

Experiments on Belief Formation in Networks*

VERONIKA GRIMM[†]

University of Erlangen–Nuremberg

FRIEDERIKE MENGEL[‡]

University of Essex

and Lund University

March 27, 2018

Abstract

We study belief formation in social networks using a laboratory experiment. Participants in our experiment observe an imperfect private signal on the state of the world and then simultaneously and repeatedly guess the state, observing the guesses of their network neighbours in each period. Across treatments we vary the network structure and the amount of information participants have about the network. Our first result shows that information about the network structure matters and in particular affects the share of correct guesses in the network. This is inconsistent with the widely used naive (deGroot) model. The naive model is, however, consistent with a larger share of individual decisions than the competing Bayesian model, while both models correctly predict only about 25 – 30% of consensus beliefs. We then estimate a larger class of models and find that participants do indeed take network structure into account when updating beliefs. In particular they discount information from neighbours if it is correlated, but in a more rudimentary way than a Bayesian learner would.

JEL-classification: C70, C91, D83, D85.

Keywords: Networks, Learning, Belief Formation, Opinion Dynamics, Experiments.

*We thank Arun Chandrasekhar, Andrea Galeotti, Jayant Ganguli, Ben Golub, Emir Kamenica, Willemien Kets, Christophe Muller, Santiago Oliveros, Juuso Välimäki, Andriy Zapechelnuk, two anonymous reviewers and seminar participants in Amsterdam, Berlin (Workshop on expectations and markets 2014), Bocconi, Brussels (CTN workshop 2014), Carlos III Madrid, Cergy Pontoise (workshop on Dynamics and Beliefs 2016), Essex (2nd European Networks meeting 2014), Glasgow, JRC Ispra, Jena (2013 socio-economic committee of the VFS), Lund, Norwich (2016 workshop on Bounded Rationality and Games), Paris School of Economics, Rome (Workshop on Economics of Limited Cognitive Resources 2015), Royal Holloway University London, Toulouse (EEA 2014) and Zuerich (ESA 2013) for helpful comments. We also thank Sandra Miltenyte, Michael Seebauer and Alexander Schneeberger for excellent research assistance. Financial support by the NWO (VENI grant 451-11-020) and the *Emerging Field Initiative* at FAU Erlangen-Nuremberg is gratefully acknowledged. This version substantially extends the working paper version published on SSRN (2013), 2361007.

[†]FAU Erlangen–Nuremberg, Lehrstuhl für Volkswirtschaftslehre, insb. Wirtschaftstheorie, Lange Gasse 20, D-90403 Nürnberg, Germany, Tel. +49 (0)911 5302-224, Fax: +49 (0)911 5302-168, email: veronika.grimm@wiso.uni-erlangen.de

[‡]Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ (UK); Department of Economics, Lund University, SE-220 07 Lund (SE); *e-mail*: fr.mengel@gmail.com

1 Introduction

Most social and economic interactions are shaped by beliefs and opinions. The most simple example are everyday consumption choices, but also investment choices depend on beliefs about future returns and political choices - like which candidate to vote for - are shaped by our beliefs and opinions about the “right” course for policy or the “right” candidate. None of these beliefs are formed by decision makers in isolation. Instead people typically communicate with others in their social network and take their experiences and opinions into account. Models of belief diffusion in social networks have been used to understand a wide range of phenomena across different areas of economics (Gale and Kariv, 2003; Golub and Jackson 2010, 2012; Mueller-Frank, 2013; Buechel et al, 2015). Among others, these models have been used to explain the emergence of political polarization (Baldassari and Bearman, 2007), but also consensus in political opinions (Katz and Lazarsfeld, 1955), to study technology adoption (Banerjee et al, 2013), compliance with the law (Drago et al, 2014), formal and informal insurance (Ambrus et al, 2014; Cai et al, 2015) or financial decisions (Hong et al, 2005; Bursztyn et al, 2014).

Existing literature has mostly relied on either of two models: Bayesian learning (Gale and Kariv (2003), Acemoglu et al (2011), Mueller-Frank (2013)) or a naive model due to deGroot (1974).¹ One of the key properties of “naive learning” is that agents completely ignore the structure of the underlying social network, while it is fully taken into account by Bayesians. Given the widespread use of these two benchmark models in the social networks literature, surprisingly little is known about how people actually form and update beliefs in networked communication and whether or not the predictions of these models match actual outcomes well.²

In this paper we provide a comprehensive experimental study (using different networks and information conditions) of social learning in networks. Our first aim is to test whether people use information about the network structure when updating beliefs, as this is one of the key differences between the naive model and more sophisticated models. We also compare how well outcomes are predicted by the Bayesian and naive model. We are interested in whether either of these models is a good predictor of behaviour in these environments in the usual “as if” sense. We are not interested in whether people are Bayesian or naive in an epistemic sense.³ Finally, we study properties of heuristics participants use to see which features of network structure participants pay attention to.

At the beginning of our experiment, participants observe an imperfect private signal about the state of the world which could be either of two colours, say “black” or “white”, with equal probability. They then simultaneously submit a binary guess about the state of the world. In all subsequent periods they observe the guesses made by their network neighbours in the previous period and submit another guess themselves. This process continues for 20 periods and is repeated 6 times with different colours (and new draws by nature).

The experiment involves a total of eleven treatments. In the initial experiment, we set up nine treatments in a 3×3 design. The first dimension varied was the network structure: we used the circle, the star, and a kite. Under the star information aggregation is centralized: one agent observes all others. In the circle information aggregation is decentralized: all agents observe some others and

¹In the literature, naive learning has also been referred to as average based updating (Golub and Jackson, 2012), best response dynamics, boundedly rational learning (de Marzo, Vayanos and Zwiebel, 2003) or myopic learning (Acemoglu and Ozdaglar, 2011). It is referred to as naive learning by e.g. Golub and Jackson (2010). We follow Golub and Jackson (2010) and refer to the deGroot (1974) dynamics as “naive learning” or “naive updating”.

²Literature is surveyed in detail below.

³These are different aims, since people might, for example, develop heuristics that induce behaviour “as if” they are Bayesian, despite not reasoning in a Bayesian manner. Moreover, there is already quite some evidence on the latter question in other environments including evidence showing that learning rules can differ across different environments (Kovarik et al, 2017).

all observe equally many agents. The kite is intermediate and was chosen because of the theoretical predictions it generates. Theoretical predictions of the Bayesian and the naive model differ across these networks both in terms of whether a consensus is reached and whether the truth is learned. The second treatment dimension varies information about the network structure. We study three information conditions: No Information (NI), Incomplete Information (II) and Complete Information (CI). The different information conditions in our design provide a clean test of one of the key tenets of the naive model: that agents ignore the network structure.

Our results show that this is not the case. The amount of information participants have about the network does affect the share of correct guesses as well as other outcomes. More information about the network structure also leads to lower consistency of individual choices (conditional on signal realizations and the history of neighbours' guesses) with the predictions of the naive model. Still, even with complete information about the network the naive model predicts individual choices better than the Bayesian model. The picture is less clear when it comes to consensus predictions. Both models predict around 25 – 30% of consensus outcomes correctly.

We then estimate learning models from a larger class. The heuristics our participants use are not too dissimilar from the naive model, but there are some crucial differences. In particular, participants do react to network structure. They place higher weight on themselves the higher their clustering coefficient and lower weight on neighbours whose clustering coefficient is high.⁴ This can be seen as a simple way to discount information from neighbours if this information is likely to be correlated (which is the case if the neighbours are neighbours themselves).

We derive an adjusted rule that deviates from the naive model in only one respect, namely in that it accounts for an agent's clustering coefficient. This rule can have fundamentally different implications than either Bayesian or naive learning. In particular, persistent disagreements are more likely under the adjusted rule than under either of the other models.⁵ We generate new experimental data in networks, where the naive, the Bayesian and the adjusted model all yield different predictions. The adjusted model is more consistent with the data than either the naive or Bayesian model.

This paper is organized as follows. In Section 2 we discuss related literature. In Section 3 we explain the experimental design. We discuss the theory in more detail and develop conjectures in Section 4. Sections 5-6 contain our results and Section 7 concludes.

2 Related Literature

We start by describing the two types of models that have been used in theoretical literature and then discuss related experimental work.

Bayesian Learning One of the dominant paradigms for modeling social learning in networks is Bayesian learning (Banerjee, 1992; Bikhchandani et al, 1992; Gale and Kariv, 2003; Acemoglu et al, 2011; Eyster and Rabin, 2010). Particularly relevant for our context are models where neighbours in a network, after initially receiving a private signal about the state of the world, *repeatedly* and *simultaneously* make guesses about the state observing their network neighbours guesses in each round. This is the setting we consider in this paper. A number of authors have shown that, under certain conditions, Bayesian learners will converge in their beliefs as to which state of the world they

⁴The clustering coefficient of an agent is the share of her first order neighbours that are neighbours themselves.

⁵If beliefs are communicated on a fine enough grid, then - as long as networks are connected - a consensus will be reached under all models. If the grid is coarser, in particular also if only actions or binary guesses are observed, as in this setting, then the adjusted model can lead to disagreement whenever clustering coefficients are high. The reason is simply that agents under the adjusted model will then place a very high weight on their own opinion.

are in and that furthermore, under additional conditions, these belief will be correct (Gale and Kariv, 2003; Mossel and Tamuz, 2010; Mueller-Frank, 2013; Mossel et al 2015).

The naive model A second class of models that has received increasing attention in recent years is based on the naive model first proposed by de Groot (1974). Under the naive model agents update beliefs in each period by taking weighted averages of their own and their network neighbours’ past beliefs. A key distinction between the naive and Bayesian models is that under naive learning agents do not make use of information about the network structure. This is one of the features of the naive model that we will test. A number of papers have established conditions under which beliefs converge in the naive model (de Groot, 1974; deMarzo et al 2003). deMarzo et al (2003) also show, under some assumptions on the updating process, that each agent’s influence is proportional to the number of direct neighbors she has, i.e. to her degree. Golub and Jackson (2010) ask not only whether a consensus will be reached, but whether agents will converge to the truth. They show that all opinions in a large society ($n \rightarrow \infty$) converge to the truth if and only if the influence of the most influential agent vanishes as the society grows.⁶ Hence, in large networks and when continuous beliefs are communicated the Bayesian and naive model tend to yield similar predictions. This is, however, not the case in some of the small networks we consider and under binary beliefs (actions). Acemoglu and Ozdaglar (2011) review some of the literature on both the Bayesian and the naive model.

Experimental Work Our paper contributes to an active experimental literature which tests this and related theory in the laboratory.⁷

Maybe the first experiment conducted on these questions is Mobius, Phan and Szeidl (2015), who study learning and belief formation in *endogenous* networks in a field experiment using the Facebook connections of Harvard undergraduates. They compared the naive model and a Bayesian model (based on Acemoglu, Bimpikis and Ozdaglar, 2014), where agents tag (link to) the source of information. They find that there is social learning, but information transmission is noisy and imperfect. When accounting for the fact that information transmission is stochastic in their setting they find some evidence for the tagged model. It is not too surprising that they find more support for the Bayesian model than we do given that the model predictions are much more straightforward in the endogenous network setting they consider.

A number of authors have tested for either or both models on exogenous networks. Corrazzini et al (2012) are focused on the naive model. They study a version of the naive model in an experiment where agents’ in-degree (how many people they observe) and out-degree (how many people they are observed by) differ. They find support for a variant of the naive model according to which social influence is proportional to an agent’s in-degree. In another study on directed networks Brandts et al (2015), however, do not find evidence for the importance of in-degree. Both of these studies test for the naive model indirectly by comparing theoretical consensus predictions to outcomes. Further, in contrast to our setting, both these studies consider *directed networks* and they study only the complete information case. Varying the amount of information about the network structure provides arguably the cleanest test of the naive model.

⁶Golub and Jackson (2012) show that homophily (a tendency of similar agents to be linked) slows down the speed of learning and hence increases the time it takes to reach a consensus. Jadbabaie et al. (2012) study a model where agents take their personal signals into account in a Bayesian way, but account for information from their neighbours in a naive way. They show that in this case agents always learn the truth. Acemoglu, et al (2010) study a version of the naive model where some “forceful” agents do not change their opinions. They study how misinformation can spread in social networks in these cases.

⁷To allow the reader to better track the history of the different experiments discussed below we provide some information on when the different experiments were conducted (whenever available). Mobius, Phan and Szeidl conducted experiments in 2004, Corrazzini et al in 2009, Chandrasekhar et al in 2010 and Mueller-Frank and Neri in 2012.

Chandrasekhar et al (2015) test the naive and the Bayesian model. They also study mixed populations allowing agents to have incomplete information about whether others are Bayesian or naive. Despite this flexibility, individuals are best described by the naive model with identical weights in their experiment. Choi et al (2012) focus uniquely on the Bayesian model studied in Gale and Kariv (2003) in three-player networks. They find that the Bayesian model fits the data quite well in these networks. Three-player networks however lack statistical power to distinguish between the naive and Bayesian model. In fact, in all of the networks they consider there are virtually no differences between the predictions of the two models (Chandrasekhar et al, 2015). Mueller-Frank and Neri (2014) conduct an experiment and show that agents rarely reach a consensus. They provide a theoretical explanation for this fact which emphasizes the role of heterogeneity.⁸ All these studies only consider the case of complete information about the network structure. The introduction of the incomplete and “no information” cases are one of the main contributions of this paper to the existing literature. While other literature on exogenous networks (particularly Chandrasekhar et al, 2015) has shown that the naive model tends to explain behaviour better than the Bayesian model, our evidence shows that one of the fundamental tenets of this model does not hold. Participants do use information about the network structure.

3 The Experimental Design

In this section we describe our experimental design. In all treatments, participants interacted in a network consisting of seven participants for six rounds of 20 periods each. The treatments differed in two dimensions: we varied the network structure (circle, star, kite, rectangle and pentagon, see figure 1) and the amount of information about the network structure that was available to participants (no info, incomplete info, or complete info). Each group of seven participants interacted in the same network structure during the six rounds of our experiment. However, participants’ positions within the network were randomized across rounds. Each **round** had the following structure:

- (1) Participants received information on the number of neighbors and (depending on the treatment) additional information about the network structure. In all treatments players were assigned “labels” (anew at the beginning of every round) so that they could follow the history of guesses of particular neighbors.
- (2) Nature drew one of two possible states $\omega \in \{B, W\}$, with commonly known probability $\frac{1}{2}$. Each state of nature represented an urn. Urn B contained four black balls and three white balls. Urn W contained four white balls and three black balls.⁹
- (3) Participants observed a private signal. If urn B was drawn, four players received a black ball and three players received a white ball. If urn W was drawn, four participants received a white ball and three participants received a black ball. Thus, if participants knew all the signals, they would know the urn for sure. In other words the distribution of signals always reflected the exact composition of the urn.¹⁰

⁸Somewhat less related are several experimental papers studying herding and informational cascades (see Anderson and Holt (1997) or Weizsaecker, 2010 for a meta-study). These settings are very different from the one studied here in that in information cascades people choose once and sequentially, while in the setting studied here they choose simultaneously and repeatedly. Also most of the cascades literature does not study networks other than the line.

⁹In the experiment we changed the colours of the urns across rounds (see Table 1). Throughout the paper we refer to black and white urns for clarity of exposition.

¹⁰Hence, unlike in much of the theoretical literature, signals in our setting are dependent. Independence is often a simplifying assumption in theory, but there is ample evidence that people don’t understand independent draws very

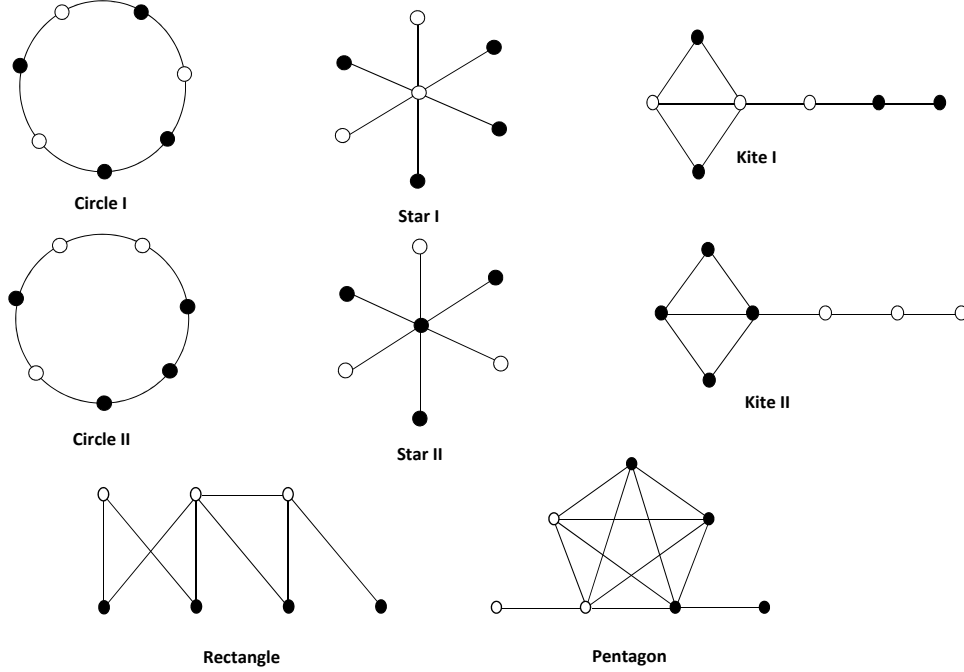


Figure 1: The Experimental Networks and Signal Distributions.

- (4) Participants had to guess the correct urn repeatedly for 20 **periods**. Each of the 20 periods consisted of two steps: First, participants stated a binary guess, B or W .¹¹ Second, all participants observed the guesses of their direct network neighbors.

We chose 20 periods of communication since the theoretical models predict convergence after at most 9 periods (see Section 4) and we wanted to give participants some more time than that. 6 rounds were chosen in order to observe mature behaviour after participants had some time to become familiar with the task.

Networks Figure 1 shows the network architectures we used. We chose networks of 7 players, because with fewer players learning about the correct urn becomes increasingly trivial and distinguishing between the naive and Bayesian model becomes impossible. Three-player networks, for example, such as those used by Choi et al (2012), lack statistical power to distinguish between the

well (see e.g. Kahnemann and Tversky, 1972). One reason for using dependent signals is to avoid biases that result from participants' incomprehension of the setting. Another advantage of using dependent signals is, as mentioned above, that we can make sure that the realized draw reflects exactly the urn composition. With independent draws there is a chance that learning is not possible.

¹¹Given that most of the theoretical literature on the Bayesian model has focused on the action (binary) setting, while most of the literature on the naive model has focused on the belief (continuous) setting, we had to make a choice here. There are four reasons we decided for a binary guess. First, in light of ample evidence that people have difficulty in communicating and reasoning about probabilities (see the research summarized in Bazerman and Moore, 2009), we decided to let participants state a binary guess instead of a probabilistic statement like “I believe the urn is white with probability 0.85”. It was important to us to minimize confusion about the environment and the task of guessing the right urn. Second, the setting with binary communication lends itself better to applications where only choices are observable. Third, the theoretical predictions of the models differ more often in this setting. (In fact with continuous signals a consensus will always be reached under both models in all networks considered (see Mueller-Frank, 2013).) And fourth, predicted convergence times are shorter.

naive and Bayesian model (Chandrasekhar et al, 2014). Networks with even more than 7 players, on the other hand, would have made learning unnecessarily hard. We chose these particular networks because the star and circle capture benchmark situations. All information is aggregated in one player in the star network, while it is evenly spread in the circle. We added the Kite network because it generates differential theoretical predictions across signal distributions, treatments and theoretical benchmark models (see Section 4). The Rectangle and Pentagon were only used in our second set of experiments and were chosen because of the theoretical predictions they entail.

Information. Treatments also differed with respect to the information our participants received about the network structure. Varying information conditions provides a clean test of the naive model, as under this model it should not matter how much information about the network structure participants have. We implemented three information conditions for each network, as follows:

- (NI) “No Information”: Players knew that there are 7 players connected in a (single-component) network as well as the number and “labels” of their first-order neighbors (i.e. their own degree). They received no other information about the network structure.
- (II) “Incomplete Information”: In addition to the information received in the NI treatments, players knew the *degree distribution* of the network. That is, all players knew how many players had how many neighbors. Note that this is equivalent to knowing the complete network structure for the Circle and the Star network, but not for the Kite, Rectangle or Pentagon networks.
- (CI) “Complete Information”: In addition to the information provided in the II treatments, players were shown a complete graphical representation of the network both in the instructions and on the screens.

