Betting on the Underdog: The Influence of Social Networks on Vote Choice*

Annika Fredén, Karlstad University

Ludovic Rheault, University of Toronto

Indridi H. Indridason, UC Riverside

Abstract

What makes people vote for an underdog? The common expectation is that people avoid wasting their vote on a party with a small probability of being elected. Yet, many voters choose to support underdogs and we still understand little about their motivations. We argue that voters gauge the support for their preferred party in the voting population from their social networks. When social networks exhibit the characteristics of echo chambers, a feature observed in real-life political networks, voters with a strong preference for an underdog tend to overestimate their chances of winning. We test this claim with voting experiments in which some treatment groups receive signals from a simulated network. We compare the effect of networks with a high degree of homogeneity against random networks. Our findings suggest that homophilic networks can generate a positive effect on the level of support for underdogs, which provides empirical evidence to back up anecdotal claims that echo chambers foster the development of fringe parties.

Accepted manuscript to appear in *Political Science Research and Methods*.

Citation: Fredén, Annika, Ludovic Rheault and Indridi H. Indridason. 2020. "Betting on the Underdog: The Influence of Social Networks on Vote Choice." *Political Science Research and Methods*, Forthcoming.

*The authors thank the Royal Swedish Academy of Sciences, Södertörn University, the Karl Staaff Foundation, the Royal Society of Arts and Sciences in Gothenburg, the Swedish Research Council (2017-02941), and the Social Sciences and Humanities Research Council of Canada for funding support, as well as Eryk Krysowski for lab assistance.

A large body of literature assumes that voters aim to affect electoral outcomes by choosing strategically among parties large enough to gain representation. This requires somewhat accurate information about the popularity of the competing parties or candidates. Many electoral contexts, however, are characterized by poor information about the likely outcome, in particular since polls are typically conducted at the national level while legislative seats are allocated at the district level. In such circumstances, or when the accuracy of polls is uncertain, voters are likely to rely on alternative sources of information when deciding how to cast their votes. Their social networks represent one such option. Because networks are often characterized by homophily, one consequence is that voters are liable to form inaccurate expectations about the popularity of parties or candidates.

We advance that distortions introduced by social networks help explain why, in many elections, a sizable share of the electorate supports a candidate or party with little hope of winning a seat—a party that could be considered an 'underdog'. Even in an established democracy like Canada, where the plurality system is expected to generate two-party races at the district level, over 21% of the votes cast during the last general election of 2015 were for a third-party candidate.[1] Similarly, in proportional representation (PR) systems, more than one in nine voters voted for parties that failed to gain representation in the 2013 Icelandic election (Indridason et al., 2017) and nearly one in six voters during the 2015 Polish election (Jasiewicz and Jasiewicz-Betkiewicz, 2016).

This paper proposes a novel explanation for the willingness of voters to support underdogs. Rather than ruling out this behavior as an anomaly, we build on the fact that voters rarely possess the information about party support needed to vote strategically. We represent the voter's dilemma as a coordination game in which players must draw inferences about whether a small party (the underdog) is viable enough to gain representation. We argue that voters rely on the signals they receive from their social networks as a cheap source of information to

[1]Third-party candidates are those not among the two front-runners in each district. Data from the Library of Parliament: https://lop.parl.ca/sites/ParlInfo/default/en_CA/ElectionsRidings/Elections (accessed December 18, 2018).

solve their coordination problem. When networks are reinforcing—that is, characterized by assortative mixing (also called homophily or echo chambers) (McPherson et al., 2001; Bakshy et al., 2015)—as is often observed in real-life networks, strong supporters tend to overestimate the chances of the underdog winning. We test this model using laboratory experiments in which participants are randomly assigned to a network, some of them receiving information about the political preferences of their peers. Our results support the view that network information can influence the decision to support underdogs, by affecting voter perceptions about their chances of success.

Exploring how social networks affect the vote is particularly relevant given the growing role of social media in politics, which has reemphasized an earlier body of work documenting the influence of peer networks on political behavior (Huckfeldt and Sprague, 1995; Mutz, 2006; Sinclair, 2012; Ahn et al., 2014). Not only have social media expanded the size of individuals' social networks, and the frequency and ease with which individuals interact, they have also made it easier to connect with like-minded individuals. There is ample evidence that online networks reproduce (or amplify) the assortative mixing observed in real-life networks (Colleoni et al., 2014; Bakshy et al., 2015; Eady et al., 2019). Empirical findings remain mixed, however, regarding the consequences of these 'echo chambers', for instance their potential impact on the decline of deliberative politics, the fomentation of extreme ideology, and political polarization (Farrell, 2012; Lee et al., 2014; Colleoni et al., 2014; Barberá et al., 2015; Flaxman et al., 2016; Bail et al., 2018; Eady et al., 2019).

