CESAR
centre for experimental studies and research

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

Heggedal, T.-R., L. Helland & K.-E. Neset Joslin

CESAR Working Paper No. 03/16

ISSN 2464-2908

NORWEGIAN BUSINESS SCHOOL

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

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

December 13, 2016

Abstract

We experimentally test the seminal platform selection model of Farrell and Saloner (1985). At the core of this model is the presence of irreversible actions and private valuations. In general, our data support the model. While complementarities in actions strongly determine follower behavior, there is a reluctance to lead not accounted for by theory. We explain observed deviations from the neo-classical equilibrium by injecting some noise in the equilibrium concept. We find that allowing cheap talk messages improves efficiency while policies aimed at insuring failed leadership or subsidizing joint choice of the challenger platform reduce efficiency.

Keywords: platform competition, networks, experimental economics
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.

[†]BI Norwegian Business School and Centre for Experimental Studies and Research (CESAR)
[‡]BI Norwegian Business School and CESAR
[§]BI Norwegian Business School and CESAR, e-mail: knut-eric.n.joslin@bi.no (corresponding author)

1 Introduction

New and superior technology platforms may fail in the market for a number of reasons. In this paper, we focus on the role of information in an environment in which switching to a new platform is irreversible. In the presence of network externalities and private valuations of available technologies, beliefs about the valuations of others become a critical determinant of whether new technologies succeed or fail. Beliefs may be too pessimistic to support switching to mutually preferred platforms. However, early switching by high valuation users may induce users with lower valuations to follow, resulting in mutually beneficial coordination of actions.

But, how large is the causal effect of private valuations on the switching probability? And to what extent does early switching by high valuation users cause users with lower valuations to follow? We address these questions by investigating the seminal network model of Farrell and Saloner (1985) in a controlled laboratory experiment.[1] In this model, the combination of a specific information structure and sequence of moves produces a unique equilibrium.[2] This equilibrium provides an unequivocal benchmark for our analysis and facilitates separate assessment of the role of strategic uncertainty and the role of complementarity in actions.[3]

In the model, players make platform choices in the presence of network externalities and incomplete information about types. The key decision is whether to "lead" and commit to a platform in the first stage or to "follow" and delay the decision to the second stage. Due to complementarity in actions, when a player leads, this creates incentives for other players to "jump on the bandwagon" and choose the same platform. However, a player may regret the decision to lead if other players fail to follow. Hence, players face a clear decision: "Should I Stay or Should I Go?" Our experiment allows us to tease out the relative importance of leader and follower behavior, and the associated efficiency consequences. In general, we find substantial support for the model in data. However, the deviations that we do observe are systematic. We show that they can be accounted for by introducing some noise in the equilibrium concept.

The dynamics that characterize the model are present in a wide range of economic interactions. The benefits of a platform to a user depend in part on idiosyncratic technological and institutional characteristics. However, the value of a platform also depends on the extent of the user base. Consider the example of the *Open Handset Alliance* (OHA). The OHA is an initiative to develop open standards for mobile devices based on the Android platform.[4] Members of the OHA include developers of applications and hardware, in addition to network operators. After the initial launch of the OHA in 2007, there have been subsequent rounds in which the OHA has expanded its membership.[5] This includes a number of companies that have deprecated or de-prioritized their own platforms in favor of the Android platform.[6] As is captured by the model, some firms switched to the OHA early on, before there was a critical mass of users, while other firms adopted the platform only after it had gained traction in the market.

[1] For textbook treatments see Shy (2001) and Belleflamme and Peitz (2015).

[2] Coordination problems are defined by the presence of multiple, Pareto-ranked equilibria. Coordination failure result 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.

[3] We follow Morris and Shin (2002) in defining strategic uncertainty as "uncertainty concerning the actions and beliefs (and beliefs about the beliefs) of others." According to Cooper and John (1988) strategic complementarities arise "when the optimal strategy of an agent depends positively upon the strategies of the other agents."

[4] The stated goal of the OHA is "to accelerate innovation in mobile and offer consumers a richer, less expensive, and better mobile experience"
`http://www.openhandsetalliance.com/`

[5] For instance, a Reuters article from December 10, 2008 reports that companies such as Sony Ericsson and Vodaphone are joining the OHA. It concludes by citing a web analyst from Sterling Market Intelligence who predicts that "...more people will jump on the bandwagon."
`http://www.reuters.com/article/us-openhandset-idUSTRE4B86M120081210`

[6] Other examples of platform technologies include payment card systems, electronic medical records, or even launch stations for space travel.

We make two main contributions. Foremost, we find that the effect of complementarity in actions is strong. If a subject takes the lead, other subjects follow with high probability. Second, we find that subjects often do not lead when it is a conditional best reponse. This reluctance to lead reduces efficiency and is not accounted for by the model. Leading carries the risk of failure; the leader might end up alone on the new platform. 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. In addition we investigate a simple extension of the basic model which permits cheap talk. We find that cheap talk boosts efficiency, as it should do in the equilibrium of Farrell and Saloner (1985).

We also analyze how a policy maker can use simple subsdies to promote technology adoption both in the context of the basic model as well as in the model with noise. Although the exact choice of subsidy does not matter, the effect of subsidies is reduced in the presence of noise. We conclude that policy-makers should be especially careful about attempts to induce technology uptake when the benefits are more uncertain. In addition, as the size of the subsidy increases, the return on investment falls.

To the best of our knowledge ours is the first experiment to address the model of Farrell and Saloner (1985). The paper closest to ours is Brindisi et al. (2014).[7] 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. This indicates that the information structure is crucial in determining the strength of bandwagon behavior in the presence of complementarities and irreversible choices. While complementarities in actions are a strong force in environments with private information about types, they appear not to be so under private information about fundamentals.

More generally, most, if not all, economic situations of interest will embody both 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.[8] 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 practise is to rely on a global games approach.[9]

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 noisey behavior impacts the equilibrium and the sixth section presents additional analyses related to efficiency. The final section concludes.

2 Model

We investigate a parameterized version of the bandwagon game with incomplete information. We consider the case with two players, $i = \{1, 2\}$, and two stages, $t \in \{1, 2\}$.[10] When referring to the generic player i, the other player will be referred to as the player's "match" and denoted by j.

[7] Brindisi et al. (2009) provides a thorough exposition of the theory.

