

Learning in social networks: Selecting profitable choices among alternatives of uncertain profitability in various networks



Bas Hofstra*, Rense Corten, Vincent Buskens

Department of Sociology/ICS, Utrecht University, Padualaan 14, 3584 CH Utrecht, The Netherlands

ARTICLE INFO

Keywords:

Social learning
Social networks
Multi-armed bandit problem
Diffusion of information

ABSTRACT

Social capital theory assumes that information is valuable. However, only rarely is this value explicitly modeled, and there are few examples of empirical tests of mechanisms that connect social network structure to valuable information. We model an individual decision problem in which individuals make choices that yield uncertain outcomes. The individuals can learn about the profitability of options from their own choices and from the network. We generate computer-simulated data to derive hypotheses about the effect of network characteristics on making profitable choices. We conduct a laboratory experiment to empirically test these hypotheses and find that, at the individual level, degree centrality has a positive effect on making profitable choices whereas betweenness centrality has no effect. At the network level, density has a positive effect on making profitable choices, whereas centralization does not have an effect.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

A great deal of research on social networks addresses the question of how different network structures affect the information diffusion in social networks and how different network positions in social networks provide benefits to individuals in terms of the valuable information they receive. In particular, [Granovetter \(1973\)](#) and [Burt's \(1992\)](#) influential theories on the strength of weak ties and structural holes specify conditions under which individuals are likely to receive novel information via social networks. In this sense, strategic social network positions are seen as a form of social capital, and much research has shown that certain network positions and structures are indeed associated with individual outcomes ([Graaf de and Flap, 1988](#); [Lin, 1999](#); [Cowan and Jonard, 2004](#); [Burt, 2004](#); [Schilling and Phelps, 2007](#); [Grund, 2012](#)). This influence extends to both individual outcomes (we refer to this as the microlevel: an actor's individual outcome) (e.g., [Burt, 2001](#)) and network performance (or the macrolevel: actors' aggregated outcomes in a network) (e.g., [Grund, 2012](#)).

We argue that the exact causal mechanisms by which individuals benefit from information received via networks are

understudied. In particular, to our knowledge, no experimental research exists that causally links the mechanism of information diffusion via network positions to better micro- and macrolevel outcomes. Empirically, most research relies on observational studies that merely establish correlations between network characteristics and outcomes. Theoretically, even formal game theory models on social capital formation ([Jackson and Wolinsky, 1996](#); [Buskens and Van de Rijt, 2008](#)) typically assume that information is valuable to actors but without specifying *how* it becomes valuable to actors. One reason for this omission might be that it is not straightforward to extrapolate from the individual outcomes typically studied in observational studies (e.g., getting a job as in [Granovetter, 1973](#)) to a more generic setting that is abstract enough to be implemented in an experiment yet preserves the core features of the diffusion mechanism assumed in the theory.

We propose that a social learning setup as proposed by [Goyal \(2007: Ch. 5\)](#), in which individuals have the opportunity to use information from their social network to decide between different actions with uncertain outcomes, provides such a setting. When faced with a decision, we often have to choose between alternatives without knowing their relative advantages ([Goyal, 2007](#)). These choices have many implications for real-life outcomes. For example, farmers often adopt new crop seeds without being fully informed about which crop seeds are most profitable (e.g., [Conley and Udry, 2010](#)). Likewise, consumers often buy laptops of different brands without being fully informed about which brand

* Corresponding author. Tel.: +31 302534748.

E-mail addresses: b.hofstra@uu.nl (B. Hofstra), r.corten@uu.nl (R. Corten), v.buskens@uu.nl (V. Buskens).

is the best (Bala and Goyal, 1998). Moreover, how do consumers decide between competing alternatives, such as iPhone versus Android (cf. Kanoria, 2012)? When making multiple decisions of this type over a longer period, actors tend to update their beliefs about which choices are most profitable (Gale and Kariv, 2003; Jackson, 2008: Ch. 8). Moreover, when making multiple decisions and experiencing their outcomes in terms of profitability, actors learn which choice is most profitable. This is how we define *learning* (i.e., intrinsically valuable information) in this study: finding profitable choices among alternatives of uncertain profitability by integrating experiences from earlier decisions to optimize further decisions.

Nevertheless, how is this information diffused between actors? As we mentioned before, an important vehicle for information diffusion is a *social network*. Actors may learn not only from their own experiences but also from the experiences of others from whom they receive information via social connections. It is not only one's own choices and payoffs that are informative but also the choices and payoffs of others, which can generate valuable information on the relative attractiveness of different choices. Actors integrate this new information with previous information, including their own, to be able to make profitable choices. We thus model, in a very general way, not only situations in which information diffuses in networks but also situations in which information is intrinsically *valuable* to actors in helping them make better individual decisions. We believe this general framework not only captures core claims of social capital theory but also applies to a range of real-life situations. Again, imagine a farmer who is uncertain which crop seed to choose; only by learning from his own crop seed profit and that of other farmers with whom he interacts (i.e., his social network) will he know which crop seeds are most profitable. Alternatively, imagine an employee who has to renew his contract with his employer. From his own past experiences in renewing his contract with the same employer and from the experiences of (ex)-colleagues, he can decide how *profitable* (in terms of wage, treatment by boss and colleagues or benefits) it is for him to renew his contract or to try to sign a contract with another employer. This leaves us with specific questions such as: At which position in a network does the employee make the most profitable choice? And what network structure helps farmers maximize the sum of profitable choices? At which network position (microlevel), whether at the boundary or in the center, and in which network structures (macrolevel), whether dense or centralized, do actors make the most profitable choices (i.e., learn)? These questions are incorporated in the main question of our study: *What is the influence of microlevel and macrolevel network characteristics on selecting profitable choices among alternatives of uncertain profitability?*

1.1. An experimental approach

Choosing between alternatives without knowing their relative advantages is also referred to as the *multi-armed bandit* problem (Robbins, 1952; Gittins et al., 2011). Mathematically modeling learning in networks by actors who update their beliefs on the profitability of choices has been a goal of theoretical economists for the past few decades (cf. Goyal, 2007; Kanoria and Tamuz, 2011). However, to our knowledge, this is the first study to examine the influence of network structure and positions on solutions to multi-armed bandit problems empirically.¹ The requirements of the data for testing predictions about network characteristics on learning behavior are high. Specifically, detailed longitudinal

data are needed on social relations and individual choices. To test hypotheses both at the micro- and macrolevel, sufficient variance and many observations at the macrolevel are needed. Moreover, when collecting field data, one must often compromise the number of observations at the macrolevel, the number of observed time points and the details of the microlevel observations (Corten and Buskens, 2010). Therefore, collecting field data with sufficient variance and many observations at the macrolevel often yields practical difficulties. While it is not impossible to collect data that are sufficient to test both microlevel and macrolevel predictions (i.e., studying network effects on individual level outcomes and complete network outcomes), these studies do remain scarce (see, for example, Grund, 2012 as an exception). Even with survey data that have properly defined networks and discrete time-points, it is difficult to distinguish learning from other phenomena that may give rise to similar outcomes (cf. Conley and Udry, 2010). Finally, it is also difficult to assess true payoffs of choice alternatives of uncertain profitability in field data setups.

We suggest laboratory experiments as a suitable alternative to test predictions about network effects at both the micro- and macrolevel. Whereas surveys generate large representative datasets, experiments provide control over information availability and incentive structures. Hereby, a clear distinction between cause and effect is possible, and causal inferences can be sustained more convincingly than in cross-sectional studies. Moreover, causal knowledge requires controlled variation, which experiments make possible (Falk and Heckman, 2009). Laboratory experiments yield precise information on behavior at every time-point. Multiple networks can be observed with varying network characteristics at the macrolevel. To generate experimentally testable hypotheses at the micro- and macrolevel, the focus in this study is on small four-person networks with different network characteristics. In essence, we vary several variables at the micro- and macrolevel to see what their effects are on the number of profitable choices actors make. At the microlevel, we vary degree centrality and betweenness centrality, and we argue that these characteristics capture, respectively, the amount of direct information and the amount of redundant information (i.e., brokerage) individual actors receive. At the macrolevel, we vary density (i.e., on average, how many actors does each actor know) and centralization, the degree to which some actors have more ties than others.

There are also disadvantages to experiments. First, the external validity of findings of abstract laboratory conditions is lower than for field data. Second, group size is often smaller than in real-life human interactions. Thus, we consider this study as an intermediate step between studies based on theoretical models on social learning and field data that are difficult to obtain.

To exploit our experimental approach, we align our experimental design and theoretical model. It is difficult to derive predictions solely by exploring informal arguments on learning because arguments and counter arguments for the same effect often exist. To solve this problem, we generate computer-simulated data within our experimental parameter space. In this manner, we can calculate for all network characteristics what their expected effect is on making profitable choices. Our hypotheses directly follow from the relations between network characteristics and profitable choices found in the simulated data.