Signal distributions and colours. Throughout the experiment each group (of seven players) interacted within the same network, but signal distributions and network positions changed across rounds. As network positions change across rounds, players don’t interact with the same neighbours too many times. One consequence of this is that it is less likely that they develop sophisticated conventions on when to experiment etc (see also footnote 18).

Moreover, we used different colors in every round in order to make it clear to participants that observations from previous rounds are not informative with respect to the right guess in a current round. We did also not tell them the true colour of the urns in previous rounds until the end of the experiment. With these design features we hoped to get mature decisions in later rounds, while at the same time minimizing the possibility of undesired spillovers across rounds.

For each network we used two fixed signal distributions, as illustrated in Figure 1. The reason to do so was to generate enough “interesting” observations, i.e. observations where (i) the learning problem is not completely trivial and (ii) where the theoretical predictions of the models differ in interesting ways, without having to run prohibitively many sessions. Participants did not know that we used different signal distributions across rounds, nor which they were. All they were told was that each participant in the network would receive one ball from the urn w/o replacement. Using the Laplace principle of insufficient reason this would lead participants to believe that the signal distribution is a random draw from a total of 35 permutations ($\frac{7!}{4!3!}$).¹² Table 1 provides details on the assignment of signal distributions and ball colors in all main treatments.

¹²If we told them the two distributions, then the problem would be trivial for several network positions, as participants could infer the signal distribution and the colour of the urn from their own signal.

Network	Round and colors					
	1 red/blue	2 green/orange	3 black/white	4 violet/yellow	5 brown/turquoise	6 grey/pink
1	1	2	2	1	2	1
2	1	1	2	1	2	2
3	2	1	2	2	1	1
4	2	2	1	2	1	1
5	1	2	1	2	1	2
6	2	1	1	1	2	2

Table 1: Signal distributions (1 or 2), urn colours and networks across rounds.

Payments. At the end of the experiment, for each player independently, we randomly selected three periods from different rounds. For each selected period the participant received Euro 6 if their guess was correct and nothing otherwise. In addition to the performance-dependent payments they received a show up fee of Euro 6. Hence participants could earn either Euro 6, 12, 18 or 24 in the experiment. On average participants earned approximately Euro 17 (all included).

Questionnaire. After the experiment participants completed an extensive questionnaire, covering emotional intelligence, cognitive reflection, as well as numeracy skills.¹³ We did not provide material incentives for correct answers in the questionnaire but emphasized that the relatively high show up fee should compensate for the additional time.

New Experiments Based on the insights from the main treatments we derived an adjusted rule that we then decided to test against new data. We hence conducted two additional treatments using the Rectangle and Pentagon networks shown in Figure 1. Table 2 summarizes our eleven treatments.

	No Info (NI)	Incomplete Info (II)	Complete Info (CI)
Star	S_NI	S_II	S_CI
Circle	C_NI	C_II	C_CI
Kite	K_NI	K_II	K_CI
Rectangle			R_CI
Pentagon			P_CI

Table 2: Treatments. In each treatment we have 5040 observations, which stem from 42 individuals observed across 120 rounds, and 6 independent observations (networks).

Further Details. The experiment took place in 2012-2013 (Star, Circle, Kite) and 2014 (Pentagon, Rectangle) at the Laboratory for Experimental Research Nuremberg (LERN). In total, 462 students from FAU Erlangen-Nuremberg participated in 22 sessions. Each session contained 3 networks of the same treatment, which can be considered independent in the absence of “static session effects” (Frechette, 2012).¹⁴ All experimental sessions were computerized.¹⁵ Written instructions were distributed at the beginning of the experiment.¹⁶ Sessions lasted between 67min-109min (including reading the instructions and answering the post-experimental questionnaire).

¹³See Appendix F for the complete set of questions.

¹⁴In other words each session contained three groups of 7 participants organized in a network. Participants never interacted across networks. Static session effects could lead to undesired within session correlation if e.g. experimenters behave differently across sessions etc.

¹⁵The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). Subjects were recruited using the Online Recruitment System ORSEE by Greiner (2004).

¹⁶The instructions for treatment K_CI, translated from German into English, can be found in Appendix I. Instructions for the remaining treatments are available upon request.

4 Theoretical Background and Research Questions

This section contains the theoretical background and the research questions we want to address with our experimental design. In Section 4.1 we derive theoretical predictions for all our networks and signal distributions. In 4.2 we discuss our research questions.

4.1 The Bayesian and the Naive Model

We index agents $i = 1, \dots, 7$ and denote i 's set of network neighbors by \mathcal{N}_i . There are two states of the world B and W according to whether the black or white urn was drawn. Agents have a common prior of $\frac{1}{2}$ on each of these states and this is common knowledge. The fact that each urn is drawn with probability $\frac{1}{2}$ is explained in the experimental Instructions which are common knowledge (read aloud at the beginning of the experiment). At the beginning of the first period all agents receive a signal $s_i \in \{0, 1\}$. If the state is W , then four agents receive signal $s_i = 0$ and three agents receive signal $s_i = 1$. If the state is B , then four agents receive signal $s_i = 1$ and three agents receive signal $s_i = 0$. Agents believe that - conditional on the state - who receives signal $s_i = 0$ and who receives $s_i = 1$ is random (see the paragraph "Signal distributions and colours" in Section 3 for a discussion of this assumption). In each period t all agents simultaneously submit guesses $g_i^t \in \{0, 1\}$ about the correct urn. We denote the vector of guesses of all agents in the network by $\mathbf{g}^t = (g_1^t, \dots, g_7^t)$.

The Bayesian Model. The Bayesian model requires assumptions on agents' priors as well as on their theory about how others behave. In line with the theoretical literature we assume that there is common knowledge of Bayesian rationality. Without this (or a similar assumption) Bayesian learning is ill-defined. Bayesian agents use their knowledge of the network, their private signal as well as the history of their own and their neighbours' guesses to update their belief in each period using Bayes rule. They choose $g_i^t = 1$ whenever their posterior is strictly above $\frac{1}{2}$ and $g_i^t = 0$ if it is strictly below $\frac{1}{2}$. Indifferences are resolved probabilistically. Note that Bayesian learning requires all agents in every period to consider the set of possible information sets of all other agents and how communication impacts the information sets of their neighbors in the subsequent period. Since in our networks the history of beliefs is not common knowledge among neighbours this quickly becomes a complex task. While this makes it seem unlikely that participants reason through the Bayesian model in its full complexity, it seems plausible that participants adopt heuristics that lead to decisions that are Bayesian in an "as if" sense. Our experiment is designed to evaluate the latter conjecture. Theoretical predictions of the Bayesian model for the CI case are summarized in Table 3 and derived in Appendix A.^{17,18}

The Naive Model. Under the naive model agents simply follow the majority. More specifically we denote by $G_i^t(\mathbf{g}^{t-1}) = \frac{g_i^{t-1} + \sum_{j \in \mathcal{N}_i} g_j^{t-1}}{|\mathcal{N}_i| + 1}$ the average of i 's and i 's neighbours past guesses. Agents

¹⁷While under CI participants observe the network structure, one may ask how participants account for the network structure under II and NI. For the circle and star networks the degree distribution (communicated in the II treatments) reveals the complete network structure, while the same is not true for the other networks. Hence, in these networks as well as in the NI condition, some assumption is needed on agent's prior over networks. Since there is no "natural" assumption for such priors we refrain from making theoretical predictions for the Bayesian model under the II and NI conditions.

¹⁸One could also consider an alternative model where Bayesian learners are not myopic, but instead choose suboptimal actions in a given round in order to extract more information from a neighbour (exploration). We do not study this model here, but it should be noted that our setting by and large does *not* offer incentives to explore. This is most obvious in the Star network where it is clearly impossible to learn faster or more by choosing suboptimal actions, as - under truthtelling - the center of the Star will know the truth for sure in period 2 (see consensus times in Table 3).

then follow this rule

$$g_i^t = \begin{cases} 0 & \text{if } G_i^t(\mathbf{g}^{t-1}) < \frac{1}{2}, \\ 1 & \text{if } G_i^t(\mathbf{g}^{t-1}) > \frac{1}{2}. \end{cases} \quad (1)$$

Here g_i^t denotes player i 's guess at time t and $|N_i|$ the cardinality of player i 's network neighborhood, i.e. her degree or the number of other players i observes excluding herself. Indifference, i.e. $G_i^t(\mathbf{g}^{t-1}) = \frac{1}{2}$ is resolved by the flip of a fair coin. Naive agents hence completely ignore the network structure when making their decisions and hence will ignore the fact that information received from two different neighbours might be correlated (see e.g. Enke and Zimmerman, 2015). Note that, as it is irrelevant how much information naive learners have about the network, the prediction of this model is the same across all three information treatments (CI, II and NI). Another property of the naive model is its “forgetfulness”. Since it takes into account only current beliefs, errors made in the past will not be recognized or accounted for. As a consequence few errors can potentially lead to very different learning outcomes in this model. Note also that $G_i^t(\mathbf{g}^{t-1})$ attaches the same weight to i 's own past guess and i 's neighbors' past guesses.¹⁹

Summary. Table 3 summarizes the theoretical predictions for the naive and Bayesian models, respectively. We ask three questions: (i) is a consensus reached? (ii) if so, do agents agree on the correct urn? and (iii) how many periods does it take to reach a steady state where no agents change their guesses anymore (convergence time)?²⁰

	<i>Naive Model</i>					
	Circle1	Circle2	Star1	Star2	Kite1	Kite2
Consensus Reached?	Yes	No	Yes	Yes	No	No
Correct Urn?	Yes	-	?	Yes	-	-
Convergence Time	4	2	≥ 2	≥ 3	2-3	1

	<i>Bayesian Model (CI)</i>					
	Circle1	Circle2	Star1	Star2	Kite1	Kite2
Consensus Reached?	Yes	Yes	Yes	Yes	Yes	Yes
Correct Urn?	Yes	Yes	Yes	Yes	No	Yes
Convergence Time	3	≥ 6	3	3	5	9

Table 3: Theoretical Predictions for all Treatments and all signal distributions. Note that the prediction for the naive model is independent of the information condition while for the Bayesian model predictions are given only for the CI case due to the multiplicity of admissible priors in the NI and II conditions. A ? should be read to say that the prediction is open - consensus could be on the correct or on the wrong urn with positive probability

4.2 Research Questions and Conjectures

In this subsection we state our research questions and outline how we will address them. Our first question asks whether people use information about the network structure when updating beliefs.

¹⁹DeMarzo, Vayanos, and Zwiebel (2003) show that the same theoretical predictions regarding whether a consensus is reached would hold as long as players attach symmetric weights to each other, but convergence time might be different. All exact derivations can be found in Appendix A.

²⁰Readers familiar with the naive model (e.g. Theorem 1 in deMarzo, Vayanos and Zwiebel, 2003) might wonder how it is possible that agents do not reach a consensus in some of our networks (Circle-2 and Kite). The difference lies in our binary communication structure. Since in our setting agents only communicate choices (or binary beliefs) it is possible that choices stop converging even if the network is connected.

Q1: Information *Do people use information about the network structure when updating beliefs?*

To address this question we will use our treatment variations CI, II and NI and ask whether outcomes differ across these three variations. If people do *not* make use of information about the network structure, as posited by the naive model, then how much information they have about the network is irrelevant. In this case we would not expect to see any differences in outcomes across the three information conditions. The outcomes we will focus on in addressing this question are (i) the share of correct guesses and (ii) the frequency of switches as measures of the quality and speed of learning.

Our second question focuses on the theoretical predictions derived in Section 4.1

Q2: Outcomes *How consistent are outcomes with the predictions of the Bayesian/Naive models, respectively?*

The answer to this question will help us evaluate how useful either model is for predicting behaviour. Even if neither model describes well how participants reason, it is still possible that the models yield quite accurate predictions. To assess this possibility we will focus on the following outcomes: (i) the share of individual choices consistent with the model predictions, (ii) the consensus beliefs contrasted with the predictions from Table 3 and (iii) the actual share of correct guesses compared to the theoretically predicted share,

The share of individual choices consistent with model predictions evaluates the models' fit at the greatest level of disaggregation. Conducting this analysis, however, requires additional assumptions especially when it comes to the Bayesian model. Remember that the Bayesian model assumes common knowledge of Bayesian rationality. In social learning environments one cannot dispense with this assumption when testing the model's explanatory power, since without a theory about how others reason Bayesian learning is not well defined. Since this assumption is used in theoretical models used to predict behaviour it is also appropriate to maintain it.²¹ This assumption, however, can lead to some problems when contrasting the model with data, because it might be contradicted by observation. In the Star treatments, for example, if the center switches at any time $t \geq 3$, then either the center cannot be Bayesian or she must believe that one or several of the spokes are not Bayesian. In either case the assumption of common knowledge of Bayesian rationality is violated. In our analysis on individual choices we assume that the decision maker reacts to such probability zero events by pretending the event did not happen.²²

We focus on two additional outcomes. Consensus beliefs evaluate the model predictions at the network level. Consensus beliefs can be seen as "long run" outcomes (in our case after 20 periods), something that theory has been particularly interested in. Our third outcome - the share of correct guesses over time also aggregates at the network level. One advantage of this measure compared to the consensus beliefs is that it allows us to study the dynamics of beliefs over time as opposed to focusing on just the last few periods.

We next present our experimental results. We start by addressing the questions above (Section 5). We then use a more explorative approach and try to understand the heuristics our participants

²¹We could come up with alternative theories of how agents reason. One might for example assume that agents are Bayesian, but believe that all others are naive. Mueller-Frank (2013) has shown, though, that if a network consists of Bayesian and non-Bayesian agents and if the updating function of each non-Bayesian agent is common knowledge, then such a network is informationally equivalent to a network consisting only of Bayesian agents. This result hinges on continuous beliefs being communicated, however.

²²This could e.g. be interpreted as the decision-maker treating such events as mistakes (and assuming that everyone else does so as well). Hence if a neighbour should choose black under common knowledge of Bayesian rationality, but chooses white, then this is treated as a mistake, i.e. as if the neighbour had chosen white.

use and which network properties they react to when incorporating information from neighbours (Section 6).

5 Main Results

In this section we present our main results. We first evaluate our first research question about the effect of our different information treatments on the share of correct guesses and the frequency of switches (Section 5.1) as well as on individual consistency with the naive model (Section 5.2). We then compare the explanatory power of the Bayesian and the naive model focusing on the outcomes discussed above (Section 5.3.2).

5.1 Information about the Network

To address our first research question we use our treatment variations CI, II and NI and ask whether outcomes differ across these three variations. If people do *not* make use of information about the network structure, as posited by the naive model, then how much information they have about the network is irrelevant. In this case, hence, we should not see differences between these three information conditions. We focus on two outcomes in this Section: (i) the share of correct guesses and (ii) the frequency of switches as measures of the quality and speed of learning.

We start by comparing the share of correct guesses across our information conditions CI, II and NI. Figure 2 shows the share of correct guesses over time in the Circle network (Panel (a)), the Star (Panel (b)) and the Kite (Panel (c)). In all networks the information condition does affect the share of correct guesses. In the Circle and Kite networks there are more correct guesses under CI compared to II and NI, which appear not significantly different. By contrast, in the Star network there are fewer correct guesses under CI compared to II and NI. Appendix D contains figures showing the share of correct guesses over time disaggregated for each combination of network type and signal distribution as well as for each matching group.

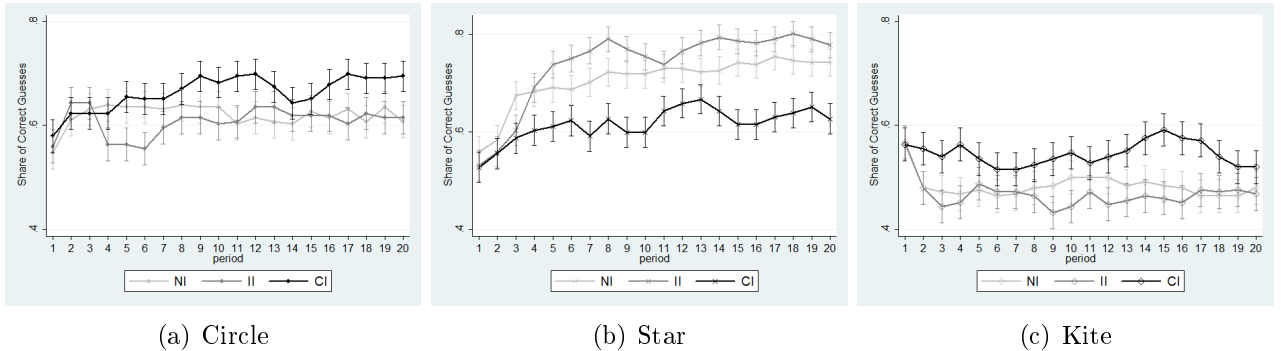


Figure 2: Share of Correct Guesses over time. Black line indicates CI, dark grey line II and light grey line NI. Error bars indicate standard errors of the mean.

To evaluate statistical significance of mean differences, we the run the following random effects OLS regressions

$$y_{ik}^t = \alpha + \beta_1 d(II)_{ik} + \beta_2 d(CI)_{ik} + \epsilon_{ik}^t \quad (2)$$

where y_{ik}^t is a binary variable indicating whether participant i in network k guessed the state correctly in t , $d(II)$ is a dummy for the II and $d(CI)$ for the CI information conditions. Standard errors are clustered at the network level and we allow for autocorrelation at the individual level. Since under the naive model the information condition should not matter we would expect $\beta_1 = \beta_2 = 0$.

Table 4 shows these regressions separately for the two signal distributions shown in Figure 1 and separately for the first and second half of the experiment (after some learning has taken place).

In the circle network, β_1 and β_2 are indeed not statistically different from zero, despite some period by period difference seen in Figure 2. This is the only network for which this is the case, i.e. for which $\beta_1 = \beta_2 = 0$ can't be rejected. In the Star network both coefficients are substantially different from zero. Interestingly the direction of the effect depends on the signal distribution. If the center of the star receives an "incorrect" signal (Star 1), more information about the network structure leads to fewer correct guesses. If the center receives a "correct" signal (Star 2), more information leads to more correct guesses. Ex post we derive the following intuition for this effect. We find evidence (in all networks) that, relative to the naive model, participants on average place too much weight on their own information (see Table 9 below). This means that also the center of a star with an "incorrect" signal will place too much weight on this signal. This overweighing can lead to an incorrect guess despite the fact that the center should know the state from round 2 onwards. Complete information in this case hurts, because spokes - knowing that the center should know the state - trust the center too much. Finally, in the Kite network $\beta_2 > 0$. Hence in this network there are more correct guesses when there is complete information about the network. This effect is particularly strong and statistically significant under signal distribution Kite2.

<i>Share of Correct Guesses - First Half of the Experiment</i>						
	(1) Circle1	(2) Circle2	(3) Star1	(4) Star2	(5) Kite1	(6) Kite2
d(II)	0.011 (0.072)	-0.015 (0.069)	-0.084 (0.075)	0.166*** (0.054)	0.048 (0.066)	-0.012 (0.080)
d(CI)	0.003 (0.072)	0.108 (0.069)	-0.217*** (0.075)	0.160*** (0.054)	0.021 (0.066)	0.159** (0.080)
Constant	0.599*** (0.051)	0.564*** (0.048)	0.730*** (0.053)	0.608*** (0.038)	0.334*** (0.047)	0.580*** (0.056)
Observations	3,780	3,780	3,780	3,780	3,780	3,780
Groups	126	126	126	126	126	126
<i>Share of Correct Guesses - Second Half of the Experiment</i>						
	(1b) Circle1	(2b) Circle2	(3b) Star1	(4b) Star2	(5b) Kite1	(6b) Kite2
d(II)	0.095 (0.074)	-0.064 (0.060)	-0.116* (0.060)	0.145** (0.072)	0.068 (0.067)	-0.039 (0.078)
d(CI)	0.023 (0.074)	0.046 (0.060)	-0.213*** (0.060)	0.196*** (0.072)	0.074* (0.047)	0.082** (0.028)
Constant	0.559*** (0.052)	0.753*** (0.042)	0.881*** (0.042)	0.618*** (0.051)	0.304*** (0.047)	0.629*** (0.055)
Observations	3,780	3,780	3,780	3,780	3,780	3,780
Groups	126	126	126	126	126	126
Robust Standard errors in parentheses						
*** p<0.01, ** p<0.05, * p<0.1						

Table 4: Random Effects regression of correct guesses on dummy variables indicating information treatments. Standard errors account for auto-correlation at the individual level and are clustered at the network (matching group) level.

