
5	0.92	0.90
Mean	0.90	0.87
Std	0.02	0.03

Table S11: $\theta_D \in [3, \theta_D^* = 6)$ vs. $\theta_N \in [3, \theta_N^* = 7.3)$

S.2.3 Logistic Regression

To demonstrate the predictive ability of the bandwagon model and to illustrate the role of complementarity, we estimate the logistic regression for individual i in repetition t with clustered errors for each individual:

$$\mathbb{P}(G_{it}^2 | a_{jt}^1) = F(\beta_0 + \beta_1 \theta_{it} + \beta_2 \mathbf{range}_{it} + \beta_3 \mathbf{treatment}_{it} + \beta_4 a_{jt}^1 + \beta_5 \mathbf{range}_i a_{jt}^1 + \beta_6 \mathbf{range}_{it} \mathbf{treatment} + \beta_7 \mathbf{treatment} a_{jt}^1 + \epsilon_{it})$$

We include three dummy variables: The **range** dummy to indicate players with types in the Follow range (relative to the Stay range), the **treatment** dummy to indicate the D treatment (relative to the N treatment), and the dummy variable a_j^1 to indicate whether the first stage action of i 's match was Go. We include the pairwise interactions between these last three variables.

G_{it}^2	Coeff. (p-value)	Std. Err.
θ_{it}	0.44 (0.00)	0.11
range	-0.44 (0.31)	0.43
treatment	0.40 (0.46)	0.53
a_{jt}^1	0.84 (0.06)	0.44
range x a_{jt}^1	3.94 (0.00)	0.41
treatment x a_{jt}^1	-0.39 (0.38)	0.44
treatment x range	0.33 (0.52)	0.52
constant	-4.50 (0.00)	0.45

Table S12: Logistic estimates of second stage Go probability

In this specification, only type θ_{it} , match action a_{jt}^1 , and the interaction between range and a_{jt}^1 are significant.⁴¹ Contrary to theory, type has an independent effect on the probability to Go and higher types are more likely to Go regardless of their match's action. However, this effect is relatively weak. The much stronger effect is the interaction between the match's action and range. In particular, when a subject is in the Follow range and their match chooses to Go in the first stage, the average probability that the test subject will Lead increases by about 70 percentage points relative to the case when their match chooses to Stay. In addition, we see that the impact of the match's action is close to zero in the case when the test subject is in the Stay range. These effects are clearly illustrated by figure 8 which plots the predicted Go probability given the match's action in both the D and N treatments.⁴²

⁴¹ Alternative specifications yielded nearly identical results with small differences in estimated coefficients and p-values.

⁴² Recall that in both treatments, the dividing line between the Stay and Follow ranges is at $\underline{\theta} = 3$.

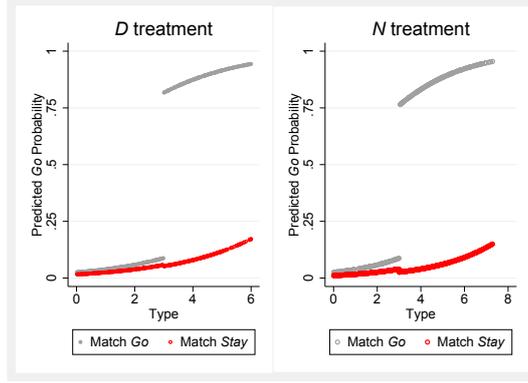


Figure S4: Predicted Leading Probabilities Conditional on Match Action

S.3 Signaling Game

Model

The equilibrium with signaling takes a form similar to the equilibrium without signals. First, observe that players will always send the message that promotes the outcome that they prefer. The message that players send therefore perfectly reveals whether a player has a type above or below θ^o . An implication is that players who send the same message will choose the same action. In addition, if players send conflicting messages, then the player who prefers outcome Go can update her belief about the type of their match. Relative to the game without signaling, the player has more information since the range of possible types for the match is truncated from θ^* to θ^o . Further, if this player choose to Stay in the first stage, then status quo prevails. This means that the upper bandwagon threshold must satisfy

$$\mathbb{P}(\theta_{-i} > \underline{\theta} | \theta_{-i} < \theta^o) \pi(G_i, G_j; \theta^*) + (1 - \mathbb{P}(\theta_{-i} > \underline{\theta} | \theta_{-i} < \theta^o)) \pi(G_i, S_{-i}; \theta^*) = \pi(S_i, S_{-i}; \theta^*).$$