The remainder of the paper is organized as follows. In Section 2, we specify our formal model of learning in social networks and discuss informal accounts of learning. Section 3 elaborates on the specific task that simulated actors and participants in the experiment face. In Section 4, we discuss the simulation and the predictions we derive, and in Section 5, we discuss the experimental procedure and the hypotheses tests based on the analyses of

¹ Conley and Udry (2010) study only social learning without involving network structures.

the experimental data. Finally, we conclude and discuss the implications of our study.

2. The model

2.1. Formalizing learning: the individual decision problem

We define learning as finding profitable choices among alternatives of uncertain profitability. Actors do so by integrating experiences from earlier decisions to optimize further decisions. We study learning a multi-armed bandit problem (Robbins, 1952; Goyal, 2007; Gittins et al., 2011) in which actors make repeated choices among different options with imperfect information on relative advantages. The updating process for selecting the profitable choices is what we call learning in this study: the closer the actor's beliefs regarding the profitability of a choice are to the true profitability, the more the actor has learned. Outcomes of choices, in this approach, are uncertain because, although one choice may, on average, yield better outcomes than another, the actually realized outcomes vary stochastically.

Formally, the multi-armed bandit can be seen as a set of k pay-off distributions $B = \{b_i \mid i = 1, \dots, k\}$. We assume that actors are fully informed about which distributions are included in B , but not specifically which b_i is which distribution. b_i consists of a given number of possible realizations b_{ij} , and j indexes the possible outcomes for each b_i . Let μ_i be the mean payoff of the payoff distribution b_i . There are N actors who each make one choice per round Y_{it} , where t is a time index and Y_{it} is a number between 1 and k representing which b_i they chose. Afterwards they observe the associated payoff. Thus, they choose a distribution b_i and obtain a random realization from that distribution b_{ijt} . They make such a choice for T times. The actor's objective is to maximize the sum of the collected payoffs $\sum_{t=1}^T b_{ijt}$. At time-point one ($t=1$), actors choose an action and see the linked payoff. At $t=2$, actors again have to make a choice. The choice problem is set up as such that the same choice yields different payoffs when chosen repeatedly. From different observations of each distribution b_i an actor tries to learn which distribution in B belongs to each b_i , and, in particular, which b_i has the highest mean payoff μ_i . Although we also vary the variances in the different distributions b_i and actors can also learn these variances, we do not expect these variances to have a major impact on the decisions of the actors.

2.2. Learning in social networks

Next, we add social networks to this decision problem. A social network consists of a set of N actors and a set G of undirected ties between those actors. We assume that when making a decision, actors use past experiences as well as the experiences of others (Goyal, 2007). Individuals see the payoffs of their own choices as well as the choices and payoffs of their neighbors: actors with whom individuals have a tie. We assume that individuals integrate knowledge on the choices and outcomes of neighbors in the network in making their own decisions, treating information from themselves and their neighbors as equally valuable. This updating process using network information such that the beliefs about the profitability of choice alternatives resembles the true profitability better is what we call *learning in social networks*.

At the individual level, we investigate two network characteristics: *degree* and *betweenness centrality*, which are linked to different theoretical mechanisms and expectations in terms of learning. First, degree centrality captures the amount of direct information one receives, as it indicates the number of neighbors an actor has. Numerous studies have found that the degree centrality of actors

has a positive influence on individual performance. For example, Sparrowe et al. (2001) find that degree centrality in advice networks is positively related to individual performance. In line with this, Ahuja et al. (2003) find that degree centrality in virtual R&D networks positively influences individual performance. Second, betweenness centrality (Freeman, 1977; Wasserman and Faust, 1994) captures the amount of non-redundant information (Burt, 2001). Numerous authors (e.g., Granovetter, 1973; Burt, 2001) argue that the amount of non-redundant information one receives positively affects individuals with respect to job outcomes or creativity (Aral and David, 2012). Moreover, relationships between two otherwise unconnected groups, yielding *structural holes*, benefit the *broker* with regard to outcomes thanks to receiving non-redundant information. In the context of our model, non-redundancy refers to the extent to which actors receive information from independent sources who do not also learn from each other.

More formally, individual network characteristics are defined as follows:

- *Degree centrality*: The number of others with whom an actor is connected:

$$\text{Degree centrality}_i = \sum_{j \neq i} g_{ij} \quad (1)$$

where i is the actor and g_{ij} is an indicator variable for whether actor i has a tie with actor j .

- *Betweenness centrality*: How well an actor is placed in a network with respect to the shortest paths between others that he lies on:

$$\text{Betweenness centrality}_k = \sum_{i \neq j \neq k} \frac{\alpha(i, k, j)}{\alpha(i, j)}, \quad (2)$$

where $\alpha(i, j)$ is the number of shortest paths from actor i to actor j , and $\alpha(i, k, j)$ is the number of those paths on which actor k resides (Freeman, 1977).

At the macrolevel, we investigate the network characteristics *density*² and *centralization*. When a network is denser, all actors in a network will receive more direct information and will receive it earlier. In denser networks, actors update expectations on profitable choices from more direct information from more actors. Furthermore, the more relationships are present in a network, the sooner the information reaches actors, such that they can learn faster and make more profitable choices. However, Goyal (2007) argues that there might be a trade-off between *speed* and learning. He states that when networks have a higher density, information might be exhausted earlier and actors tend to stop experimenting with choice alternatives sooner. As a result, these networks might perform worse because actors conform to the same, but not necessarily the most profitable, choice earlier.

Where density captures the overall connectedness of a network, centralization focuses on the extent to which some actors in a network are more central than others (Freeman, 1978/1979; Snijders, 1981). Networks where individuals observe few others but are observed by many others often tend to perform worse than networks without these *common observers* (Chamley, 2004). Bala and Goyal (1998) offer a mathematical proof and Grund (2012) gives empirical evidence for the proposition that star-shaped networks perform worse than less centralized networks.

² Because we do not vary network size in our simulations or in the experiment, we operationalize density via the average degree. This facilitates comparisons with the variable degree centrality.

Table 1
Number of cards available and their value in the decks.

Value	-200	-100	-50	50	100	200	Average card value	SD
Deck 1	3	1	2	1	1	2	-25	156.79
Deck 2	1	4	1	2	1	1	-25	120.76
Deck 3	1	1	2	1	4	1	25	120.76
Deck 4	2	1	1	2	1	3	25	156.79

Formally, our macrolevel network characteristics are defined as:

- **Density:** The extent to which actors in a network are connected with one another:

$$\text{Density} = \frac{\sum_i \text{Degree}_i}{N}, \tag{3}$$

where Degree_i is the degree centrality of actor i , and N is the total number of actors in a network.

- **Centralization:** The degree to which some actors have more neighbors than others, measured as the standard deviation in the actors' degrees (Snijders, 1981):

$$\text{Centralization} = \sqrt{\frac{\sum_i (\text{Degree}_i - \text{Density})^2}{N}} \tag{4}$$

3. A Decision task for learning in social networks

We further specify the decision problem for our experiment and simulation by using an adaptation of a multi-armed bandit problem, known in psychology as the Iowa Gambling Task (Bechara et al., 1994). In this task, there are four choice alternatives that can be selected. An actor chooses one card from one of four different decks of cards, representing the payoff distributions b_i , in every period. Table 1 shows these decks and the distributions of payoffs. For example, in Deck 1, there are three cards of -200 points, and in Deck 3, there are 4 cards with a payoff of 100 points. There are ten cards in every deck, and each of them is drawn with a 10% chance. Cards are drawn *with* replacement, and cards within decks are shuffled after every draw, such that the result of any draw is independent of all previous draws. Actors, both simulated and human, draw one card per period from one of the four decks and see the payoff of that card in the next period. The actors (both simulated and participants in the experiment) have full knowledge about which cards are available in the decks and with what frequency a priori, but actors do not know which deck is which. That actors know the distribution of different payoffs in the four decks is a variation on Bechara et al. (1994), where participants are unaware of the cards available. Our variation is intended to prevent participants in the actual experiment from having different expectations about which cards are available in the decks, while we as researchers are unaware of these differences between participants. For example, actor x could think that there are cards between -10,000 and +10,000 in the decks, while actor y expects cards between -1000 and +1000. Decks 1 and 2 are not as profitable as Decks 3 and 4: the average expected value of a card drawn from Decks 3 and 4 is higher. Therefore, we label Decks 3 and 4 as profitable. Nevertheless, due to chance it is possible to draw, for example, in five subsequent rounds a -200 card from Deck 4. Thus, the income that deck choices yield is uncertain. While, on average, some decks are better than others, in different periods, the same deck choice can yield different payoffs. The actors in our simulation and experiment try to determine which two of the four decks are on average more profitable by drawing one card from one of the four decks in each period. This means the actors face the problem of having to determine which of the 24 equiprobable permutations of the four distributions is

Table 2
Network characteristics for the four actors (micro) and networks (macro).

