



Are bridging ties really advantageous? An experimental test of their advantage in a competitive social learning context

Eva Vriens*, Rense Corten

Department of Sociology/ICS, Utrecht University, The Netherlands



ARTICLE INFO

Article history:

Available online 3 February 2018

Keywords:

Social learning
Bridging ties
Clustering
Competition
Experiment

ABSTRACT

Despite the widespread acceptance of the claim that bridging ties help to obtain profitable outcomes, its underlying mechanisms remain understudied. Starting from a multi-armed bandit problem, we tested the bridging tie hypothesis experimentally by studying the outcomes of social learning for different network positions (in terms of local clustering and closeness centrality) with and without competition. We found a positive effect of bridging ties, but only within one's direct network (i.e., when local clustering is lower), in competitive contexts, and for choices characterized by higher uncertainty. This stresses the importance of outlining more clearly the scope in which the bridging tie hypothesis applies.

© 2018 Elsevier B.V. All rights reserved.

Introduction

Social networks form an important source of information on which individuals base their opinions and actions. They have, for instance, been argued to convey news about new products and technologies (Conley and Udry, 2010; Duflo and Saez, 2003; Kremer and Miguel, 2007) and about job openings (Granovetter, 1995; Lin, 2001). In particular, lots of scholarly attention has been devoted to the hypothesis of Granovetter (1973) and Burt (1992) that actors with ties to otherwise distant social cliques obtain better individual outcomes (e.g., higher paid jobs).

This hypothesis, known as the bridging tie hypothesis, is widely accepted due to its intuitive appeal. Whether its underlying mechanisms apply, however, has up until recently rarely been questioned. The available support is largely based on observational studies that merely established correlations between network characteristics and individual outcomes (de Graaf and Flap, 1988; Lin, 2001). Only recently did scholars start making efforts to expose the underlying causal mechanisms and the evidence is mixed: while some studies (e.g., Conley and Udry, 2010; Mason and Watts, 2012) indeed find support for network effects, others, both observational (e.g., Mouw, 2003) and experimental (Choi et al., 2004; Hofstra et al., 2015; Rutten, 2014), could not relate bridging ties to better individual outcomes.

These results signify the importance of more thoroughly investigating the relationship between bridging ties and opportunities for receiving novel information by studying the exact underlying mechanisms as explicitly as possible, as well as the conditions under which these mechanisms are likely to operate. In this article, we do so by incorporating the impact of competition. In most situations described by seminal studies such as Granovetter (1973) and Burt (1992), the advantage of receiving novel information from one's network is discussed relative to the position of others. That is, the advantage is obtained only if the actor obtains it before others do. This interdependency in actors' behavior is likely to influence their learning strategies (Denrell and March, 2001): Rather than aiming to learn the *best* outcome as soon as possible, actors first and foremost have to obtain a *better* one.

To illustrate, for a scholar to benefit from his network in his pursuit to add to theory development, he should act quickly when he hears a colleague present about a promising new approach on a conference. If he immediately incorporates and improves this approach, he might be able to gain a competitive advantage over other scholars. If he waits too long, the first colleague will have published the results and the knowledge will have become more widespread. More scholars start working on this project, and it is more difficult to still gain a competitive advantage (Burt, 2004). We investigate how this competitive aspect of learning better alternatives before others do affects the relation between network positions and the likelihood of obtaining beneficial outcomes.

Altogether, we more rigorously test the mechanisms underlying the relationship between network positions and individual outcomes. We aim to answer the following questions: To what extent do bridging ties facilitate the actor's goal to obtain the best possi-

* Corresponding author at: Department of Sociology/ICS, Utrecht University, Padualaan 14, 3584 CH Utrecht, The Netherlands.
E-mail address: e.vriens@uu.nl (E. Vriens).

ble outcome? And does the competitiveness of the learning context matter? We investigated these questions in a laboratory experiment, as that enables studying mechanisms of social learning in isolation (Falk and Heckman, 2009) and thereby to identify causal effects of social network positions.

We presented 200 subjects with a variant of the Iowa Gambling Task (Bechara et al., 1994), a decision task in which subjects have to learn the most profitable action among alternatives of uncertain profitability. They made multiple decisions over time and could infer new information (i.e., learn) from the outcomes of their own earlier actions and those of their neighbors. By varying network positions and whether or not the task is competitive, we exposed the conditions that facilitate or hinder the chance of making a profitable choice.

Theory

To understand to what extent information from one's network is used to make decisions we use the model of a multi-armed bandit problem (Robbins, 1952), named after the dilemma a gambler faces when deciding which of K slot machines (one-armed bandits) to play. This model reflects a decision task with uncertainty. Uncertainty, in this case, means that the actor is confronted with a range of possible actions (e.g., slot machines to play) and has to pick one without knowing in advance which generates the highest payoff (Bubeck and Cesa-Bianchi, 2012).

Multi-armed bandit problems are characterized by the stochastic nature of the actions' outcomes. Similar to slot machines, a suboptimal action might return multiple winnings purely by chance and an optimal action might return several failures. Therefore, actors are assumed to learn through reinforcement (Camerer, 2003, chap. 6), as multiple trials are needed to learn which action is more profitable (Auer et al., 1995).

Formally, we consider a range of K different actions $a \in A$ and an unknown state of the world θ , a random, exogenous variable that determines which action is most profitable (denoted as a^*) in terms of payoff maximization (Bala and Goyal, 1998). Although the actor is not informed about the exact value of θ , he does hold certain private beliefs x_i about it, composed of θ and a margin of error ε_i —the latter normally distributed with a mean of 0 and a standard deviation of 1 (DeMarzo et al., 2003). Based on these private beliefs x_i the actor chooses the action that supposedly yields the highest possible profit. When the decision task is repeated over multiple time periods, indexed in discrete steps by $t \in \{1, 2, \dots\}$, the actor can integrate information from earlier experiences to update prior beliefs x_i , which decreases the difference between x_i and θ and therefore increases his chances of learning a^* (Goyal, 2012).

Learning strategies

To solve a multi-armed bandit problem, a payoff-maximizing actor aims to strike a proper balance between two learning strategies: exploration and exploitation (Auer et al., 1995; Mason and Watts, 2012). Exploration involves experimenting with new actions to gather new information to develop alternative solutions (Baum et al., 2000; March, 1991). To illustrate, a clinical researcher applying this strategy would try different treatments to see which provides the best results in terms of curing a disease.

The strategy is rewarding for it creates a large variety of information, which increases the chances of gaining insight into θ and therefore to learn about a^* (Fang et al., 2010). Nonetheless, it is often considered costly, because its outcomes are uncertain beforehand (Denrell and March, 2001). Each new action could both result in better and worse payoffs, meaning that the trial of new actions for future gain could go to the expense of current profit.

To avoid this actors often opt for exploitation instead; a learning strategy that reuses and refines existing knowledge and known solutions (Gupta et al., 2006). Exploiting actors stick to the action that previously generated the highest payoff. The clinical researcher, for instance, cannot endlessly experiment with new treatments on patients at random, hoping to one-day find the perfect cure. Instead, he might opt for the treatment that yielded the best results in earlier years.