In this study, we examine whether reinforcing networks can lead voters to make biased inferences about the viability of small parties. The proposed mechanism sheds light on the purported link between echo chambers and the support for fringe parties. Our findings represent a novel contribution to the literature on social networks and political behavior (see, e.g., Sokhey and McClurg, 2012; Bond et al., 2012; Santoro and Beck, 2017). Many studies of political networks have focused on two-party systems (see, e.g., Zuckerman et al., 1994; Huckfeldt and Sprague, 1995),[2] where opportunities for strategic voting rarely arise, or examined network

[2]A notable exception is Beck (2002), which examined support for Perot in the 1992 presidential election.

effects on other types of outcomes such as political participation (McClurg, 2003, 2006; Großer and Schram, 2006; Battaglini et al., 2008; Gil de Zúñiga et al., 2012; Tufekci and Wilson, 2012; Boulianne, 2015), the impact of elite communication on the formation of opinions (Ahn et al., 2014; Siegel, 2009, 2013), and "correct voting" (Ryan, 2011; Pietryka, 2016). We expand on this literature by focusing on vote choice, specifically in multi-party systems where voters face a different challenge in the form of coordination problems.

Theoretical Considerations

We focus our attention on a setting in which voters have narrowed their possible choices to two parties. In doing so, we seek to zero in on situations where strategic voting is possible: voters facing a choice between a preferred party with limited chances of gaining representation (the underdog), and a second preference that is viable (the safe option). This is a scenario that often occurs, e.g., in elections under proportional representation where one of two ideologically similar parties is at risk of not reaching the threshold for representation. For example, if a Portuguese voter sincerely prefers the newly-formed Alliance Party but the party is at risk of not gaining seats in her district, the voter faces a dilemma between supporting the underdog or rallying around the safer option, the Social Democratic Party. As in many real-life campaigns, we assume that each voter is uncertain about the preferences of other voters.

This decision problem can be represented as a n-player coordination game. The underdog has a chance of winning if enough voters coordinate their efforts on their sincere preference. On the other hand, voting for the underdog when other voters fail to do the same implies a wasted vote—the viable second-preferred party would have done better had everyone voted for it. Table 1 presents the payoff structure for this problem. We assume that voters differ in their preference for the underdog. We denote the strength of this preference with the variable x_i, and assume it is drawn randomly from a uniform distribution. Casting a vote for a party that gains representation yields a constant payoff of c; in effect this is the value of not wasting one's vote. When the underdog receives enough votes to win representation (at least equal to some threshold T), the voter's total payoff is the sum $c + x_i$. A voter choosing the underdog when it fails to

Table 1: Payoff structure

	Underdog does not gain representation	Underdog gains representation
Vote Underdog (U)	x_i	$c + x_i$
Vote Safe Option (S)	c	$c + x_i$

win representation only receives a payoff of x_i. Thus, in choosing the underdog, a voter has a chance of getting the maximum payoff of $c + x_i$, but also risks only receiving a payoff of x_i. The voter's other option, which we refer to as the *safe option*, is to vote for the viable party, which guarantees a minimum payoff of c. We assume voters assign greater value to not wasting their vote than expressing a sincere preference, i.e., $x_i < c$.

This game resembles the *stag hunt*, a coordination game with two pure-strategy equilibria: one in which the voters coordinate on the risky option for a higher payoff ($x_i + c$), and one in which the voters choose the safe option for a certain but lower payoff (c). The voters in our model face the same type of coordination problem. It departs from a pure spatial model of candidate competition, in that it incorporates an element of expressive support while retaining the essential features of the coordination problem in the stag hunt. We do so for two reasons. First, expressive support strikes us as a reasonable motivation for voters and, in particular, supporters of parties that might be considered underdogs. Part of the rationale for supporting underdogs is to protest against the more mainstream alternatives, and there has been an upsurge in research on expressive voting of late (for an overview, see Hamlin and Jennings, 2019). Second, these payoffs provide the simplest possible formulation for the coordination problem faced by supporters of underdog parties. A first equilibrium arises in which all voters choose the safe option; switching the vote to the underdog reduces the voter's payoff from c to x_i. In a second type of equilibrium, at least T voters choose the underdog, where T is the threshold, i.e., the number of votes needed for representation. In such scenarios, no voters can improve their payoff by modifying their decision, and players achieve a Pareto-efficient outcome $c + x_i$. Experimental evidence on stag hunt games suggest a tendency for players to coordinate on the safe option (Skyrms, 2013).