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

[9] A few prominent examples include bank runs (Goldstein and Pauzner, 2005); speculative currency attacks (Morris and Shin, 1998); and political revolutions (Egorov and Sonin, 2011).

[10] The model can be generalized to the case with n players and n stages. The essential conclusions translate to that setting.

In this game, players choose to either stay with an incumbent platform x or switch to an entrant platform y. We denote the actions by x_i^t and y_i^t where the superscript specifies the stage in which the action is chosen and the subscript indicates the player. Payoffs in the game only depend on the outcome at $t = 2$. For expositional purposes, we refer to action y and x as respectively *Go* and *Stay.* In addition, we refer to action y^1as a decision to *Lead* and action $y_i^2|y_j^1$ (choosing y in the second period conditional on the match choosing y in the first period) as a decision to *Follow.*

The game has the following timeline:

0. Prior to first stage, nature draws a type θ_i for each player. Type draws are i.i.d from a uniform distribution: $\theta_i \sim U[0, 10]$. Each player's type is private information and is revealed to the player but not to the player's match. θ parameterizes preferences, with higher realizations associated with higher payoffs from platform y.

1. In the first stage, players simultaneously select action x^1 or action y^1. The choice of y^1 also commits a player to y^2. Players observe the stage 1 action of their match at the conclusion of the stage. A player who observes that their match has chosen y^1 thus knows that action y^2 is played by their match in the second stage.

2. In the second stage, players who chose x^1 in stage 1 choose between x^2 and y^2. Players who chose y^1 do not make a decision in the second stage. If both players chose x^1, second stage decisions are taken simultaneously.

A key feature of our implementation is the choice of simple, linear payoff functions: $\pi_{\theta_i}(x_i^2, x_j^2) = 7; \pi_{\theta_i}(y_i^t, x_j^2) = \alpha\theta_i; \pi_{\theta_i}(x_i^2, y_j^t) = 5; \pi_{\theta_i}(y_i^t, y_j^t) = \theta_i + 2$. In our two main treatments, D and N, $\alpha = 1$ and $\alpha = 1/2$, respectively.[11]

A strategy in the bandwagon game specifies an action in both stages of the game. Formally, a strategy for player i consists of a mapping from θ_i to a first stage action and a mapping from θ_i and j's first stage action (x_j^1 or y_j^1) to a second stage action.

In the unique symmetric equilibrium of the game, players use a *bandwagon strategy* governed by the thresholds $\underline{\theta}$ and θ^*.[12] These thresholds divide the players into three strategic regions: A *Stay* range $[0, \underline{\theta})$ in which players use strategy $s_1 = (x_i^1, x_i^2)$, a *Lead* range $[\theta^*, 10]$ in which players use strategy $s_2 = y_i^1$, and a *Follow* range $[\underline{\theta}, \theta^*)$ in which players use strategy $s_3 = \left(x_i^1, (x_i^2|x_j^1; y_i^2|y_j^1)\right)$.[13] Notice that players who use strategy s_3 condition their second stage action on the first stage action of their match. In addition to the bandwagon thresholds, we identify the thresholds $\bar{\theta}$ and θ°. A player with a type below θ° prefers a joint choice of platform x while a player with a type above θ° prefers a joint choice of platform y. The threshold $\bar{\theta}$ defines the region in which *Lead* is a dominant strategy; a player with a type greater than $\bar{\theta}$ prefers a unilateral choice of *Go* rather than a joint decision to *Stay.* This threshold is$\bar{\theta} = 7$ in the D treatment but 10 in the N treatment. This means that no players in the N treatment have a dominant strategy to *Lead.*

[11]The parameterizations of the treatments and the associated predictions are presented in table 1. Linear payoff functions are a special case. The equilibrium described exists and is unique for any $\pi_{\theta_i}(y_i^t, y_j^t)$ and $\pi_{\theta_i}(x_i^2, y_j^t)$ that are continuous and strictly increasing in θ, as long as $\pi_{\theta_i}(y_i^t, y_j^t) - \pi_{\theta_i}(x_i^2, x_j^2)$ is monotone in θ.

[12]This is an *ex-ante* strategy. It prescribes a strategy for each possible type realization; that is, prior to that agent observing their type. We may refer to the strategy after the agent observes their type as an *interim* strategy. In a slight overburdening of language, we also refer to this as a strategy.

[13]There are only three combinations of actions that need to be compared: s_1, s_2, and s_3. We may disregard strategies (x_i^1, y_i^2) and $\left(x_i^1, (y_i^2|x_j^1; x_i^2|y_j^1)\right)$ as they are dominated by other combinations of actions. In the bandwagon equilibrium, when a player chooses s_3 it induces a joint switch to platform y if their match has a type in the range $(3, 6]$. In contrast, when a player chooses strategy (x_i^1, y_i^2), the player forgoes the opportunity to induce a joint switch to y. Moreover, $\left(x_i^1, y_i^2|x_j^1; x_i^2, x_i^2|y_j^1\right)$ can never be an optimal strategy because it guarantees that players choose different platforms.

The two strategically relevant thresholds, $\underline{\theta}$ and θ^*, are defined by indifference conditions. $\underline{\theta}$ is defined as the point of indifference between staying on x alone and a joint switch to y, that is $\pi_{\underline{\theta}}(x_i^2, y_j^t) = \pi_{\underline{\theta}}(y_i^t, y_j^t)$. Given our specification, $\underline{\theta} = 3$ in all treatments. Because players with types less than $\underline{\theta}$ prefer to stay on x rather than a joint move to y, s_1 is strictly dominant for players in the *Stay* range.

Although all players face strategic uncertainty, only those with types in the interval $\left(\underline{\theta}, \overline{\theta}\right)$ face genuine trade-offs. Players in this range have conditional best responses. In particular, players with types in the interval $\theta_i \in \left(\theta^\circ, \overline{\theta}\right)$ prefer a joint move to platform y but would choose x if they knew that their match was going to select x.[14] These are players that *ex post* regret a decision to *Lead* if their match chooses to *Stay*. These players therefore assess the expected benefits of s_2 against their best alternative strategy s_3. Such a player balances the benefits of leading, in the hope of promoting a joint move to platform y, with the cost of possibly ending up isolated on y. In the unique equilibrium of the bandwagon game, θ^* denotes the point at which a player is indifferent between leading and following. Players with types greater than θ^* should therefore *Lead*.