It is easy to see that as long as the best solution has not been found, exploitation ultimately is not as profitable as exploration (March, 1991). However, embeddedness in a social network, with others facing the same or similar choice alternatives, shifts this balance to some extent (Mason and Watts, 2012). Within networks, another means of exploitation is to reuse and refine existing information arising from the experiences of others (Lazer and Friedman, 2007). In this respect, exploiting actors use opportunities for social learning, where they improve their decision by inferring information about θ from how others behave and/or from the payoff they receive (Gale and Kariv, 2003; Goyal, 2012). In our example, the clinical researcher might learn that a colleague experimented with a treatment that provided promising results and could opt to use the same treatment on his own patients.

By enabling social learning, exploitation also offers the actor more information than he otherwise would have had (Gupta et al., 2006). However, not only profitable actions spread through the network; information about suboptimal actions might spread as well, certainly when no one has learned a^* . So even though social learning might provide the actor with new information there is no guarantee that the integration of all currently available information helps him to ultimately grasp θ . To sustain learning about a^* , actors facing a multi-armed bandit problem should therefore find a proper balance between exploration and exploitation (Vermorel and Mohri, 2005).

In finding that balance, we follow Bala and Goyal (1998) and assume actors to be boundedly rational decision makers. They integrate information from their own experiences and those of others to increase knowledge about θ , but do not use this information to infer what these actors must have learned from their connections. Furthermore, they use this information to determine their own best course of action, but do not consider how their own actions could influence the behavior of others. Deviating from full rationality, these assumptions simplify the model in terms of tractability and enable translation to real-world settings (DeMarzo et al., 2003; Mobius et al., 2015).

In behavioral terms, it means that we expect the boundedly rational actor to favor social learning over independent exploration (Baum et al., 2000). He follows a neighbor in situations wherein this neighbor repeatedly received a higher payoff and only explores new alternatives when neither he nor his neighbors obtained profitable outcomes. When a single action repeatedly generates profitable payoffs, he sticks to exploiting this action (Lazer and Friedman, 2007).

Learning outcomes in different network positions

To predict the outcome of using social exploitation as a learning strategy, we have to take the network structure into account. That is, the probability that social exploitation enables an actor to learn a^* depends largely on how well connected he is—with connectedness determined by his position within the network (Mason and Watts, 2012). All actions taken by neighbors provide the actor with information, but some neighbors provide more valuable information than others.

To see how, consider a single component network g , composed of $N \geq 3$ actors connected in such a way that $e_{ij} \in g$ means an undirected tie exists between actors i and j (Mobius and Rosenblat,

2014). That is, if information can flow from i to j , it also flows from j to i . We assume information passes through observation and actors observe all actions and corresponding payoffs of their neighbors. Actors cannot withhold information and value information from all neighbors equally, without taking into account differences in connectedness. Along the lines of the previously specified behavioral assumptions, these assumptions restrict possibilities for strategic decision-making. This allows us to hypothesize on differences in the opportunities provided by different network positions over and above individual actions.

More specifically, we hypothesize on opportunity differences originating from the presence or absence of bridging ties: ties that connect the actor to otherwise distant social cliques. It is often argued that actors with one or several bridging ties have better access to new and relevant information, because these ties pass along information that none of the actor's other neighbors have access to (Burt, 1992; Granovetter, 1973). Related to the multi-armed bandit problem, the general mechanism is that the more bridging ties an actor has, the more likely he is to obtain more varied information about θ , which in turn increases his chances of learning a^* . To explain how this works, we distinguish two mechanisms, one operating locally and one in reference to the entire network.

Locally, clustering describes the presence or absence of connections between the actor's neighbors. It is higher when more neighbors are connected to each other and highest when the actor and his neighbors form a complete clique (Watts and Strogatz, 1998). Thus, locally a bridging tie implies lower clustering for it connects the actor to a neighbor that is not connected to his other neighbors. Network simulation models have proven that the higher the degree of local clustering, the slower the pace by which new information reaches the actor (Buskens, 2002, chap. 4; Buskens and Yamaguchi, 1999).

Moreover, it has theoretically been proven that the higher the degree of local clustering, the faster social learning induces convergence of behavior (Ellison and Fudenberg, 1993; Goyal, 2012, chap. 5). The reason is that in structures of high local clustering, the actors' neighbors also learn from each other's experiences, which makes it more likely that they behave similarly. Given that the likelihood of imitation increases when more neighbors take the same action (Broere et al., 2017; Centola, 2010; Granovetter, 1978), it follows that structures with higher local clustering increase the likelihood that actors converge on taking the same action.

In terms of learning a^* , the combination of these two mechanisms yields the prediction that higher local clustering increases the chance of premature convergence. Premature convergence describes a situation where the chosen action, although it provided the highest expected utility given the information available at the time, is suboptimal when related to θ (Fang et al., 2010).

Therefore, it can be deduced that compared to subjects with lower local clustering, those with higher local clustering are less likely to learn a^* : when new and relevant information reaches them in later stages they are more likely to get stuck on suboptimal actions. In an earlier experiment, Mason and Watts (2012) indeed found that it took actors with higher local clustering longer to find the optimal action, if they found it at all. Therefore, our first hypothesis reads:

H1. The higher the actor's local clustering, the poorer his chances of making profitable choices.

A second mechanism through which bridging ties influence the actor's learning outcome is via his indirect connections to other parts of the network. It might be that it is not actor i 's neighbor j that obtained a high payoff, but j 's neighbor k . Since we assume information to spread through observation, actor i will only learn about this via social learning when actor j infers that actor k 's action

was more profitable than his own, chooses the same action in the next decision round, and also obtains a profitable outcome.

Thus, the more intermediary actors are located on the path between actors i and k , the longer it takes for this information to reach actor i (if it ever reaches him at all). From this, we can derive that i is more likely to gather all relevant information present in the network when he is connected to all others through the shortest possible path length, since that means fewer intermediaries are involved that may or may not transmit information about possibly profitable actions (Freeman, 1979; Leavitt, 1951).

This benefit is captured by the actor's closeness centrality, which is higher the shorter the paths to all other actors in the network (Freeman, 1979). Path lengths are shortened considerably by the presence of bridging ties. Hence, the more bridging ties, the higher the actor's closeness centrality. This in turn increases the chance that all relevant knowledge about θ reaches the actor through the behavior of his neighbors.

Network simulations have proven that this mechanism should at least hold for differences related to the number of actors that are connected to actor i in at most two steps (Buskens, 2002, chap. 4; Buskens and Yamaguchi, 1999). Moreover, support for the positive relation between closeness centrality and problem solving derives from an earlier experiment on social learning (Mason and Watts, 2012). We thus propose the hypothesis that:

H2. The higher the actor's closeness centrality, the better his chances of making profitable choices.