Our main contribution is to consider social networks as a mechanism for equilibrium selection. If voters possess information on the preferences x_j of individuals with whom they share connections in the network, they may use this information to infer the likelihood of coordination on the risky option. The network signals reproduce the opinions that voters in a real-life setting would observe on social media, for instance, by reading comments from their friends or followers. As discussed earlier, social networks are usually not random, which may distort perceptions of the likelihood that voters will coordinate on the underdog. Our central claim is that network information is a key mechanism used by voters to decide between risky and safe choices in elections.[3]

Building upon this discussion, we offer two hypotheses. The first concerns the influence of homophilic networks on equilibrium selection. In presence of network signals revealing the preferences of other voters (x_j), we expect that a voter is more likely to coordinate on the underdog when she observes a strong preference for the candidate in her network, compared to the control condition (Hypothesis 1). In contrast, signals coming from a random network provide no cues that help solve the coordination problem. Thus, we expect no difference in behavior between voters receiving signals from a random network and where no network information is available (Hypothesis 2).

Experimental Design

We designed a laboratory voting experiment in which participants choose between two parties competing in simulated elections: Party S (the Safe Option) and Party U (the Underdog). The experiment was conducted in a computer lab on May 29, 2018. We recruited 96 participants in total, 24 for each session. The participants were randomly assigned into subgroups of 6 voters for *each* election. We informed participants that the Underdog must receive at least 5 out of 6 votes to get elected. Before each election, each voter was assigned a random number (from a discrete uniform distribution ranging from 1 to 9) representing the strength of preference for the

[3] Our setup defines an underdog as a party that has some theoretical chance of achieving representation, albeit a marginal one. The model could be extended to situations where an underdog does not actually have a theoretical chance to win, even when all of its supporters vote sincerely.

Underdog (x_i). Thus, as the distribution of preferences and subgroups of six voters change each round, the voters face a new electorate every time they vote. We set the constant invoked in the theory section to $c = 10$. The participants voted in 20 consecutive elections.

To examine the effects of social networks on coordination, we considered two network treatments along with a control condition. The participant was randomly assigned to one condition. In the control group, the participants receive no information about other voters. Participants are informed about the existence of randomized private values, but they do not observe the private values of other participants. In the first treatment (random network) group, each voter observes the private values of two randomly selected peers among the six voters participating in the election. In the second treatment group (homophilic network), the network ties are reinforcing. We partition the voters so that the three participants with the highest x_i values are grouped together, with the other three participants forming another group. Voters observe the private values of the voters within their network who are similar to their own (see Figure 1). This simulates the phenomenon of assortative mixing, whereby voters who share similar attributes are more likely to develop network ties. The online appendix provides additional information about the experimental design, along with descriptive statistics and balance checks.