As an alternative outcome and indicator of learning, we also studied the frequency of switches, i.e. how likely participants are to switch their guess between two rounds. Figure 3 does not show a visible difference in switching frequencies over time across the three information conditions. It does show, however that switching decreases over time in all networks.²³ Table 5 shows the results of regressions analogous to those presented in Table 4, but with switching as main outcome. The table shows some statistically significant differences between the NI and II (CI) cases and in each of these

²³Appendix D contains a figure showing the share of correct guesses over time disaggregated for each combination of network type and signal distribution.

$\beta_1 > 0$ ($\beta_2 > 0$). More information about the network structure seems to lead to somewhat more switching.

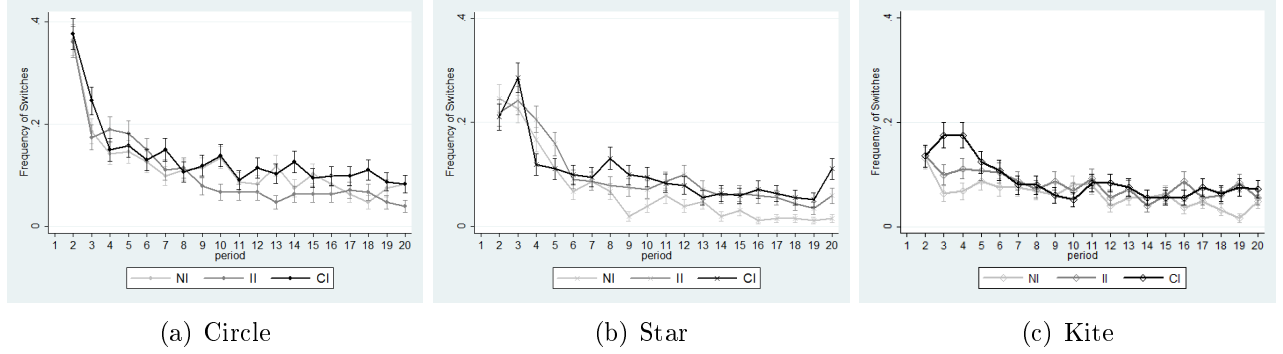


Figure 3: Share of Switches over time. Error bars indicate standard errors of the mean.

<i>Frequency of Switching - First Half of the Experiment</i>						
	(1)	(2)	(3)	(4)	(5)	(6)
	Circle1	Circle2	Star1	Star2	Kite1	Kite2
d(II)	0.067*	-0.001	0.070***	-0.025	0.044*	-0.010
	(0.038)	(0.029)	(0.025)	(0.030)	(0.025)	(0.024)
d(CI)	-0.009	0.023	0.019	0.029	0.046*	0.014
	(0.038)	(0.029)	(0.025)	(0.030)	(0.025)	(0.024)
Constant	0.160***	0.122***	0.054***	0.109***	0.070***	0.063***
	(0.027)	(0.021)	(0.018)	(0.021)	(0.017)	(0.017)
Observations	3,591	3,591	3,591	3,591	3,591	3,591
Number Participants	126	126	126	126	126	126
<i>Frequency of Switching - Second Half of the Experiment</i>						
	(1b)	(2b)	(3b)	(4b)	(5b)	(6b)
	Circle1	Circle2	Star1	Star2	Kite1	Kite2
d(II)	0.018	0.004	0.032	0.011	0.051**	-0.009
	(0.030)	(0.036)	(0.023)	(0.027)	(0.024)	(0.020)
d(CI)	-0.010	0.063*	0.016	0.044*	0.025	0.025
	(0.030)	(0.036)	(0.023)	(0.027)	(0.024)	(0.020)
Constant	0.098***	0.110***	0.063***	0.068***	0.067***	0.043***
	(0.021)	(0.026)	(0.016)	(0.019)	(0.017)	(0.014)
Observations	3,591	3,591	3,591	3,591	3,591	3,591
Number of Participants	126	126	126	126	126	126

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 5: Random Effects regression of switches on dummy variables indicating information treatments. Standard errors account for auto-correlation and are clustered at the network (matching group) level.

Result 1 *Information about the network structure affects behaviour and in particular the share of correct guesses.*

5.2 Information and Consistency with the Naive Model

We next ask how consistent individual choices are with the naive model. Table 6 shows how consistent participants are with the naive model depending on the information treatment. The table shows that the amount of information does not affect the share of decisions consistent with the naive model in the Circle networks. This is different in the Star and Kite networks. Here we do see a substantial drop in the share of choices consistent with the naive model particularly under CI. Regression Analysis can be found in Table 7. The regression is analogous that that in equation (2), except that outcome y

	<i>Consistency with Naive Model</i>		
	Circle	Star	Kite
NI	0.90 ($n=4788$)	0.90 ($n=4788$)	0.92 ($n=4788$)
II	0.91 ($n=4788$)	0.88 ($n=4788$)	0.91 ($n=4788$)
CI	0.88 ($n=4788$)	0.87 ($n=4788$)	0.86 ($n=4788$)

Table 6: Percentage of individual decisions (across periods $t = 2, \dots, 20$ and all rounds) consistent with the predictions of the naive model conditional on the history of guesses ($\mathbf{g}^1, \dots, \mathbf{g}^{t-1}$) and depending on information conditions and network type. Indifference ($G_i^t(\mathbf{g}^{t-1}) = \frac{1}{2}$) is coded as 50% success.

now indicates consistency with the naive model. The table shows that complete information leads to lower consistency with the naive model (compared to NI) under all treatments and signal distributions across the second half of the experiment ($\beta_2 < 0$). The difference is statistically significant at the 5% level in Star-2, Kite-1 and Kite-2. This is hence another way in which information about the network structure affects play.

	<i>Consistency with Naive Model - First Half of the Experiment</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
	Circle1	Circle2	Star1	Star2	Kite1	Kite2
d(II)	0.049 (0.033)	0.004 (0.026)	-0.054* (0.032)	0.038 (0.028)	-0.009 (0.032)	0.014 (0.031)
d(CI)	-0.024 (0.033)	-0.020 (0.026)	-0.009 (0.032)	-0.007 (0.028)	-0.034 (0.032)	-0.010 (0.031)
Constant	0.859*** (0.023)	0.879*** (0.018)	0.875*** (0.023)	0.836*** (0.019)	0.883*** (0.023)	0.893*** (0.022)
Observations	3,780	3,780	3,780	3,780	3,780	3,780
Number Participants	126	126	126	126	126	126
	<i>Consistency with Naive Model - Second Half of the Experiment</i>					
	(1b)	(2b)	(3b)	(4b)	(5b)	(6b)
	Circle1	Circle2	Star1	Star2	Kite1	Kite2
d(II)	-0.000 (0.030)	0.006 (0.031)	-0.030 (0.026)	0.010 (0.025)	-0.033 (0.033)	-0.003 (0.023)
d(CI)	-0.003 (0.030)	-0.058* (0.031)	-0.041 (0.026)	-0.060** (0.025)	-0.075** (0.033)	-0.048** (0.023)
Constant	0.883*** (0.021)	0.892*** (0.021)	0.895*** (0.018)	0.898*** (0.018)	0.898*** (0.023)	0.837*** (0.016)
Observations	3,780	3,780	3,780	3,780	3,780	3,780
Number of Participants	126	126	126	126	126	126

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 7: Random Effects regression of dummy indicating consistency with naive model on dummy variables indicating information treatments. Standard errors account for auto-correlation and are clustered at the network (matching group) level.

Result 2 *Complete information about the network structure leads to lower consistency with the naive model in the Star and Kite networks compared to “no information” (NI).*

5.3 The Naive vs the Bayesian Model

This subsection compares consistency of the data with the naive model and the bayesian model outlined in Section 4. This section only uses data from the complete information (CI) treatments.

5.3.1 Individual Decisions

We start by comparing the share of individual choices explained by the Bayesian and naive models. The percentage of decisions consistent with the Bayesian model ranges from just below 70% to 85%

across treatments. The percentage of decisions consistent with the naive model is even above 85% in most treatments. Note that, since choices are binary, random (uniform) choice would be consistent with a given model 50 % of the time. To distinguish the two models we should be most interested in what happens when the predictions of the two models differ. Depending on the network, this happens about 20-40 % of the time (least often in the Circle, most often in the Star). If predictions differ, then decisions are consistent with the naive model in more than 78% of all cases in all treatments, but consistent with the Bayesian model in less than 22% of all cases.²⁴

5.3.2 Network Level Results

Next we study to which extent aggregate outcomes, in particular consensus beliefs and the share of correct guesses, are in line with theoretical predictions.

Consensus Beliefs Table 8 shows the share of networks consistent with the consensus predictions of the Bayesian and Naive model in the last period of each round. Overall 28 % of networks are consistent with the Bayesian model and 26 % with the naive model. If either model predicts consensus on the right urn, then this happens in about a third of the time in the experiment (27% for the Bayesian model and 32% for the naive model). If we allow for a mistake by one participant out of seven, i.e. we require at least 6 participants to agree on the correct urn (right panel), then about half the networks are in line with the theoretical prediction (44% in the cases where the Bayesian model predicts a consensus on the correct urn and 53% in the case of the naive model). The difference between the two models is not statistically significant in any of these cases (two-sided ranksum test, $p > 0.1$). If the Bayesian model predicts a consensus on the wrong urn, this happens in 33% of the cases (56% if we allow for one mistake).

Consistent with...	Bayes	Naive		Bayes	Naive
Overall	0.28	0.26	Allowing 1 mistake	0.47	0.54
<i>Conditional on model predictions</i>			<i>Conditional on model predictions</i>		
Cons. Right Urn (7 correct)	0.27 (0.29,0.30,0.27)	0.32 (0.38,0.30,-)	Cons. Right Urn (6-7 correct)	0.44 (0.41,0.49,0.39)	0.53 (0.61,0.49,-)
No Consensus (4 correct)	-	0.24 (0.17,-,0.31)	No Consensus (3-5 correct)	-	0.69 (0.73,-,0.61)
No Consensus (2 correct)	-	0.08 (-, -,0.08)	No Consensus (1-3 correct)	-	0.33 (-, -,0.33)
Cons. Wrong Urn (0 correct)	0.33 (-, -,0.33)	-	Cons. Wrong Urn (0-1 correct)	0.56 (-, -,0.56)	-

Table 8: Share of Networks consistent with consensus predictions of Bayesian and Naive model overall and conditional on model predictions. Hence each entry shows which percent of those networks where a model predicts x agents to guess correctly ($x = 0, 2, 4, 7$) reach a state in which this is the case in period 20. In brackets below separate shares for the three networks (circle, star, kite). Three rightmost columns allow for one mistake, i.e. require only 6 agents to agree for a consensus. Data from the last period of each round.

Share of correct guesses Figure 4 contrasts the share of correct guesses over time observed in the data with theoretical predictions. The figure yields a similar insight as the consensus analysis. The theoretical models seem to have something to say about the data, but neither model is a very good predictor of behaviour in all treatments. In the Kite treatments, for example, data points are

²⁴Indifference ($G_i^t(\mathbf{g}^{t-1}) = \frac{1}{2}$) is not very common in the data. The naive model is indifferent for 5.2% of observations and the Bayesian model for 3.8% of observations.

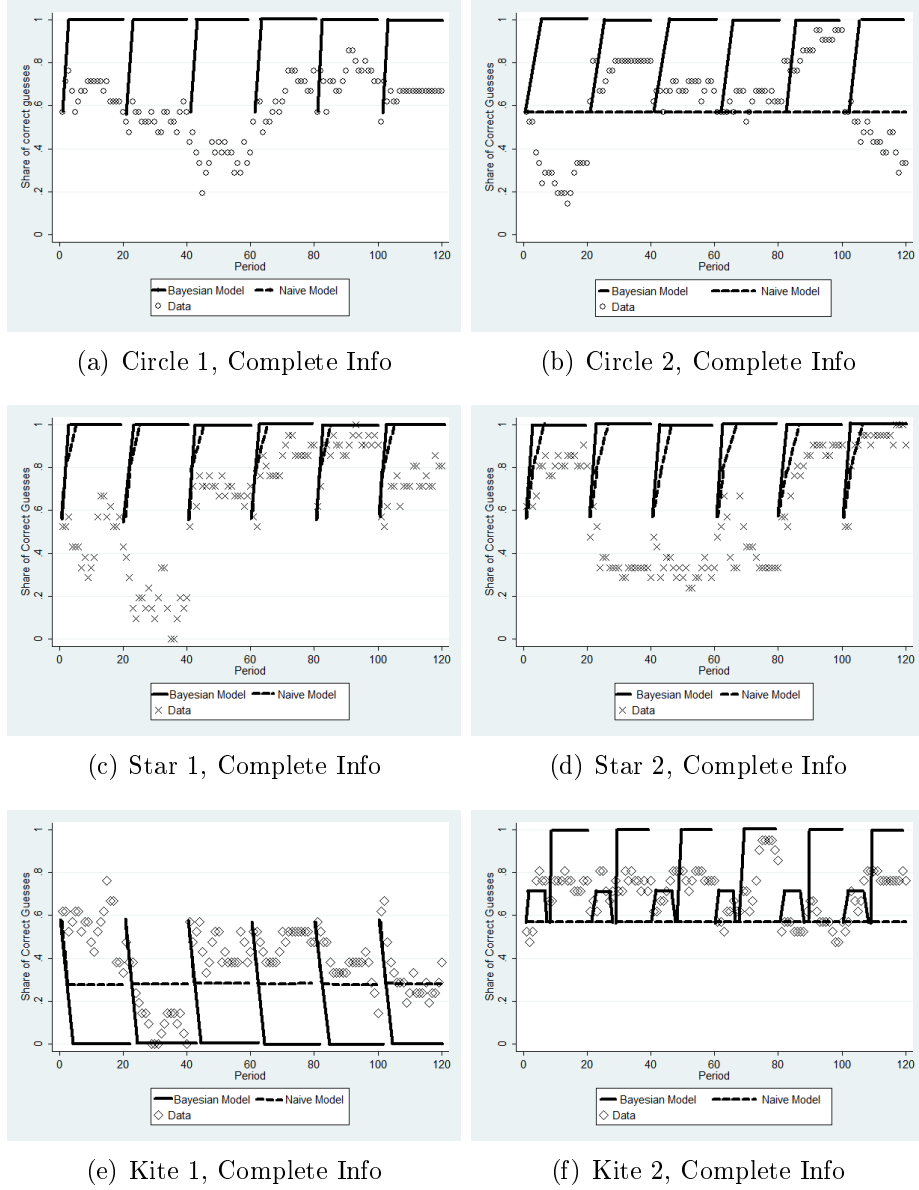


Figure 4: Share of Correct Guesses over time. The solid lines indicate the predicted share of correct guesses in the Bayesian Model, the dashed lines in the naive model and data points are scattered.

clustered around high shares of corrected guesses in Kite-2 and around low shares of correct guesses in Kite-1 as predicted by both models, but the dynamics is quite different and a substantial gap between theoretical predictions and data averages remains. In the Star and Circle networks, data points are very close to theory in some rounds (e.g. round 5 in Circle-2, round 4-5 in Star-1 or rounds 5-6 in Star-2), but they seem disconnected from theory in other rounds (round 3 in Circle-1, round 1 in Circle-2, round 2 in Star-1 etc.). Generally data seem closer to theory in the second half of the experiment, i.e. after some learning has occurred. In terms of model comparison neither model seems to clearly outperform the other.

Summary To sum up, we have seen that while the naive model seems to be able to explain individual choices well, it does far worse at predicting network level outcomes, such as consensus beliefs or the share of correct guesses. In fact, at the network level both models seem to predict about equally well. How can we explain these patterns? First, these findings suggest a substantial

impact of few deviations from the model on aggregate properties. Note that the naive model in particular is vulnerable to such deviations because of its “forgetfulness”. For example in Circle-1 one mistake in period 2 (a player with a white signal communicating white) will imply the network fails to reach a consensus. Hence, in this example consensus predictions of the naive model would fail despite 99.2% individual consistency (only one mistake). Second, comparison of our information treatments suggest that participants might be using rules of thumb that, while not Bayesian, are less “naive” than the naive model would suggest. One possible explanation, hence is that the instances where agents fail to be consistent with the naive model are “crucial” instances where opinions are divided. Indeed, if we focus on early periods, where such disagreements tend to be more likely, then individual-level consistency with the naive model drops (Table 6). If we focus on instances where consistency with the naive model requires switching, consistency even drops to below 60% under CI, i.e. about 2-3 out of 7 participants make decisions which are inconsistent with the naive model in these cases (though note that this is a selected sample).²⁵

Result 3 *Both models correctly predict a consensus in about 25 – 30% of the cases, which increases to $\approx 50\%$ under a relaxed definition of consensus. The naive model is a better predictor of individual choices than the Bayesian model.*

6 Results: Empirical Properties of Learning

We now move a step beyond our two benchmark models to understand what are the properties of heuristics participants use in our experiment. We study this questions within the following class of updating processes.

$$g_i^t = \begin{cases} 0 & \text{if } G_i^t(\mathbf{g}^{t-1}, \lambda) < \frac{1}{2}, \\ 1 & \text{if } G_i^t(\mathbf{g}^{t-1}, \lambda) > \frac{1}{2}, \end{cases} \quad (3)$$

where $G_i^t(\mathbf{g}^{t-1}, \lambda) = \lambda_{ii}(t)g_i^{t-1} + \sum_{j \in \mathcal{N}_i} \lambda_{ij}(t)g_j^{t-1}$. As before g_i^t denotes player i ’s guess at time t and \mathcal{N}_i is player i ’s network neighborhood, i.e. the set of other participants i observes. λ_{ii}^t is the weight player i attaches to her own past guess at time t and λ_{ij}^t the weight player i assigns to the guess of neighbour j at time t , where $\lambda_{ii}^t + \sum_j \lambda_{ij}^t = 1, \forall t$. Equation (3), hence, describes a more general class of models, where participants could ignore information from others ($\lambda_{ij} = 0$), overweigh their own information, change weights over time or have different weights for different neighbours depending on the network structure. Equation (3) nests the naive model ($\lambda_{ij}^t = \frac{1}{|\mathcal{N}_i|+1} \forall i, j, t$). It does not nest the Bayesian model, since the weights λ are not history dependent. However, time dependent weights allow us to pick values of λ_{ij}^t such that the path of a Bayesian learner is simulated. Hence, while given our results above we certainly don’t expect participants to be Bayesian, the class of models (3) precludes neither very naive learners (as in the naive model) nor very sophisticated learners (as in the Bayesian model).²⁶

In the following we will try to understand how the weights λ depend on participants’ network positions. To this end we estimate the following linear OLS model for each network position i in network k at each time $t = 2, \dots, 20$.

²⁵Population heterogeneity (between Bayesian and naive learners) cannot by itself explain these results. In fact, Mueller-Frank (2013) has shown for the case with continuous signals that if a network consists of Bayesian and non-Bayesian agents and if the updating function of each non-Bayesian agent is common knowledge, then such a network is informationally equivalent to a network consisting only of Bayesian agents. Population heterogeneity which allows for other types could potentially explain some of these results (see also Mueller Frank and Neri, 2014).

²⁶It should be noted that - as above - we refer to the Bayesian model with common knowledge of Bayesian rationality (see Section 4.1).

$$g_{ik}^t = \lambda_{ii;k}^t g_{ik}^{t-1} + \sum_{j \in N_i} \lambda_{ij;k}^t g_{jk}^{t-1} + \epsilon_{ik}^t. \quad (4)$$

Since we estimate $\lambda_{ij;k}^t$ separately for each time period, network and network position this means that some estimates (in particular those for the hub in the star and most kite positions) would be based on only few observations. For these positions only we estimate instead $\lambda_{ij;k}^\tau$ for three or four adjacent time periods, where $\tau = 1$ for $t = 2, 3, 4$ ($\tau = 2$ for $t = 5, 6, 7, \dots$, $\tau = 6$ for $t = 17, 18, 19, 20$). As a consequence each $\lambda_{ij;k}^t$ is estimated based on between 18 to 48 observations.