Given the uniform distribution of types on the interval $[0, 10]$ and our parameterization of the payoff functions, this reduces to

$$\frac{\theta^o - \underline{\theta}}{\theta^o} (\theta^* + 2) + \frac{\underline{\theta}}{\theta^o} \theta^* = 7,$$

where $\underline{\theta} = 3$, and $\theta^o = 5$. Relative to the D treatment, the upper bandwagon threshold increases slightly from 6 to 6.2 in the game with signaling.

Figures

We present the first-stage results in figure S5. Comparison of the D and S treatments demonstrates that communication improves coordination of actions whenever subjects have the same preferred outcome.

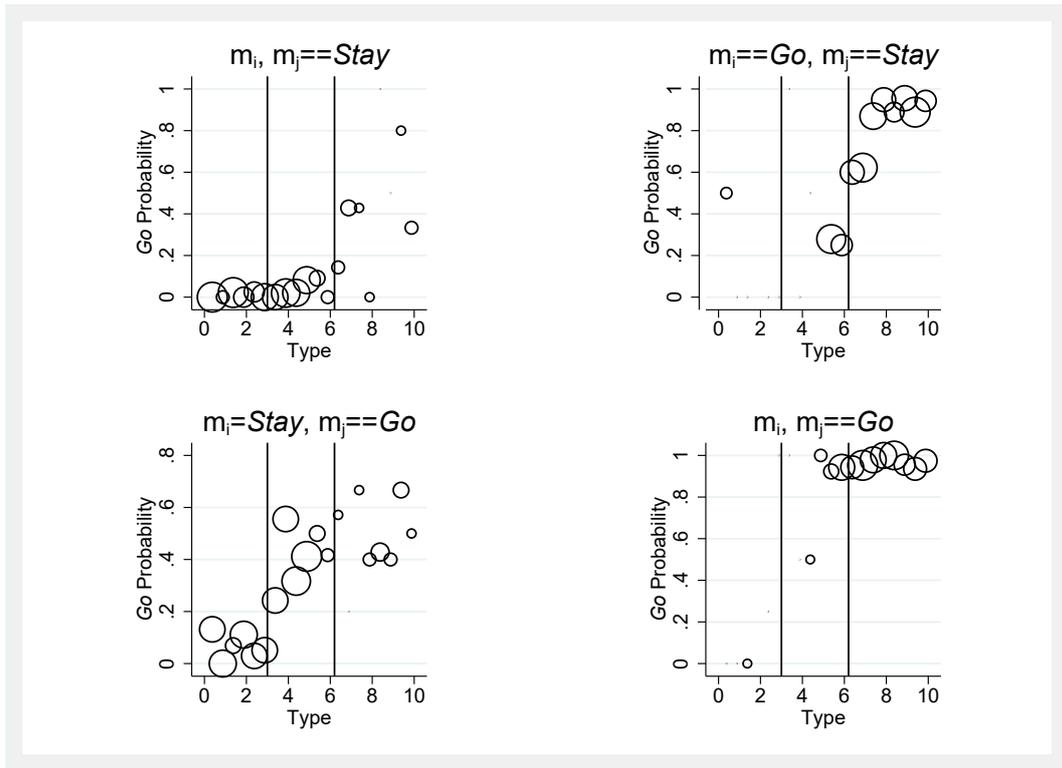


Figure S5: First Stage GO by Signal

On the main diagonal, we see that when players send the same message, they overwhelmingly choose the same action in the first stage. When both players signal Go or both signal Stay, the coordination success of subjects is substantially higher than what is observed in the D and N treatments. On the off-diagonal, we see the instances in which subjects send conflicting signals. In these cases, the subjects should play bandwagon strategies similar to the D treatment, with the exception that $\theta^* = 6.2$. As expected, we observe that subjects with dominant strategies (respectively $\theta < \underline{\theta}$ and $\theta > \bar{\theta}$) behave as they should. However, consistent with the other treatments, there is an over-eagerness for subjects just below θ^* to Lead.