Network	a 	b 	c 	d 	e 	f 
Degree actor # 1	0	3	2	2	3	1
Degree actor # 2	0	3	2	2	1	1
Degree actor # 3	0	3	2	1	1	1
Degree actor # 4	0	3	2	1	1	1
Betweenness actor # 1	0	0	0.5	0.8	1	0
Betweenness actor # 2	0	0	0.5	0.8	0	0
Betweenness actor # 3	0	0	0.5	0	0	0
Betweenness actor # 4	0	0	0.5	0	0	0
Density	0	3	2	1.5	1.5	1
Centralization	0	0	0	0.5	.87	0

a, EMPTY; b, COMPLETE; c, CIRCLE; d, LINE; e, STAR; f, TWO DYADS. For numbering see Fig. 1.

the true one. The values of cards in the four decks differ also in variance or standard deviation. Decks 2 and 3 have a standard deviation equal to 120.76, while Decks 1 and 4 have a standard deviation of 156.79. Thus, the variance in Decks 1 and 4 is higher than that in Decks 2 and 3. This might matter for human subjects if they are, for instance, risk averse. We do not incorporate that factor in our theoretical analysis but only test empirically whether subjects have a preference for one or the other profitable deck (Deck 3 or 4). By having all possible cards at least once in all decks, the actors will never know for sure which deck is where.

3.1. Network conditions

Fig. 1 shows the different network conditions (experimental treatments) we use in the simulation and in the experiment. These conditions are labeled EMPTY, COMPLETE, CIRCLE, LINE, STAR and TWO DYADS. This generates the actor types ISOLATE, DYAD, CENTER IN LINE, PERIPHERY IN LINE, CIRCLE, CENTER IN STAR, PERIPHERY IN STAR and FULL. The rationale for choosing this specific set of networks is that within a set of four actors and with the network characteristics used, these networks (1) ensure a large variance between networks and actor types (and, hence, network characteristics); and at the same time (2) they include the most symmetric networks, thus ensuring enough observations per type of actor in these networks. Taking into account other networks does not add more variance with regard to network types and actor types. We model our individual decision problem within these six network conditions. Actors perform the task with three other actors in all conditions. The arrows in Fig. 1 represent the relationships participants have. Next to their own choice and payoff, actors see the deck choices and payoffs of their neighbors after each period. Therefore, neighbors are, next to the actors' own information, sources of information for actors to update their beliefs. Actors have their own four decks of cards, which are reshuffled after every period. In essence, actors know from which deck neighbors draw, but the draws are independent from their own draws.

Table 2 provides values for the network characteristics for the networks used in the simulation and experiment. The actors are numbered as displayed in Fig. 1.³

³ One could also think of larger networks and varying more network characteristics, such as closeness centrality. In this set of networks, however, closeness centrality and degree centrality are too highly correlated to take the former into account.

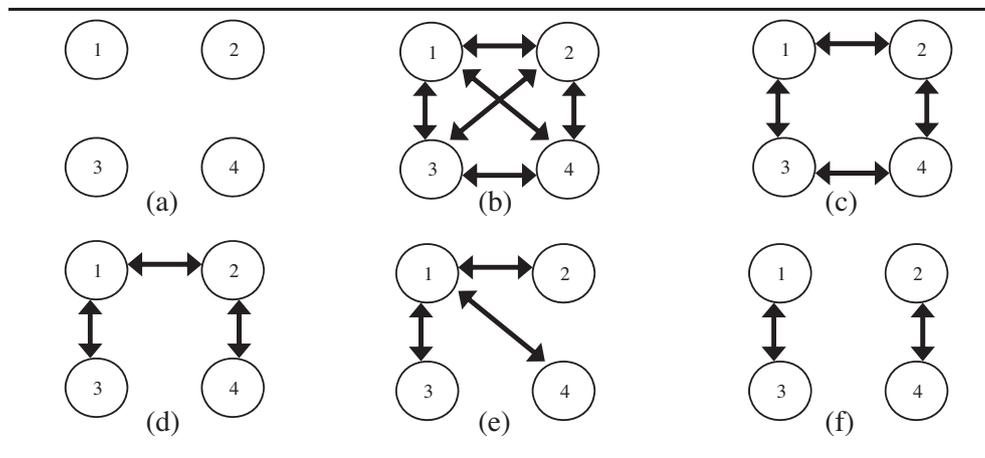


Fig. 1. Network conditions used in the simulation and experiment: EMPTY (a), COMPLETE (b), CIRCLE (c), LINE (d), STAR (e) and TWO DYADS (f).

3.2. Individual and network level outcomes

We study the outcomes of the learning processes at the individual (micro) level and the network (macro) level, which we specify as follows:

- *Individual profitable choices*: a binary variable for whether an actor made a profitable choice per period, or, in essence, whether an actor chose a distribution b_i with the highest mean payoff μ_i in period t .
- *Network profitable choices*: a variable indicating the number of individual profitable choices in a network per period, or, in essence, a count of the actors in a network who chose a b_i with the highest mean payoff μ_i in period t .

4. Bayesian updating and simulation

Because informal accounts on learning easily lead to contradictory arguments, as illustrated above, we simulate actors in our decision problem and derive predictions based on these calculations. We propose that actors in the simulation use rational Bayesian rules for updating beliefs on the profitability of choice alternatives. There are four payoff distributions (decks) b_i that are randomly assigned to a position, and this position is fixed for several periods but is the same for every actor. There are 24 equiprobable ways to order the four decks (permutations). Actors are aware that every deck has an equal chance of being at each position. Because two decks have expected payoffs of -25 and two decks have expected payoffs of 25, the expected payoff of each position at the start of the process is 0. Actors try to find out which b_i is placed at which position (i.e., which of the 24 equiprobable permutations the true one is) by drawing one realization (card) of one of the decks in every period. By drawing these cards, actors update beliefs about which position is the most profitable to select. Because simulated actors know that every possible deck occurs exactly once among the choice options, information about one position also implicitly provides information about the other positions. Updating proceeds by the following set of assumptions adapted from [Bala and Goyal \(1998\)](#) and [Goyal \(2007\)](#):

- Actors are aware of the set of payoff distributions for the choice alternatives a priori;
- Actors update the expected payoffs for each choice in a Bayesian manner using information from their own choices and their neighbors' choices;

- Actors choose an option that has the highest expected payoffs given their updated beliefs;
- Actors are risk neutral and thus do not take the variances in outcomes of decks into account.

Experimenting with choice alternatives at the beginning could be beneficial for long-term payoffs, but there is no straightforward formalization that, in this case, determines how much experimentation would be beneficial for the actors. Therefore, the simulated actors do not experiment but choose every time the option that has the highest expected benefit given their current beliefs. This relatively simple learning mechanism might be a bit more demanding cognitively than other adaptive learning or reinforcement learning mechanisms (e.g., [Charness and Levin, 2005](#)). Still, our Bayesian updating rule implies that decks with good outcomes have a greater likelihood of being chosen in the future. While additional parameters need to be specified for other reinforcement types of learning mechanisms to determine the extent to which certain decks become more likely to be chosen, this is endogenous in the Bayesian updating mechanism. Therefore, while the Bayesian updating may be more demanding in the cognitive sense, it is more parsimonious in the model sense. Therefore, we stick with this mechanism in the present study.

4.1. Simulation results

We simulate the learning process in social networks using the model as described above in four-person networks. By modeling our simulation using the same parameters as in our experiment, we are able to derive predictions directly for our experimental conditions (cf. [Corten and Buskens, 2010](#)). To optimize the fit between our theoretical model and experimental conditions, we use, as much as possible, the same statistical techniques on our simulation data as on our empirical data. In our simulation, each of the six network conditions is played 500 times by 4 actors for 20 periods, resulting in a total of 240,000 observations (i.e., choices).