6.1 Network Properties

In Section 5 we have seen that information about the network structure matters. We now ask which information about the network structure participants use and how.

To do so, we first ask how the weight participants attach to themselves depends on their network position. To do this we estimate the following random effects OLS regression, where the coefficients estimated from the equation above are now the endogenous variable.²⁷ We start by focusing on the diagonal weights λ_{ii} estimating regressions of the following form

$$\text{rel}\lambda_{ii,k}^t = \alpha + \beta \text{rel}\mathbf{X}_{ik} + z_k + \epsilon_{ik}^t, \quad (5)$$

where $\text{rel}\lambda_{ii}^t = \frac{\lambda_{ii}^t}{\bar{\lambda}_{ij}^t}$ is the weight an agent places on herself (λ_{ii}^t) relative to the weight she places on her neighbours on average ($\bar{\lambda}_{ij}^t$). \mathbf{X} are network characteristics of interest (degree, clustering, eigenvector centrality) with $\text{rel}\mathbf{X}_{ik} = \frac{\mathbf{X}_{ik}}{\bar{\mathbf{X}}_{jk}}$ the corresponding “relative network characteristics”.²⁸ The advantage of using a formulation based on *relative* weights is that the theoretical predictions of the naive model do not depend on the network structure. In particular under the identical weights model we should expect $\alpha = 1$ and $\beta = 0$ in all our networks. We standardize all relative variables $\text{rel}\mathbf{X}_{ik}$ to mean zero and standard deviation one and we include network-type (circle, star, kite) fixed effects z_k in all regressions.

Table 9 (top panel) shows the results of this regression focusing on the second half of the experiment.²⁹ Column (1) simultaneously includes relative degree, eigenvector centrality and clustering in the regression. Degree measures how many neighbours a node has, eigenvector centrality is a measure of how “central” a node is in the network and clustering measures how many of a node’s neighbours are neighbours themselves.³⁰ Columns (2)-(4) include the three network characteristics separately. Columns (5)-(7) include interactions with the star and kite networks. Note that within the circle there is no variation in any of these characteristics. As a consequence columns (5)-(7) include only

²⁷We assume random effects as this model permits consistent estimation also of coefficients on time-invariant regressors, such as the network characteristics of different nodes. Most of the time the weights estimated in equation (4) add up to one. If this does not come out of the regressions automatically, then we normalize the estimated weights s.t. they do sum to one.

²⁸Tables C.2 and C.3 in Appendix C also show regressions with X_{ik} as exogenous variable with very similar results.

²⁹Table C.1 in Appendix C shows results for the first half. There we detect few statistically significant effects. The clearest effect is that participants seem to increase the weight on themselves the higher their degree. This effect seems to come predominantly from the Kite network, as column (5) in Table C.1 illustrates.

³⁰We define $\mathbf{A} = [\mathbf{a}_{ij}]$ as the adjacency matrix of a network, where $a_{ij} = 1$ if there is a link between agents i and j and zero otherwise. The clustering coefficient c_i of agent i is defined as follows: $c_i = \frac{\sum_{j < k} a_{ij} a_{ik} a_{jk}}{\sum_{j < k} a_{ij} a_{ik}}$. Eigenvector Centrality is defined as $EC_i = \frac{1}{\lambda} \sum_{j \in N_i} EC_j = \frac{1}{\lambda} \sum_{j \in G} a_{ij} EC_j$. These network characteristics are explained in more detail in Appendix B.

interactions with the Star and Kite, respectively. All regressions allow for autocorrelation at the individual level. Standard errors are clustered at the network level.

The first thing to notice is that $\alpha \neq 1$. Participants place substantially higher weight on themselves compared to their average neighbour. This is inconsistent with the identical weights version of the naive model, but would still be in line with some of the modifications used in the literature (see e.g. DeMarzo et al (2003)). More fundamentally, however, we also see that $\beta \neq 0$ in a number of cases, i.e. that relative weights depend on network positions. This cannot be accommodated by any of the standard versions of the naive model.

Degree, clustering and eigenvector centrality all have a positive influence on the weight i places on herself when considered separately, as in columns (2)-(4). When entered together, though, the effect of clustering and possibly eigenvector centrality seem to dominate, though the latter is very imprecisely estimated and not statistically significant (column (1)). An F-test fails to reject the hypothesis that degree and eigenvector centrality are simultaneously equal to zero at the 5% level ($F = 2.39; p = 0.0928$). A one standard deviation increase in clustering, though seems to lead to an $\approx 10\%$ increase in the weight i places on herself.

We next ask how off-diagonal weights λ_{ij} , ($j \neq i, j \in \mathcal{N}_i$) depend on network characteristics. To answer this question we run the following regression

$$\text{rel}\lambda_{ij,k}^t = \alpha + \beta \text{rel}\mathbf{X}_{jk} + z_k + \epsilon_{ik}^t, \quad (6)$$

where $\text{rel}\lambda_{ij}^t = \frac{\lambda_{ij}^t}{\lambda_{i(-j)}^t}$ is the weight an agent places on neighbour j relative to the average weight placed on other neighbours (including herself). $\text{rel}\mathbf{X}_{jk} = \frac{\mathbf{x}_j}{\bar{\mathbf{x}}_{(-j)_k}}$ is the relative value of network characteristic \mathbf{X} for neighbour j compared to all neighbours of i on average. Under the (identical weights) naive model we would again expect $\alpha = 1$ and $\beta = 0$.

We now find that $\alpha < 1$ which is again due to the fact that participants place more weight on themselves compared to their neighbours. For the first half of the experiment (Table C.1, bottom panel) we cannot reject the hypothesis that $\beta = 0$. All coefficients are relatively small and statistically not significant.

This changes when we study data from the second half of the experiment (Table 9, bottom panel). Here neighbour j 's clustering coefficient has a negative impact on how much weight i places on j . In particular, one standard deviation increase in j 's relative clustering leads to an $\approx 22\%$ decrease in the weight i places on j . An F-test cannot reject the hypothesis that degree and eigenvector centrality are simultaneously equal to zero ($F = 0.20; p = 0.8175$) in column (1). If we focus on the star network, where there is no variation in clustering, we find a marginally significant effect of degree (column (5)). Agents tend to place higher relative weight on the center (compared to themselves) than the center places on spokes (relative to other neighbours). This is in line with the information effect identified in Section 5.1, where we argued that a higher weight on the center under complete information can lead to more or less correct guesses depending on the signal the center receives.

Result 4 *The weights participants place on neighbours' guesses depend on their and their neighbours' network characteristics.*

This analysis confirms our earlier result (Section 5.1) that participants do take network structure into account when forming beliefs. In particular, we found that a higher clustering coefficient induces participants to increase the weight on themselves, while a higher clustering coefficient of a neighbour decreases the relative weight placed on that neighbour. Both of these effects have an intuitive interpretation. Remember that an agent's clustering coefficient measures the share of her first-order neighbours who are neighbours themselves. It can hence be seen as a crude measure of how

	<i>Relative Diagonal Weights λ_{ii}</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
rel degree	-0.073 (0.458)	0.330* (0.190)					
rel clustering	0.364** (0.158)		0.320** (0.157)				
rel centrality	0.448 (0.425)			0.336* (0.176)			
kite×rel degree					1.110 (0.728)		
star×rel degree					0.273 (0.197)		
kite×rel centrality						0.419 (0.310)	
star×rel centrality						0.296 (0.214)	
kite×rel clustering							0.320** (0.157)
Constant	3.303*** (0.205)	3.302*** (0.206)	3.187*** (0.197)	3.305*** (0.205)	4.635*** (0.496)	4.635*** (0.496)	4.635*** (0.495)
Observations	404	404	404	404	404	404	404
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.046	0.032	0.035	0.034	0.035	0.034	0.035
	<i>Relative Off-Diagonal Weights λ_{ij}</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
rel degree	0.030 (0.127)	-0.069 (0.080)					
rel clustering	-0.177*** (0.040)		-0.185*** (0.038)				
rel centrality	0.056 (0.101)			-0.115* (0.061)			
kite×rel degree					-0.130 (0.242)		
star×rel degree					0.062* (0.035)		
kite×rel centrality						-0.176* (0.089)	
star×rel centrality						0.061 (0.084)	
kite×rel clustering							-0.185*** (0.038)
Constant	0.796*** (0.078)	0.731*** (0.077)	0.801*** (0.075)	0.727*** (0.075)	0.291* (0.161)	0.291* (0.161)	0.291* (0.160)
Observations	1,423	1,423	1,423	1,423	1,423	1,423	1,423
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.021	0.005	0.021	0.007	0.005	0.008	0.021

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 9: Relative Weights regressed on relative network positions and network-type interactions. **Second Half of the Experiment.** The top panel focuses on diagonal weights, the bottom panel on off-diagonal weights. 540 observations (9 positions across 10 rounds (11,...,20) and 6 repetitions) for diagonal weights and 1440 for off-diagonal weights, respectively. Missing observations are due to the average weight on neighbours being zero.

correlated an agent’s neighbours’ opinions are among each other (ignoring correlation caused by having more distant neighbours in common).³¹ Participants then decrease the weight on neighbours whose information is correlated (and in turn increase the weight on themselves).

6.2 An adjusted Rule

In the previous subsection we have seen that the weight participants attach on information from neighbours reflects their own as well as their neighbours’ network positions. In particular clustering seems to matter. While the effect of clustering is intuitive, one downside is that it is only identified off the Kite network, as both the Circle and Star networks have zero clustering.

In this section we augment the naive model to allow weights to vary with clustering and then generate new data (from networks with variation in clustering) to test whether the augmented rule can indeed explain these new data better than the naive (or the Bayesian) model.

To define the adjusted rule we build on the naive model. We denote by c_i the clustering coefficient of agent i and, as above, by $|\mathcal{N}_i|$ the number of neighbours of i , i.e. the cardinality of her first-order neighborhood. The following rule corresponds to the naive rule, but modifies it in one important manner: agents increase the weight placed on themselves proportionally to their clustering coefficient c_i .

$$g_i(t) = \begin{cases} 0 & \text{if } G_i^t(\mathbf{g}^{t-1}, \mathbf{c}) < \frac{1}{2}, \\ 1 & \text{if } G_i^t(\mathbf{g}^{t-1}, \mathbf{c}) > \frac{1}{2}, \end{cases} \quad (7)$$

where $G_i^t(\mathbf{g}^{t-1}, \mathbf{c}) = \left(\frac{1+\psi c_i}{(|\mathcal{N}_i|+1)} \right) g_i^{t-1} + |\mathcal{N}_i|^{-1} \sum_{j \in \mathcal{N}_i} \left(1 - \left(\frac{1+\psi c_i}{(|\mathcal{N}_i|+1)} \right) \right) g_j^{t-1}$, i.e. where agents increase or decrease the weight on themselves with their own clustering coefficient. We hence opted for a minimal change of the naive model. Alternative models could account for clustering of i as well as of neighbours or include other network properties as well.

If $\psi = 0$, then this rule corresponds exactly to the naive rule. The higher ψ , the more strongly agents react to correlated information from their neighbours by increasing the weight on themselves. This rule can have fundamentally different implications than either Bayesian or naive learning. In particular, persistent disagreements are more likely under the adjusted rule than under either of the other models. Note also that under the adjusted rule agents will discount neighbour’s opinions both in the cases where it does not reflect new information, but also in the cases where it does. In that sense it is a very rudimentary way to account for correlated information. Below we will assume $\psi > \frac{1}{2}$ which will ensure that the adjusted rule makes different predictions from the naive rule in the networks we consider.

In our existing networks (circle, star, kite) the adjusted rule offers differing predictions only in the Kite, where clustering coefficients are not all zero. Here, it predicts no consensus just as the naive model. The adjusted rule, however, predicts a higher share of correct guesses than the naive model in both Kite networks. Figure E.1 in Appendix E.2 compares the three models in the Kite network. In the existing data the adjusted rule seems to do somewhat better than either the Bayesian or naive models. Since, however, we have derived this rule by using these data, we will compare all three models in data generated from two new networks to test the adjusted rule out of sample.

Our new networks are the “Rectangle” and “Pentagon” networks with signal distributions as illustrated in Figure 1. In the “Rectangle” network all models converge to a consensus. The Bayesian and Adjusted Model (equation (7)) agree that agents should learn the truth. Under the naive model,

³¹This is somewhat similar in spirit to the idea behind the Newey-West estimator of the covariance matrix (Newey and West, 1987), where in this case correlation decreases with network distance (estimated by the clustering coefficient for first-order neighbours and assumed zero for larger network distances).

however, agents are predicted to agree on the wrong urn. To be fair to the naive model, the latter prediction hinges on the fact that agents stick to their own signal in case of indifference, which seems not unreasonable given the evidence seen above. In the “Pentagon” network the naive and Bayesian models agree. Both predict that a consensus will be reached and that agents will learn the truth. The adjusted model (with $\psi > \frac{1}{2}$), however, predicts that no consensus will be reached.

Consistent with...	Bayes	Naive	Adjusted		Bayes	Naive	Adjusted
Overall	0.44	0.21	0.33	Allowing 1 mistake	0.51	0.25	0.53
<i>Conditional on model predictions</i>				<i>Conditional on model predictions</i>			
Cons. Right Urn (7 correct)	0.44 (0.55,0.33)	0.33 (-,0.33)	0.55 (0.55, -)	Cons. Right Urn (6-7 correct)	0.51 (0.66,0.36)	0.36 (-,0.36)	0.66 (0.66,-)
No Consensus (4 correct)	-	-	0.11 (-, 0.11)	No Consensus (3-5 correct)	-	-	0.39 (-,0.39)
Cons. Wrong Urn (0 correct)	-	0.08 (0.08, -)	-	Cons. Wrong Urn (0-1 correct)	-	0.13 (0.13, -)	-

Table 10: Share of Networks consistent with consensus predictions of Bayesian and Naive model conditional on model predictions. Hence each entry shows which percent of those networks where a model predicts x agents to guess correctly $x = 0, 2, 4, 7$ reach a state in which this is the case in period 20. In brackets below separate shares for the two networks (rectangle, pentagon). Three rightmost columns allow for one mistake, i.e. require only 6 agents to agree for a consensus. By definition, in this case, all networks are consistent with “no consensus”. Data from the last period of each round.

Table 10 shows the share of networks that did reach a consensus in the last period of each round. In the “Rectangle” network, the naive model does very poorly with only 8 % of networks (13% under a relaxed definition of consensus) converging to a consensus on the wrong urn as the model would predict. Around 55 % of networks (66% under a relaxed definition) reach a consensus on the truth as the adjusted and Bayesian model would predict. In the “Pentagon” network only a third of the networks end up agreeing on the truth, which increases to 36% under the relaxed definition.

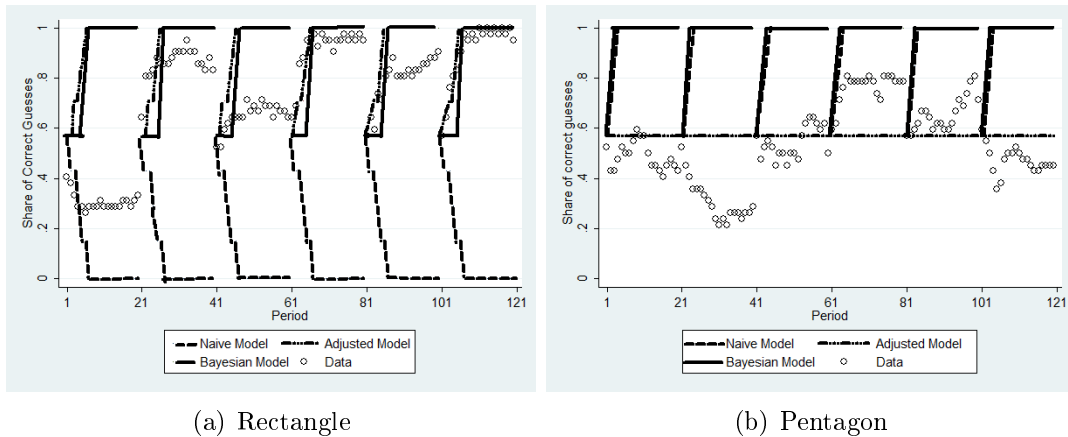


Figure 5: The share of correct guesses over time according to the Bayesian model (solid line), the naive model (dashed line) and the adjusted rule (mixed dashes) as well as data points.

Figure 5 shows the share of correct guesses over time in the two networks. In the “Rectangle” network, behaviour looks most in line with the naive model in round 1. Starting from round 2, however, the data track the theoretical prediction of the Bayesian and adjusted model remarkably well and in fact almost perfectly by the last round. A round by round t-test rejects the hypothesis that shares equal theoretical predictions (at the 5% level) in all six rounds for the naive model and in

three out six rounds for the adjusted and the Bayesian model. In the “Pentagon” the share of correct guess doesn’t track any of the three models perfectly. It oscillates between 0.2 and 0.8. In rounds 1, 3, 5 and 6 the share remains quite close to the predicted share of the adjusted model (0.57), but in round 2 it seems well below and in round 4 well above that share. A round by round t-test rejects the hypothesis that shares equal theoretical predictions (at the 5% level) in all six rounds for the naive and the Bayesian model and in two out six rounds for the adjusted model.

To sum up, the model adjusting for clustering is more consistent with the data in terms of its consensus predictions and the share of correct guesses across the two networks than either the naive or the Bayesian model. Accounting for more network characteristics (with more degrees of freedom) has the potential to improve the model’s fit and prediction power.

Result 5 *A model adjusting for network characteristics explains the share of correct guesses and consensus predictions in new data better than the naive or Bayesian models.*

In the previous two subsections we analyzed whether and how participants take network position into account in the CI treatments. Network position can presumably explain less in the II and especially also NI treatments, where participants are only partially aware of network structure. In Appendix E.3 we explore the possibility that participants react to the frequency with which their neighbours switched in the past between two guesses. There we regress weights on past frequency of switches with the strong caveat that switching is an endogenous outcome as switching of j can depend on how much weight i places on j . Table E.4 in Appendix E.3 shows the results of these regressions. Across all information conditions, weights are positively related to neighbours’ switching. This effect is particularly strong and only statistically significant in treatment NI, where a one standard deviation increase in the frequency of switching is associated with an ≈ 38 percentage point increase in the weight.

7 Conclusions

We performed two different, but complementary exercises in this paper. First we confronted the two main theories of belief formation in networks (Bayesian learning and naive learning) with experimental data. Varying the amount of information available about the network structure, we found that a key tenet of the naive model does not hold: participants *do* make use of information about the network structure when forming and updating beliefs. Still, the naive model outperforms the Bayesian model in explaining individual decisions, while both models are only partially successful in explaining aggregate patterns.

We then estimated a larger class of updating rules and found that heuristics our participants use have properties that the naive model does not capture. In particular participants increase or decrease the weight placed on neighbours depending on their network position as well as their past frequency of switching. We derive a new updating rule that allows for weights to depend on network position (in particular on clustering) and find that it does somewhat better than the two existing models in explaining a different set of data.

Future experimental work could augment the naive rule with a variety of network characteristics and contrast the performance of rules accounting in more or less sophisticated ways for various network characteristics. Future research in theory is needed to see what are the properties of the long run dynamics if agents use rules that do well in the laboratory. Another important direction for future research is to understand the role of heterogeneity in these models and to see which rules are successful in the presence of naive types in the population, if these types cannot easily be identified.