Competition in the learning process

Several macro level factors could influence the actor's learning process. An important factor to consider is competition (Denrell and March, 2001). In competitive situations, actors have to outperform each other to obtain the profitable outcome, as the benefit is obtained only if the action was learned before most others have. Competition thus alters the basic assumption behind the social learning model that although the actors' goals are interdependent their payoff for a certain action does not depend on the actions taken by others (Gale and Kariv, 2003).

Instead, in competitive situations payoffs are zero-sum, meaning that an increase in actor i 's payoff goes to the expense of the payoff actor j (Aumann, 1961; Burt, 1992). The industry serves as a clear example. When one or a few actors (companies) innovate their production process, they might benefit from a relative profit gain. This benefit is maintained until the point that all other actors in the network adopted the same approach. As the competitive advantage disappeared, relative payoffs once again stabilized.

This scarcity in profitable payoffs is likely to alter the actor's learning goal. Whereas in noncompetitive situations the actor wants to learn the action providing the most profitable outcome in general, in competitive situations the actor's main interest is to obtain an outcome that is at least as good and preferably better than that of all other actors in the network. In these circumstances, exploitation, with its promise of immediate benefits, is likely to increase in attractiveness further, for it enables a direct comparison with one's neighbors.

To illustrate, if actor i observes actor j obtaining a higher payoff than his own, i would want to find an action that allows him to catch up with j . Although exploration potentially results in such an outcome, one cannot know for sure. If a poor outcome were obtained instead, it would give j an even larger lead. Imitating j 's action is more attractive, because it is more certain to end actor j 's comparative advantage. In that regard, competition invokes imitation learning (Camerer, 2003, chap. 6) on top of general processes of social learning, meaning that the actor has even more incentive

to incorporate information from one's network than he would have had otherwise.

In general, exploitation strategies are considered profitable under competitive payoffs. In fact, [Duersch et al. \(2012\)](#) have proven that the imitation strategy is unbeatable in many zero-sum games, precisely because it influences not only on actor i 's payoffs, but also those of the imitated actor j . Nonetheless, this increased reliance on exploitation strategies resulting from competition can be expected to affect the payoff obtained differently depending on one's network position. Recall that in terms of social learning, actors are more likely to stick to exploiting one action the more of their neighbors do the same. This tendency of behavioral convergence is likely to become even stronger if, due to competition, actors are more focused on the behavior and outcomes of others around them. They will be more likely to use social exploitation strategies already when only one of their neighbors chose an action that resulted in a better payoff—let alone when multiple neighbors did. Add to this the higher risk of premature convergence the sooner convergence takes place ([Goyal, 2012](#)) and it can be inferred that:

H3. Competition reinforces the negative relationship between the actor's local clustering and his chances of making profitable choices.

The same logic can be applied to the effects of indirect connections to all others in the network. When competition drives actors to exploit the information inferred from their neighbors' behavior sooner than they otherwise would have done, information will spread throughout the entire network at a faster pace. Compared to other actors, the actor that is connected to all others through a shorter path length will then also be more likely to observe all relevant information available within the network sooner, which in turn increases his chances to understand θ and thus to gain a competitive advantage over other actors. Actors with higher closeness centrality are connected to all others through a shorter path length. Therefore, they have a competitive advantage over their counterparts when information spreads faster, as their chances to be one of the first to learn about the better outcome through their network increase. Translated into our final hypothesis, we thus expect that:

H4. Competition reinforces the positive relationship between the actor's closeness centrality and his chances of making profitable choices.

Learning task

To study social learning and assess the proposed hypotheses we translated the multi-armed bandit problem to the laboratory context by adapting the Iowa Gambling Task (IGT) ([Bechara et al., 1994](#)), a task originally developed in Psychology for research on emotion and cognition. In the IGT, subjects are asked to draw a card from one of four card decks that differ from each other in terms of average profitability and the distribution of cards in the decks. The subjects' goal is to find the most profitable card deck, since that yields them the highest sum of payoffs. Learning the position of this card deck is possible by introducing repetition: The subject repeatedly draws cards from the different decks and thereby increases insight into each deck's profitability ([Robbins, 1952](#)). This updating process between each decision is what we refer to as learning.

Design

For the basic design applied in this study we follow the set up used by [Hofstra et al. \(2015\)](#). Contrary to the original IGT ([Bechara et al., 1994](#)), this set up gives subjects full knowledge about the division of card types over the decks. This prevents unobserved differences between subjects' expectations about the type of cards

Table 1
Overview of type of cards available in each deck.

Card deck	-100	-50	-25	25	50	100	Average card value	Variation
Deck 1	11	7	6	4	4	8	-12.5	High
Deck 2	7	7	10	8	4	4	-12.5	Low
Deck 3	4	4	8	10	7	7	+12.5	Low
Deck 4	8	4	4	6	7	11	+12.5	High

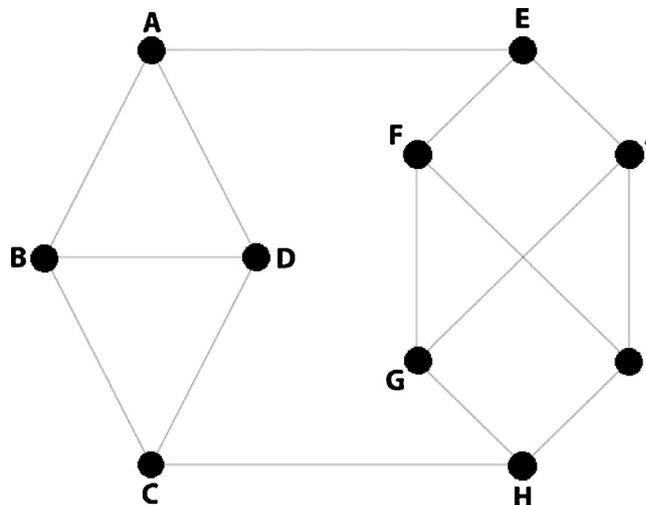


Fig. 1. Network structure used in experiment.

available. The four decks consisted of 40 cards each, with values -100, -50, -25, 25, 50, and 100.

The number of cards of each type differed per deck (see [Table 1](#) for the distribution). Two decks were profitable, providing an average payoff of +12.5, and two unprofitable, returning -12.5 on average. Similar to [Hofstra et al. \(2015\)](#), we introduced differences in variance in the decks' distribution to increase uncertainty about each deck's profitability and to control for differences in subjects' risk preferences. One profitable and one unprofitable deck had higher internal variation; meaning that they contained more extreme value cards than the other two.

Before the start of each treatment, the four decks were randomly placed in positions A–D and for each subject the 40 cards within the decks were independently and randomly shuffled. Thus, although the decks were in the same position for all subjects in the same treatment, different subjects would not necessarily draw the same card if they chose the same deck. Furthermore, cards were drawn with replacement and cards within the deck were shuffled after each draw. Therefore, profitable card decks could also return poor outcomes several periods in a row. This created an uncertainty in choice alternatives that supposedly invoked the exploration-exploitation trade-off, where subjects face the dilemma of exploiting decks that did well in the past or exploring decks that might return better payoffs in the future.