Figure 1: Information in Homophily Treatment

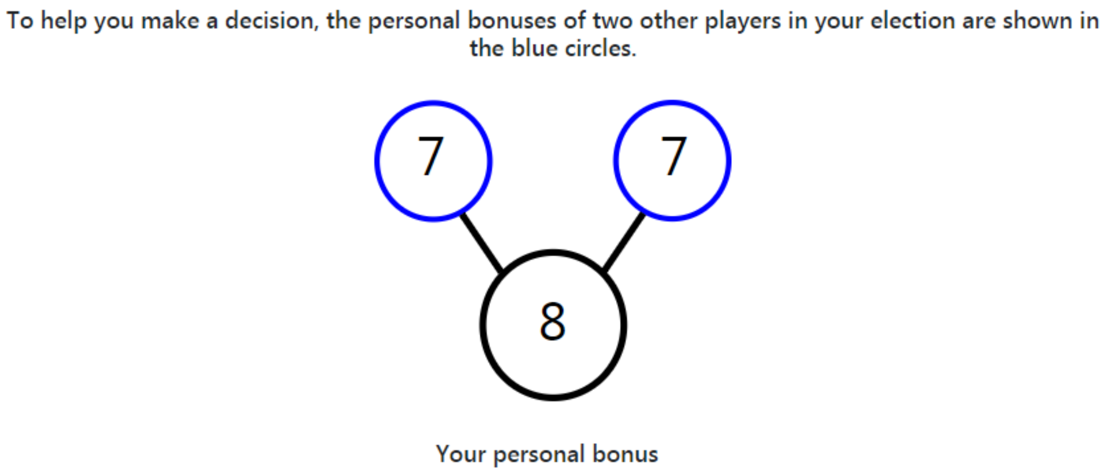

Findings

We start by examining the baseline rates of support for the underdog party across experimental groups. Table 2 reports shows that the underdog (Party U) is chosen roughly 20% of the time by participants assigned to the control group. This proportion increases to about 25% in the homophilic treatment ($p = 0.12$; bootstrapped cluster-robust p-value). However, the comparison obscures the actual effect of homophily as our treatment generates two groups of participants. The fourth and fifth rows report the same cross-tabulation, this time by contrasting the participants whose network contained the three lowest payoffs for selecting the underdog (which we label "Low Signal") and those with the high payoffs ("High Signal"). When exposed to a network of voters who have a strong preference for the underdog, the tendency to vote for that party doubles, to 43%.

Table 2: Cross-tabulation of the vote for the underdog, by experimental group

	Vote Choice	
	Safe Option (S)	Underdog (U)
Control	79.58%	20.42%
Random Network	78.75%	21.25%
Homophilic Network	74.79%	25.21%
Low Signal	92.29%	7.71%
High Signal	57.29%	42.71%
All	76.98%	23.02%
Observations	1,920	

Conversely, when the signal received from the network indicates weaker support, participants are much more likely to select the safe option, in which case the overall support for the underdog drops below 8%. Once private signals (x_i) are controlled for, however, a network of players with weak preferences for the underdog does not behave differently from the control group (see Table 3 and discussion below). Our experimental design also affords us the opportunity to contrast reinforcing networks with random networks. The underdog vote share in the random network treatment is 21%, which is statistically indistinguishable from the control group ($p = 0.80$). This is consistent with our second hypothesis.

The results support the principal contention made in this paper. When real-life social networks are composed of individuals who think alike, voters with a strong preference for an underdog party receive a signal that may overstate the overall strength of support for that option. In turn, this signal increases the likelihood of voting for the underdog. Our claim is that this phenomenon is a key mechanism explaining the paradox of voters who "fail" to defect from non-viable candidates. In game-theoretic parlance, social networks appear to serve as a tool for equilibrium selection. However, as illustrated by the strikingly different patterns between the homophilic and random treatments, the network itself provides no useful information for coordination when network connections are random. Signals need to be one-sided, as one would expect inside networks with the characteristics of echo chambers.

Table 3 reports logistic regressions of the binary decision to vote for the underdog, with cluster-robust standard errors (where clusters are the individual participants). The models include the treatment variables, in addition to a time trend. As rounds progress, the overall share of participants selecting the risky option decreases, suggesting that participants adjust their behavior after observing that the underdog rarely wins. The second model controls for the individual preference for the underdog x_i (the private signal). The third model adds a control variable measuring the general risk preference of respondents measured on a 0-10 scale (Dohmen et al., 2010). These models support the conclusions outlined above. The likelihood of selecting the underdog is greater under the homophilic treatment with high signals (using the control group as a base category), everything else being equal, a result that is statistically significant at conventional levels.

The results illustrate why homophilic networks tend to benefit underdogs in the aggregate. Once the individual payoffs are taken into account (Models 2 and 3 from Table 3), the treatment effect is statistically significant only for the subgroup receiving high signals from the homophilic network. In other words, voters are equally likely to choose the safe option when their own sincere preference for the underdog is weak, regardless of whether or not they are informed of the reinforcing preferences of their network. This finding is consistent with the observed tendency of players to choose the safe option in previous stag hunt game experiments (Skyrms, 2013). On the other hand, when the network brings together voters with a strong preference

Table 3: Treatment Effects (Logistic Regressions with Cluster-Robust Standard Errors)

	Vote Choice (Underdog=1)		
	Model 1	Model 2	Model 3
Homophilic Network (Low Signal)	-1.178^{***}	-0.001	0.078
	(0.284)	(0.307)	(0.291)
Homophilic Network (High Signal)	1.167^{***}	0.647^{*}	0.751^{**}
	(0.209)	(0.270)	(0.260)
Random Network	0.054	0.177	0.352
	(0.284)	(0.329)	(0.305)
Private Signal		0.539^{***}	0.559^{***}
		(0.052)	(0.051)
Tolerance to Risk			0.255^{***}
			(0.060)
Round	-0.118^{***}	-0.148^{***}	-0.157^{***}
	(0.014)	(0.016)	(0.018)
Constant	-0.252	-3.180^{***}	-4.737^{***}
	(0.197)	(0.359)	(0.548)
Observations	1,920	1,920	1,920
Participants (Clusters)	96	96	96
Nagelkerke's R^2	0.221	0.402	0.433