S.4 Efficiency

Wilcoxon rank-sum tests identify a statistically significant difference with respect to the *realized efficiency* (empirically observed payoff as a percent of the maximum possible payoff) between each pair of treatments, between S and D ($W = 1.4, p = 0.087$), between S and N ($W = 2.6, p = 0.005$), and between D and N ($W = 2.6, p = 0.005$).

Session	S	D	N
1	97.0	95.4	89.4
2	95.6	95.0	88.3
3	94.8	94.9	88.3
4	94.6	93.9	88.0
5	94.1	93.2	88.0
Mean	95.2	94.1	88.4
Std	0.01	0.01	0.01

Table S13: Empirical efficiency as percent of maximum by treatment–session,

S.5 Equilibrium with Noise

Regarding the estimation of the AQRE, there are a few features on which it is worth commenting. Because the type space is continuous, in our estimation we discretize the type space into B equally spaced bins with types corresponding to the mid-point in the bin. Since all actions are played with some non-zero probability in an AQRE, the expected payoff for a player in bin i depends on how likely it is to get a match j in each of the $1 \dots B$ bins and the probability that the match will choose to Go conditional on the actions taken so far. In particular, in the second stage, players update their beliefs about their match’s type based on whether their match chose to Stay or Go in the first stage. Although we present estimates based on 20 bins in the paper, we estimated versions with up to 100 bins. Since increasing the number of bins did not change any conclusions—even delivering the same estimate of the noise parameter—we choose to present the simpler version.

Given the discretization, the estimation involves two stages. In the first stage, we estimate a fixed point for the vector of first stage Go probabilities taking the second stage Go probabilities as given. Since agents are forward looking, they anticipate how likely it is that their match will Go in the first stage and how their own first stage action will affect the second stage action of their match. In the second stage, agents that chose to Stay are in one of two possible situations: Either their match chose to Go or their match chose to Stay. In both cases, we must estimate a QRE for the second stage Leading probability for each of the B types. Moreover, in the case when the match chose to Stay, the second stage estimation depends on the first-stage probability estimates because agents need to update their beliefs about how likely each type is because a high type will be more likely to Go in the second stage. The first and second stage decisions are thus interlinked because the first-stage decisions depend on the anticipated second stage probabilities and the second stage decisions depend on the updated beliefs generated in the first stage.

The actual estimation proceeded by looping through the first and second stage, using the estimated probabilities from the previous iteration of the procedure as beliefs. To efficiently estimate the model, we vectorize the computations. For example, in the first stage we compute the payoff from choosing Go for all the types $i = 1, \dots, B$ from the matrix multiplication

$$\frac{1}{B} \mathbf{1} (\mathbf{p}_1^T \pi_i(\mathbf{G}_i^1, \mathbf{G}_j^1) + (\mathbf{1}^T - \mathbf{p}_1^T) \mathbf{p}_{2, \text{Go}}^T \pi_i(\mathbf{G}_i^1, \mathbf{G}_j^2) + (\mathbf{1}^T - \mathbf{p}_1^T) (\mathbf{1}^T - \mathbf{p}_{2, \text{Go}}^T) \pi_i(\mathbf{G}_i^1, \mathbf{S}_j^2))$$