References

- [1] Acemoglu, D. and K. Bimpikis and A. Ozdaglar (2014), Dynamics of Information Exchange in Endogenous Social Networks, *Theoretical Economics* 9(1), 41-97.
- [2] Acemoglu, D., M.A. Daleh, I. Lobel and A. Ozdaglar (2011), Bayesian Learning in Networks, *Review of Economic Studies* 78, 1201-1236.
- [3] Acemoglu, D., A. Ozdaglar and A.ParandehGheibi (2010), Spread of (Mis)-information in Social Networks, *Games and Economic Behavior* 70, 194-227.
- [4] Acemoglu, D. and A. Ozdaglar (2011), Opinion Dynamics and Learning in Social Networks, *Dynamic Games and Applications* 1: 3-49.
- [5] Alevy, J.E., M.S. Haigh and J.A List (2007), Information Cascades: Evidence from a Field Experiment with Financial Market Professionals, *Journal of Finance* 62(1), 151-180.
- [6] Ambrus, A., M. Mobius and A. Szeidl (2014), Consumption risk-sharing in social networks, *American Economic Review* 104(1), 149-182.
- [7] Anderson L. and C. Holt (1997), Information Cascades in the laboratory, *American Economic Review* 87, 847-862.
- [8] Baldassari, D. and P.S. Bearman (2007), Dynamics of Political Polarization, *American Sociological Review* 72, 784-811.
- [9] Bandiera, O. and I. Rasul (2006), Social Networks and Technology Adoption in Northern Mozambique, *Economic Journal* 116(514), 869-902.
- [10] Banerjee, A., A.G. Chandrasekhar, E. Duflo and M.O. Jackson (2013), The Diffusion of Microfinance, *Science* 341
- [11] Bazerman, M.H. and D.A. Moore (2009), *Judgement in Managerial Decision Making*, 7th edition, Wiley and Sons.
- [12] Brandts, J., A.E. Giritligil and R. Weber (2015), An Experimental Study of Persuasion Bias and Social Influence in Networks, *European Economic Review*, forthcoming.
- [13] Buechel, B., T. Hellman and S. Kloessner (2015), Opinion Dynamics and Wisdom under Conformity, *Journal of Economic Dynamics and Control*, forthcoming.
- [14] Bursztyn, L., F. Ederer, B. Ferman and N. Yuchtman (2014), Understanding Mechanisms Underlying Peer Effects: Evidence from a Field Experiment on Financial Decisions, *Econometrica* 82(4), 1273-1301.
- [15] Cai, J., A.D. Janvry and E. Sadoulet (2015), Social networks and the decision to insure, *American Economic Journal: Applied Economics* 7(2), 81-108.
- [16] Chandrasekhar, A., H. Larreguy and J.P. Xandri (2015), Testing models of social learning on networks: Evidence from a lab experiment in the field, NBER working paper 21468.
- [17] Choi, S. and D. Gale and S. Kariv (2012), Social Learning in Networks: A Quantal Response Equilibrium Analysis of Experimental Data, *Review of Economic Design* 16(2), 135-157.
- [18] Conley, T.G. and C.R. Udry (2010), Learning about new technology: Pineapple in Ghana, *American Economic Review* 100(1), 35-69.
- [19] Corrazini, L., F. Petrovesi, B. Petrovich and L. Stanca (2012), Influential listeners: An experiment on persuasion bias in social networks, *European Economic Review* 56(6), 1276-1288.

- [20] de Groot, M. (1974), Reaching a Consensus, *Journal of the American Statistical Association* 69 (345), 118-121.
- [21] deMarzo, P.M., D. Vayanos and J. Zwiebel (2003), Persuasion Bias, Social Influence and Unidimensional Opinions, *The Quarterly Journal of Economics* 118(3), 909-968.
- [22] Drago, F., F. Mengel and C. Traxler (2014), Compliance Behaviour in Networks: Evidence from a Field Experiment, mimeo.
- [23] Enke, B. and F. Zimmermann (2015), Correlation Neglect in Belief Formation, mimeo.
- [24] Frechette, G. (2012), Session effects in the laboratory, *Experimental Economics* 15, 485-498.
- [25] Gale, D. and S. Kariv (2003), Bayesian learning in social networks, *Games and Economic Behavior* 45, 329-346.
- [26] Goeree, J. K., T. Palfrey, B. Rogers and R. McKelvey (2007), Self-Correcting Information Cascades, *Review of Economic Studies* 74, 733-762.
- [27] Golub, B. and M.O. Jackson (2010), Naive Learning in Social Networks and the Wisdom of Crowds, *American Economic Journal: Microeconomics* 2(1), 112-149.
- [28] Golub, B. and M.O. Jackson (2012), How Homophily affects the speed of learning and best response dynamics, *The Quarterly Journal of Economics* 127 (3), 1287-1338.
- [29] Hong, H., J. Kubik and J. Stein (2005), Thy neighbor's portfolio: Word-of-mouth effects in the holdings and trades of money managers, *Journal of Finance* 60, 2801-2824.
- [30] Jadbabaie, A., P. Molavi, A. Sandroni and A. Tahbaz-Salehi (2012), Non-bayesian social learning, *Games and Economic Behavior* 76, 210-225.
- [31] Kahneman, D. and A. Tversky (1972), Subjective probability: A judgement of representativeness, *Cognitive Psychology* 3(3), 430-454.
- [32] Katz, E. and P.F. Lazarsfeld (1955), *Personal Influence*, New York: Free Press.
- [33] Kovarik, J., F. Mengel and J.G. Romero (2017), Learning in Network Games, *Quantitative Economics*, forthcoming.
- [34] Mobius, M., T. Phan and A. Szeidl (2015), Treasure Hunt: Social Learning in the Field, CEPR working paper 10493.
- [35] Mossel, E. and O. Tamuz (2010), Efficient Bayesian Learning in Social Networks with Gaussian Estimators, arXiv:1002.0747.
- [36] Mossel, E. and A. Sly and O. Tamuz (2015), Strategic Learning and the Topology of Social Networks, *Econometrica* 83(5), 1755-1794.
- [37] Mueller Frank, M. (2013), A general framework for rational learning in social networks, *Theoretical Economics* 8(1), 1-40.
- [38] Mueller Frank, M. and C. Neri (2014), Social Learning in Networks: Theory and Experiment, SSRN-id2328281.
- [39] Newey, W.K. and K.D. West (1987), A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55(3), 703-708.
- [40] Newman, M.E.J. (2003), Mixing Patterns in Networks, *Physical Review E* 67, 026126.

- [41] Petrides, K.V. (2009), Psychometric Properties of the Trait Emotional Intelligence Questionnaire (TEIQue). In C. Stough, D.H. Saklofske, and J.D.A. Parker (Eds.), *Advances in the Measurement of Emotional Intelligence*. New York: Springer.
- [42] Ward, E. (1982), Conservatism in human information processing. In D. Kahnemann, P. Slovic, and A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- [43] Weizsaecker, G. (2010), Do we follow others when we should? A simple test of rational expectations, *American Economic Review* 100(5), 2340-2360.

Online Appendix

Contents

A Proofs of the Theoretical Predictions	1
A.1 Bayesian Learning	2
A.2 The Naive Model	4
B Network Characteristics	7
C Additional Tables	8
C.1 Additional Tables Weights Estimations	8
D Additional Figures	11
D.1 Share of Correct Guesses by Combination of Network and Signal Distribution	11
D.2 Switching by Combination of Network and Signal Distribution	11
D.3 Share of Correct Guesses per Matching Group	12
E Additional Results	16
E.1 Time to Convergence	16
E.2 Adjusted Model	18
E.3 Frequency of Switching	19
F Questionnaire	20
G Questionnaire Results	23
H Screenshots	25
I Instructions Treatment K_CI	26

A Proofs of the Theoretical Predictions

This appendix contains the derivations of our theoretical predictions. Recall from section 4 that at the beginning of each round (consisting of 20 periods) the correct urn (BLACK or WHITE) is determined by a random draw. If urn Black is drawn, four participants receive a black signal (B) and three participants receive a white signal (W), if urn WHITE is drawn, four participants receive a white signal (W) and three participants receive a black signal (B). Conditional on the state who receives a black and who receives a white signal is random. We denote a player's signal by $s_i \in \{B, W\}$ and label network positions $i = 1, \dots, 7$ as illustrated in Figure A.1.

We denote by \mathcal{N}_i the set of neighbors of player i , by $k = (k_i)_{i=1, \dots, 7}$ the degree distribution and by $\mathbf{g}^t = (g_1^t, \dots, g_7^t)$ the vector of guesses made at time t by the 7 participants.

Note first, that — since agents get paid for three randomly drawn periods from different rounds and since their payoffs are independent of what their neighbours communicate — all agents should only be interested in communicating the correct colour in as many periods as possible. This implies that all agents have and incentive to communicate their signal at $t = 1$. In later periods, the agents' perception about the correct colour depends on the way they update their beliefs. We start by discussing the results of the Bayesian model and then afterwards prove the results relating to the naive model as outlined in Section 4.

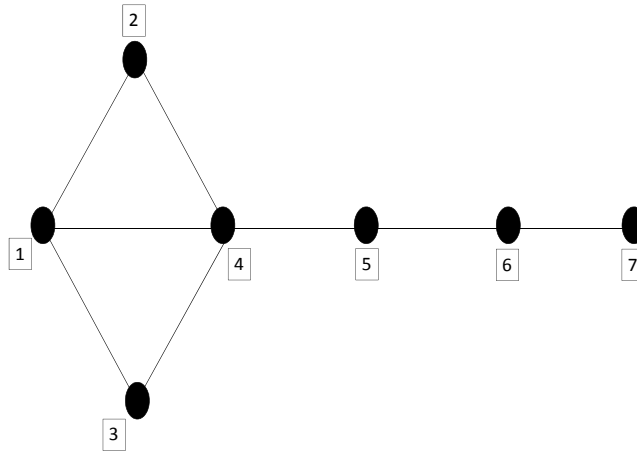


Figure A.1: Labelled Network Positions in the Kite network. In the Star the center is labelled number 1 and the spokes are 2,..7. In the circle all network positions are the same and are labelled 1,...,7 consecutively.

A.1 Bayesian Learning

In the Bayesian model, given a uniform prior on the two urns and common knowledge of rationality, agents only need to ask whether they have more evidence of four black or four white balls in the network. The answer to this question, then determines their guess.

A.1.1 Star

Note that in the star network complete and incomplete information are the same, since the star network can be uniquely deduced from its degree distribution.

Star — Bayesian Model: (i) **Star1.** *Under Initial Condition $(s_1, ..s_7) = (W, B, B, B, W, B, W)$ there is consensus on B after 3 rounds.*

(ii) **Star2.** *Under Initial Condition $(s_1, ..s_7) = (B, B, W, B, W, B, W)$ there is consensus on B after 3 rounds of communication.*

Proof. It is common knowledge that the center of the star has all the information at $t = 2$. Hence starting from $t = 3$ everybody should communicate what the center communicated at $t = 2$. The center on the other hand will know B starting from $t = 2$. \square

A.1.2 Circle

In the circle, as in the star, the network can be deduced uniquely from the degree distribution and hence complete and incomplete information are the same.

Circle — Bayesian Model: (i) **Circle1.** *Under Initial Condition $(s_1, ..s_7) = (W, B, W, B, W, B, B)$ there is consensus on B after 3 rounds of communication.*

(ii) **Circle2.** *Under Initial Condition $(s_1, ..s_7) = (W, B, W, W, B, B, B)$ there is consensus on B after 6 or more rounds.*

Proof. (i) **Circle1.** At $t = 2$ all agents will follow the majority of guesses they observed at $t = 1$. Hence at $t = 2$ agents communicate $g^2 = (B, W, B, W, B, B, B)$. At $t = 3$ participants 6 and 7 have evidence of three B-signals (their own, their neighbour's observed at $t = 1$ and one they can infer from the fact that their neighbour with a W-signal switched at $t = 2$), but only evidence of one W-signal (the one observed at $t = 1$). Hence they will guess B. Participants 2 and 4 have evidence of two W-signals and

three B-signals (since they can infer that each of their neighbours must have observed two B-guesses at $t = 1$ (one of them being themselves)). Hence participants 2 and 4 will choose B. Likewise participants 1 and 5 have evidence of three B-signals (two observed at $t = 1$ and one inferred from the fact that one of their neighbours did not swap at $t = 2$), but only evidence of two W-signals. Hence at $t = 3$ all agents will guess B.

- (ii) **Circle2.** At $t = 2$ all agents will communicate their signal, except for agents 1 and 2. At $t = 3$ all agents communicate again the same as at $t = 2$. (Agent 1 knows (W, B, W, \dots, B, B) , since she saw player 2 switch and player 7 not switch. Agent 2 knows (W, B, W, W, \dots, B) since she saw Player 1 switch and Player 3 not switch.) At $t = 4$ player 2 knows $(W, B, W, W, \dots, B, B)$ since she saw Player 1 stick at $t = 3$. Player 1 has now the same knowledge. All other participants know less. The only person participants 1 and 2 do not know about is player 5. However they know that player 5 has one neighbour choosing W and one choosing B. But both player 5 choosing W or B at $t = 2$ is consistent with everything observed so far. Hence participants 1 and 2 are now indifferent.

If player 2 resolves indifference by choosing B then we get consensus to B. The reason is as follows. Player 3 knows (W, B, W, W, \dots) and she knows that player 2 knows (W, B, W, W, \dots) . Now if Player 2 resolves the indifference by choosing B, then Player 3 will deem it more likely that she has evidence for four B rather than 3 and will switch to B. Because player 4 subsequently switches to B, player 3 knows that player 4 must have one B neighbour and hence sticks. Because player 3 sticks, player 2 knows that Player 5 had signal B and hence will stick forever to B.

The reverse logic doesn't work. If player 1 resolves indifference towards W, then player 6 will *not* switch to W even if player 7 switches. The reason is player 6 knows that if player 7 switches only after $t = 3$ it is not possible that all four other participants are W. Hence at least one of participants 1, 2, 3, 4 must have had a B signal, but player 6 knows three B signals for sure and hence will stick. But then player 7 can infer that 5 must be B etc. At this point consensus to B will be reached. (Note that, since our experiment has 20 rounds only, there is a small probability $((\frac{1}{2})^{20})$ that no consensus will be reached.)

□

A.1.3 Kite

Kite — Bayesian Model :

- (i) **Kite1.** Under Initial Condition $(s_1, \dots, s_7) = (W, B, B, W, W, B, B)$ there is consensus on W after 5 rounds.
- (ii) **Kite2.** Under Initial Condition $(s_1, \dots, s_7) = (B, B, B, B, W, W, W)$ there is consensus on B after 9 rounds of communication.

Proof. (i) **Kite1.** At $t = 2$ agents will communicate $(?, W, W, W, W, B, B)$. At $t = 3$ agent 1 won't be indifferent anymore, because s/he knows that player 5 must have communicated W at $t = 1$. (Else 4 would have communicated B instead of W). Hence all participants 1, 2, 3, 4 have evidence for 3 W and 2 B (the signals of participants 1, ..., 5). Player 5 has evidence of 2B and 2W (the signals of participants 4, 5, 6 and the signal of player 7 inferred from the fact that player 6 communicated W at $t = 2$). From the fact that 4 chose W at $t = 2$ player 5 can make inference on the eight possible signal combinations among participants 1, 2 and 3, where a higher expected number of W signals are consistent with player 4 communicating W at $t = 2$. Hence at $t = 3$ participants communicate $g^3 = (W, W, W, W, W, B, B)$. From the fact that Player 5 did not switch at $t = 3$ and knowing that player 5 must know at least 2B, player 6 infers that player 5 must know of at more than three W signals in expected terms. Hence at $t = 4$ participants communicate $g^4 = (W, W, W, W, W, W, B)$ and at $t = 5$ $g^5 = (W, W, W, W, W, W, W)$.

- (ii) **Kite2.** Player 4 knows that the true urn is B, because s/he observes four B at $t = 1$. Hence player 4 will never switch. From the fact that player 4 does not switch at $t = 2$ player 5 infers that she can

have observed at most one W apart from player 5 herself. Player 5 hence has evidence for 3B and 2W. Player 7 can be W or B with equal probability. Hence the overall evidence of player 5 is in favour of B and at $t = 3$ we have the following communication pattern (B, B, B, B, B, W, W) . At $t = 4$ player 6 can infer that player 5 must have evidence of at least three B. However s/he also has evidence of three W. Overall the probability that participants 1,2,3,4 are all B is lower than the probability that three of them are B and one W. Hence player 6 has more evidence of W and sticks. At $t = 4$ the following is communicated (B, B, B, B, B, W, W) . Now at $t = 5$ player 5 observes that player 6 did not switch, from which she can infer that player 7 must be W. Hence now player 5 has the same evidence as Player 6 in the previous round and switches back to W. At $t = 5$ the following is communicated (B, B, B, B, W, W, W) . At $t = 6$ player 4 does not switch because s/he *knows* that there are 4B. Now at $t = 7$ player 5 knows that player 4 saw him switch at $t = 5$ and did still not switch at $t = 6$. Hence player 5 knows that player 4 knows that 5 must have evidence of 3W. Given that 4 did not switch in light of this evidence player 5 knows now that player 4 must know 4W. Hence at $t = 7$ player 5 switches again and we observe (B, B, B, B, B, W, W) . Now the process unravels and at $t = 9$ we will have consensus on B. □

A.1.4 Pentagon

Pentagon — Bayesian Model *Under the initial condition depicted in figure 1 there is consensus on B after 4 rounds.*

Proof. Note that the player with degree 5 (call him player 5) knows the truth and will hence always stick to B. In period one the W player linked to one W and three B signals (say player 3) will switch from W and B. Her W-neighbour knows that player 3 does not have any information she does not have herself, but she knows that player 5 observes a player she does not observe. From the fact that player 5 did not switch in period 2 she infers that that player must have communicated B in period 1. Hence she has observed 4 B signals and hence knows the truth. She will switch to B. In period 4 her neighbour (with degree 1) will switch to B as well. □

A.1.5 Rectangle

Rectangle — Bayesian Model *Under the initial condition depicted in figure 1 there is consensus on B after 5 rounds.*

Proof. The rightmost player with a W-signal in Figure 1 knows the truth after 2 periods of communication. She observes two B signals herself and knows that her neighbour on the left observes two W signals. Upon seeing that neighbour switch to B, she knows that she must have observed 3 B signals out of which she herself only sees one. Hence she has evidence of four B signals after 2 rounds of communication. Her direct neighbours will learn the truth in period 3, her second-order neighbours in period 4 and her third-order neighbours in period 5 by which time the process will have converged to all B. □

A.2 The Naive Model

In this section we derive the predictions for the naive agents as described in Section 4.

A.2.1 Star

(Star — The naive model) :

- (i) **Star1.** *Under Initial Condition $(s_1, ..s_7) = (W, B, B, B, W, B, W)$ there is consensus on B or W after 3 or more rounds of communication.*

- (ii) **Star2.** *Under Initial Condition $(s_1, ..s_7) = (B, B, W, B, W, B, W)$ there is consensus on B after 4 or more rounds of communication.*

Proof. (i) Consensus can be on either B or W since all agents that have a “true signal” are indifferent in period 2. Whether or not agents converge to the truth or not depends on how indifferences are resolved. (ii) The center of the star will know the truth and choose it from $t = 2$ (since always at least 3 of the spokes will choose B). All other agents either choose B or are indifferent. How long convergence takes depends on how indifferences are resolved. \square

A.2.2 Circle

(Circle — The naive model) :

- (i) **Circle1.** *Under Initial Condition $(s_1, ..s_7) = (W, B, W, B, W, B, B)$ there is consensus on B after 4 rounds of communication.*
- (ii) **Circle2.** *Under Initial Condition $(s_1, ..s_7) = (W, B, W, W, B, B, B)$ there is no consensus.*

Proof. (i) **Circle1.** At $t = 2$ all W agents communicate B. All B agents communicate B except for participants 2 and 4 (that had two W neighbours at $t = 1$). At $t = 3$ all agents communicate B except for player 3, both of whose neighbours communicated W at $t = 2$. Hence at $t = 3$ all agents have both of their neighbours communicating B. There is consensus at $t = 4$ with everyone communicating B.

- (ii) **Circle2.** At $t = 2$ player 1 will communicate B and Player 2 will communicate W. All other participants will communicate again their signal. Hence communication at $t = 2$ has the following structure: (B, W, W, W, B, B, B) which is the same structure (up to permutation of player labels) as under (i). Hence there is no consensus. \square

A.2.3 Kite

(Kite — The naive model) :

- (i) **Kite1.** *Under Initial Condition $(s_1, ..s_7) = (W, B, B, W, W, B, B)$ there is no consensus.*
- (ii) **Kite2.** *Under Initial Condition $(s_1, ..s_7) = (B, B, B, B, W, W, W)$ there is no consensus.*

Proof. (i) **Kite1.** At $t = 2$ agents will communicate $g^2 = (?, W, W, W, W, B, B)$ and at $t = 3$ they will communicate (W, W, W, W, W, B, B) which is what the communication process converges to.