To compute θ^*, assume that players use the bandwagon strategy with thresholds $\underline{\theta}$ and θ^*. Beyond θ^*, players lead while players in the region $[\underline{\theta}, \theta^*)$ choose to follow. Next, let θ^* denote the point of indifference above which players use s_2 and below which they use s_3:

$$\mathbb{E}\left[\pi_{\theta^*}(s_2)\right] = \mathbb{E}\left[\pi_{\theta^*}(s_3)\right]$$
$$\mathbb{P}(\theta_j > \underline{\theta})\pi_{\theta^*}(y_i^t, y_j^t) + (1 - \mathbb{P}(\theta_j > \underline{\theta}))\,\pi_{\theta^*}(y_i^t, x_j^2) = \mathbb{P}(\theta_j > \theta^*)\pi_{\theta^*}(y_i^t, y_j^t) + (1 - \mathbb{P}(\theta_j > \theta^*))\,\pi_{\theta^*}(x_i^2, x_j^2)$$

Using the payoff functions from the D treatment (i.e. $\alpha = 1$), this equation reduces to:

$$\frac{(10 - \underline{\theta})}{10}(\theta^* + 2) + \frac{\underline{\theta}}{10}\theta^* = \frac{(10 - \theta^*)}{10}(\theta^* + 2) + \frac{\theta^*}{10}7$$
$$\theta^{*2} - 5\theta^* - 2\underline{\theta} = 0$$

The only positive root of this equation is $\theta^* = 6$. The bandwagon strategy $(\underline{\theta} = 3, \theta^* = 6)$ is thus a unique best response to itself. Moreover, any equilibrium strategy must have the threshold form: Regardless of a player's beliefs, the benefits of leading are non-decreasing in the player's type θ. As a consequence, if it is optimal for a player of type $\theta^{'}$ to *Go* in the first period, then it is optimal for any players $\theta > \theta^{'}$ to also *Go* in the first period. All types above θ^* will therefore switch in the first period, and the bandwagon strategy $(\underline{\theta} = 3, \theta^* = 6)$ completely characterizes the equilibrium. An analogous calculation establishes that $\theta^* = 7.3$ for treatment N in which $\alpha = 1/2$.

The model produces two varieties of inefficiencies, Pareto and Kaldor-Hicks. The first type of inefficiency arises when both players have types in the range $\theta^o < \theta < \theta^*$. In this range, both players use strategy s_2. Although the players would prefer a joint switch to platform y, neither player "gets the bandwagon rolling," and they remain on platform x. We refer to this as Pareto inefficiency because both players would be better off if they could arrange a joint switch to y.

The second type of inefficiency arises when the players prefer different platforms and end up on the platform that yields lower social surplus. We refer to this as Kaldor-Hicks inefficiency because agents would choose a different platform if they were able to negotiate side-payments. Given our specification of payoff functions, whenever the sum of the players' types is greater than 10, efficiency requires a joint switch to y. However, it is possible for players to remain on x even though the sum of type draws is greater than 10 and for the players to switch to y even though the sum of type draws is less than 10. The first situation will occur if a player just below θ^* meets a player just below θ°—in which case neither player initiates a switch—and the second situation will occur if one player has a type just above θ^*—and initiates a switch—while the other player follows but has a type just above $\overline{\theta}$.

[14]Symmetrically, players in the interval $\theta_i \in (\underline{\theta}, \theta^\circ)$ prefer to stay jointly on platform x but will choose y if their match chooses y.

We also investigate a version of the bandwagon game with communication (a formal analysis of the signaling equilibrium is provided in the supplementary materials). The game is identical to the basic model but with the addition of a cheap talk stage after the agents have observed their type θ_i but prior to the first stage. The message space is $m_i = \{x_i, y_i\}$. This allows players to announce their preference for one of the platforms. Importantly, this message is non-binding and the players know this.

In equilibrium players with types below θ° prefer x and will send the message x while types above θ° prefer y and will send the message y.[15] Based on this message, players can partially update their beliefs about their match's type. If both players send the same message, this indicates that they have aligned preferences. In equilibrium, this eliminates Pareto inefficiency as players that send the same message will choose the same action. In the case in which agents send conflicting signals, the agents use a threshold strategy analogous to in the bandwagon game without communication. However, relative to the basic model, the probability that a match will follow, conditional on giving signal x, is lower than the unconditional probability.[16] As a consequence, the threshold θ^* is higher in the game with communication.

3 Design and procedures

Design The design is organized around the comparison of three treatments: Two treatments without pre-play communication, D and N, and a single treatment with pre-play communication, S. Table 1 summarizes the parameterization and predictions for these treatments. The payoff functions are identical in the D, S, and N treatments except that the payoff of unilaterally switching is reduced by half in the N treatment. The fourth and fifth columns in Table 1b indicate the regions in which the best response to lead is, respectively, not dominant and dominant.

Payoffs stage one

	$\pi_{\theta_i}(x_i^2, x_j^2)$	$\pi_{\theta_i}(x_i^2, y_j^t)$	$\pi_{\theta_i}(y_i^t, x_j^2)$	$\pi_{\theta_i}(y_i^t, y_j^t)$
D	7	5	θ_i	$\theta_i + 2$
N	7	5	$\theta_i/2$	$\theta_i + 2$
S	7	5	θ_i	$\theta_i + 2$

(a)

Predictions

	Thresholds			Best Response $= y_i^1$	
	$\bar{\theta}$	θ°	θ^*	Conditional	Dominant
D	3.0	5.0	6.0	$\theta_i \in [6.0, 7.0)$	$\theta_i \in [7.0, 10.0]$
N	3.0	5.0	7.3	$\theta_i \in [7.3, 10.0]$	$\theta_i \in \emptyset$
S	3.0	5.0	6.2	$\theta_i \in [6.2, 7.0)$	$\theta_i \in [7.0, 10.0]$

(b)

Table 1: Payoffs (1a) and predictions (1b) for D, N, and S ($i = \{1, 2\}\,; t = \{1, 2\}$)

[15] It should be clear why truthful signaling is optimal: Because the message is costless, players should send the message that promotes their highest payoff.

[16] The computations are analogous to those presented for the basic model, but take into account the partial updating that results from observing the message of the match.