where all vectors are denoted in bold, are of length $1 \times B$, and transposes are indicated by a T . The vector \mathbf{p}_1 denotes the probability of match $j = \{1, \dots, B\}$ going in the first stage, $\mathbf{p}_{2, \mathbf{Go}}$ denotes the probability of match $j = \{1, \dots, B\}$ going in the second stage conditional on i choosing to Lead, and π_i denotes payoffs to a type $i = \{1, \dots, B\}$ that depends on the outcome realized in the second stage. The first term in the entire product, $\frac{1}{B} \mathbf{1}$, is the probability of meeting each of the B types while the second term (everything inside the outer parenthesis) is a $B \times B$ matrix that in each i, j -cell gives the expected payoff to a type i of meeting a type j . Notice that the second factor is composed of three outer products that respectively give the payoff from (1) both choosing to Lead, (2) i leading and j following in the second stage, and (3) i leading while j chooses to Stay in the second stage. The product thus produces a vector of length $1 \times B$ that in each position gives the expected payoff to a type i that results from the sum of payoffs from meeting all the $j \in \{1, \dots, B\}$ types. Analogous computations were carried out for the first stage payoff of choosing to Stay, as well as for the second stage payoffs that also depend via updated beliefs on the first stage action of the match. QRE probability estimates in each stage could then be computed as the ratio of payoffs from Go relative to the sum of the payoffs from Go plus the payoffs of Stay as shown in section 5. The QRE probabilities from the first and second stage for all B types then characterize the behavior profile. The estimation terminated when the estimates converged sufficiently that the maximum distance between the previous best estimates and the current best estimates fell below the tolerance.

How to select the level of noise and fit the AQRE is an open question. Foremost, the question is whether the noise should be jointly estimated across the two stages of the model and whether the noise should be jointly estimated across treatments. To discipline our analysis against over-fitting, we chose to constrain noise to be the same across both the first and second stage of the estimation. However, this does create some issues. Specifically, the second-stage decisions are more “simple” than the first stage decisions in the sense that the payoff consequences of making a mistake are more clear and more stark. Thus, even with low levels of noise, the second stage behavior is sharper. In turn, this has consequences for the first-stage estimates.

In figure 9, we present the same plots as in figure 1 but with the noise parameter estimated jointly for both treatments. In the top panel we present the data and fitted AQRE for all the periods while in the bottom panel the same information for the last ten periods.