Network positions

Subjects were not only aware of their own history of choices and corresponding card payoffs, but were also informed on the choices and payoffs of others. We opted for a network of $N=10$ subjects connected to $k=3$ neighbors each (the latter to rule out confounding influences of degree centrality). A total of 19 different network structures comply with these restrictions. As most of these networks have either little to no variation in the network positions of interest or a very high correlation between these positions, we opted to use a single network structure over all sessions and treatments. This network (displayed in [Fig. 1](#)) was chosen for

it had the highest variation in terms of local clustering ($SD = .401$) and closeness centrality ($SD = .327$), while at the same time the correlation between the two was minimized as much as possible. This increases the possibility of disentangling the separate network effects (Valente et al., 2008).

Subjects were not informed about the network structure other than its total size and the constant of $k=3$ neighbors per subject. In line with our behavioral assumptions, this feature made it impossible that actors take into account that some of the information they receive might be from redundant sources—they do not know whether this is the case. Preventing any noise resulting from differences in skills and strategic behavior resulting from such knowledge, this provides a clean test on the effect of differences in the opportunities provided by different network positions.

Competition

We made two variations on the learning task: one competitive, the other noncompetitive. In the noncompetitive treatment, subjects could gather information by observing their neighbors' actions and payoffs, but each subject's payoff was determined independent of his neighbors' actions (cf. Hofstra et al., 2015). In the competitive treatment, contrarily, the outcomes of all network members were interdependent. Subjects were primed to compare their own performance to that of the other network members, as they were informed about their own position relative to the rank of the other network members.

Moreover, before the start of this treatment, subjects were informed that after the final round was played, the three subjects with the best scores would obtain another 400, 200, and 100 points, whereas the three worst performing subjects would lose 400, 200, and 100 points. This division was chosen for it entailed a clear and substantial gain or loss on top of the total payoff. To illustrate, the best performing subject would earn a bonus worth four times the highest card type. Given that even averaged over the two profitable decks the probability of drawing a 100-point card is only $9/40$, this bonus creates substantial incentive to follow neighbor j if he did obtain this score.

Data and procedure

We translated the IGT into a computerized laboratory experiment using z-Tree (Fischbacher, 2007) and conducted the experiment in February 2016 in the Experimental Laboratory for Sociology and Economics (ELSE) at Utrecht University. Subjects were invited to participate through the web-based ORSEE recruitment system (Greiner, 2004). We conducted 10 sessions with $N=20$ subjects each, resulting in $N=200$ subjects in total. Of these subjects, 177 (88.5%) were students, 126 (63%) were female, and 104 (52%) were Dutch. Non-Dutch participants came from all over the world, most notably from Greece (13), Italy (9), Brazil (6), China (5), Russia (5), and Germany (4). Finally, subjects were between 18 and 51 years old, with an average of 23 years ($SD=4.18$, $Mdn=22$).

All subjects participated in both the competitive and the non-competitive treatment, each consisting of 20 decision-making rounds. The sequence of the two treatments was systematically varied across the 10 sessions to avoid differences between the treatments resulting from increased familiarity with the learning task over the course of the experiment. Before the start of each treatment, the decks were randomly placed on positions A to D. Additionally, subjects were randomly divided into two networks of $N=10$ and randomly assigned a position within the network.

There were separate instructions for each treatment, which were handed out prior to starting that treatment. These instructions, including figures of the experiment's screen layout, can be found in the Supplementary materials (S1). They explain the learn-

ing task, the payoff determination, and the composition of card decks. After reading the instructions subjects were asked several control questions to ensure everyone understood the instructions correctly. In case a wrong answer was given, the screen would show additional clarification.

During the treatments, the screen displayed a table that recorded the previous deck choices and yielded payoffs for both the subject and his or her three neighbors. By documenting the history of events this way, unobserved differences between subjects related to their ability to memorize previous events were minimized. In the competitive treatment, the screen also contained a table that ranked the subject's own performance against that of the other network members (without showing which decks they chose).

After completion of the two treatments, subjects had to fill out a small questionnaire. This questionnaire assessed the subject's risk preferences and contained several questions on socio-demographics. Sessions lasted between 50 and 65 min in total. Subjects were paid according their performances and were informed on this matter beforehand. They started each treatment with 1000 points and earned or lost points with each decision depending on the value of the card drawn. The amount of points earned over the treatments was translated into a monetary payoff with 200 points worth €1.00. Subjects earned between €3.50 and €20.00, with an average of €11.22 ($SD=2.92$) per subject. By paying subjects according to their performance, we assured that they all had the same aim during the experiment: to obtain the highest possible monetary reward (Falk and Heckman, 2009).

Measures

For each decision round, it was recorded whether the deck chosen by the subject was in fact deck 3 or deck 4—i.e., one of the profitable decks. Per subject, we calculated the *sum of profitable choices* per treatment by counting over all rounds how often the subject chose deck 3 or deck 4.

Local clustering was calculated by dividing the number of relations among the subject's neighbors by his degree centrality. As degree centrality was fixed to $k=3$ neighbors and the network was comprised of a single component, the clustering coefficient ranged from 0 (when none of the neighbors were connected to one another) to .667 (when two relations were present among the subject's neighbors). We unit standardized the variable so that the transformed minimum and maximum are 0 and 1, thereby simplifying comparisons between positions.

The subject's *closeness centrality* was measured as the inverse of the average length of the shortest paths to all other subjects in the network (Freeman, 1979). The higher the coefficient, the more shortest paths cross the actor. In the network used in the experiment, the subjects' closeness centrality ranged from .053 to .059. For interpretative purposes we again applied unit standardization based on the observed minimum and maximum score.

To measure *competition*, finally, a dichotomous variable indicated on the treatment level whether or not the treatment was competitive.

Descriptive statistics for these four variables are reported in Table 2.

Analytical strategy

Hypotheses were tested twofold. First, we used one-sample and paired sample t -tests to obtain overall tests of whether the effects hold on the network/treatment level. To test whether the sum of profitable choices depends on the position within the networks we calculated for each of the $g=40$ networks (20 networks per treatment) the Pearson correlation coefficients between the subjects'

Table 2
Descriptive statistics of the dependent and independent variables.

	<i>N</i>	<i>M</i>	<i>SD</i>	Min	Max
Individual/treatment level					
Clustering centrality	400	.300	.401	0	1
Closeness centrality	400	.600	.327	0	1
Sum of profitable choices	400	13.953	4.636	0	20
Treatment level					
Competition	20	.500		0	1

network position and the sum of profitable choices over the course of the treatment. By aggregating the data this way, we take care of the interdependencies of repeated choices within subjects as well as the interdependencies between subjects within the same networks. Subsequently, we tested whether the mean of these correlations deviated significantly from 0 using one sample *t*-tests. Significant negative correlations would signal support for H1, while significant positive correlations would be in line with H2.

To test interaction H3 and H4, we compared for each network position the average correlation in the noncompetitive context to that in the competitive context. We used paired sample *t*-tests, with observations paired on the network level. Expressing the difference in terms of the added effect of competition (i.e., Difference = Competition – No competition), we expected a negative difference for H3 and a positive difference for H4.