Running the simulation, we obtained the microlevel descriptive results provided in [Table 3](#). [Table 3](#) shows the mean number of profitable choices over periods 2 through 20 for different network positions. [Table 3](#) reveals that when degree centrality increases, the number of profitable choices increases accordingly. We observe a similar relationship for betweenness centrality. [Fig. 2](#) shows the proportion of profitable choices made by the different actor types per period. The actors have equivalent positions in the networks, with the exception of the LINE and STAR networks, in which some actors are in the center and others in the periphery. In accordance

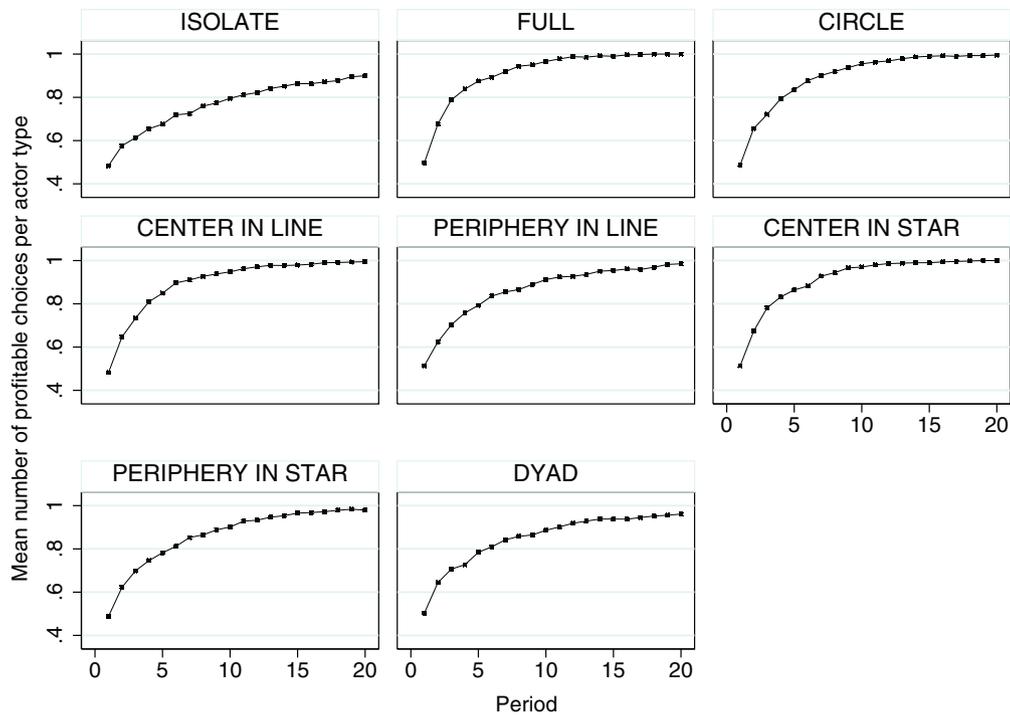


Fig. 2. Simulation: network positions and profitable choices per period.

Table 3

Descriptive simulation results on the mean number of profitable choices per actor in a condition.

Actors	Mean	SD	Observations
Degree = 0	14.889	4.825	2000
Degree = 1	16.654	3.531	4500
Degree = 2	17.456	2.407	3000
Degree = 3	17.775	2.145	2500
Betweenness = 0	16.503	3.764	8500
Betweenness = .5	17.445	2.374	2000
Betweenness = .8	17.478	2.471	1000
Betweenness = 1	17.764	2.135	500

with Table 3, actors in the positions CENTER IN STAR, FULL and CIRCLE perform best, while ISOLATE performs worst. Finally, actors seem to learn fast in the first few periods, but the pace of learning decreases with time.

Table 4 provides descriptive results from the simulation on the macrolevel and provides network parameters and the mean number of profitable choices for periods 2 through 20 per network. The mean number of profitable choices peaks when the density is

Table 4

Descriptive simulation results on mean number of profitable choices per network.

Condition	Density	Centralization	Mean	SD	Observations
a 	0	0	59.556	10.054	500
b 	3	0	71.110	8.558	500
c 	2	0	69.780	7.958	500
d 	1.5	0.5	68.526	8.888	500
e 	1.5	0.87	68.090	9.300	500
f 	1	0	65.990	10.757	500

a, EMPTY; b, COMPLETE; c, CIRCLE; d, LINE; e, STAR; f, TWO DYADS.

highest. When the density decreases, the mean number of profitable choices decreases accordingly. When centralization is zero, averaged over EMPTY, COMPLETE, CIRCLE and TWO DYADS, the mean number of profitable choices is 65.526. When centralization rises, the mean number of profitable choices increases accordingly. Fig. 3 provides the mean number of profitable choices for all actors in a network per period. We see that COMPLETE performs best, followed by CIRCLE and STAR. LINE, TWO DYADS and EMPTY perform worst. We observe again that the speed of learning decreases with time.

Table 5 provides the random effect logistic regression results for the microlevel network characteristics. The response variable is a dichotomous variable of whether an actor made a profitable choice in a given period for periods 2 through 20. In the first period, actors have not yet received information and are not yet capable of updating beliefs about profitability from either their own information or the information of neighbors (i.e., in the first round, a deck is chosen randomly). Therefore, we exclude the first period. Because observations are clustered within actors and groups, we add a random term for the group level. We do not need a random term at the actor level because all actors are programmed to behave similarly in the simulation, and therefore, there is no unobserved heterogeneity at this level. We control for the logarithm of time, as the descriptive results showed a slower pace of learning over time. It is reasonable to assume that in the first few periods, actors learn faster than in later periods. Indeed, we see that this predictor is significant (OR = 5.199; $p < .001$). The model fit statistic indicates that at least one predictor differs significantly from zero (Wald $\chi^2(3) = 18587.650$; Prob. $> \chi^2 = .000$) and that the network model does fit better than the baseline model (LR-test: $\chi^2(3) = 1018.580$, $p < .000$). We also see that the Intra-Class-Correlation (ρ) is .458, and therefore, we can state that 45.8% of the propensity to choose a profitable choice can be attributed to networks rather than individual actors. We see that when the degree increases by one unit, the odds of choosing a profitable choice increases by 117.5%, all else constant (OR = 2.175; $p < .001$). When betweenness centrality increases by one unit, the odds of choosing a profitable choice decreases by 33.2%, ceteris paribus (OR = .668; $p < .000$). This seems to contradict

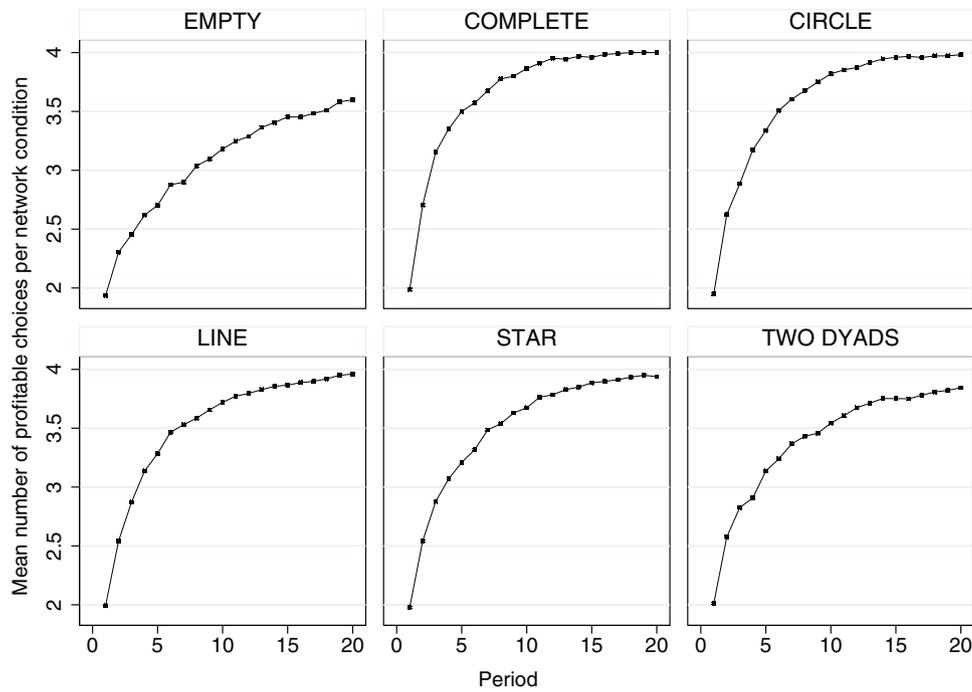


Fig. 3. Simulation: network conditions and profitable choices per period.

Table 5

Microlevel simulation results: random effect logistic regression on profitable choice per period.^a

	Baseline model			Network model		
	Odds ratio	SE	ρ	Odds ratio	SE	ρ
Fixed						
Constant	0.685	(0.29)	0.000	0.231	(0.013)	0.000
Degree				2.174	(0.078)	0.000
Betweenness				0.668	(0.045)	0.000
Log. time	5.159	(0.063)	0.000	5.199	(0.064)	0.000
Random						
σ^u	1.834	(0.032)		1.688	(0.029)	
ρ	0.506	(0.009)		0.458	(0.009)	
Log likelihood	-58288.257			-57778.968		
Wald χ^2 (df)	18078.050 (1)			18587.650 (3)		
Prob. > χ^2	0.000			0.000		

^a Level-1 observations = 228,000; level-2 observations = 3000; likelihood ratio test baseline model versus network model: $\chi^2(2) = 1018.580$, $p = .000$.

the observation in Table 3, which is due to the correlation between degree and betweenness centrality.