Our key comparison is between behavior in the D and N treatments. The strategic difference between these treatments is that $\theta^* = 6.0$ in the D treatment and $\theta^* = 7.3$ in the N treatment. The other bandwagon threshold, $\underline{\theta}$, is the same in both treatments. Because the only strategically relevant threshold that changes is θ^*, a comparison of D and N provides a clean test of the model. The basic analysis compares the behavior of players in the strategic ranges *Stay*, *Follow*, and *Lead*. The model predicts that behavior in the same range will be identical across treatments.

In the first stage, the crucial decision is whether to *Lead*. However, although the model predicts that all players with types greater than θ^* will lead, the relevance of beliefs is distinctly different in the D and N treatments. In the D-treatment, the decision to lead is dominant for players with types greater than $\bar{\theta} = 7.0$ (see column five in table 1b). In contrast, the decision to lead in the N treatment is always predicated on beliefs. Comparison of subjects in the D and N treatment thus facilitates a test of the behavioral impact of beliefs on leadership.

In the second stage, the bandwagon equilibrium predicts that players in the *Follow* range will choose the same action as their match chose in the first stage. We refer to this as a complementarity effect. In particular, due to irreversibility, when a player's match chooses y^1, this resolves strategic uncertainty in the second stage of the game. We therefore expect that leadership will powerfully determine behavior because it eliminates the need for beliefs.

Our second stage analysis compares behavior of subjects conditioned on the first stage choice of their match. Of special interest is the difference between the behavior of the subjects in the *Stay* and *Follow* ranges: When a player's match chooses to *Stay*, players in the two ranges should behave identically. However, when a player's match leads, players in the two ranges should make opposite choices. Hence, a direct measure of the complementarity effect may be computed by comparing the difference of the *Go* frequency in the presence and absence of a leader in the *Stay* and *Follow* ranges.

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. We compute overall measures of efficiency and also report the prevalence of specific varieties of inefficiency.

Moreover, in contrast to most other studies, we examine the role of complementarities in the presence of conflicts of interest. For instance, players who prefer to remain on platform x will send the signal that they intend to stay on this platform. However, if they are in the *Follow* range, they will *Go* if their match leads. This preference flipping highlights the strength of complementarities in our setting.

Experimental procedures All sessions were conducted in 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, 2004). In the D and N treatments we ran five sessions per treatment with between 16 and 20 subjects per session. These data are supplemented with one session of 20 subjects for the S treatment. 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, subjects were randomly allocated to cubicles in the lab 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. Subjects were asked to choose either shape *Circle* (that is, *Stay*) or shape *Square* (that is, *Go*). A goal was to avoid prioritizing one of the actions as a default option.[17] Sample instructions and screen shots are provided in the supplementary materials.

[17]For instance, if subjects were asked "would you like to stay with the incumbent technology or switch to the entrant technology?" this may have affected subject decisions due to connotations associated with technology or as a consequence of incumbency biases.

Each session of the experiment began with two non-paying test games in which subjects could 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.[18] In each game, subjects were matched with one other subject, their "match", according to a highway protocol. Every subject thus met every other subject once and only once.[19] In total our data consists of 1913 unique 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 (ECU). 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 is formulated in the following fashion: At the beginning of each each new game, each subject receives a private number drawn from a uniform distribution on the interval $(0, 10)$ with two decimal points of precision. This number corresponds to the subject's type θ_i. A dedicated screen is used to display this information. Thereafter, subjects observe a 2×2 matrix with their own payoffs from the four possible combinations of outcomes and a button to choose a first-stage action. The first stage concludes when both subjects in the match have made their decisions. If both subjects choose to *Go* in the first stage, they continue directly to the feedback and bypass the second stage.

The second stage begins with a screen that reveals the first stage actions of both subjects in the match. Subjects who chose to *Stay* in the first stage make a decision. If a subject's match chose *Go* in the first stage, then the subject observes a truncated 2×1 matrix in which the payoffs conditioned on the match choosing *Stay* is removed. This reflects the fact that the subject's match has committed to *Go*.

After all second stage decisions are resolved, the subjects move to a feedback screen. The feedback displays payoffs from the current game as well as a history of type draws, choices, and outcomes in all previous periods in which the subject participated.

The signal treatment S includes an additional stage between the type draw stage and the first stage action choice. In this stage, subjects choose between two messages "I choose circle" or "I choose square". Next, the message is 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. Figure 1 shows first stage behavior of subjects in the D and N treatments. On the horizontal axis is a set of twenty bins, 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 $\theta \in [0.5, 1)$, etc. Each of the bins shows the proportion of subjects in the given range who chose to *Go* in the first period. We interpret this 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.

The theoretical prediction for the first stage behavior is a step function at θ^*: In the equilibrium of the model, players with types below θ^* *Stay* in the first period while those above θ^* *Go*. For each treatment, this threshold is indicated by a stapled line. The plots in figure 1 illustrate that subjects with low types tend to *Stay* in the first period while subjects with high types tend to *Go*. Moreover, the frequency of switching increases steeply in the vicinity of θ^* in both treatments. This is consistent with the use of bandwagon strategies.

[18] Hence, in a treatment with 20 participants, each participants played 19 rounds with payoffs.

[19] This protocol eliminates certain dynamic problems, such as strategic teaching and reciprocity (see Fréchette (2012) for a discussion).