The *t*-tests provide a clean test of an overall effect, but do suffer from low power due to the low amount of observations on which the correlation variables are based (*N*=10) and the low amount of correlations to be compared on the network * treatment level (*N*=20). To get a more precise measure of the magnitude of the effects, we also conducted multilevel OLS regressions on the subject/treatment level (*N*=200 subjects * 2 treatments = 400), with treatment scores nested in subjects and cross-classified in two networks. The advantage of this method is that it allows for an estimation of the difference in the amount of profitable choices associated with different network positions or treatment conditions. A disadvantage is that the analysis method with cross-classification in two different networks does not allow to also control for a higher-level nesting in different sessions. Therefore, not all interdependencies in behavior are accounted for.

The analyses were performed in Mplus (Muthén and Muthén, 2017), using a Bayesian estimator with noninformative priors (default in Mplus). The models (one for the main effects and another in which the interaction terms were added) were estimated using 50,000 iterations and a 50% burn-in phase. Convergence was assessed by the Gelman-Rubin criterion for a cutoff value of .01 (Gelman et al., 2004) and by manual inspection of whether the chains in the trace plots converged to the same target distribution (all trace plots are added in the Supplementary materials [S2] of this paper).

Finally, by means of robustness check we performed separate tests for the number of times deck 3 (profitable with low internal payoff variation) and deck 4 (profitable with high internal payoff variation) were chosen.

Results

Descriptive results

With two out of four card decks considered profitable, we would expect subjects to make a profitable choice in 50% of the decisions if choices were random and no learning occurred. As can be seen from Fig. 2, this was indeed the case for the first round, when decisions were made without prior knowledge about the type of cards available in each deck. In later rounds, a learn-

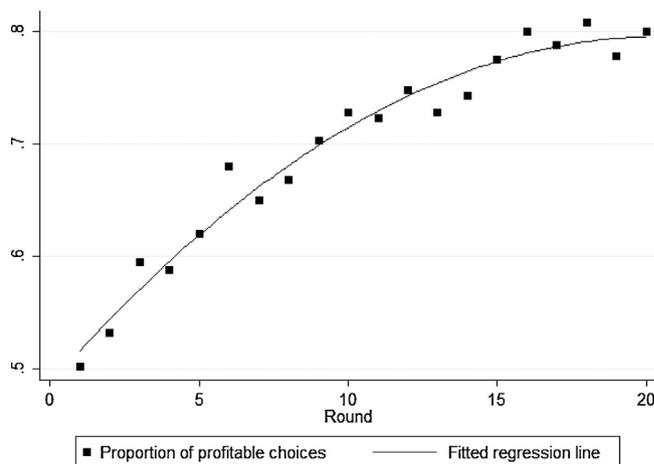


Fig. 2. Proportion of profitable choices made per decision round.

Table 3
Paired and One Sample *T*-tests for the average correlation between network positions and the number of profitable choices, compared over treatments.

	ρ_M	ρ_{SD}	<i>T</i>	<i>df</i>	<i>p</i>
ρ clustering, profitable choice					
All treatments	.026	.463	.359 ^a	39	.721
Noncompetitive treatments	.167	.358	2.085 ^a	19	.051
Competitive treatments	-.114	.520	-.981 ^a	19	.339
Difference	-.281		-1.989 ^b	38	.054
ρ closeness, profitable choice					
All treatments	-.010	.366	-.172 ^a	39	.865
Noncompetitive treatments	-.093	.303	-1.370 ^a	19	.187
Competitive treatments	.073	.411	.794 ^a	19	.437
Difference	.166		1.452 ^b	38	.155

^a Test statistic for One Sample *t*-test.

^b Test statistic for Paired Sample *t*-test.

ing curve becomes apparent, as subjects increasingly drew cards from the more profitable decks 3 and 4. Deck 3—the profitable deck with low internal variation—was chosen 2834 times (35.42%) and deck 4—the profitable deck with high internal variation—2747 times (34.34%). Together, this resulted in 5581 (69.76%) profitable choices, whereas decks 1 and 2—unprofitable on average—were only chosen 1270 (15.88%) and 1149 times (14.36%), respectively.

Fig. 3 plots the sum of profitable choices made by subjects as a function of the two network positions and the treatment’s competitiveness. The figure suggests that in general the sum of profitable choices does not differ for subjects in positions of low or high local clustering (left panel) or low or high closeness centrality (right panel). However, if the relation is plotted separately for the two treatment types differences do arise. Rather than strict negative effects of local clustering, the figure hints at positive effects of local clustering in a context without competition (left panel). Only for competitive treatments do we observe the hypothesized negative relationship between higher local clustering and the probability of making profitable choices. The same goes for the effect of closeness centrality (right panel): the hypothesized positive relation seems to exist only in the competitive context. In the noncompetitive treatments, the relation appears to be negative.

Hypotheses tests

Fig. 3 suggests that the effect of network positions on the sum of profitable choices depends crucially on the context in which it takes place. One sample *t*-tests of the general network effects support this observation (Table 3). The average Pearson correlation between the sum of profitable choices and local clustering was *M* = .026, with

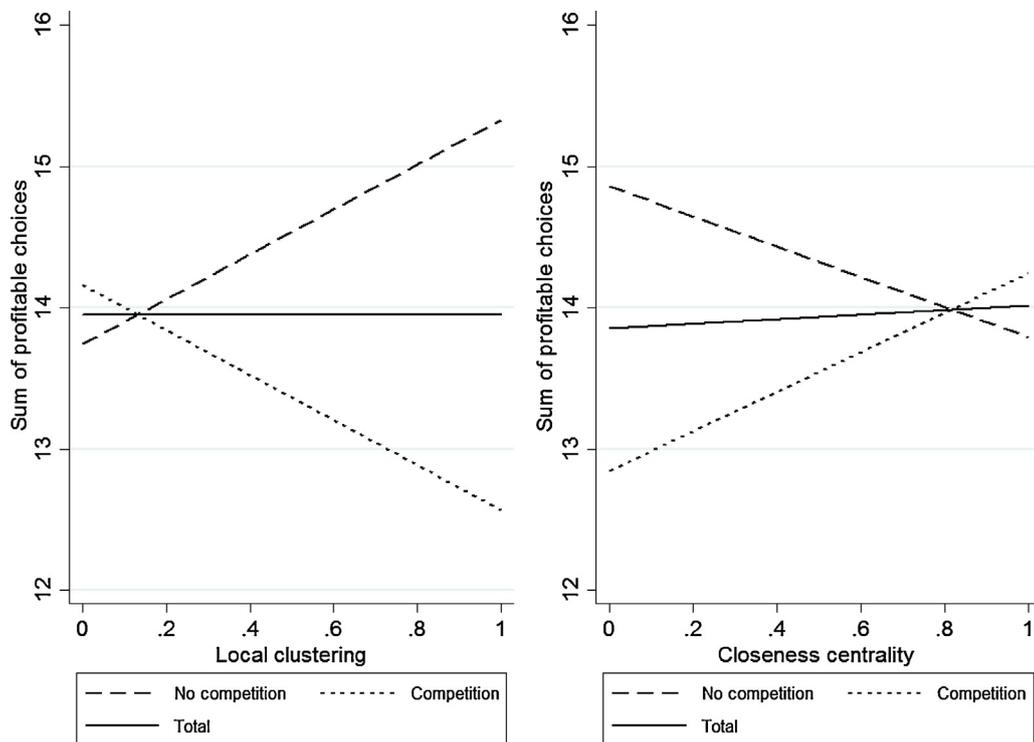


Fig. 3. Sum of profitable choices as a function of network positions for the different treatments.