Finally, when we pairwise compare different actor types as dummies in the same analysis as Table 5 (random effect logistic regression with a random term for the group level with a Bonferroni correction for multiple comparisons), we see, consistent with the degree centrality effect, that FULL performs best, followed by center in STAR, CIRCLE, CENTER in LINE, PERIPHERY in STAR, PERIPHERY in LINE, DYAD and ISOLATE ($p < .05$, except for PERIPHERY in LINE versus DYAD, PERIPHERY in STAR versus DYAD, CIRCLE versus CENTER in LINE, PERIPHERY in STAR versus PERIPHERY in LINE and CENTER in STAR versus CIRCLE: see Appendix A for the specific coefficients and confidence intervals). More specifically, we compare FULL (betweenness = 0) with CENTER in STAR (betweenness = 1) and CIRCLE (betweenness = .5) with CENTER in LINE (betweenness = .8) to examine betweenness effects more precisely (degree centrality is constant between these actor types). FULL performs better than CENTER in STAR, and CIRCLE performs better than CENTER in LINE, as we also would expect from our negative odds ratio for betweenness centrality from the analyses found in Table 5.

Table 6 provides random effect generalized least squares regression (GLSR) results.⁴ When we run a Breusch-Pagan/Cook-Weisberg test for heteroscedasticity (using ordinary least squares regression (OLSR), we see that the null hypothesis of homoscedasticity (constant variance) is rejected ($\chi^2(1) = 13266.33$, $p < .001$). With a random effect GLSR, we reduce the likelihood of statistically inefficient results or even misleading inferences. In our random effect GLSR, we estimate the effects of network characteristics on the number of profitable choices made by all actors in a network in a period, excluding period one. We add a random term for the group level to account for the multiple observations of the same group. At least one of the covariates differs significantly from zero (Wald $\chi^2(3) = 16904.600$; Prob. > $\chi^2 = .000$), and 31.5% of the variance in

⁴ Alternatively, we could use logistic regression for grouped data (e.g., Corten and Buskens, 2010) in a multilevel framework. This approach accounts for the number of successes within a group and thereby adjusts standard errors. Considering the already complex data structure and extensive analyses, we chose to use random effect GLSR in both the simulation and empirical analyses.

Table 6
Macrolevel simulation results: random effect generalized least squares regression on profitable choices per period.^a

	Baseline model			Network model		
	Coef.	SE	<i>p</i>	Coef.	SE	<i>p</i>
Fixed						
Constant	2.252	(0.014)	0.000	1.923	(0.021)	0.000
Density				0.202	(0.010)	0.000
Centralization				0.112	(0.025)	0.000
Log. time	0.576	(0.005)	0.000	0.576	(0.005)	0.000
Random						
σ^u	0.503			0.467		
σ^e	0.688			0.688		
ρ	0.348			0.315		
Wald χ^2 (df)	16464.900 (1)			16904.600 (3)		
Prob. > χ^2	0.000			0.000		
R^2 within	0.000			0.234		
R^2 between	0.000			0.128		
R^2 overall	0.158			0.200		

^a Level-1 observations = 57,000; level-2 observations = 3,000; running this model using maximum likelihood estimation does not lead to different results.

the response variable is explained by the group level random effect. Our macrolevel descriptive results suggested a slower pace of learning with the passage of time. Therefore, we include the logarithm of time as a control, and we see that this variable indeed has a significant effect ($B = .576$, $p < .000$). The results imply that when the density increases by one unit, the number of profitable choices by actors in a network per period increases by .202 units and that when centralization increases by one unit, the number of profitable choices by actors in a network per period increases by .112 units ($B = .202$, $p < .000$; $B = .112$, $p < .000$). Finally, the overall variance explained in the response variable by this analysis is 20%.

Finally, when we pairwise compare different network conditions as dummies in a random effect GLSR (corrected for multiple comparisons), we see, consistent with the density effect, that COMPLETE performs best, followed by CIRCLE, STAR, LINE, TWO DYADS and EMPTY ($p < .05$, except for CIRCLE versus LINE, STAR versus LINE, STAR versus CIRCLE and COMPLETE versus CIRCLE; see Appendix B for the specific coefficients and confidence intervals). When we contrast STAR (centralization = .87) with LINE (centralization = .5), we do not see that these two networks differ significantly from one another in making profitable choices ($p > .05$).

4.2. Hypotheses

While the informal arguments about effects of network characteristics on making profitable choices in networks were still ambiguous, the following four hypotheses for our experimental set-up follow directly from our behavioral assumptions and the simulation results in Tables 5 and 6.

- **H1.** The higher the degree centrality of actors, the more profitable choices actors make.
- **H2.** The higher the betweenness centrality of actors, the less profitable choices actors make.
- **H3.** The higher the density of a network, the more profitable choices actors in a network make.
- **H4.** The higher the centralization of a network, while controlling for the density of a network, the more profitable choices actors in a network make.

5. Experiment

Participants in the experiment face the gambling task as described before. Participants start with 1200 points and lose or gain points depending on their draws in the experiment. They receive monetary rewards (1000 points equals 1.6 Euros) based

on performance after the experiment. They observe their own and their neighbors' draws and outcomes on the screen. If no relationship is present between them and a member in the group, a question mark is displayed on the computer screen. Participants can keep track of decisions and payoffs of neighbors thanks to a table on their screens that includes this information for every period played thus far. As a result, we reduce the unobserved differences between subjects related to their ability to memorize previous events. Participants do not receive information about the structure of the networks, in essence, the relations between other group members. Therefore, individual and network level outcomes are only the results of learning processes, without confounding these processes with different expectations about information received from well-connected or peripheral neighbors. The decks are randomly placed per treatment under labels A, B, C or D, and these positions are equal for every participant in a treatment. Actors gather information from their neighbors about where the profitable card decks are placed on the screen (with the exception of EMPTY) by drawing one card from one of the four decks in every period. Every participant plays every network for 20 rounds and receives 1200 points before every network "game" (a total of $6 \times 20 = 120$ rounds are played by each participant). To control for sequence effects, the six network conditions were ordered differently at each session (see Appendix D).⁵ Participants were randomly assigned to a group before each treatment and were randomly placed at a network position within this group. The group and position were fixed for the 20 periods of a treatment.

Data were collected between December 2012 and January 2013. The experiment was programmed and conducted using the z-Tree software (Fischbacher, 2007). Subjects were recruited among students at Utrecht University using the Internet recruitment system ORSEE (Greiner, 2004). Six sessions took place at the Experimental Laboratory for Sociology and Economics (ELSE), where a total of $N = 144$ participants took part in the experiment. There were 62 (43.1%) male participants and 82 (56.9%) female participants. Furthermore, 66 (45.8%) participants were undergraduate students, 51 (35.4%) were graduate students and 27 (18.8%) came from various occupations. Participants received written instructions in English at the beginning of the experiment. These instructions were the same for every participant (instructions with screenshots are found in

⁵ To ensure that there were indeed no order effects, we controlled for order in our analyses. While there does seem to be an order effect, where more profitable choices are made for networks later in the experimental sessions, it did not qualitatively change our results.

Table 7
Descriptive statistics for the variables used in the analyses.

	Min.	Max.	Mean	SD	Obs.
Microlevel dependent variables					
Actor prof. choices in period	0	1	0.710	0.453	15960
Microlevel independent variables					
Degree	0	3	1.514	1.011	15960
Betweenness	0	1	0.190	0.312	15960
Macrolevel dependent variables					
Group prof. choices in period	0	4	2.843	1.006	3990
Macrolevel independent variables					
Density	0	3	1.514	0.922	3990
Centralization	0	0.87	0.235	0.343	3990

Table 8
Descriptive empirical results on the number of profitable choices per actor.

Actors	Mean	SD	Observations
Degree = 0	12.347	5.051	144
Degree = 1	13.123	5.059	300
Degree = 2	14.097	4.702	216
Degree = 3	14.356	4.640	180
Betweenness = 0	13.296	4.988	588
Betweenness = .5	14.278	4.573	144
Betweenness = .8	13.740	4.962	72
Betweenness = 1	13.361	5.357	36

variables at the micro- and macrolevel. A correlation matrix of the independent variables in the empirical data is found in Appendix E.

Table 8 provides descriptive results for the mean number of profitable choices made over periods 2 through 20 by actors based on their degree and betweenness centrality. When degree centrality increases, the number of profitable choices also increases. However, we see more ambiguous results for betweenness, where an increase in betweenness does not show an increase or decrease in the number of profitable choices. Fig. 4 provides the proportion of profitable choices made by different types of actors in the network conditions per period. Although the relations in Fig. 4 are less clear than in Fig. 2, FULL and CIRCLE seem to perform better than other network positions, while ISOLATE performs worse than most others.