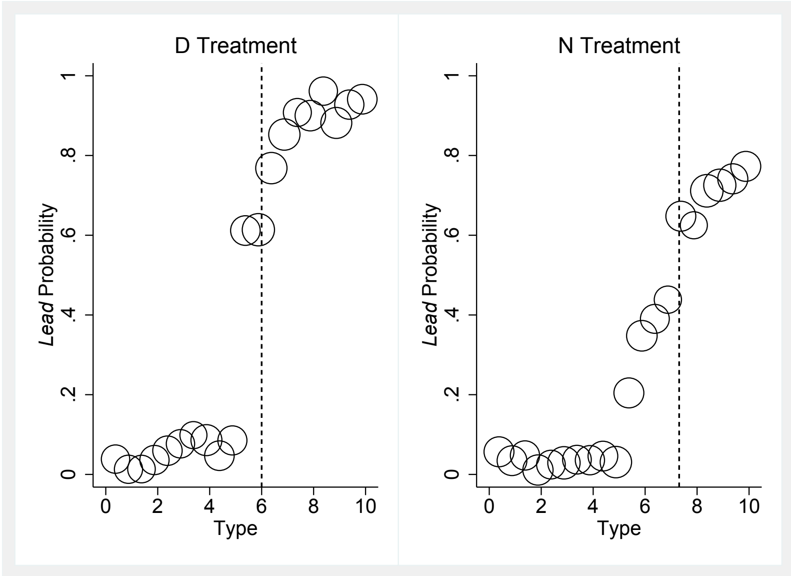

Figure 1: First Stage Behavior

To formally assess the predictions of the model, we compare behavior across treatments using Wilcoxson Rank Sum (WRS) tests. Throughout, we denote tables and figures in the supplementary materials with the prefix "S-." Using session level data, between treatment comparisons find no significant differences in behavior over the D and N treatments in either the *Stay range* or the *Follow range* (tables S1 and S2). In contrast, we reject the equality of treatments in the *Go range.* In the range $[\theta^*, 10]$, subjects are significantly more likely to *Go* in the D treatment than the N treatment (table S3).

Our results suggest that beliefs have a behavioral impact on the decision to lead. While there is a range of subjects in the D treatment for whom the decision to lead is strictly dominant, there is no such region in the N treatment (see table (1b)). Hence, beliefs matter for high types in the N treatment but not for high type in the D treatment. Although the model predicts that subjects in both treatments will lead, subjects in the D treatment with types above $\bar{\theta}$ *Go* at a much higher rate than subjects in the N treatment with equally high type draws (see table S4). Specifically, subjects *Go* 92 percent of the time in the D treatment compared to 71 percent of the time in the N treatment. We attribute this to the fact that test subjects in the D treatment have an unconditional best response that does not depend on beliefs. Furthermore, in the region in which decisions are conditioned on beliefs (*above* θ^* but *below* $\bar{\theta}$*),* we do not identify a significant difference between treatments. We can not reject equality of behavior in the vicinity of θ^*—regardless of whether implemented by a test around or just beyond this threshold (tables S5 and S6)—or if we compare the D treatment subjects in the region $[\theta^*, \bar{\theta})$ with subjects anywhere in the *Lead* range of the N treatment.[20] Thus, when beliefs are relevant, we find no difference in behavior.

We conclude that leadership is undermined when it is predicated on beliefs. When beliefs are relevant for actions, subjects tend to be more tentative and adopt a "wait-and-see" approach.[21] Below, we use the (agent form) quantal response equilibrium—in which beliefs are relevant everywhere—to rationalize observed deviations from the (subgame perfect) Bayesian Nash equilibrium of the model.

[20]This holds for a comparison of the range $\theta_i^D \in (6, 7)$ with the range $\theta_i^N \in (7.3, 10.0)$, as well as for a restricted comparison using ten base points above the *Go* treshold in both treatments to get a balanced set of data (i.e $\theta_i^D \in (6, 7)$ vs. $\theta_i^D \in (7.3, 8.3)$).

[21]Duffy and Ochs (2012) observe a similar "wait-and-see" dynamic in a study of binary entry games.

Relative to the model, subjects with types above θ^* do not choose y^1 often enough. In doing so, subjects forgo an opportunity to induce their favored platform if their match is in the Follow range.[22] This is costly. When we compare the payoff consequences a decision to *Stay* for players who have a conditional best response to *Lead,* we find that on average players in the D treatment earn 3.6 ECU less while players in the N treatment earn 1.6 less.[23]

Second Stage Behavior Figure 2 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 stayed in the first stage while the right panel presents the second stage behavior for subjects whose match chose to *Go* in the first stage. Data from the D treatment are presented as hollow bubbles and data from the N 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 are marked by vertical lines: $\underline{\theta}$, which is identical for both treatments, is identified by a short dashed black line, while θ^* is denoted by a short dashed black line for the D treatment ($\theta^* = 6$) and a long dashed gray line for the N treatment ($\theta^* = 7.3$).

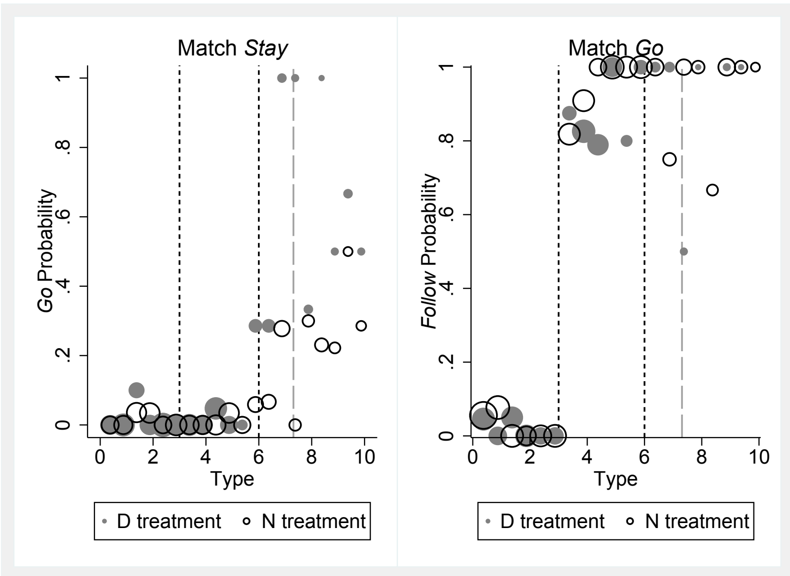

Figure 2: Second Stage Behavior by Match Action

As is evident from figure 2, complementarity has a strong effect on the eventual outcomes in the bandwagon game. When a subject in the *Follow* range has a match that stays in the first stage, the subject also stays with high probability: In 0.92 of cases in the D treatment and in 0.93 of cases in the N treatment. In contrast, when a subject in the follow range has a match that switches in the first stage, the subject also tends to switch with a high probability: In 0.90 of cases in the D treatment and in 0.87 of cases in the N treatment. Subjects in the *Follow* range essentially mirror the first stage action of their match. This is evidence of a strong complementarity effect. Furthermore, WRS tests do not identify

[22] Because types are distributed uniformly, subjects are in the *Follow* range 30% of the time in the D treatment and 43% of the time in the N treatment. The actual rates realized in the treatments were 35% and 43%.