- (ii) **Kite2.** All agents will keep communicating their signals. \square

A.2.4 Pentagon

Pentagon — The naive Model *Under the initial condition depicted in figure 1 there is consensus on B after 4 or more rounds.*

Proof. In period 2 all participants stick to their signals except for the player with a W-signal with three B-neighbours who switches to B. In period 3 her neighbour with a W-signal then observes 4 neighbours communicating B and switches from W to B. In period 4 or later the remaining W player switches to B (she is indifferent between B and W). \square

A.2.5 Rectangle

Rectangle — The naive Model *Under the initial condition depicted in figure 1 there is consensus on W after 5 or more rounds.*

Proof. Label the participants with W signals 1,2,3 from left to right and those with B signals 4,5,6,7 also from left to right. Since agents follow the majority in their neighborhood, in period 2 the following guesses will be communicated $g^2 = (B, B, W, W, W, W, ?B?)$, where $?B?$ means that player 7 is indifferent, but may stick to their own signal with high probability. In period 3 $g^3 = (W, W, W, B, B, W, ?B?)$, in period 4 guesses will be $g^4 = (B, W, W, W, W, W, ?B?)$ and in period 5 $g^5 = (W, W, W, W, W, W, ?B?)$. The time to convergence is determined by the time it takes player 7 to resolve her indifference. \square

B Network Characteristics

In this subsection we summarize network properties of our different networks and initial conditions. We outline three indicators that are prominent in the literature. Table B.1 summarizes those network characteristics for the different network positions. In the following, we refer to $\mathbf{A} = [\mathbf{a}_{ij}]$ as the adjacency matrix of a network, where $a_{ij} = 1$ if there is a link between agents i and j and zero otherwise.

- (i) **Degree.** The Degree of an agent is given by the number of his neighbors.
- (ii) **Clustering.** The clustering coefficient is the fraction of neighbours of i who are neighbours themselves. The clustering coefficient c_i of agent i is defined as follows:

$$c_i = \frac{\sum_{j < k} a_{ij} a_{ik} a_{jk}}{\sum_{j < k} a_{ij} a_{ik}}.$$

The clustering coefficient is of particular interest in our setting, because one way in which the Bayesian and the naive model differ is that the former accounts for dependencies of information (beliefs) in clustered neighborhoods while the latter does not.

- (iii) **Eigenvector Centrality.** Eigenvector Centrality (EC) is one of several measures that determines the relative importance of a node within a network. This measure matters for convergence to a consensus under the naive model of learning (deGroot, 1974; de Marzo, Vayanos and Zwiebel, 2003). The measure assigns relative scores to all nodes in the network, assuming that connections to high-scoring nodes contribute more to the score of the node in question than equal connections to low-scoring nodes. Eigenvector Centrality is defined as

$$EC_i = \frac{1}{\lambda} \sum_{j \in \mathcal{N}_i} EC_j = \frac{1}{\lambda} \sum_{j \in G} a_{ij} EC_j$$

The equality can be rewritten as the eigenvector equation $AEC = \lambda EC$. Newman (2006) shows that only the highest λ satisfies the requirement of entirely positive entries of the vector EC and thus, Eigenvector Centrality of agent i is uniquely determined as the i^{th} entry of the respective eigenvector EC .

	Circle	Star Hub	Star Spokes	Kite 1	Kite 23	Kite 4	Kite 5	Kite 6	Kite 7
degree	2	6	1	3	2	4	2	2	1
clustering	0	0	0	0.66	1	0.33	0	0	0
EV centrality	0.14	0.29	0.11	0.22	0.17	0.25	0.11	0.04	0.01

Table B.1: Characteristics of different network positions.

C Additional Tables

C.1 Additional Tables Weights Estimations

	<i>Relative Weights λ_{ii} - First Half of the Experiment</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
rel degree	0.630 (0.887)	0.564* (0.302)					
rel clustering	-0.230 (0.350)		0.176 (0.263)				
rel centrality	-0.0726 (0.849)			0.040 (0.284)			
kite×rel degree					0.359*** (0.121)		
star×rel degree					0.366 (0.310)		
kite×rel centrality						0.432 (0.521)	
star×rel centrality						0.397 (0.339)	
kite×rel clustering							0.176 (0.263)
Constant	5.201*** (0.354)	5.103*** (0.341)	4.947*** (0.331)	5.047*** (0.338)	3.104*** (0.759)	3.104*** (0.766)	3.104*** (0.766)
Observations	451	451	451	451	451	451	451
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.031	0.028	0.021	0.025	0.042	0.025	0.021
	<i>Relative Weights λ_{ij} - First Half of the Experiment</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
rel degree	0.075 (0.066)	0.010 (0.042)					
rel clustering	-0.007 (0.021)		-0.014 (0.020)				
rel centrality	-0.063 (0.051)			-0.023 (0.032)			
kite×rel degree					-0.102 (0.125)		
star×rel degree					0.024 (0.0450)		
kite×rel centrality						-0.074 (0.0466)	
star×rel centrality						0.024 (0.0446)	
kite×rel clustering							-0.014 (0.0202)
Constant	0.634*** (0.041)	0.629*** (0.040)	0.630*** (0.039)	0.621*** (0.039)	0.508*** (0.084)	0.508*** (0.084)	0.508*** (0.084)
Observations	1,282	1,282	1,282	1,282	1,282	1,282	1,282
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.021	0.020	0.020	0.020	0.020	0.022	0.020

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table C.1: Relative Weights regressed on relative network positions and network-type interactions. **First Half of the Experiment**. The top panel focuses on diagonal weights, the bottom panel on off-diagonal weights.

	<i>Relative Weights λ_{ii} - First Half of the Experiment</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
degree	0.164** (0.670)	0.420 (0.280)					
clustering	0.280 (0.481)		0.355 (0.296)				
centrality	-0.156* (0.821)			-0.016 (0.287)			
kite\times degree					0.518 (0.541)		
star\times degree					0.384 (0.328)		
kite\times centrality						-0.253 (0.333)	
star\times centrality						0.660 (0.564)	
kite\times clustering							0.355 (0.296)
Constant	4.787*** (0.397)	5.016*** (0.334)	5.040*** (0.341)	4.931*** (0.338)	3.104*** (0.766)	3.104*** (0.766)	3.104*** (0.765)
Observations	451	451	451	451	451	451	451
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.037	0.025	0.023	0.020	0.025	0.024	0.023
	<i>Relative Weights λ_{ij} - First Half of the Experiment</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
degree	0.033 (0.062)	0.006 (0.036)					
clustering	0.018 (0.035)		0.006 (0.027)				
centrality	-0.028 (0.053)			-0.001 (0.027)			
kite\times degree					-0.027 (0.060)		
star\times degree					0.024 (0.044)		
kite\times centrality						-0.010 (0.031)	
star\times centrality						0.032 (0.058)	
kite\times clustering							0.006 (0.027)
Constant	0.603*** (0.058)	0.626*** (0.039)	0.618*** (0.050)	0.626*** (0.039)	0.508*** (0.084)	0.508*** (0.084)	0.508*** (0.084)
Observations	1,282	1,282	1,282	1,282	1,282	1,282	1,282
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.020	0.020	0.020	0.020	0.020	0.020	0.020

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table C.2: Relative Weights regressed on network positions (absolute measures) and network-type interactions. **First Half of the Experiment**. The top panel focuses on diagonal weights, the bottom panel on off-diagonal weights.

	<i>Relative Weights λ_{ii} - Second Half of the Experiment</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
degree	0.964*	0.319*					
	(0.414)	(0.174)					
clustering	0.722**		0.340**				
	(0.288)		(0.172)				
centrality	-0.921*			0.284			
	(0.498)			(0.173)			
kite\times degree					0.399		
					(0.324)		
star\timesdegree					0.286		
					(0.207)		
kite\times centrality						0.220	
						(0.197)	
star\timescentrality						0.493	
						(0.356)	
kite\times clustering							0.340**
							(0.172)
Constant	2.921***	3.274***	3.095***	3.286***	4.635***	4.635***	4.635***
	(0.249)	(0.201)	(0.204)	(0.204)	(0.496)	(0.496)	(0.495)
Observations	404	404	404	404	404	404	404
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.048	0.033	0.034	0.031	0.033	0.032	0.034
	<i>Relative Weights λ_{ij} - Second Half of the Experiment</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
degree	0.254*	-0.052					
	(0.191)	(0.068)					
clustering	-0.314***			-0.194***			
	(0.101)			(0.052)			
centrality	-0.060		-0.193***				
	(0.067)		(0.052)				
kite\times degree					-0.035		
					(0.116)		
star\timesdegree					-0.061		
					(0.084)		
kite\timescentrality						-0.226***	
						(0.059)	
star\timescentrality						-0.080	
						(0.110)	
kite\timesclustering							-0.194***
							(0.052)
Constant	0.810***	0.750***	0.975***	0.749***	0.291*	0.291*	0.291*
	(0.111)	(0.075)	(0.097)	(0.074)	(0.161)	(0.160)	(0.160)
Observations	1,423	1,423	1,423	1,423	1,423	1,423	1,423
NW type fixed effects	YES	YES	YES	YES	YES	YES	YES
R-squared	0.021	0.005	0.014	0.014	0.005	0.015	0.014
Standard errors in parentheses							
*** p<0.01, ** p<0.05, * p<0.1							

Table C.3: Relative Weights regressed on network positions (absolute measures) and network-type interactions. **Second Half of the Experiment**. The top panel focuses on diagonal weights, the bottom panel on off-diagonal weights.

D Additional Figures

D.1 Share of Correct Guesses by Combination of Network and Signal Distribution

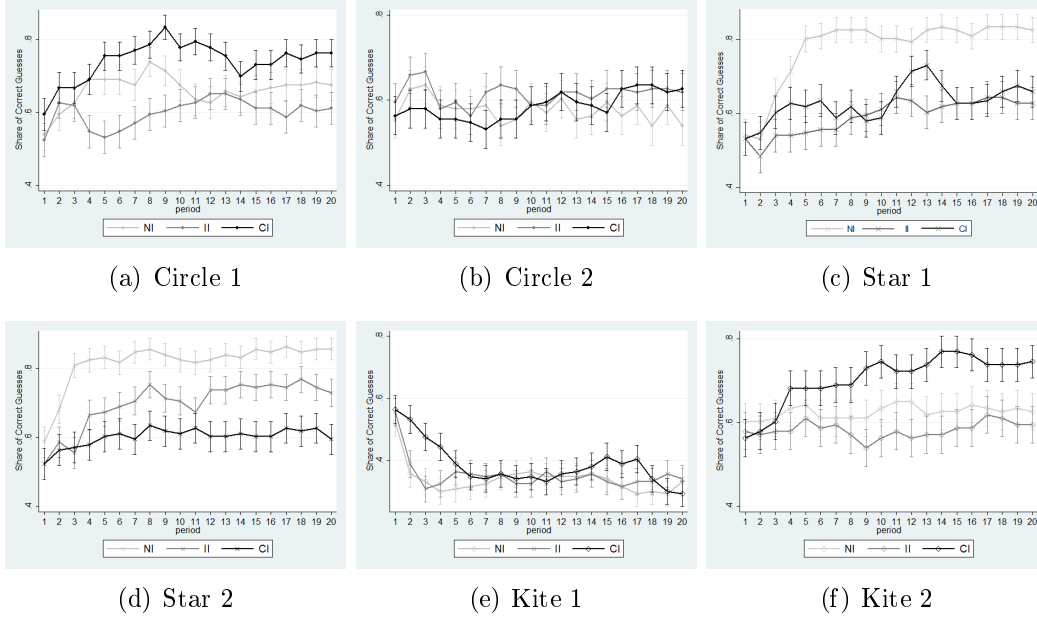


Figure D.1: Share of Correct Guesses over time for each network and signal distribution.

D.2 Switching by Combination of Network and Signal Distribution

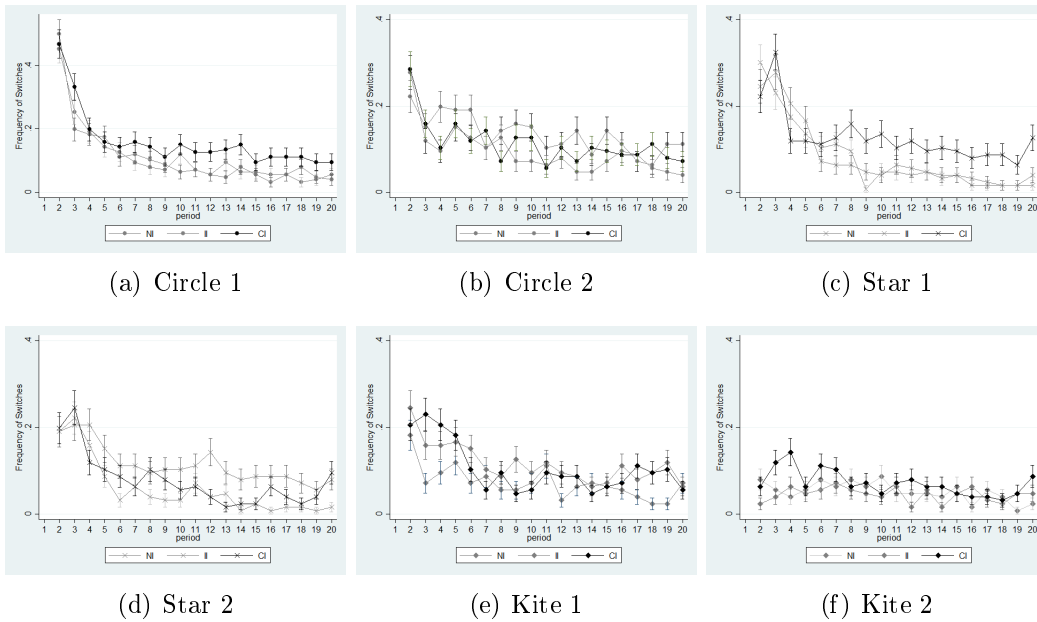


Figure D.2: Frequency of switches over time for each network and signal distribution.

D.3 Share of Correct Guesses per Matching Group

Figures D.3 - D.10 show the share of correct guesses over time for different matching groups.

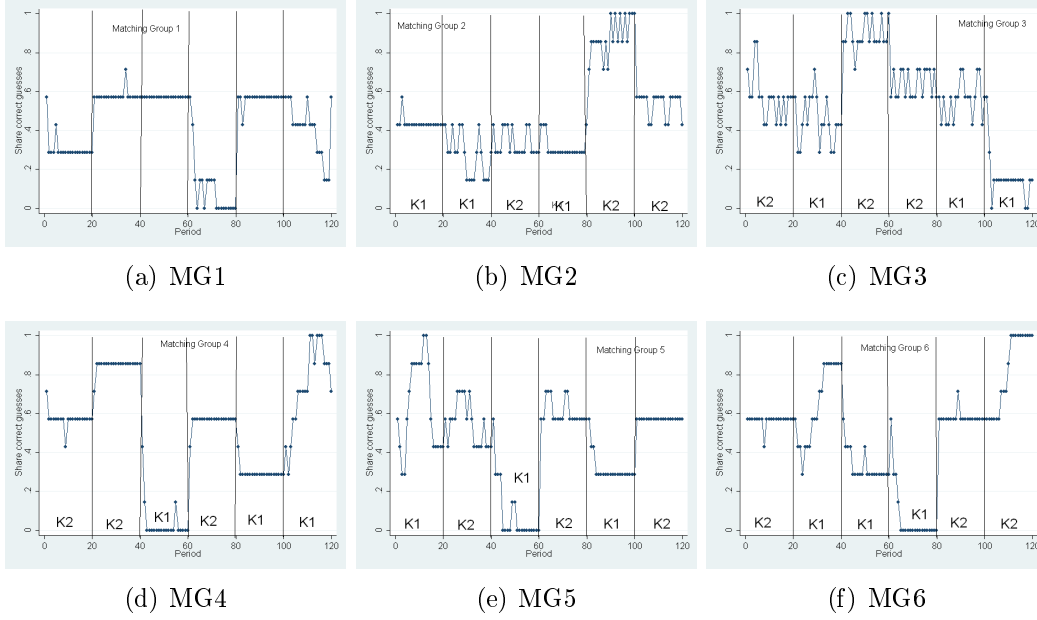


Figure D.3: Share of Correct Guesses over time. K-NI. K1 and K2 indicate initial conditions.

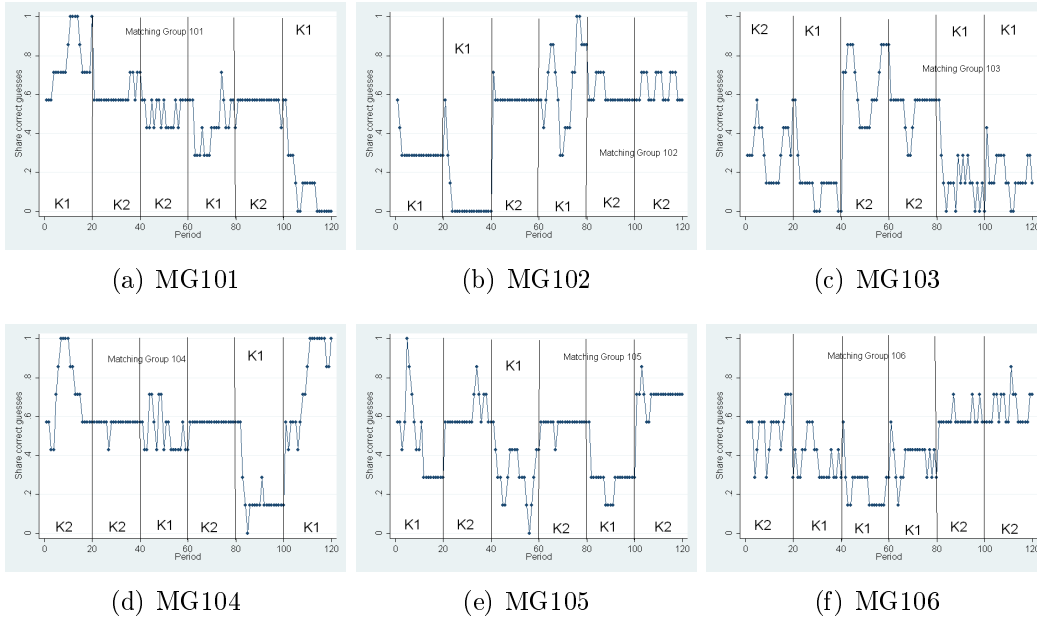


Figure D.4: Share of Correct Guesses over time. K-II. K1 and K2 indicate initial conditions.

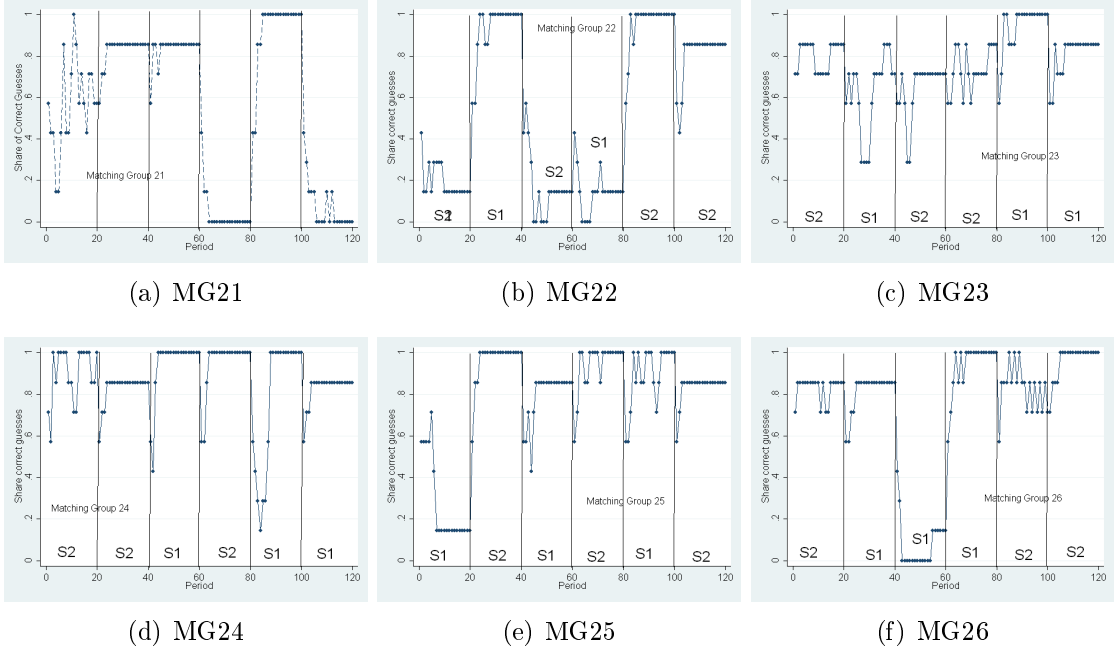


Figure D.5: Share of Correct Guesses over time. Treatment 2 (Star, No Info). S1 and S2, respectively, indicate initial conditions.