Table 9 shows the descriptive results on the macrolevel and provides network parameters and the mean number of profitable choices of rounds 2–20 for networks. When density is highest, the mean number of profitable choices is at a peak and when density decreases, the mean number of profitable choices decreases accordingly. When centralization is 0, averaged over EMPTY, COMPLETE, CIRCLE and TWO DYADS, the mean number of profitable choices is 54.428. When centralization is .5, the mean number of profitable choices is 52.778. When centralization is highest, the number becomes 53.694. Thus, there is no clear relationship between centralization and the number of profitable choices. Fig. 5 provides

Appendix C). Each participant played three practice periods before the actual experiment started to allow them to get used to the computer interface. A questionnaire was presented to the participants after the experiment to obtain demographic information as well as attitudes toward risk and trust. Participants earned between a minimum of 6.5 Euros and a maximum of 18 Euros with a mean of 12.8 Euros for 75 minutes of their time. A computer crash in the second session resulted in a forgone first treatment (TWO DYADS) for that session. Finally, the participants in the experiment made a total of 16.800 decisions (120 periods*120 participants + 100 periods × 24 participants).

5.1. Variables and descriptives

As in our simulation, we use one dependent variable at the microlevel from our experimental data: a binary variable of whether a profitable choice is made per period, after the first period. Likewise, the dependent variable at the macrolevel resembles the dependent variable in the simulation. We use a variable that is a count of the number of profitable choices by four actors in a network per period.

Overall, 11,782 decisions (71%) were profitable, and 5018 (29%) were non-profitable. Table 7 shows the descriptive statistics for the

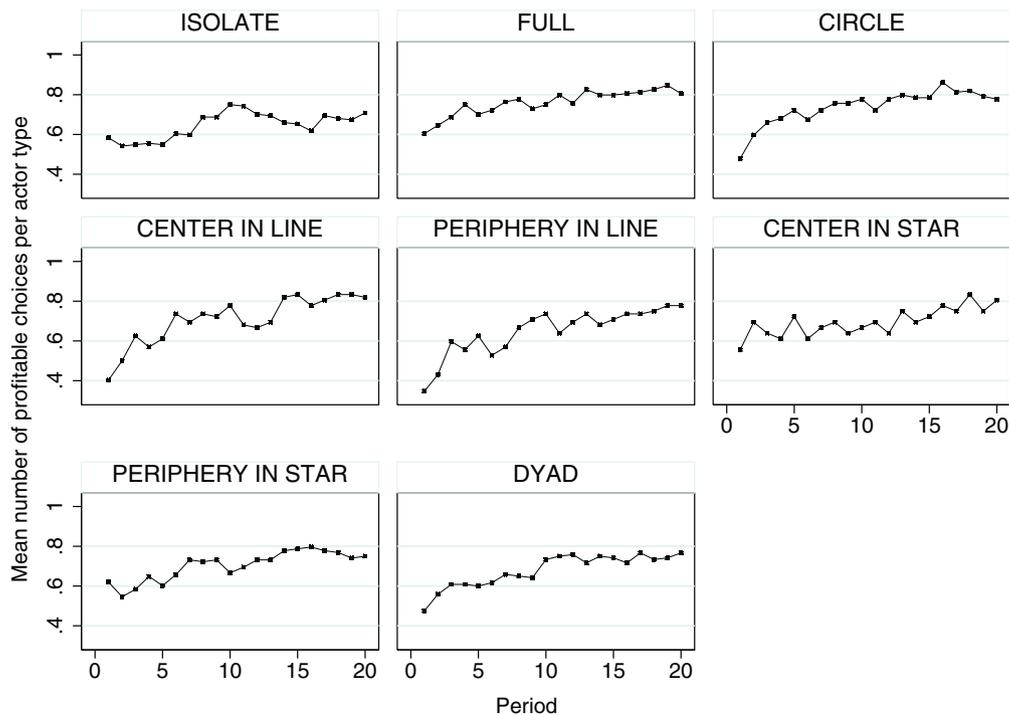


Fig. 4. Empirical: network positions and profitable choices per period.

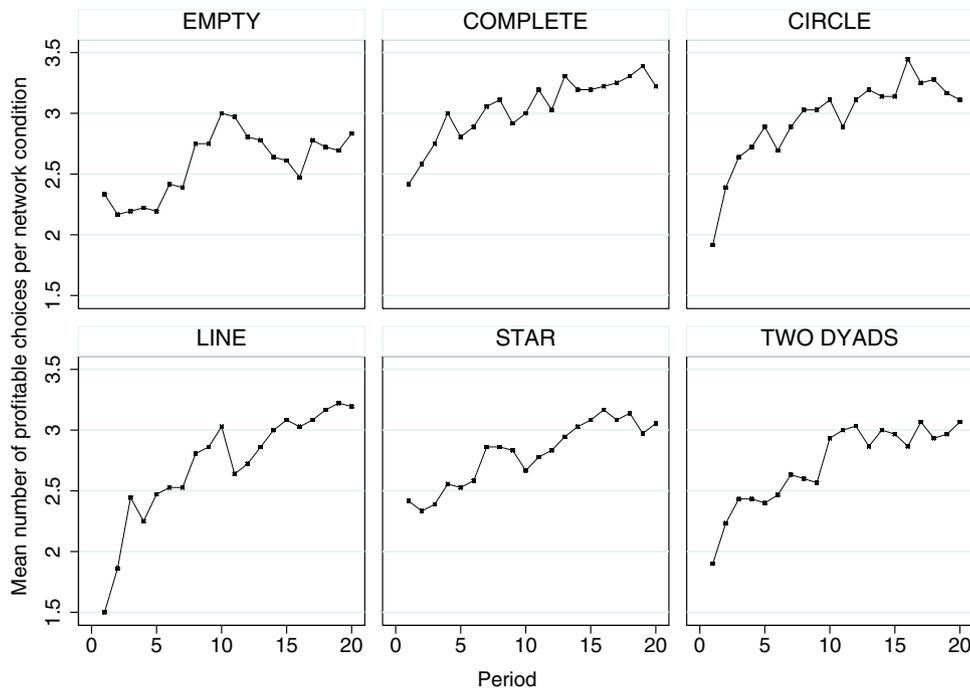


Fig. 5. Empirical: network conditions and profitable choices per period.

Table 9

Descriptive empirical results on mean number of total profitable choices per network condition.

Condition	Density	Centraliz.	Mean	SD	Observations
a 	0	0	49.389	8.794	36
b 	3	0	58.417	10.432	36
c 	2	0	57.111	12.483	36
d 	1.5	0.5	52.778	13.357	36
e 	1.5	0.87	53.694	10.796	36
f 	1	0	52.467	11.383	30

a, EMPTY; b, COMPLETE; c, CIRCLE; d, LINE; e, STAR; f, TWO DYADS.

the mean number of profitable choices for each network condition per period. COMPLETE seems to perform best, while EMPTY seems to perform worst. As already suggested in the simulation’s descriptive results, with the exception of EMPTY, Fig. 5 suggests a learning effect where the mean profitable choice increases in the first couple of rounds and then wears off when time increases.

5.2. Hypotheses tests

Table 10 provides the multilevel logistic regression models for the effect of individual network characteristics on making or not making a profitable choice per period. The analyses on our empirical data in Table 10 resemble the analyses on our simulation data of Table 5. We added a random term for groups and subjects. In contrast to the simulation analysis, we also need to control for subject specific tendencies in choosing and learning because we cannot assume that all subjects choose according to the same principles, while for the simulations we know that every actor does. We control for the logarithm of time and see that this predictor is indeed significant in the baseline model as well as in our model with the

network variables included (OR = 1.886, $p < .001$).⁶ The model fit statistic for our network model predicts that at least one predictor differs significantly from zero (Prob. $> \chi^2 = .000$). When we look at the network model, we see that when degree centrality increases by one unit, the odds of making a profitable choice increase by 32.9% (OR = 1.329, $p < .01$). We also see that the network model fits the data significantly better than the baseline model (LR-test: $\chi^2(2) = 12.72, p < .01$). The effect of betweenness centrality does not differ significantly from zero.

Finally, when we pairwise compare different actor types in a random effect logistic regression (Bonferroni correction for confidence intervals), we see that FULL performs best, followed by CIRCLE, CENTER IN LINE, PERIPHERY IN STAR, DYAD, CENTER IN STAR, PERIPHERY IN LINE and ISOLATE. Although 15 of the 28 coefficients for the comparisons are statistically insignificant ($p > .05$), the order is in line with what we would expect from our simulation results (contrasts between actor types and their confidence intervals are found in Appendix F). However, CENTER IN STAR performs worse in the empirical results compared to the simulation results. When we contrast FULL with CENTER IN STAR (degree centrality is constant), FULL performs significantly better than center in star ($p < .05$). Furthermore, when we contrast CIRCLE with CENTER IN LINE (degree is constant), CIRCLE performs significantly better. Both these contrasts are in line with what we found in the simulation results. While we do not find a significant result for betweenness centrality in Table 10, these specific network position comparisons provide evidence for the negative effect of betweenness centrality that was hypothesized.