[23] These 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.

differences in the behavior across treatments (see tables S8 and S9). This indicates that the second stage actions are dictated by the bandwagon strategy.

These effects can be neatly summarized using a logistic regression (see figure S3). We find that the strongest predictor of behavior is the interaction between range (*Stay* or *Follow*) and the first stage action of the match. In both the D and N treatments, if a player is in the follow range, the probability that the player will switch increases by about 80 percentage points if the match chooses to Go in the first stage relative to the case when the match chooses to Stay. Moreover, although type is statistically significant—thus implying that higher types Go more often regardless of other factors—this effect is dwarfed by the effect of the match's action.

Although the model predicts that all subjects with types above θ^* will *Go* in the first period, some fail to do so. These are subjects who have made "mistakes" relative to the equilibrium. Nevertheless, the optimal second stage behavior for these subjects is easy to characterize: When a subject's match chooses to *Go*, all subjects above $\underline{\theta} = 3$—including those above θ^*—should choose to *Go* in the second stage. This is due to the complementarity effect. Otherwise, if a subject's match chooses to *Stay* in the first stage, then only subjects who have dominant strategies should *Go* in the second stage.[24] Since only players in the D treatment have dominant strategies, we have distinct predictions for the D and N treatments: When a subject's match chooses to *Go,* behavior should be identical in the treatments. However, when a subject's match chooses to *Stay*, only subjects in the D treatment with types above $\bar{\theta}$ should *Go*. These are subjects who have a dominant strategy, but made a mistake in the first round.

Consistent with these predictions, when a subject's match has chosen to *Go*, subjects with types above $\underline{\theta}$ *Go* with high probability. In contrast, when a subject's match chooses to *Stay*, behavior is less clearly determined. However, consistent with the predictions, subjects with high types in the D treatment *Go* in the second stage with higher probability than analogous subjects in the N treatment.[25] This difference is consistent with subjects in the D treatment realizing that they have a dominant strategy to *Go*.

Signal Results for the signal treatment are based on a single session of 20 students. In the signal treatment, 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. We present the first-period results in figure 3.

The opportunity for pre-play communication enables players to update their beliefs about their match's type (section S.4 shows the computation of the bandwagon thresholds in this case). If both subjects signal the same action, then both subjects should choose that action in the first period. These results are shown on the diagonal. In the top left panel, we see that when subjects with types below $\theta < \theta^\circ = 5$ meet each other they nearly always signal *Stay* and then *Stay* in the first period. Similarly, on the bottom right, we see that when subjects with types $\theta > \theta^\circ = 5$ meet each other they nearly always send the signal of *Go* and then *Go* in the first period. Relative to the D and N treatments, the coordination success of subjects in the signal treatment is meaningfully higher.

On the off diagonal, we see the instances in which subjects sent conflicting signals. In this case, the subjects should play bandwagon strategies similar to the D treatment, with the exception that $\theta = 6.2$ in the S treatment. The pattern of behavior is similar to the D treatment: subjects with a dominant strategy switch as predicted (respectively $\theta < \underline{\theta}$ and $\theta > \overline{\theta}$) and there is an over eagerness for subjects just below θ^* to *Go*. Comparison of the D and S treatments suggests, therefore, that the information effect from communication improves coordination substantially whenever subjects have the same preferred platform.

[24]There are of course beliefs that can support second stage switching. However, such beliefs require a large proportion of subjects to switch in the second stage. This is not consistent with rational behavior nor with the observed switching behavior. It should therefore be difficult for such beliefs to survive.

[25]see the left panel figure 2 in which players with dominant strategies correct their mistake in the second stage even if their match choose to stay.

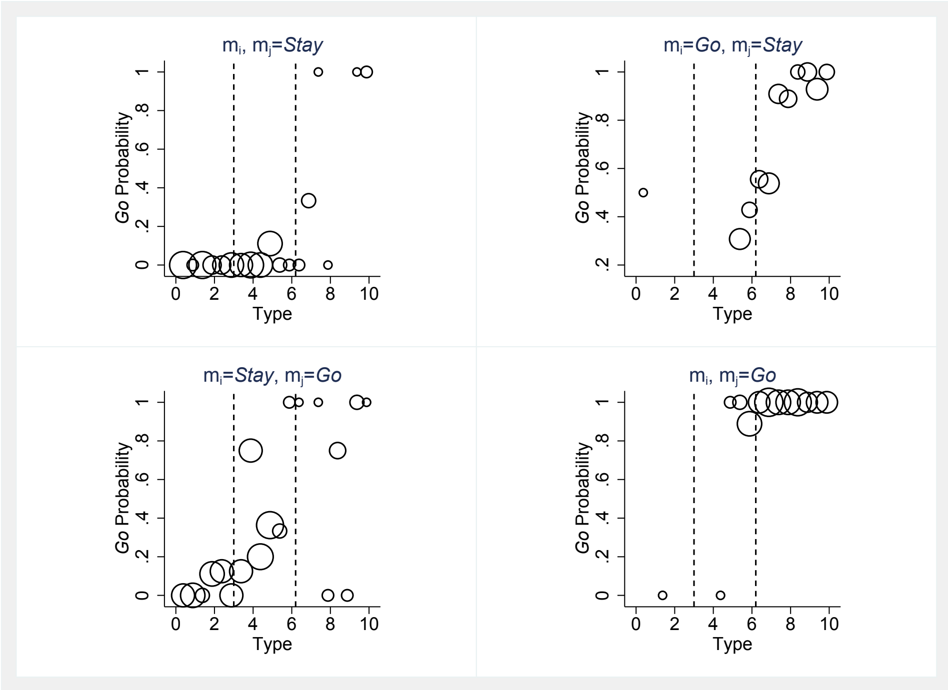

Figure 3: First Stage GO by Signal

Summary of behavior We find that behavior is indistinguishable in the *Stay* and *Follow* ranges across treatments. Behavior is also indistinguishable when *Go* is a best response but not a dominant action. However, contrary to the prediction of the model, subjects in the *Lead* ranges who have a dominant strategy switch at a much higher rate than subjects in the same range who do not have a dominant strategy. We conclude that the model organizes the data well, but the need to condition the decision to lead on beliefs has a significant and substantially negative effect on leadership.