substantial variation over the different networks ($SD = .463$). As already suggested by the straight line in Fig. 3, the correlation does not significantly deviate from 0 ($T(39) = .359, p = .721$). H1 is not supported. Likewise, the average Pearson correlation between the sum of profitable choices and closeness centrality was $M = -.010$ ($SD = .366$), implying that closeness centrality does not significantly affect the subject's sum of profitable choices either ($T(39) = -.172, p = .865$). H2 is not supported.

Splitting the average correlations for the two treatments, we find that the average correlation between local clustering and the sum of profitable choices was positive in the noncompetitive context ($M = .167, SD = .358$), while it was negative in the competitive context ($M = -.114, SD = .520$). This corresponds to the diverging plot lines in Fig. 3 (left panel). The difference in outcomes between the two treatments approaches conventional significance ($T(38) = -1.989, p = .054$). This signals that subjects in positions of higher local clustering indeed perform worse in competitive contexts compared to noncompetitive contexts: they make fewer profitable choices. Since the main effect is not significant, strictly speaking competition does not reinforce the effect of local clustering. Instead, the results suggest that the negative effect of local clustering is conditional on the treatment type. Nonetheless, the result hints at support for the moderation effect posited in H3.

For closeness centrality, the average correlation was negative over the noncompetitive treatments ($M = -.928, SD = .303$), while it was close to zero for the competitive treatments ($M = .073, SD = .411$). The difference between the two conditions is not significant ($T(38) = 1.452, p = .155$), indicating that H4 is not supported by the data.

Table 4 reports the results of the multilevel regression analyses. Performed on the individual/treatment level, these analyses increase the number of observations and thereby statistical power, allowing for a more detailed interpretation of the effects. Both the model with main effects (Model 1) and the model with interaction effects (Model 2) have Bayesian posterior predictive p -values close to .5, indicating good model fit. Model 2 (which includes the

interaction terms) has a lower score on the hierarchical Deviance Information Criterion.

Looking at the parameter results, Model 1 again signals the lack of overall network effects. In Model 2, we observe a significant positive effect of local clustering without competition. Those in the highest clustering position are estimated to make 3.328 more profitable choices than those in the lowest clustering position ($p = .012$). The interaction of local clustering with competition is negative ($B = -5.638, p = .003$). This means that in competition, instead, those in position of high local clustering are estimated to make 2.334 (i.e., $5.638 - 3.328$) fewer profitable choices than those in position of low local clustering. The outcome is worse than in the noncompetitive treatment, and thereby supports H3. The main and interaction effects of closeness centrality are not significant.

Robustness

To test the robustness of our findings, we also analyzed the effects of network positions and competition for the two profitable decks separately. Although we did not speculate on potential differences in network effects related to the amount of uncertainty surrounding the action's profitability, Fig. 4 hints at different learning curves for locating deck 3 (profitable with low internal card variation) and deck 4 (profitable with high internal card variation). The learning curve for deck 3 is steeper, suggesting that its location was easier to discover.

In relation to our hypotheses, separate t -tests for deck 3 indicate that there are no network or network*competition effects on the number of times deck 3 was chosen by subjects (see Supplementary materials [S3]). Instead, the pattern observed in Table 3 seems to derive solely from the likelihood of learning the location of deck 4. For this deck, the average correlation with the extent of local clustering was $M = .118$ ($SD = .367$) in the noncompetitive context, compared to $M = -.217$ ($SD = .472$) in the competitive context. The difference is statistically significant in the hypothesized direction ($M = .335, T(38) = -2.503, p = .017$). For closeness centrality, the

Table 4
Results of Multilevel Cross-classified OLS regressions (Bayes estimator) of the treatment and network effects on the number of profitable choices.

	Model 1		Model 2	
	<i>B</i>	(95% PPI)	<i>B</i>	(95% PPI)
Subject/treatment level				
Local clustering	.472	(−1.624, 2.603)	3.328*	(.396, 6.293)
Closeness centrality	.739	(−1.826, 3.377)	2.505	(−1.083, 6.131)
Clustering* competition			−5.638**	(−9.812, −1.589)
Closeness* competition			−3.461	(−8.589, 1.570)
Network/treatment level				
Competition	−.537	(−2.033, .986)	3.230	(−1.151, 7.696)
Intercept	13.639***	(11.197, 15.999)	11.718***	(8.535, 14.857)
Residual variance				
Subject/treatment level	17.095***	(14.211, 20.323)	16.725***	(13.890, 19.908)
Subject level	1.254***	(.030, 3.754)	1.132***	(.042, 3.636)
Network level	4.040***	(1.955, 7.888)	4.081***	(1.991, 8.035)
Model fit indices				
Bayesian <i>ppp</i>	.475		.465	
<i>DIC (pD)</i>	2323.419	(55.112)	2319.492	(59.884)

* $p < .05$.

** $p < .01$.

*** $p < .001$.

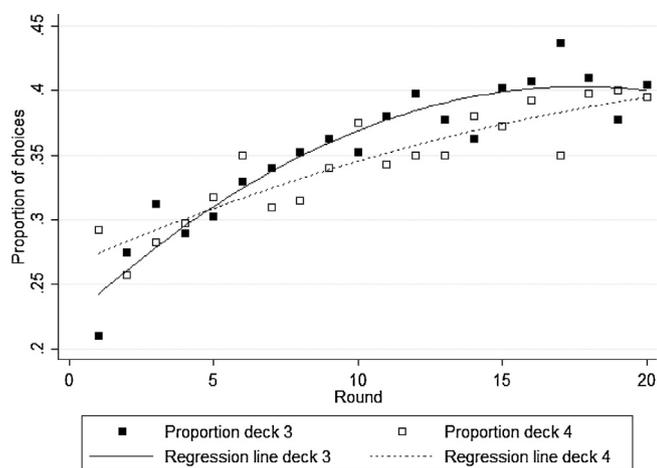


Fig. 4. Proportion of profitable choices made per decision round, separated by deck choice.

average correlation was close to zero in the noncompetitive context ($M = -.024$, $SD = .357$) and positive in the competitive context ($M = .138$, $SD = .357$). Just like for the general effect, this difference is not statistically significant ($M = -.152$, $T(38) = 1.466$, $p = .151$).