The table reports the output of logistic regression models computed with bootstrapped, cluster-robust standard errors, where clusters represent the participants. The base category for treatment effects is the control group. Standard errors are in parentheses. $^{*}p<0.05$; $^{**}p<0.01$; $^{***}p<0.001$

for the underdog, networks have a significant impact on the decision to coordinate on the risky option. In that case, we estimate the (conditional) average treatment effect, using the difference in predicted probabilities between the homophilic treatment with high signal and the control group, to be around 14 percentage points (see the appendix for an extended discussion). The fact that homophilic networks induce stronger coordination effects in the high-signal group explains why, when considering the homophilic treatment group as a whole (i.e., both those receiving high and low signals), the support for the underdog is higher overall than in the other two comparison groups.

Concluding Remarks

We examined the impact of social networks on the support for underdogs using laboratory experiments. We expected that network information would affect evaluations of the underdog's chance of winning, in particular when a voter belongs to a reinforcing network where other voters share strong preferences for the underdog. We find clear evidence supporting the existence of such an effect. In other words, a promising explanation for people choosing to 'waste' their vote on underdog candidates in real-world elections is that their social network may lead them to overestimate vote intentions for that alternative. In our experiments, the tendency to vote for underdogs is significantly higher for voters receiving signals from a reinforcing network of strong supporters, when compared to scenarios without network information and where network connections are random. Our results thus suggest that social networks have important effects on strategic voting, but these effects only arise when networks have the characteristics of echo chambers. A substantive implication for the study of democracy is that echo chambers can foster support for fringe parties, consistent with recent claims suggesting that social media may spur the growth of extremist ideologies (see, e.g., Flaxman et al., 2016).

References

Ahn, TK, Robert Huckfeldt, and John Barry Ryan (2014). *Experts, activists, and democratic politics: Are electorates self-educating?* Cambridge: Cambridge University Press.

Bail, Christopher A, Lisa P Argyle, Taylor W Brown, John P Bumpus, Haohan Chen, MB Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences 115*(37), 9216–9221.

Bakshy, Eytan, Solomon Messing, and Lada A. Adamic (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science 348*(6239), 1130–1132.

Barberá, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science 26*(10), 1531–1542.

Battaglini, Marco, Rebecca B Morton, and Thomas R Palfrey (2008). Information aggregation and strategic abstention in large laboratory elections. *American Economic Review 98*(2), 194–200.

Beck, Paul A (2002). Encouraging political defection: The role of personal discussion networks in partisan desertions to the opposition party and Perot votes in 1992. *Political Behavior 24*(4), 309–337.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler (2012). A 61-million-person experiment in social influence and political mobilization. *Nature 489*(7415), 295.

Boulianne, Shelley (2015). Social media use and participation: A meta-analysis of current research. *Information, Communication & Society 18*(5), 524–538.

Colleoni, Elanor, Alessandro Rozza, and Adam Arvidsson (2014). Echo chamber or public sphere? Predicting political orientation and measuring political homophily in Twitter using big data. *Journal of Communication 64*(2), 317–332.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review 100*(3), 1238–60.

Eady, Gregory, Jonathan Nagler, Andy Guess, Jan Zilinsky, and Joshua A Tucker (2019). How many people live in political bubbles on social media? Evidence from linked survey and twitter data. *SAGE Open 9*(1), 1–21.

Farrell, Henry (2012). The internet's consequences for politics. *Annual Review of Political Science 15*(1), 35–52.

Flaxman, Seth, Sharad Goel, and Justin M Rao (2016). Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly 80*(S1), 298–320.

Gil de Zúñiga, Homero, Nakwon Jung, and Sebastián Valenzuela (2012). Social media use for news and individuals' social capital, civic engagement and political participation. *Journal of Computer-Mediated Communication 17*(3), 319–336.

Großer, Jens and Arthur Schram (2006). Neighborhood information exchange and voter participation: An experimental study. *American Political Science Review 100*(2), 235–248.