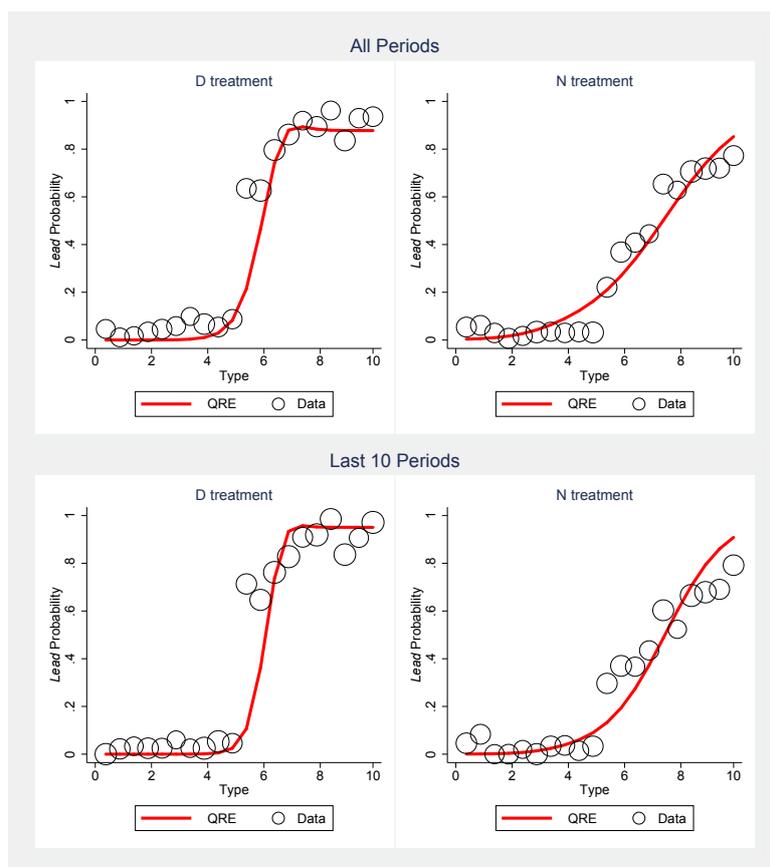


Figure S6: AQRE, Joint estimate of lambda

We investigated several procedures to select the level of noise: Maximum likelihood, Euclidean distance, and absolute distance. In all cases, our fit was based on the closeness of the first-stage estimates to the data. All three procedures produced similar results, although maximum likelihood and Euclidean distance estimated somewhat higher levels of noise. The maximum likelihood estimates were strongly affected by the fact that for the lower half of the types, the predicted first-stage behavior is close to zero. Because actual behavior was somewhat greater than zero—even for lowest types—the maximum likelihood requires a high level of noise. The second feature that created issues was that subjects in the bin $[5, 5.5)$ were leading at a higher rate in this bin than in the bin $[5.5, 6)$. This non-monotonic behavior is strongly punished by the quadratic distance and led to a higher noise for the Euclidean distance. Given these issues and our interest in using the model for prediction, we fit the AQRE based on the absolute distance.

Goodness of fit

To assess the goodness of fit of the AQRE model, we use a χ^2 test. Observe that under a given model of behavior, each type of player is associated with a probability of choosing Go and the distribution of Go will follow a binomial distribution. The binomial distribution approaches the normal distribution asymptotically. If we group players in narrow bins, the joint distribution of the probability of Go will therefore approach a multivariate normal distribution. A χ^2 test can then be used to evaluate whether

the AQRE frequencies are consistent with the frequencies expected under an alternative hypothesis.

Consider the case of testing the AQRE model against the hypothesis of random behavior. Let there be q bins indexed by j , each with n_j observations. Denote the AQRE predictions by \hat{x}_j and observe that the prediction of random behavior is $p = 0.5$ for all j . Under the null hypothesis of random behavior, $\mathbf{x} = ((\hat{x}_1 - n_1 0.5) / \sqrt{n_1 0.5(1 - 0.5)}, \dots, (\hat{x}_q - 0.5 n_q) / \sqrt{n_q 0.5(1 - 0.5)})'$ will approach a multivariate normal with the zero vector in expectation and the identity matrix as the covariance matrix. Furthermore, $\mathbf{x}^T \mathbf{x} = \sum_{j=1}^q (\hat{p}_j - 0.5 n_j)^2 / n_j 0.5(1 - 0.5)$ will (asymptotically) be distributed as χ_q^2 under the null hypothesis. A χ^2 test can thus be carried out to comparing the AQRE prediction with the anticipated probability of Go.

A similar type of test can be used to compare the AQRE model with the Nash prediction. However, because the Nash equilibrium predicts zero probabilities for certain types, we use the AQRE predictions as the null hypothesis (to avoid zeros in the denominator).

For both random and Nash behavior, we can strongly reject the null hypothesis. The χ^2 statistics are large in each case. The AQRE is inconsistent with the predictions given by these other models of behavior. For the case of random behavior, the test statistics are 1643 and 1474 for the D and N treatments respectively. This is far in excess of the threshold of about 38 required for 1% significance. Similarly, compared with Nash, test statistics for both treatments were in excess of 159.

S.6 Appendix. Instructions for baseline treatment

INSTRUCTIONS

You are going to participate in an experiment financed by the Department of Economics at BI and the Norwegian Research Council.

You will earn money. How much you earn depends on the decisions you make, as well as on the decisions made by other subjects.

All interactions are anonymous and are performed through a network of computers. The administrators of the experiment will not be able to observe your decisions during the experiment.

There are 20 subjects participating in the experiment. All participants are in this room. They have all been recruited in the same way as you and are reading the same instructions as you are for the first time. It is important that you do not talk to any of the other subjects in the room until the experiment is over.