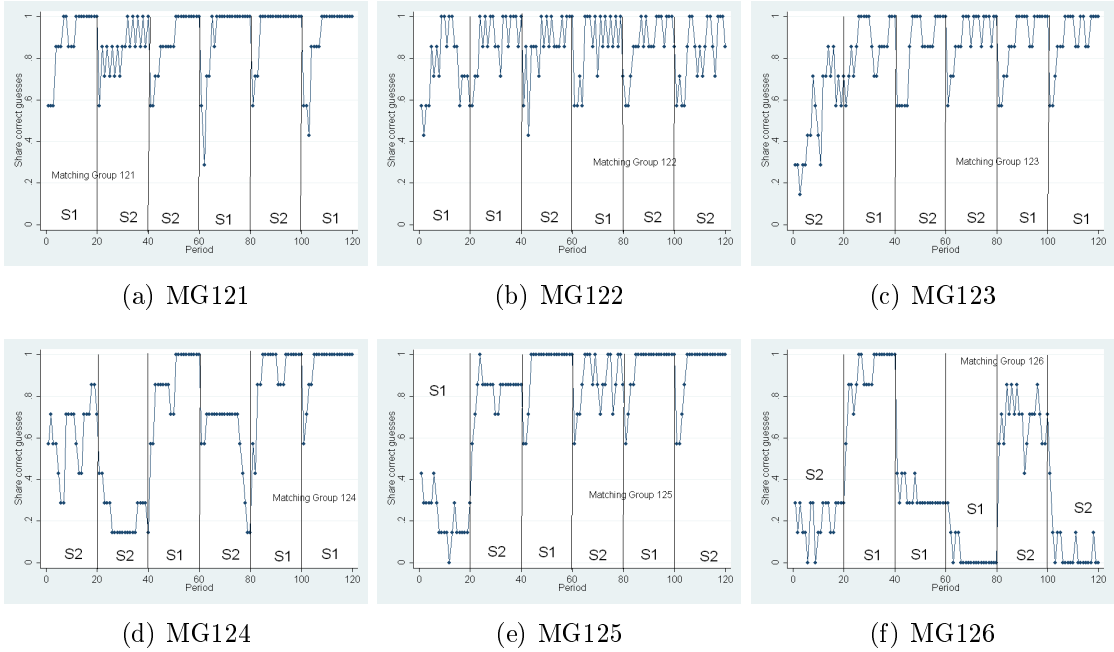


Figure D.6: Share of Correct Guesses over time. Treatment 12 (Star, Incomplete Info). S1 and S2, respectively, indicate initial conditions.

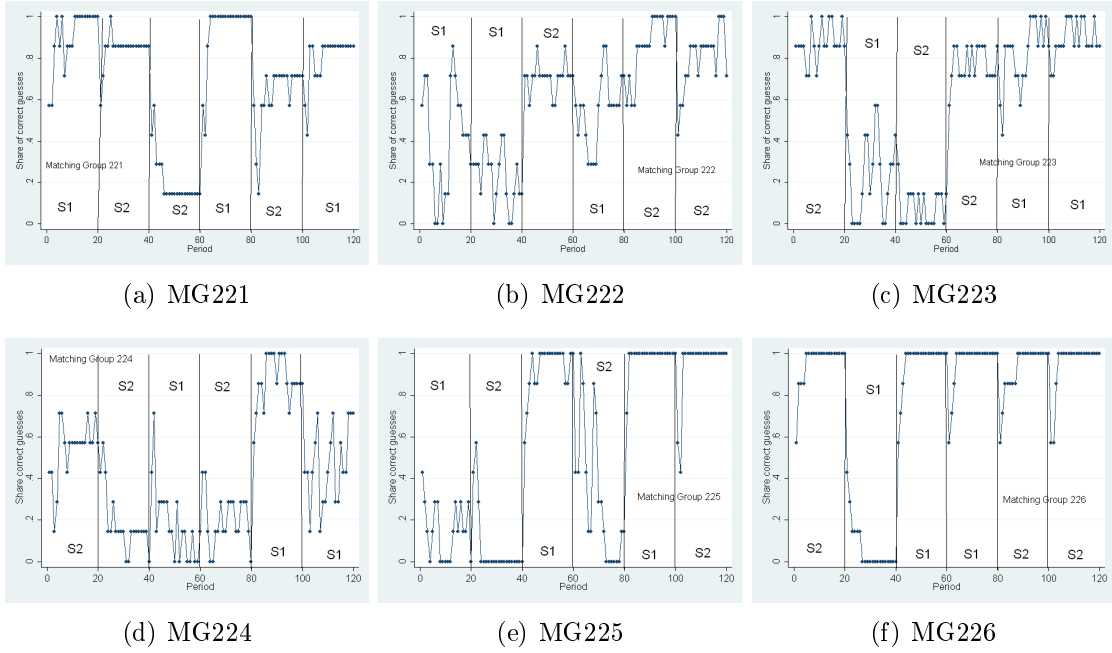


Figure D.7: Share of Correct Guesses over time. Treatment 22 (Star, complete Info). S1 and S2, respectively, indicate initial conditions.

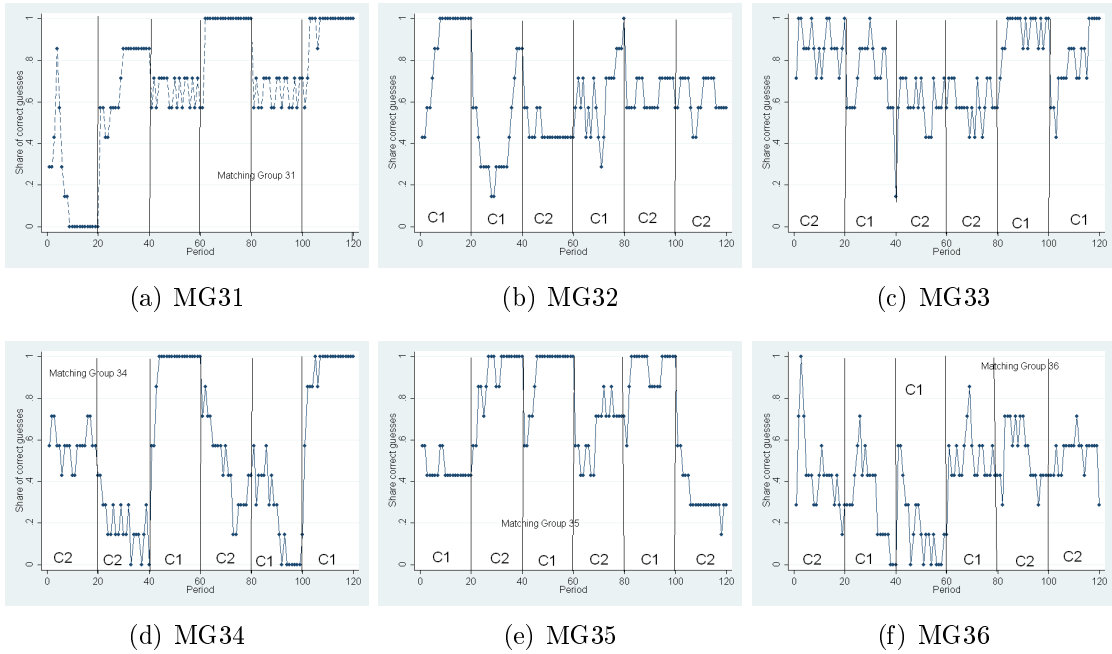


Figure D.8: Share of Correct Guesses over time. Treatment 3 (Circle, No Info). C1 and C2, respectively, indicate initial conditions.

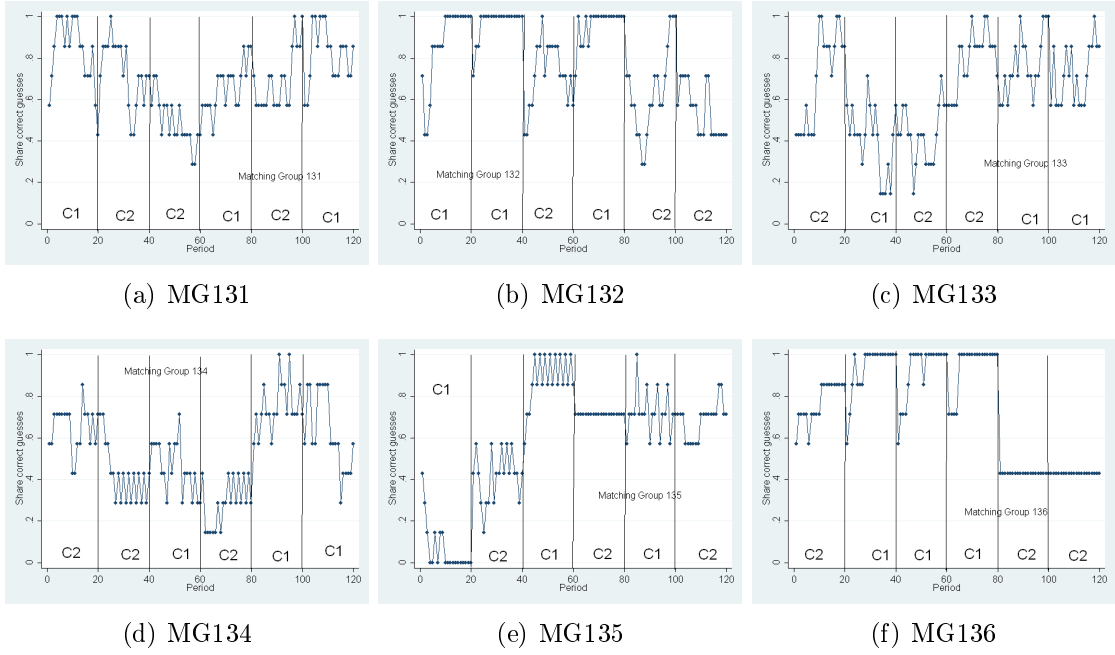


Figure D.9: Share of Correct Guesses over time. Treatment 13 (Circle, Incomplete Info). C1 and C2, respectively, indicate initial conditions.

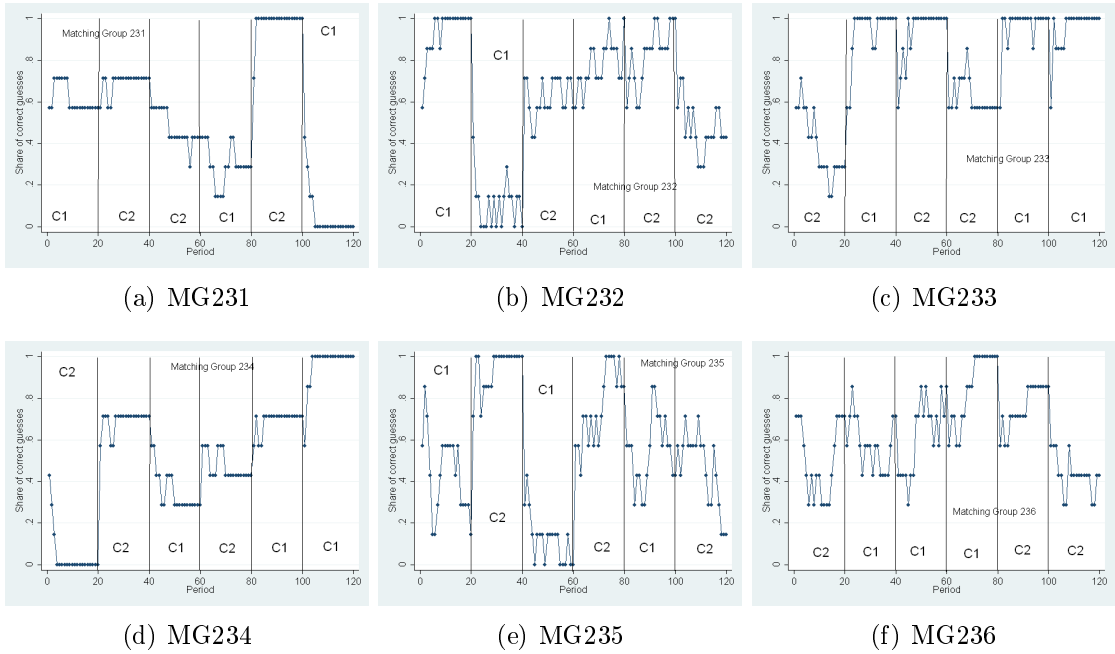


Figure D.10: Share of Correct Guesses over time. Treatment 23 (Circle, Complete Info). C1 and C2, respectively, indicate initial conditions.

E Additional Results

E.1 Time to Convergence

In this subsection we collect some additional evidence on convergence times. Table E.1 measures the time to convergence. The measure used is relatively strict in that it requires for convergence to obtain in period τ that no player switches not even once at any period $t > \tau$. Despite the measure being quite strict the table shows that convergence is obtained in several networks well before the end of a round.

	Round 1	Round 2	Round 3	Round 4	Round 5	Round 6	Mean No of Switches
Kite-NI	9.5 (6;10)	17 (19;16)	13.5 (11;17)	8.5 (5;13)	7 (4;10)	17 (20;11)	12.2 (13.2; 14.2)
Kite-II	18 (12;20)	19 (19;19)	18 (19;17)	15.5 (20;11)	12 (12;8)	18.5 (20;18)	14.6 (19.4; 9.8)
Kite-CI	19 (19;15)	19.5 (19;20)	13.5 (15;12)	15 (11;20)	16 (16;16)	19.5 (20;19)	15.8 (18; 13.6)
Star-NI	15 (10;15)	4.5 (8;4)	7 (8;6)	10 (8;12)	8 (8;5)	4.5 (7;4)	13.2 (14.8;11.6)
Star-II	20 (20;20)	18.5 (18;19)	11 (9;16)	19 (7;19)	14.5 (14;20)	12.5 (8;19)	17 (15; 19)
Star-CI	20 (20;20)	13.5 (20;6)	19.5 (19;20)	20 (4;20)	16.5 (20;16)	13.5 (19;4)	17.6 (21; 14.2)
Circle-NI	15 (9;20)	18 (18;12)	14 (6;20)	19 (19;19)	19 (19;19)	16 (7;19)	19.6 (18.6; 20.6)
Circle-II	15.5 (10;20)	19.5 (8;20)	19.5 (19;17)	13 (7;20)	19 (20;18)	19.5 (20;14)	22 (24.8;19.2)
Circle-CI	13.5 (9;16)	11 (13;7)	17.5 (18;17)	12.5 (14;10)	16.5 (15;18)	12.5 (5;19)	18.6 (19.2;17.6)

Table E.1: Time to Convergence. Median Number of per τ s.t. there are no switches by any player at any $t > \tau$. In brackets separate values for each signal distribution. The last column shows the average number of switches per network and round.

Table E.2 then relaxes the previous criterium and allows for one switch (by one player) in every period. Relaxing the criterium in this manner reduces convergence times drastically, illustrating that long convergence times are often due to one player displaying somewhat unstable behaviour while all others have converged.

	Round 1	Round 2	Round 3	Round 4	Round 5	Round 6
Kite-NI	4 (6;1)	1 (1;1)	1.5 (2;1)	1.5 (2;1)	1 (1;1)	6 (16;1)
Kite-II	16 (12;20)	1.5 (17;1)	3.5 (17;1)	12 (16;11)	1.5 (3;1)	13 (16;2)
Kite-CI	18 (18;6)	11 (9;13)	9 (3;15)	7 (5;9)	5 (7;1)	9.5 (13;4)
Star-NI	4.5 (7;1)	3 (3;3)	4.5 (3;5)	3 (2;3)	4.5 (4;5)	3.5 (5;3)
Star-II	12.5 (10;15)	3 (3;3)	5 (4;6)	4.5 (5;4)	4 (3;12)	3 (4;2)
Star-CI	11.5 (16;7)	11.5 (20;4)	2.5 (1;3)	3 (3;3)	4.5 (8;4)	5 (7;3)
Circle-NI	7.5 (6;9)	15.5 (4;19)	9 (3;15)	13.5 (14;13)	9.5 (14;2)	4.5 (4;5)
Circle-II	11.5 (7;16)	12 (3;18)	8 (4;9)	5 (5;5)	15.5 (13;15)	12.5 (18;3)
Circle-CI	13.5 (9;16)	11 (13;7)	7.5 (8;7)	4 (4;4)	6.5 (5;8)	6.5 (5;8)

Table E.2: Time to Convergence. Median Number of per τ s.t. there is at most one switch at any $t > \tau$. In brackets separate values for each signal distribution. The last column shows the average number of switches per network and round.

Another prediction we evaluate refers to the time it takes to reach a consensus. The time it takes to reach a consensus should be higher the more homophily there is in the network. The intuition is based on a result by Golub and Jackson (2012) who have shown that homophily (the fact that people with the same

beliefs tend to be linked up with increased probability) slows down convergence under the naive model. In our experiment it is the initial conditions Kite2, Star2 and Circle2 which present higher homophily (or lower signal dispersion) compared to initial conditions 1 (Kite1, Star1, Circle1).

	Newman Assort.	Time to reach Consensus		
		NI	II	CI
Circle1	-0.05	12.05	10.83	12.05
Circle2	0.30	16.27	15.77	15.41
Δ_{2-1}		4.22***	4.94***	3.36***
Star1	-0.20	12.16	9.33	8.38
Star2	0.47	11.27	11.22	10.38
Δ_{2-1}		-0.89	1.89*	2.00**
Kite1	0.09	15.05	13.8	11
Kite2	0.86	18	18.66	15.16
Δ_{2-1}		2.95***	4.86***	4.16***

Table E.3: Newman’s Assortativity and Time to Consensus. Time to consensus is measured by the first period in each round in which all agents agree irrespective of whether the agreement breaks down later on. No agreement is counted with the value 20. Statistical significance of difference between initial conditions 1 and 2 (Δ_{2-1}) is determined from random effects OLS regression of variable indicating a (correct guess; correct urn; time to consensus) on dummy for signal distribution (standard errors clustered by matching group). ***, **, * significance at 1,5,10 percent level.

Table E.3 reports Newman’s assortativity coefficient (our measure of signal dispersion or homophily) and the average time to consensus for all our networks, information conditions and initial signal distributions.¹ In the table time to consensus is measured by the first period in each round in which all agents agree irrespective of whether the agreement breaks down later on. No agreement is counted with the value 20. Table E.3 shows that indeed convergence times are much slower under initial conditions 2 with higher homophily. The differences are strongest and most significant in the Kite and Circle networks. These results are consistent with the naive model.

¹The definition of Newman’s assortativity can be found in Appendix B. We only exploit the fact, however, that homophily (or assortativity) is higher under initial conditions Circle-2, Star-2 and Kite-2. The results do not depend on the exact numerical values and hence on the exact measure of homophily used.

E.2 Adjusted Model

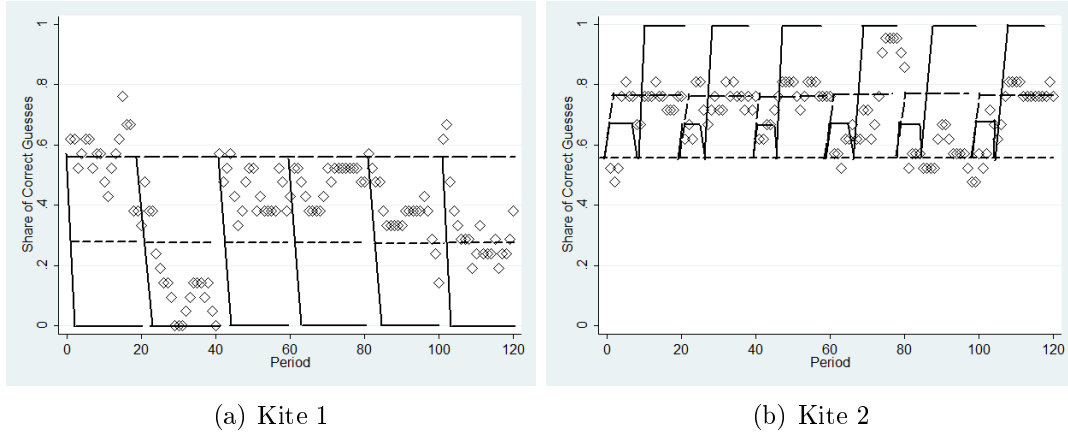


Figure E.1: The share of correct guesses over time according to the Bayesian model (solid line), the naive model (dashed line) and the adjusted rule (long dashes) as well as data points.