Second stage behavior testifies to the strong influence that sequential interaction has on coordination. If a match chooses *Go* in the first period, then nearly all subjects with types above $\underline{\theta}$ choose to follow. Conversely, if neither subject in a match choose to *Go* in the first period, it is unlikely that they will *Go* in the second period.

Finally, cheap talk significantly promotes subjects ability to coordinate actions on mutually beneficial outcomes. We discuss this further in section 6.

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 bandwagon equilibrium. In particular, the pattern of errors is asymmetric in the vicinity of θ^*. By "errors," we mean deviations from the prediction of the model. For example, if the model predicts that subjects *Go* with certainty, an 80% frequency of *Go* is a 20% error rate. The asymmetry is evident in figure S2 which plots the distribution of errors along with a smoothed trend. 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 one of the platforms, rarely do. This is consistent with the core intuition for a quantal response equilibrium. We therefore estimate the agent quantal response equilibrium (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 game 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 totally mixed behavior profiles. Letting $\lambda \to \infty$ identifies a subset of the set of sequential equilibria as limiting points (McKelvey and Palfrey (1998)). Thus, when noise vanishes one is back in standard equilibrium theory. On the other hand, and for a given game, moderate noise can get amplified in an AQRE, resulting in substantial deviations from standard equilibrium theory.

We estimate the logit AQRE on 20 equally sized bins (i.e. the empirically observed switching frequency in that range) for the three decision nodes: The first stage action and two second stage stage actions that relate to 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 QRE estimates. Figure 4 presents the best fit for each treatment individually. We choose to present the individually estimated logit AQREs because the treatments are quite different, both in terms of the costs of ending on y alone (which are higher for the N treatment) and in terms of the complexity of the environment (in the D treatment the majority of the players have a dominant strategy whereas the majority of players in the N treatment have only a conditional best response).[26]

The QRE reproduces key features of the data. Crucially, it captures the asymmetry around θ^*: The AQRE correctly predicts that types just below θ^* deviate from the model to a greater extent than types just about this cut-off. The AQRE also identifies key features such as the stable level of switching for high types in the D treatment.[27] Between treatments, the AQRE correctly predicts that there should be

[26]Jointly estimated logit AQRE are presented in the supplementary materials figure S4. 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 and Kosenok (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.

[27]The flat (and even declining for high noise) *Lead* probability for players in the D treatment with high types is the outcome of the subgame structure. For players with high types, if they fail to switch in the first period, there is still a high probability of switching in the second (since they prefer y alone). The payoff consequence is therefore about the same for all players in this range: It is the size of the payoff externality from not inducing the preferred platform. This predicts similar behavior for these players. In addition, when behavior is noisy, lower types are less likely to correct their mistakes

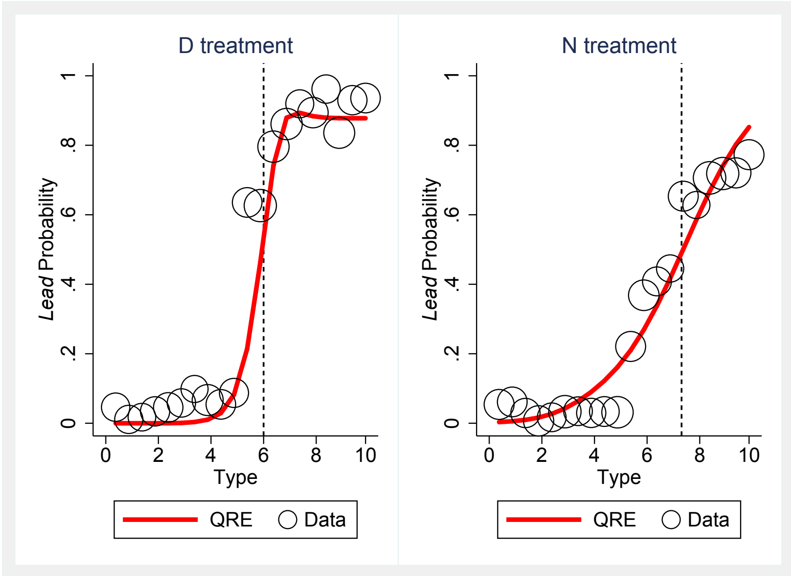

Figure 4: AQRE and Data by Treatment and Periods

a 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 modelling actual behavior.

The main feature not captured by the AQRE is the rate of *Lead* in the bin $\theta \in [5, 5.5)$. Subjects in this range *Lead* at a higher rate than predicted. We conjecture that this is due to fact that 5 defines the halfway point in the type draws and is the threshold above which a subject prefers a joint switch to y rather than jointly sticking to x.

6 Efficiency and policy

The different treatments affect the incentives and ability of subjects to achieve efficient outcomes. Figure 5 presents the empirical and theoretical efficiency of each treatment, computed as the fraction of the maximum possible payoffs.[28] The realized efficiency is highest in the S treatment and lowest in the N treatment. In addition, although players in all treatments earn less than predicted by the model, the biggest difference between expected and realized earnings is in the N treatment.

The first comparison is between the S treatment, the treatment with communication, and the D treatment. These treatments have an identical parameterization and differ only in regard to the presence of the cheap talk stage. As can be seen in figure 3, subjects in the S treatment coordinate effectively in the first period when they send the same message. But how much better off are these subjects relative to

in the second period than higher types. This can explain why for lower levels of noise it is actually types in the vicinity of $\bar{x}$ for whom a error to not *Lead* is most costly.

[28]The *maximum* payoff is computed as the payoff that would be realized if a social planner chose the subjects' actions to maximize total payoff; the maximum payoff is *nearly* identical across treatments. The *theoretical maximum* payoff is computed as the payoff that would be realized if subjects played the bandwagon equilibrium. The *empirical* payoff is the realized payoff from actual play.