In the experiment you will earn points. At the end of the experiment, you will be paid in Norwegian Kroner (NOK) based on your total earnings in points from all the games of the experiment. The exchange rate from points to NOK is:

$$1 \text{ point} = 1,25 \text{ NOK}$$

The more points you earn, the more cash you will receive.

DESCRIPTION OF THE EXPERIMENT

The experiment consists of 19 games. In each game, you are randomly matched with another participant (your “match”). You will be matched with each participant once and only once.

Each game has an identical structure and consists of two stages:

- **First stage:** In the first stage, you select either action CIRCLE or action SQUARE. If your action is SQUARE in the first stage, then your action is SQUARE in the second stage as well. Your match faces an analogous first-stage decision between action CIRCLE and action SQUARE. After you and your match have selected actions, you observe your match’s action and your match observes your action. This concludes the stage.
- **Second Stage:** In the second stage, there are two possibilities:
 - If your action in the first stage is CIRCLE then you again select between action CIRCLE and action SQUARE.
 - If your action in the first stage is SQUARE then you do not make a second-stage selection: Your second stage action is SQUARE.

Your match has the same options: If your match selected CIRCLE in the first stage, then he/she selects again. Otherwise, your match's second stage action is SQUARE.

At the conclusion of the second stage, you observe your match's action and your match observes your action. Points are allocated based on which of the four combinations of actions occurs. The four combinations are:

1. Your action is CIRCLE and your match's action is CIRCLE
2. Your action is CIRCLE and your match's action is SQUARE
3. Your action is SQUARE and your match's action is CIRCLE
4. Your action is SQUARE and your match's action is SQUARE

The payoff structure is discussed in the next section. After points are allocated, the game is concluded.

Payoff structure

At the beginning of each game, the software in the machine draws a random number between 0.00 and 10.00. Each number in the interval 0.00 to 10.00 has an equal probability of being drawn.

This random number determines your payoffs in that game and that game only. You observe this number before you choose any actions.

Specifically, your payoff structure takes the following form:

	Your payoffs (in points)
Your action is CIRCLE and your match's action is CIRCLE:	7
Your action is CIRCLE and your match's action is SQUARE:	5
Your action is SQUARE and your match's action is CIRCLE:	Your random number
Your action is SQUARE and your match's action is SQUARE:	2+Your random number

The software in the machine also draws a separate random number between 0.00 and 10.00 for your match. Again, each number in the interval 0.00 to 10.00 has an equal probability of being drawn. This random number determines the payoffs for your match. Your match observes this number before he/she makes any decisions.

Your match's payoff structure takes the following form:

	Your match's payoffs (in points)
Your action is CIRCLE and your match's action is CIRCLE:	7
Your action is CIRCLE and your match's action is SQUARE:	Your match's random number
Your action is SQUARE and your match's action is CIRCLE:	5
Your action is SQUARE and your match's action is SQUARE:	2+Your match's random number

In other words, your payoffs and the payoffs of your match only differ due to differences in your random numbers.

Information

Your payoff structure is revealed to you at the beginning of the game. You observe the points associated with each possible combination of actions before making a selection in the first stage.

In addition, prior to making your choice in the second stage, you observe the first-stage action of your match.

However, you never observe your match's payoffs.

Your match has similar information: Your match observes his/her own payoffs at the beginning of the game and observes your first-stage action prior to the second stage. But your match never observes your payoffs directly.

At the end of each game, historical statistics are presented. These statistics show the decisions made in each of the games that have been played. It also shows your profits from each game, as well as your total accumulated profits.

EARNINGS

After the last game is completed, your earnings in points are converted to Norwegian Kroner at the stated exchange rate. Your earnings will be paid in cash as you exit the lab.

TEST GAMES

We run two "test games" before we start the experiment. You do not earn points in these test games. However, the test games allow you to familiarize yourself with the screens used in the experiment.

Are there any questions?