Conclusion and discussion

We studied whether bridging ties are advantageous by closely examining the hypothesis proposed by Granovetter (1973) and Burt (1992). In an experiment involving a multi-armed bandit problem we highlighted several conditions that supposedly give rise to bridging tie benefits (i.e., fixed degree centrality and a competitive learning task). We contributed to existing knowledge on the workings and scope of the bridging tie benefit by disentangling, both theoretically and empirically, the bridging tie effect in one's direct connections (local clustering) from that related to the entire network (closeness centrality) and by exploring empirically characteristics of the learning task that give rise to differences resulting from competition and different network positions.

Our main conclusion is that the benefit is not as clear-cut and all-encompassing as the hypothesis suggests. Characteristics of both the context and the profitable action itself seem to matter for the type of network position that is deemed beneficial. As for the con-

text, first of all, in line with Burt's (1992) arguments our results hint at the importance of competition to reap the benefit of bridging tie positions. Specifically, the only robust significant finding was that subjects with more direct bridging ties (i.e., lower degrees of local clustering) made more profitable choices in competitive contexts than in noncompetitive contexts. That is, a lack of local clustering increases the benefits of social learning under competition. In line with the theory, this suggests that competition might more strongly invoke learning strategies of social exploitation, which, in positions of high local clustering, increases the likelihood of premature convergence on a suboptimal action.

Although the results of the experiment also hint at a benefit of indirect bridging ties in competitive contexts over noncompetitive contexts, the resulting difference was not statistically significant. Therefore, the results of the experiment suggest important limits to the scope in which the bridging tie hypothesis applies: it restricts the benefits to direct bridges and to competitive settings. In this light, our findings might also explain why earlier experiments (Choi et al., 2004; Hofstra et al., 2015; Rutten, 2014) did not find an advantage of bridging tie positions. In none of these studies was the task to be solved of a competitive nature, which might actually be crucial.

With respect to characteristics of the profitable action itself, secondly, sensitivity analyses suggested further limits to the scope of the bridging tie hypothesis. We found that the uncertainty surrounding the task's profitability seems to determine whether differences related to network position become apparent. Most subjects learned the position of the profitable card deck with low internal payoff variation in early stages of the experiment. They needed more time to recognize the profitability of the profitable card deck with high internal payoff variation, potentially because the occasional high negative card returns discouraged subjects from repeatedly choosing that deck.

The fact that more trials were needed to learn that this deck was also profitable might explain why the effects of network positions and treatment types stood ground only with respect to the number of times this deck was chosen. To recognize the profitability of the low-variation deck, subjects might not have needed the additional information provided by their neighbors. For the high-variation deck, on the other hand, the additional information provided by the actor's neighbors' trials appeared to have been crucial to recognize its profitability, at least as long as actors actually rely on social exploitation strategies. As these strategies seem to be invoked by a

competitive context, it is in this setting that bridging ties helped to learn the profitability of more complex learning tasks.

These results are in line with research on group-level collective problem-solving, where it is generally argued that networks containing a number of cross-group connections (i.e., moderately efficient networks) perform best when problems are more ambiguous or complex (Lazer and Friedman, 2007; Mason et al., 2008). For clearly beneficial problems, on the other hand, regional clusters are thought to benefit the diffusion of the optimal solution (Gibbons, 2004).

Conclusions with respect to this characteristic should be drawn with caution, though, as the current design inhibits us from grasping what exactly characterizes uncertainty regarding the profitability of the optimal action. With four alternative actions and two of them profitable, the variation in uncertainty is too limited to say what kind of tasks ask for bridging tie positions in order to be solved. At what level of uncertainty can we expect local clustering to become detrimental?

Future research could clarify this matter by making the uncertainty of the learning task the key element of study. Different tasks could be compared, varying in the number and type of (suboptimal) choice alternatives. This could potentially increase insight in the turning point after which actors in positions of low local clustering gain a competitive advantage over their counterparts in highly clustered parts of the network.

Some other issues are at stake as well. For instance, although we improved on previous experiments by increasing the network size to $N = 10$, it remains problematic to properly disentangle the separate network effects. Minor increases in network size would already enable building structures with lower correlation between our two network measures, thereby providing cleaner tests of their separate effects. Moreover, such an increase would introduce structures with more variation on the closeness centrality measure, as the amount of network members that can only be reached indirectly increases. It might be that more indirect connections are needed for information differences resulting from closeness centrality to become apparent.

Second, power issues might have affected the results of the analyses, especially those related to the paired sample *t*-tests. For the separate treatments the number of observations—20 correlations taken over 10 observations each—is rather low, especially considering that the standard deviation of the outcome measure (the number of profitable choices) is by definition inflated by the stochastic nature of the learning task at hand. As profitable decks could also have returned several unprofitable cards, on the individual level it might be that in some instances the profitability of the optimal deck was never revealed. More observations on the treatment level might have been necessary to cancel out this noise.

As for the experiment's external validity, we should recognize that we only captured the mechanism of information opportunities underlying different network positions. In real life several confounding factors might be at play. For instance, we did not consider the control aspect of occupying a bridging tie position, which is another dimension central to Burt's (1992) theory.

The extent to which an actor's network position enables him to control the flow of information throughout the network is usually captured by betweenness centrality (Freeman, 1979). The higher the betweenness centrality, the higher the control over information flows. The more an actor can control the information flows, the more he might be able to use the available information to his own advantage. Future experiments could incorporate this feature by giving each subject the choice whether or not to share his choices and payoffs with others. Assuming there is no need to exercise such control in noncompetitive situations, perhaps that reveals stronger reinforcement of the effects of betweenness centrality in competitive situations.

Additionally, in our experiment subjects did not know to whom they were connected and what the network structure looked like beyond their direct ties. In real-world settings the subjective value actors place on different neighbors can be expected to exert substantial influence. This could mean that, in line with our hypotheses, more weight is placed on neighbors that have a better network position. In that case we underestimated the effect of closeness centrality and overestimated that of local clustering.

Other reasons to attach different weights to the information of different network members, not related to their structural position, can be thought of as well. For example, one may place more weight on the experiences of a family member, without there being a relation to the relative experience this person has or the type of information he receives from his network.

Regardless of such potentially confounding factors in real life, we have undoubtedly established the importance of low local clustering in terms of opportunities to learn novel information. We advocate for future research to translate these findings to non-laboratory settings. Is it possible to find the same effects in the field, now that we know more about the type of settings in which they should occur? Can we, for example, find an effect of local clustering in a network of producers in a competitive market, each looking for the best method to increase output under a limited cost range? Many such situations can be thought of to study whether the interrelation between competitiveness, network positions, and task uncertainty is present in real life as well.

Acknowledgements

We would like to thank Ineke Maas, Anne-Rigt Poortman, and Wojtek Przepiorka, who provided valuable feedback on earlier drafts of this paper. We are also thankful for the help of Vincent Buskens, Bas Hofstra, and Charlotte Rutten in designing the experiment. Finally, we thank the two anonymous Social Networks reviewers for their elaborate reviews and suggestions.

Supplementary data

Supplementary data associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.socnet.2018.01.007>.