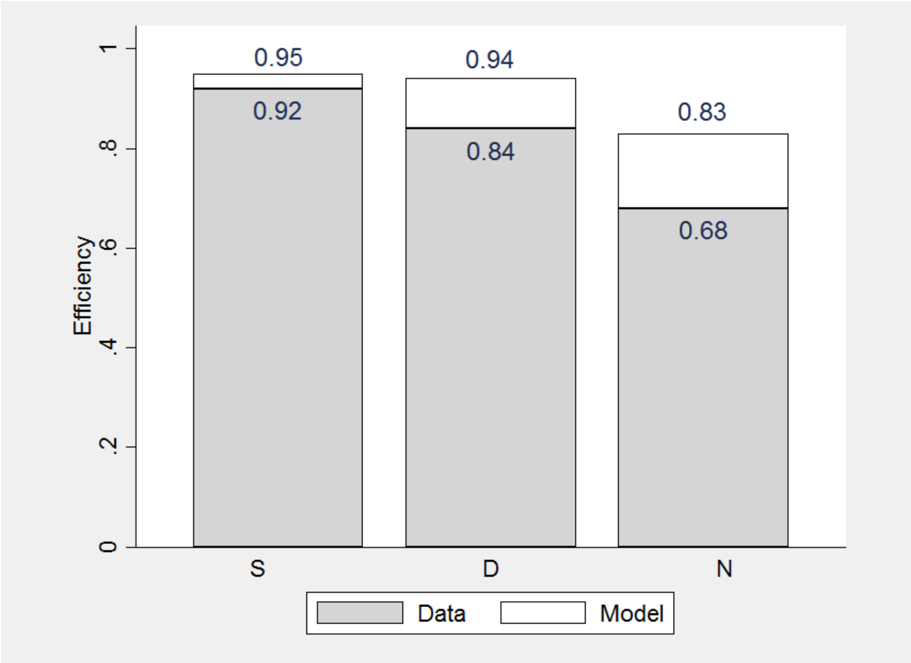

Figure 5: Efficiency

the same subjects in the D treatment? To assess this, we compare the realized efficiency of subjects in the S treatment who send the same signal with subjects in the D treatment who, *if they had an opportunity to send a message*, would send the same message. We find that the ability to send a message increases payoffs by 6 percentage points, from 89% to 95% of the maximum. Communication therefore yields a meaningful increase in payoffs.

The second comparison is between the D treatment and the N treatment. The sole difference between these treatments is that the payoff from unilaterally choosing y is halved in treatment N. Payoffs in the N treatment are thus smaller than in the D treatment because payoffs of failed leadership are reduced and because θ^* is higher and there are fewer joint moves to y. However, in addition, the realized outcomes in the N treatment are relatively less efficient than in the D treatment as measured by a proportion of the expected payoffs from the model. This appears to be due to the impact of beliefs: In the D treatment, beliefs are sharper and this facilitates joints moves to y.

A question is what scope there is for policy and whether the effects of policy interventions are different in the equilibrium of the standard model compared to the AQRE.

A conclusion from the signal treatment is that communication has a strong impact on realized payoffs. Policies facilitating communication therefore make sense. In addition, we perform a number of policy exercises to investigate the role of subsidies. In these excercises we consider a planner who faces the same information asymmetry as agents but is permitted to adjust the payoffs from the different outcomes by a fixed amount.[29] Thus, the planner can influence the incentives of agents but not pick actions on their behalf. We consider the effect of subsidies on behavior in the equilibrium of the standard model as well as in the AQRE. Specifically, we analyze two policies: Subsidizing players for leading contingent on the final

[29] We do not allow the planner to condition the size of the subsidy on type, only on the realized outcome. In addition, in the equilibrium with noise, we let the prize be awarded when a specific outcome arises—for example, the prize for a joint switch to y is paid even if it occurs due to both players switching in the second period.

outcome being (y^t, y^t), or subsidizing players for choosing y^t contingent on the outcome being (y^t, x^2).[30] We denote these as a "prize" and an "insurance," respectively.

We find that the nature of the subsidy does not matter, even in the presence of noise: Prizes and insurance have the same impact on payoffs per unit of financing.[31] Quantitatively, we find that each unit of financing increases payoffs by less than 1 and is declining as the size of the subsidy increases. In addition, we find that subsidies are relatively less effective in the presence of noise: To achieve the same increase in payoffs requires a larger expenditure when agents make mistakes. We conclude that policy makers should be wary about using such instruments because they become less cost effective in the presence of noise and as the size of the subsidies increase.

7 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 most pronounced when a leader's best response is not a dominant strategy. 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 leader switches he or she produces a strong incentive to follow for a match with moderate valuation of the challenger platform. This is because the leader resolves all uncertainty on behalf of potential followers. We find that these complementarities in actions strongly determine 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.

In terms of policy we demonstrate that cheap talk improves players' ability to coordinate actions on mutually beneficial outcomes. We also find that policies aimed at insuring failed leadership or subsidizing joint choice of the challenger platform does not promote efficiency. Furthermore, noisy decisionmaking exacerbates these problems.

[30] The criteria for "leading" is that players choose y without first observing a choice of y by their match. This allows for the case in which players switch in the second period.

[31] The policy equivalence of the different instruments is the result of expected utility maximization and uniform type draws.

References

Belleflamme, P. and M. Peitz (2015). *Industrial organization: markets and strategies.* Cambridge University Press.

Brindisi, F., B. Çelen, and K. Hyndman (2009). On the role of information and strategic delay in coordination games: Theory and experiment. *Working Paper*.

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*, 264–281.

Carlsson, H. and E. Van Damme (1993). 12 equilibrium selection in stag hunt games. *Frontiers of game theory*, 237.

Cooper, R. and A. John (1988). Coordinating coordination failures in keynesian models. *The Quarterly Journal of Economics*, 441–463.

Duffy, J. and J. Ochs (2012). Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior 76*(1), 97–116.

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 G. Saloner (1985). Standardization, compatibility, and innovation. *The RAND Journal of Economics*, 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.

Goldstein, I. and A. Pauzner (2005). Demand–deposit contracts and the probability of bank runs. *the Journal of Finance 60*(3), 1293–1327.

Greiner, B. (2004). The online recruitment system orsee 2.0-a guide for the organization of experiments in economics. *University of Cologne, Working paper series in economics 10*(23), 63–104.

Haile, P. A., A. H. and G. Kosenok (2008). On the empirical content of quantal response equilibrium. *American Economic Review 91*(1), 180–200.

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*, 587–597.

Morris, S. and H. S. Shin (2002). Measuring strategic uncertainty.

Shy, O. (2001). *The economics of network industries.* Cambridge University Press.

Turocy, T. (2010). Computing sequential equilibria using agent quantal response equilibria. *Economic Theory 42*, 255–69.