We describe the theoretical predictions under the adjusted model, starting with the pentagon network (Figure 1).

Pentagon — The adjusted model *Under the initial condition depicted in figure 1 there is no consensus with agents sticking forever to the initial condition.*

Proof. There is only one player who in the first period observes a majority of guesses that contradict her own signal. This is the player with W signal who observes one neighbour guessing W and three neighbours guessing B. This player, however, has a clustering coefficient of one. Hence, under the adjusted rule, she will place a weight of $\frac{2}{5}$ on herself. Given that she has a W signal herself, she would then need at least $\frac{5}{6} * 4$ of her neighbours to communicate B. Since she observes only 3 B guesses she will stick to W. Given this all agents will stick to their signals in the first period and hence in all periods. \square

Rectangle — The adjusted model *Under the initial condition depicted in figure 1 there is consensus on B after 6 rounds of communication.*

Proof. Label the agents with W signals 1,2,3 from left to right and those with B signals 4,5,6,7 also from left to right. Since agents follow the majority in their neighborhood, in period 2 the following guesses will be communicated $g^2 = (B, B, W, W, W, B, ?)$. In contrast to the naive model, player 6 here will not follow the majority but her own signal of B. The reason is that her clustering coefficient is one and hence she places a weight of $\frac{2}{3}$ on her own opinion. Given this, she will follow her own signal no matter what the neighbours say. In period 3, then, guesses will be $g^3 = (W, W, B, B, B, B, ?)$ with all agents following the majority, player 6 sticking to B and player 7 potentially being indifferent (or saying B). In period 4 guesses will be $g^4 = (B, B, B, W, W, B, ?)$ with again all agents following the majority. In period 5 guesses will be $g^5 = (W, B, B, B, B, B, ?)$ with all agents having followed the majority and in round 6 there will be consensus on B (with potentially player 7 being still indifferent). \square

E.3 Frequency of Switching

In the previous two subsections we analyzed whether and how participants take network position into account in the CI treatments. Network position can presumably explain less in the II and especially also NI treatments, where participants are only partially aware of network structure. One possibility that we explore in this subsection is that participants react to the frequency with which their neighbours switched in the past between two guesses.

Ex ante it is not clear how participants should react to switching. It can be indicative of confusion or it can be indicative of being able to react to information. Both types of behaviours have also been widely documented as violations of Bayesian rationality: (i) base-rate neglect, which leads to overweighing of sampled information (Kahnemann and Tversky, 1972) can lead to excessive switching and (ii) conservatism, which leads to underweighing or even ignoring the sample (Ward, 1982) can induce participants not to switch enough.

To understand whether and how participants in our experiment react to switching, we run the following regression:

$$\text{rel}\lambda_{ij,k} = \alpha + \beta S_{jk} + \epsilon_{ik}, \quad (8)$$

where S_{jk} is the (standardized) frequency of switches of neighbour j compared to other neighbours k across all periods 1-19. Under the naive model, we would again expect $\beta = 0$ as weights in this model are fixed and do not depend on neighbours' behaviour. It should be noted that switching is an endogenous outcome as switching of j can depend on how much weight i places on j .

	(1) NI	(2) II	(3) CI
S_{jk}	0.383*** (0.094)	0.093 (0.105)	0.129 (0.108)
Constant	0.413*** (0.039)	0.378*** (0.068)	0.379*** (0.105)
Observations	322	321	323
R-squared	0.017	0.001	0.001
Robust standard errors in parentheses			
*** p<0.01, ** p<0.05, * p<0.1			

Table E.4: Relative weight placed on neighbour j depending on how often that neighbour has switched guesses in the previous periods of that round. Data from the last period of each round. Standard errors clustered at the network (matching group) level. Missing observations to 324 are missing because of no variation among neighbours weights.

Table E.4 shows the results of these regressions. Across all information conditions, weights are positively related to neighbours' switching. This effect is particularly strong and only statistically significant in treatment NI, where a one standard deviation increase in the frequency of switching is associated with an ≈ 38 percentage point increase in the weight.

F Questionnaire

Our questionnaire consisted of three parts which we list in turn.

Emotional Intelligence (TEIQue-SF)

There are seven possible responses to each statement ranging from “Completely Disagree” (number 1) to “Completely Agree” (number 7).

- Expressing my emotions with words is not a problem for me.
- I often find it difficult to see things from another person’s viewpoint.
- On the whole, I’m a highly motivated person.
- I usually find it difficult to regulate my emotions.
- I generally don’t find life enjoyable.
- I can deal effectively with people.
- I tend to change my mind frequently.
- Many times, I can’t figure out what emotion I’m feeling.
- I feel that I have a number of good qualities.
- I often find it difficult to stand up for my rights.
- I’m usually able to influence the way other people feel.
- On the whole, I have a gloomy perspective on most things.
- Those close to me often complain that I don’t treat them right.
- I often find it difficult to adjust my life according to the circumstances.
- On the whole, I’m able to deal with stress.
- I often find it difficult to show my affection to those close to me.
- I’m normally able to “get into someone’s shoes” and experience their emotions.
- I normally find it difficult to keep myself motivated.
- I’m usually able to find ways to control my emotions when I want to.
- On the whole, I’m pleased with my life.
- I would describe myself as a good negotiator.
- I tend to get involved in things I later wish I could get out of.
- I often pause and think about my feelings.
- I believe I’m full of personal strengths.
- I tend to “back down” even if I know I’m right.
- I don’t seem to have any power at all over other people’s feelings.

- I generally believe that things will work out fine in my life.
- I find it difficult to bond well even with those close to me.
- Generally, I'm able to adapt to new environments.
- Others admire me for being relaxed.

Cognitive Reflection Test

1. A bat and a ball cost Euro 1.10 in total. The bat costs Euro 1.00 more than the ball. How much does the ball cost?
2. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?
3. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?

Probability Biases

- Lisa is 38 and pregnant for the first time. She is worried about birth defects such as Down Syndrome. Her doctor tells her the probability is actually only 1 out of 1000 that a woman of her age has a baby with Down Syndrome. Nevertheless, she takes a test. The test is moderately accurate. When a baby has Down Syndrome the test delivers a positive result 86 percent of the time. There is however a small “false positive” rate. 5 percent of babies produce a positive result despite NOT having Down Syndrome. Lisa takes the test and obtains a positive result. What are the chances that her baby has Down Syndrome.
 - 0-20 %
 - 21-40 %
 - 41-60 %
 - 61-80 %
 - 81-100 %
- Which of the following instances appears most likely, which appears second most likely?
 - Drawing a red marble from a bag containing 50 percent red and 50 percent white marbles
 - Drawing a red marble seven times in succession, with replacement (i.e. a selected marble is put back in the bag before the next marble is selected), from a bag containing 90 percent red marbles and 10 percent white marbles
 - Drawing at least one marble in seven tries, with replacement, from a bag containing 10 percent red marbles and 90 percent white marbles.
- Which of the following sequences (of heads H and tails T) arises most likely from flipping a fair coin 5 times?
 - HHHHH
 - HTHTH
 - HHTTT
 - All equally likely

- B and C equally likely (and more likely than A)
- I am drawing three cards from a deck of 32 (16 of which are red and 16 black). Now I am holding those three cards in my hands. You randomly draw one of my three cards and it is black. What is the probability that out of the three cards I drew exactly two are black?
 - 0
 - $1/4$
 - $1/3$
 - $1/2$
 - $2/3$
 - $3/4$
 - 1

Risk Numeracy

- Imagine that we roll a fair, six-sided die 1000 times. Out of 1000 rolls, how many times do you think the die would come up as an even number?
- The chance of getting a viral infection is .0005. Out of 10 000 people, about how many of them are expected to get infected?
- In a contest, the chance of winning a car is 1 in 1000. What percent of tickets of that contest win a car?
- Imagine we roll an unfair six-sided die, where the probability of rolling a six is twice as high than the probability of rolling any other number. How often should we expect to see a six if we roll the die 105 times?
- Out of 1000 people in a small town 500 are members of a choir. Out of these 500 members in a choir 100 are men. Out of the 500 inhabitants that are not in a choir 300 are men. What is the probability that a randomly drawn man is a member of the choir?
- A drawer contains red and black socks. When two socks are drawn at random, the probability that both are red is 0.5. How many socks must the drawer contain at least?

G Questionnaire Results

In this last section we report evidence from our questionnaire. Apart from general statistics on gender, age and nationality we collected our participant’s responses to questions in 4 areas that we conjectured might be related to how well they are able to solve the problem of belief updating in networks. The first area is emotional intelligence which is measured by the TEIQue (Petrides, 2009), the second is participant’s willingness to reflect (measured by the cognitive reflection test (CRT); Fredericks, 2005), the third area asks 5 questions relating to typical biases in probability (such as neglect of base rates) and the fourth area asks questions relating to risk numeracy. From each of these areas we then construct a variable that indicates which share of questions a participant answered correctly (in the case of CRT, probability biases and risk numeracy) or which share of responses indicates high emotional intelligence in the case of the TEIQue questionnaire.

	Emo.Intelligence	CRT	Prob. Biases	Risk Numeracy
Circle NI	0.72 [0.5,0.9]	0.48 [0,1]	0.37 [0,0.6]	0.54 [0,1]
Circle II	0.74 [0.5,0.8]	0.51 [0,1]	0.36 [0,0.8]	0.54 [0.2,1]
Circle CI	0.75 [0.5,0.9]	0.45 [0,1]	0.37 [0,0.8]	0.48 [0,0.8]
Star NI	0.72 [0.5,0.8]	0.57 [0,1]	0.35 [0,1]	0.6 [0,1]
Star II	0.75 [0.6,0.9]	0.5 [0,1]	0.39 [0,0.8]	0.55 [0,1]
Star CI	0.75 [0.5,0.9]	0.5 [0,1]	0.33 [0,0.8]	0.53 [0,0.8]
Kite NI	0.73 [0.5,0.9]	0.59 [0,1]	0.30 [0,0.8]	0.56 [0,1]
Kite II	0.72 [0.5,0.9]	0.59 [0,1]	0.30 [0,0.8]	0.52 [0,1]
Kite CI	0.73 [0.5,0.8]	0.48 [0,1]	0.36 [0,0.8]	0.55 [0,1]

Table G.1: Summary Statistics on Questionnaire Indicators. Mean Share of Correct Answers and in brackets range.

Table G.1 shows the mean and range of these four variables in all of our treatments. The means are similar across treatments and in most cases (except for the TEIQue) the range covers almost the entire interval $[0,1]$. We then ask whether any of these measures can explain how well participants do in our experiment. There are reasons to conjecture that any of these might be able to do so. The risk numeracy indicator summarizes how well participants understand probabilistic statements and uncertain environments. The probability biases indicator checks for some well known mistakes in dealing with probabilities and in particular also with Bayesian updating. Participants who score better in these test hence should have less problems with Bayesian updating and hence might be better at incorporating information from neighbours into their assessments. The cognitive reflection test (CRT) measures how willing people are to reflect carefully when making a decision. The emotional intelligence test is the odd one out in that it doesn’t test for reasoning abilities. Instead it measures how good participants are in assessing their own and others emotional responses. If there is a lot of heterogeneity, decisions are very noisy and following different heuristics, then - rather than reasoning abilities - a good intuition for neighbours’ decisions might be important for participants.

Table G.2 shows the marginal effects of an OLS regression where we regress a binary variable indicating whether a participant had above median payoffs (1) or not (0) on the four questionnaire indicators. The table shows that, by and large, the questionnaire indicators are insignificant. Interestingly, the emotional intelligence indicator does best at explaining success in our experiment. It has a significantly positive effect in the Star-II, Star-CI and Kite-NI treatments.

	Emo.Intelligence	CRT	Prob. Biases	Risk Numeracy
Circle NI	−0.76	0.01	−0.83	0.02
Circle II	−1.08	0.33	0.28	−0.09
Circle CI	−1.52	−0.22	−0.52	1.31***
Star NI	−0.35	0.11	−0.28	0.51
Star II	0.78*	0.32	0.02	−0.31
Star CI	2.31**	−0.29	0.51	0.18
Kite NI	2.14**	−0.13	−0.02	−0.19
Kite II	0.40	0.14	−0.13	−0.19
Kite CI	0.39	0.15	−0.60	−0.20

Table G.2: Marginal Effects from OLS regression of a binary variable indicating whether a participant had above median payoffs (1) or not (0) on questionnaire indicators. Standard errors clustered by matching group.

H Screenshots

Figures (H.1) and (H.2) show screenshots from the “Rectangle” and “Pentagon” treatments, respectively. The screens are the first period where the participant sees the network (no previous decisions by neighbours because it is period 1) and her signal and is asked to make a guess on the colour of the urn at the bottom.

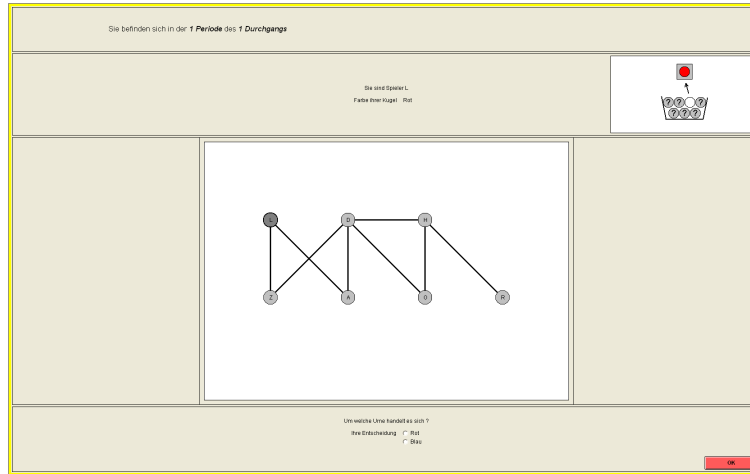


Figure H.1: Screenshot of “Rectangle” Complete Information treatment. First period decision.

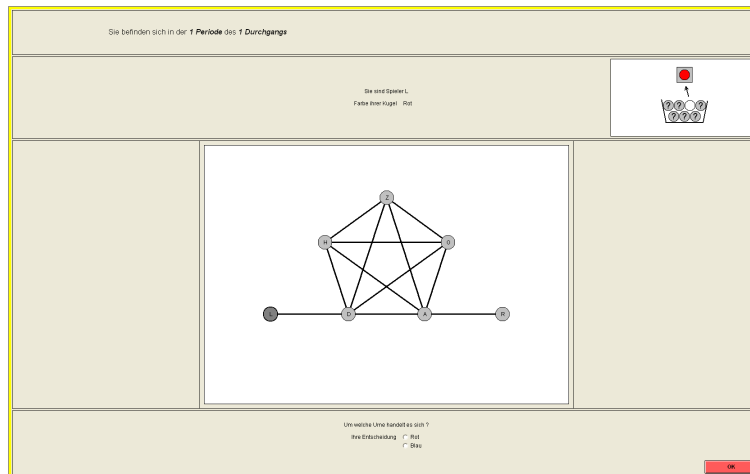


Figure H.2: Screenshot of “Pentagon” Complete Information treatment. First period decision.

I Instructions Treatment K_CI

Instructions²

Welcome and thank you for participating in this experiment. Please read these instructions carefully.

Should you have any questions, please raise your hand. An experimenter will come to you and answer your questions. From now on all communication with other participants in the experiment is forbidden. If you don't follow these rules we will have to exclude you from the experiment. Please do also switch off your mobile phone now.

You will receive 6 Euros for showing up on time to this experiment and for answering a questionnaire at the end of the experiment. During the course of the experiment you can earn more. How much depends on your and others' decisions. All your decisions will be treated confidentially.

The Experiment The experiments consists of 6 rounds in which the same game will be played. The 6 rounds are independent of each other, meaning that your decisions from previous rounds do not affect future rounds in any way.

In every round you will interact with seven randomly selected participants from this session. All your decisions will remain anonymous for all other participants.

At the beginning of each round There are two urns. Each urn contains balls of different colours. One urn (we refer to it as the "BLACK" urn) contains four black balls and three white balls. The other urn (urn "WHITE") contains four white balls and three black balls. At the beginning of each round the computer randomly selects one urn. Both urns are equally likely. You will not know which urn was chosen. Your goal in each round will be to guess which urn was chosen.

After the urn was chosen each player receives one ball from the urn, the colour of which she will see. This means.

- If urn BLACK was chosen four group members will receive a black ball and three group members a white ball.
- If urn WHITE was chosen four group members will receive a white ball and three group members a black ball.

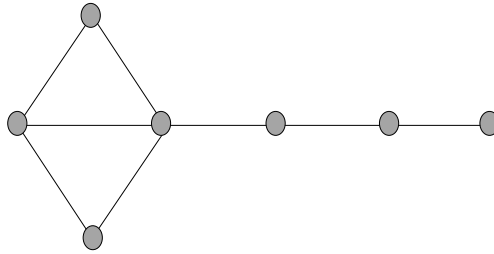
Each group member hence receives information about only one ball from the urn that was chosen. If you knew all colours of all the balls given to the group members, then you would know exactly which urn was chosen. At the beginning of a round each group member, however, knows only the colour of their own ball.

The Network The seven members of your group (including yourself) get a pseudonym in each period: L,D,H,Z,A,O or R. All seven members of your group are directly or indirectly linked via information channels as illustrated in the following graph.

In each round hence, there is

- one group member with one "neighbour"
- four group members with two "neighbours" each,
- one group member with three "neighbours" and
- one group member with four "neighbours".

²Translation from german.



“Neighbours” are other group members who you transmit information directly and from whom you receive information. In the graph shown above, neighbours are connected directly with a line. You will learn below what type of information is communicated between neighbours and how this is done.

At the beginning of each round you will learn your position in the network shown above. The network will appear on your screen and your position within the network will be marked black. Hence you will also be able to see how many neighbours you have in that round. You will also learn you own pseudonym and that of your neighbours. Just like you, all the other group members will also learn their position in the network, the number of neighbours they have, their pseudonym and that of their neighbours.

Network positions as well as pseudonyms are assigned anew at the beginning of each round. Hence both your pseudonym and those of your neighbours will change in each round. Also the positions of all group members change each round. You will have a new position and you will interact with different participants as your neighbours.

The Decision Phase After all group members have received information about their network positions, their pseudonyms, the pseudonyms of their neighbours as well as their private signals (either a black or white ball) each round continues as follows.

Each round contains 20 periods. Each period contains the following two steps.

1. Your decision: You guess whether the urn the computer chose at the beginning of the round is BLACK or WHITE.
 - (a) You say BLACK if you believe it is more likely that there are four black and three white balls in the urn.
 - (b) You say WHITE if you believe it is more likely that there are four white and three black balls in the urn.
2. After all participants have made their guess information will be exchanged simultaneously between direct neighbours in the network: You learn which guess your neighbours made and then learn about your guess. On your screen the positions of your direct neighbours will appear in the colour that corresponds to their guess.

Keep in mind that your neighbours might receive information not only from you but also from other group members. (That depends on how many neighbours they have apart from you.)

Your earnings At the end of the experiment three periods from different rounds will be randomly drawn. If your decision in the chosen period was correct (if you guessed the right urn in that period) you will receive 6 Euros for that period. If your decision was incorrect you receive nothing for that period.

This means that, if your decision was correct in all three selected periods, you receive $3 * 6 = 18$ Euros, which, together with the base payment of 6 Euros, gives you earnings of 24 Euros in total. If you are wrong in all three selected periods you will receive only the base payment of 6 Euros. If you are correct in one period you receive 12 Euros in total and if you are correct in two periods you receive 18 Euros.

Different rounds are independent of each other. As already mentioned you will participate in 6 different rounds. All these 6 rounds are independent of each other. This means in particular that

- your neighbours and your position in the network will change between rounds
- the colours of the urns and balls will change in each round (black and white were only used as examples and will not be used in the actual experiment)
- a new urn will be chosen at the beginning of each round.

Summary The following is a graphical summary of a round.