References

- Auer, P., Cesa-Bianchi, N., Freund, Y., Schapire, R.E., 1995. Gambling in a rigged casino: the adversarial multi-armed bandit problem. In: *Proceedings of the 36th Annual Symposium on Foundations of Computer Science (FOCS '95)*, IEEE Computer Society Press, Los Alamitos, CA, pp. 322–331.
- Aumann, R.J., 1961. Almost strictly competitive games. *J. Soc. Ind. Appl. Math.* 9, 544–550.
- Bala, V., Goyal, S., 1998. Learning from neighbors. *Rev. Econ. Stud.* 65, 595–621.
- Baum, J.A.C., Li, S.X., Usher, J.M., 2000. Making the next move: how experiential and vicarious learning shape the locations of chains' acquisitions. *Adm. Sci. Q.* 45, 766–801.
- 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.
- Broere, J., Buskens, V., Weesie, J., Stoof, H., 2017. Network effects on coordination in asymmetric games. *Sci. Rep.* 7, 1–9.
- Bubeck, S., Cesa-Bianchi, N., 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Found. Trends Mach. Learn.* 5, 1–122.
- Burt, R.S. (Ed.), 1992. *Structural Holes: The Social Structure of Competition*. Harvard University Press, Cambridge, MA.
- Burt, R.S., 2004. Structural holes and good ideas. *Am. J. Sociol.* 110, 349–399.
- Buskens, V., Yamaguchi, K., 1999. A new model for information diffusion in heterogeneous social networks. *Sociol. Methodol.* 29, 281–325.
- Buskens, V., 2002. *Social Networks and Trust*. Kluwer Academic Publishers, Dordrecht.
- Camerer, C.F. (Ed.), 2003. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton.
- Centola, D., 2010. The spread of behavior in an online social network experiment. *Science* 329, 1194–1197.

- Choi, S., Gale, D., Kariv, S., 2004. Behavioral aspects of learning in social networks: an experimental study. *Adv. Appl. Microecon.* 13, 25–61.
- Conley, T.G., Udry, C.R., 2010. Learning about a new technology: pineapple in Ghana. *Am. Econ. Rev.* 100, 35–69.
- de Graaf, 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, 452–472.
- DeMarzo, P.M., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and unidimensional opinions. *Q. J. Econ.* 118, 909–968.
- Denrell, J., March, J.G., 2001. Adaptation as information restriction: the hot stove effect. *Organ. Sci.* 12, 523–538.
- Duersch, P., Oechssler, J., Schipper, B.C., 2012. Unbeatable imitation. *Games Econ. Behav.* 76, 88–96.
- Duflo, E., Saez, E., 2003. The role of information and social interactions in retirement plan decisions: evidence from a randomized experiment. *Q. J. Econ.* 118, 815–842.
- Ellison, G., Fudenberg, D., 1993. Rules of thumb for social learning. *J. Polit. Econ.* 101, 612–643.
- Falk, A., Heckman, J.J., 2009. Lab experiments are a major source of knowledge in the social sciences. *Science* 326, 535–538.
- Fang, C., Lee, J., Schilling, M.A., 2010. Balancing exploration and exploitation through structural design: the isolation of subgroups and organizational learning. *Organ. Sci.* 21, 625–642.
- Fischbacher, U., 2007. z-Tree: Zurich Toolbox for ready-made economic experiments. *Exp. Econ.* 10, 171–178.
- Freeman, L., 1979. Centrality in social networks: conceptual clarification. *Soc. Netw.* 1, 215–239.
- Gale, D., Kariv, S., 2003. Bayesian learning in social networks. *Games Econ. Behav.* 45, 329–346.
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 2004. *Bayesian Data Analysis*. Chapman & Hall/CRC, London.
- Gibbons, D.E., 2004. Network structure and innovation ambiguity effects on diffusion in dynamic organizational fields. *Acad. Manag. J.* 47, 938–951.
- Goyal, S. (Ed.), 2012. *Connections: An Introduction to the Economics of Networks*. Princeton University Press, Princeton.
- Granovetter, M.S., 1973. The strength of weak ties. *Am. J. Sociol.* 78, 1360–1380.
- Granovetter, M.S., 1978. Threshold models of collective behavior. *Am. J. Sociol.* 83, 1420–1443.
- Granovetter, M.S., 1995. *Getting a Job: A Study of Contacts and Careers*. University of Chicago Press, Chicago.
- Greiner, B., 2004. The online recruitment system orsee 2.0: a guide for the organization of experiments in economics. Working Paper Series in Economics, vol. 10. University of Cologne, Cologne.
- Gupta, A.K., Smith, K.G., Shalley, C.E., 2006. The interplay between exploration and exploitation. *Acad. Manag. J.* 49, 693–706.
- Hofstra, B., Corten, R., Buskens, V., 2015. Learning in social networks: selecting profitable choices among alternatives of uncertain profitability in various networks. *Soc. Netw.* 43, 100–112.
- Kremer, M., Miguel, E., 2007. The illusion of sustainability. *Q. J. Econ.* 122, 1007–1065.
- Lazer, D., Friedman, A., 2007. The network structure of exploration and exploitation. *Adm. Sci. Q.* 52, 667–694.
- Leavitt, H.J., 1951. Some effects of certain communication patterns on group performance. *J. Abnorm. Soc. Psychol.* 46, 38–50.
- Lin, N., 2001. *Social Capital: A Theory of Social Structure and Action*. Cambridge University Press, Cambridge, MA.
- March, J.G., 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* 2, 71–87.
- Mason, W.A., Watts, D.J., 2012. Collaborative learning in networks. *Proc. Natl. Acad. Sci.* 109, 764–769.
- Mason, W.A., Jones, A., Goldstone, R.L., 2008. Propagation of innovations in networked groups. *J. Exp. Psychol. Gen.* 137, 422–433.
- Mobius, M., Rosenblat, T., 2014. Social learning in economics. *Annu. Rev. Econ.* 6, 827–847.
- Mobius, M., Phan, T., Szeidl, A., 2015. *Treasure Hunt: Social Learning in the Field*.
- Mouw, T., 2003. Social capital and finding a job: do contacts matter? *Am. Sociol. Rev.* 68, 868–898.
- Muthén, L.K., Muthén, B.O., 2017. *Mplus User's Guide*, Eighth ed. Muthén & Muthén, Los Angeles, CA.
- Robbins, H., 1952. Some aspects of the sequential design of experiments. *Bull. Am. Math. Soc.* 55, 527–535.
- Rutten, C., 2014. *Network Entrepreneurial Personality, Network Constraint and Social Learning: An Experimental Test of Burt's Theory*. Utrecht University, Utrecht.
- Valente, T.W., Coronges, K., Lakon, C., Costenbader, E., 2008. How correlated are network centrality measures? *Connections* 28, 16–26.
- Vermorel, J., Mohri, M., 2005. Multi-armed bandit algorithms and empirical evaluation. 16th Eur. Conf. Mach. Learn.
- Watts, D.J., Strogatz, S.H., 1998. Collective dynamics of small-world networks. *Nature* 393, 440–442.