S.7 Screenshots

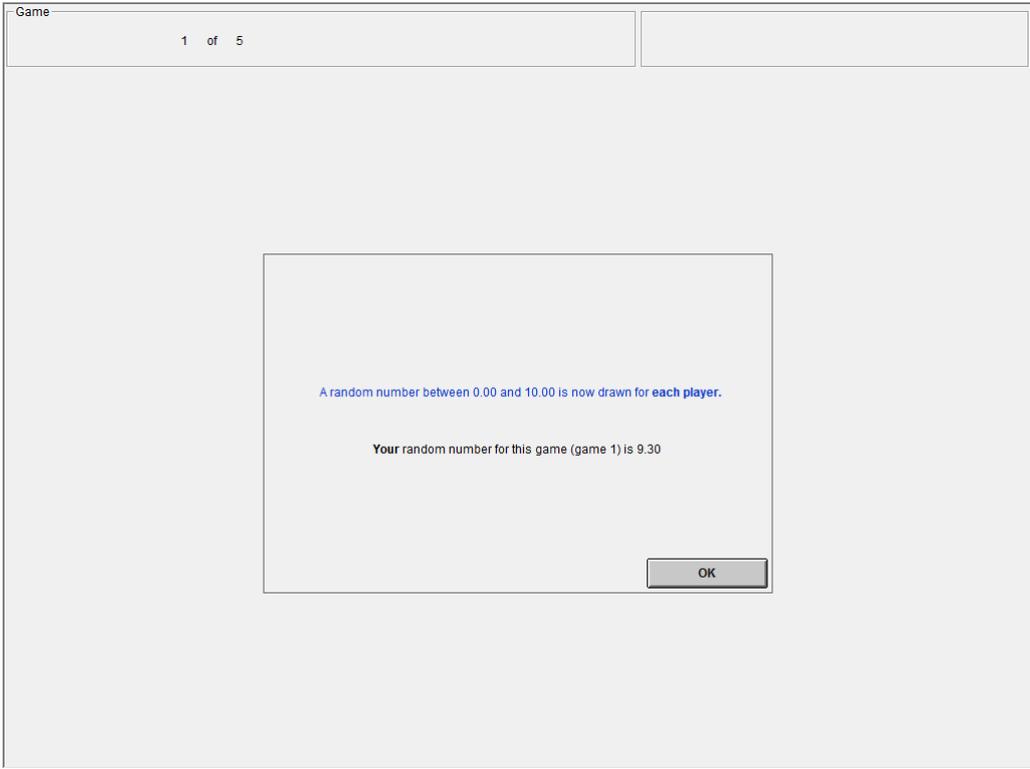


Figure S7: Stage 0. Random Number

Game

1 of 5

Your payoff matrix	Your match choose CIRCLE	Your match choose SQUARE
You choose CIRCLE	7.00	5.00
You choose SQUARE	9.91	11.91

This is stage ONE

Do you choose CIRCLE or SQUARE?

Figure S8: Stage 1. First Decision

Game

1 of 5

Your payoff matrix	Your match choose CIRCLE	Your match choose SQUARE
You choose CIRCLE	7.00	5.00
You choose SQUARE	9.30	11.30

This is stage TWO

In stage one your match chose CIRCLE

Do you choose CIRCLE or SQUARE?

Figure S9: Screen 2. Both Stay

Game

2 of 5

Your payoff matrix	Your match choose SQUARE
You choose CIRCLE	5.00
You choose SQUARE	7.88

This is stage TWO

In stage one your match chose **SQUARE**

Do you choose CIRCLE or SQUARE?

CIRCLE

SQUARE

Figure S10: Screen 2. Match Go

Game		4 of 5			
Your accumulated payoff (in points) 30.53					
<input type="button" value="OK"/>					
Game	Your stage ONE action	Your stage TWO action	Stage ONE action of your match	Stage TWO action of your match	Your points
1	CIRCLE	SQUARE	CIRCLE	CIRCLE	9.30
2	SQUARE	-	CIRCLE	SQUARE	9.47
3	SQUARE	-	SQUARE	-	4.75
4	CIRCLE	CIRCLE	CIRCLE	CIRCLE	7.00

Figure S11: Screen 3. Feedback