Based on these results, we can state that we have found evidence for both microlevel hypotheses. The higher the degree centrality of actors, the more profitable choices they make. Contrasting actors with the same degree centrality, actors with larger betweenness centrality always perform worse.⁷

⁶ A likelihood ratio test indicates that similar models with dummies for the period variable do not fit the data significantly better than the model with the logarithm of time.

⁷ We also ran a random effect logistic regression model (level-1 = period, level-2 = subject, level-3 = network) only for those who made a profitable choice. We then

Table 10
Microlevel experimental results: random effect logistic regression on profitable choices per period.^a

	Prediction	Baseline model			Network model		
		Odds ratio	SE	<i>p</i>	Odds ratio	SE	<i>p</i>
Fixed							
Constant		1.011	(0.108)	0.917	0.669	(0.104)	0.010
Degree	+				1.329	(0.115)	0.001
Between	–				0.909	(0.228)	0.704
Log. time		1.886	(0.062)	0.000	1.886	(0.062)	0.000
Random							
Level-2		1.677	(0.067)		1.674	(0.067)	
Level-3		0.721	(0.101)		0.667	(0.103)	
Log likelihood		–7691.967			–7685.605		
Wald χ^2 (df)		373.630 (1)			385.030 (3)		
Prob. > χ^2		0.000			0.000		

As a robustness analysis, we estimated the models with controls for woman, English and risk. Woman is a binary variable for participants being female (1) or not (0), English is a variable to measure the extent to which participants had difficulty answering the questionnaire after the experiment in English (0 not difficult – 6 very difficult) and risk is a variable that measures the mean of seven statements indicating whether participants are risk averse (6) or risk seeking (0) (e.g., one statement is “I am not willing to take risks when choosing a job or company to work for”; Cronbach’s Alpha of seven variables = .801). While the direction and significance of our network variables do not change after adding these controls, we do find that woman and English are negative predictors for choosing a profitable deck ($p < .01$). Risk does not influence choosing a profitable deck. A likelihood ratio test for the model against its counterpart with woman, English and risk included, and it shows that the models with these controls included do not fit significantly worse than the second model in Table 10 ($p < .01$).

^a Level-1 observations = 15,960; level-2 observations = 840; level-3 observations = 210; Likelihood ratio test against the baseline model; likelihood ratio test baseline model versus network model: $\chi^2(2) = 12.720, p = .000$.

Table 11
Macrolevel experimental results: random effect generalized least squares regression on profitable choices per period.^a

	Prediction	Baseline model			Network model		
		Coef.	SE	<i>p</i>	Coef.	SE	<i>p</i>
Fixed							
Constant		2.022	(0.060)	0.000	1.854	(0.020)	0.000
Density	+				0.167	(0.044)	0.000
Centralization	+				–0.065	(0.119)	0.588
Log. time		.368	(0.019)	0.000	0.368	(0.019)	0.000
Random							
σ^u		0.582			0.563		
σ^e		0.787			0.787		
ρ		0.353			0.339		
Wald χ^2 (df)		350.030(1)			374.550 (3)		
Prob. > χ^2		0.000			0.000		
R^2 within		0.000			0.087		
R^2 between		0.000			0.066		
R^2 overall		0.055			0.079		

^a Level-1 observations = 3990; level-2 observations = 210; running this exact same model with maximum likelihood estimation does not lead to different results.

Table 11 provides random effect GLSR results. In accordance with our simulation analyses, a Breusch-Pagan/Cook-Weisberg test for heteroscedasticity (using OLSR) shows that the null hypothesis of homoscedasticity (constant variance) is rejected ($\chi^2(1) = 17.44, p < .001$), and a (random effect) GLSR is suitable. We estimated the effect of network characteristics on the number of profitable choices made by actors in a network in a period, excluding the first period. We added a random term for the group level. The

checked whether risk aversion positively affected making a profitable choice with a lower standard deviation (Deck 3, see Table 1), as drawing from a deck with a higher standard deviation implies a larger risk than drawing from a deck with the same mean and a lower standard deviation. Risk aversion does not affect the likelihood of making a profitable choice from a deck with a lower standard deviation ($p > .05$). In later stages of the experiment, all participants more often made profitable choices from the deck with a lower standard deviation (Deck 3), while those who found it difficult to answer the questionnaire in English more often made profitable choices from the deck with a higher standard deviation (Deck 4).

analysis on the empirical data of Table 11 resembles the analysis on the simulation data of Table 6. We see that at least one of the predictors significantly differs from zero (Wald $\chi^2(3) = 374.55$; Prob. > $\chi^2 = .000$). Again, we control for the logarithm of time and observe that this predictor is indeed significant ($B = .368, p < .001$). The explained variance of our model in the response variable that can be attributed to the second level term is 33.9%, while the overall variance explained in the response variable by this model is 7.9%. We see that when the density increases by one unit, the number of profitable choices by actors in a network per period increases by .167 units ($B = .167, p < .001$). We do not find an effect of centralization on making profitable choices in networks ($B = -.065, p > .05$).

When we pairwise compare different network types (adjusted for multiple comparisons) and their performance, we see that COMPLETE performs best, followed by CIRCLE, STAR, LINE, TWO DYADS and EMPTY. Although the ordering is completely in line with the density hypothesis and the order found in the simulation, the performance

of COMPLETE differs significantly from EMPTY (exact coefficients and confidence intervals found in Appendix G). In addition, as we also found in our simulation, when varying centralization and keeping the density constant, the performance in STAR does not differ significantly from that in LINE.

With these results, we found evidence in support of hypothesis 3: the higher the density of a network, the more profitable choices actors in a network make. We did not, however, find evidence in support of hypothesis 4: a higher centralization of a network, while controlling for density, does not seem to affect whether actors in a network make profitable choices.⁸

6. Conclusions & discussion

We developed and tested hypotheses explaining the influence of network characteristics on selecting profitable choices through learning. We formalized a decision problem where actors repeatedly make choices of uncertain profitability in small networks. To move beyond informal intuitive arguments on the possible effects of network characteristics on individual and network performance, we derived predictions from computer-simulated data that we tested using data gathered from laboratory experiments.

First, at the microlevel, we derived the hypothesis that the number of relationships an actor has positively influences making a profitable choice (degree centrality: H1). Second, we hypothesized that receiving less redundant information negatively affects making profitable choices (betweenness centrality: H2). Both effects were corroborated in our experiment, which provides evidence for earlier claims based on observational studies that degree centrality is positively related to individual performance (e.g., Sparrowe et al., 2001; Ahuja et al., 2003). However, contrary to previous findings (e.g., Granovetter, 1973; Burt, 2001), we obtained evidence that being in a network position where one receives less redundant information (i.e., betweenness centrality) negatively affects making profitable choices. This last finding needs to be interpreted with some caution because the contrasted network positions that provide evidence for this hypothesis could also be explained by the fact that individual actors in denser networks make more profitable choices than actors in less dense networks. However, because the individual network characteristic betweenness centrality can only be high in networks that are not too dense, these two effects cannot be empirically distinguished. Moreover, density can also be seen as a measure for redundancy, and a positive effect of density on individual performance thus provides evidence that redundancy promotes performance in our experiment.

At the macrolevel, we obtained the hypotheses that both the average number of relationships in a network (density: H3) and the degree to which some actors in a network are more central than others (centralization: H4) positively influences learning and making profitable choices. We only found evidence for the predicted positive effect of density on network performance. The literature is not unanimous concerning the effect of density and centralization on group performance. Some authors argue that both density and

centralization negatively affect network performance (e.g., Bala and Goyal, 1998; Goyal, 2007; Grund, 2012). However, from arguments based on information availability, one could also argue that both measures positively influence performance. The latter effect for density dominates in our set-up both theoretically and empirically.

Before we move to some broader theoretical and practical implications of our study, we first briefly discuss some shortcomings of our experimental approach. First, we investigated exogenous networks, while real-life social networks are (sometimes) endogenous. We chose to keep networks fixed to not confound the effect of network characteristics on learning with decision processes of subjects on changing relationships. For subsequent studies, we suggest keeping average degree in the network conditions constant while varying betweenness centrality as well as other network characteristics, such as closeness centrality. The reason is that the effects of degree centrality and average degree are so strong that they seem to overwhelm the effects of other network characteristics. Furthermore, we suggest examining larger networks to see whether the effects of network characteristics are similar in larger networks.

Overall, we find that direct information is the most convincing mechanism for learning, both at the micro- (degree centrality) and macrolevel (density). Actors update their expectations on profitability of choices mostly with their own information and information obtained directly from their neighbors. Furthermore, based on these first experimental results, we can tentatively state that we find no causal evidence that learning in networks benefits from receiving less redundant information, as, for example, argued by Granovetter (1973) and Burt (2001). In our specific setup of contrasting specific actor types, more unique information might even work against performance. However, we have to keep in mind that four-person networks may be too small to find evidence for these arguments and that the context we examine (learning) is very specific. Future research should examine larger networks to convincingly support these statements. The argument by Goyal (2007), that actors in a network stop experimenting and perform less well in denser networks, is not supported by our findings.

While we believe that our setup addresses some of the core claims of social capital theory about the impact of network structure on the diffusion of valuable information, other aspects of this theory may be less well captured by the model. In particular, the arguments by Granovetter (1973) and Burt (2001) were developed for a context that is *competitive* (e.g., the job market) and in which actors may benefit from receiving certain information earlier than others or in which they can strategically broker information. This aspect is not present in our model, although it could be incorporated in our experimental framework relatively easily. Nevertheless, we stress that the core claims of social capital theory about information diffusion (such as the redundancy argument) are typically formulated such that they do not rely on competition processes. With our relatively simple setup, we show that these claims are at least not generally true.

Finally, we propose suggestions for future research in this area. In general, scholars should focus more on gathering empirical data on multi-armed bandit settings within various networks and on gathering survey data. Because choices concerning uncertain profitability are so common in real-life, it is important to gain knowledge of how social networks could influence the optimization of these choices. Furthermore, investigating larger networks and allowing more variation in network characteristics would increase the generalizability of the results. Scholars may also incorporate learning tasks in experiments and vary the average values of decks to examine whether more difficult or easier learning tasks lead to the same results and conclusions. In addition, scholars may test network hypotheses with learning tasks where decks are chosen more specifically such that influences of risk preferences can be studied better. Given that network relations clearly have

⁸ Because we found statistical evidence for the effect of the logarithm of time on making profitable choices, we tested interactions between the network characteristics and the logarithm of time (both micro- and macrolevel). One could expect that there is variation in the *dynamics of learning* across network positions and network types. We added interactions between the network characteristics and the logarithm of time for both the simulation data and the empirical data analyses to explore whether the dynamics of learning vary across measures. In the empirical-data analyses, none of the interactions differ significantly from zero. In the simulation-data analyses, having a higher degree centrality becomes more important in later stages of the decision task. Regarding the macrolevel network characteristics, both density and centralization become less important for the number of profitable choices actors in a network make in later stages of the decision task.

value in our set-up, an obvious follow-up question is how actors themselves would construct these networks. Some experimental research exists on investment in social networks for solving other problems, such as trust problems (Frey et al., 2012; Raub et al., 2013); this line of work might be extended to learning in networks as well.

Acknowledgements

First, we are thankful for the helpful comments and insights of Anne-Rigt Poortman, Ineke Maas and participants in the SaSR-10 and SaSR-11 seminars. Second, we thank Jeroen Weesie for advice on methodological issues. Finally, we thank two anonymous reviewers from *Social Networks* for their helpful suggestions to improve earlier drafts of this paper.

Appendices A–G. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.socnet.2015.04.011>

References

- Ahuja, M.K., Galetta, D.F., Carley, K.M., 2003. Individual centrality and performance in virtual R&D groups: an empirical study. *Manage. Sci.* 49 (1), 21–38.
- Aral, S., David, V., 2012. The anatomy & dynamics of vision advantages. In: 33rd International Conference on Information Systems, Available at SSRN: <http://ssrn.com/abstract=2172485> or <http://dx.doi.org/10.2139/ssrn.2172485>
- Bala, V., Goyal, S., 1998. Learning from neighbors. *Rev. Econ. Stud.* 65 (3), 595–621.
- Bechara, A., Damasio, A.R., Damasio, H., Anderson, S.W., 1994. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* 50, 7–15.
- Burt, R.S., 1992. *Structural Holes: The Social Structure of Competition*. Harvard University Press, Cambridge.
- Burt, R.S., 2001. Structural holes versus network closure as social capital. In: Lin, N., Cook, K.S., Burt, R.S. (Eds.), *Social Capital: Theory and Research*. Transaction Publishers, Piscataway, pp. 31–56.
- Burt, R.S., 2004. Structural holes and good ideas. *Am. J. Sociol.* 110 (2), 349–399.
- Buskens, V., Van de Rijt, A., 2008. Dynamics of networks if everyone strives for structural holes. *Am. J. Sociol.* 114 (2), 371–407.
- Chamley, C.P., 2004. *Rational Herds: Economic Models of Social Learning*. Cambridge University Press.
- Charness, G., Levin, D., 2005. When optimal choices feel wrong: a laboratory study of Bayesian updating, complexity, and affect. *Am. Econ. Rev.* 95 (4), 1300–1309.
- Conley, T.G., Udry, C.R., 2010. Learning about a new technology: pineapple in Ghana. *Am. Econ. Rev.* 100 (1), 35–69.
- Corten, R., Buskens, V., 2010. Co-evolution of conventions and networks: an experimental study. *Soc. Netw.* 32 (1), 4–15.
- Cowan, R., Jonard, N., 2004. Network structure and the diffusion of knowledge. *J. Econ. Dynam. Control* 28 (8), 1557–1575.
- Falk, A., Heckman, J.J., 2009. Lab experiments are a major source of knowledge in the social sciences. *Science* 326, 535–538.
- Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10 (2), 171–178.
- Freeman, L.C., 1977. A set of measures of centrality based on betweenness. *Sociometry* 40 (1), 35–40.
- Freeman, L.C., 1978/1979. Centrality in social networks conceptual clarification. *Soc. Netw.* 1 (3), 215–239.
- Frey, V., Corten, R., Buskens, V., 2012. Networks and information effects in the emergence of efficient conventions: theory and experimental findings. *Rev. Netw. Econ.* 11 (3), 1–28.
- Gale, D., Kariv, S., 2003. Bayesian learning in social networks. *Games Econ. Behav.* 45, 329–346.
- Gittins, J., Klazebrook, K., Weber, R., 2011. *Multi-Armed Bandit Allocation Indices*. John Wiley & Sons Ltd, Essex.
- Goyal, S., 2007. *Connections: An Introduction to the Economics of Networks*. Princeton University Press, Princeton.
- Graaf de, N.D., Flap, H.D., 1988. With a little help from my friends: social resources as an explanation of occupational status and income in West Germany, The Netherlands, and the United States. *Soc. Forces* 67 (2), 452–472.
- Granovetter, M.S., 1973. The strength of weak ties. *Am. J. Sociol.* 78 (6), 1360–1380.
- Greiner, B., 2004. The online recruitment system ORSEE 2.0: a guide for the organization of experiments in economics. Working Paper Series in Economics, 10. University of Cologne, Germany.
- Grund, T.U., 2012. Network structure and team performance: the case of English Premier League soccer teams. *Soc. Netw.* 34 (4), 682–690.
- Jackson, M.O., 2008. *Social and Economic Networks*. Princeton University Press, Princeton.
- Jackson, M.O., Wolinsky, A., 1996. A strategic model of social and economics networks. *J. Econ. Theory* 71, 44–74.
- Kanoria, Y., Tamuz, O., 2011. Tractable Bayesian Social Learning. Working Paper. Stanford University.
- Kanoria, Y., 2012. Learning in Social and Economic Networks: Research Statement. Working Paper. Stanford University.
- Lin, N., 1999. Social networks and status attainment. *Annu. Rev. Sociol.* 25, 467–487.
- Robbins, H., 1952. Some aspects of the sequential design of experiments. *Bull. Am. Math. Soc.* 58 (5), 527–535.
- Schilling, M.A., Phelps, C.C., 2007. Interfirm collaboration networks: the impact of large-scale network structure on firm innovation. *Manage. Sci.* 53 (7), 1113–1126.
- Raub, W., Buskens, V., Frey, V., 2013. The rationality of social structure: cooperation in social dilemmas through investments in and returns on social capital. *Soc. Netw.* 35 (4), 720–732.
- Snijders, T.A.B., 1981. The degree variance: an index of graph heterogeneity. *Soc. Netw.* 3 (3), 163–174.
- Sparrowe, R.T., Liden, R.G., Wayne, S.J., Kraimer, M.L., 2001. Social networks and the performance of individuals and groups. *Acad. Manage. J.* 44 (2), 316–325.
- Wasserman, S., Faust, K., 1994. *Social Network Analysis*. Cambridge University Press, Cambridge.