Hamlin, Alan and Colin Jennings (2019). Expressive voting. In B. G. Roger D. Congleton and S. Voigt (Eds.), *The Oxford Handbook of Public Choice*, Volume 1, Chapter 16, pp. 333–350. Oxford: OUP.

Huckfeldt, R Robert and John Sprague (1995). *Citizens, politics and social communication: Information and influence in an election campaign*. Cambridge: Cambridge University Press.

Indridason, Indridi H., Eva Heida Önnudóttir, Hulda Thórisdóttir, and Ólafur Th. Hardarson (2017). Re-electing the culprits of the crisis? Elections in the aftermath of a recession. *Scandinavian Political Studies 40*(1), 28–60.

Jasiewicz, K. and A Jasiewicz-Betkiewicz (2016). Poland. *European Journal of Political Research 55*, 204–215. Political Data Yearbook.

Lee, Jae Kook, Jihyang Choi, Cheonsoo Kim, and Yonghwan Kim (2014). Social media, network heterogeneity, and opinion polarization. *Journal of Communication 64*(4), 702–722.

McClurg, Scott D (2003). Social networks and political participation: The role of social interaction in explaining political participation. *Political Research Quarterly 56*(4), 449–464.

McClurg, Scott D (2006). The electoral relevance of political talk: Examining disagreement and expertise effects in social networks on political participation. *American Journal of Political Science 50*(3), 737–754.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology 27*(1), 415–444.

Mutz, Diana C (2006). *Hearing the Other Side: Deliberative Versus Participatory Democracy*. Cambridge: Cambridge University Press.

Pietryka, Matthew T (2016). Accuracy motivations, predispositions, and social information in political discussion networks. *Political Psychology 37*(3), 367–386.

Ryan, John Barry (2011). Social networks as a shortcut to correct voting. *American Journal of Political Science 55*(4), 753–766.

Santoro, Lauren Ratliff and Paul A. Beck (2017). Social networks and vote choice. In J. N. Victor, A. H. Montgomery, and M. Lubell (Eds.), *The Oxford Handbook of Political Networks*, pp. 383. Oxford: Oxford University Press.

Siegel, David A (2009). Social networks and collective action. *American Journal of Political Science 53*(1), 122–138.

Siegel, David A (2013). Social networks and the mass media. *American Political Science Review 107*(4), 786–805.

Sinclair, Betsy (2012). *The Social Citizen*. Chicago: University of Chicago Press.

Skyrms, Brian (2013). *The Stag Hunt and the Evolution of Social Structure*. Cambridge: Cambridge University Press.

Sokhey, Anand Edward and Scott D. McClurg (2012). Social networks and correct voting. *Journal of Politics 74*(3), 751–764.

Tufekci, Zeynep and Christopher Wilson (2012). Social media and the decision to participate in political protest: Observations from Tahrir square. *Journal of communication 62*(2), 363–379.

Zuckerman, Alan S, Nicholas A Valentino, and Ezra W Zuckerman (1994). A structural theory of vote choice: Social and political networks and electoral flows in Britain and the United States. *Journal of Politics 56*(4), 1008–1033.

Supplementary Materials

Additional Information on Experimental Design

We conducted the experimental sessions in the computer lab of WULABS at the Vienna University of Economics and Business, Austria on May 29, 2018. The experimental module was programmed using the oTree library, a web-based platform using the Python Django framework (Chen et al., 2016). The project received ethics approval from the Institutional Review Board of the University of California, Riverside (protocol HS-17-131).

In total, we ran four sessions of 24 players per session. Participants were recruited from the lab's student pool and consisted mainly of first-year business school students. We used the laboratory's standard recruitment procedures, inviting 32 individuals per session by email. The participants who showed up were greeted by a research assistant and randomly assigned a card corresponding to a computer upon arrival, in order to avoid selection biases such as early or late show ups. As the experiment was designed for 24 participants, invited individuals who were not selected to participate were awarded a show-up fee of 5 Euros. The selected participants gave their consent by signing an electronic form that stated the purpose and main procedures of the game. Each session took approximately 45 minutes to complete.

Each experimental session began with an introduction to the electoral scenario. We informed participants that two parties are facing each other (called A and B during the experiment), and that the election results would be determined by the choices made by sub-groups of six participants. The payoff structure was explained with a concrete example. Note that in the manuscript and below, we relabel the parties S and U for simplicity, to emphasize which one represents the 'safe option' and which one is the underdog.

Consistent with the structure introduced in Table 1 of the main text, the lowest payoff occurs when a participant votes for the Underdog and the party receives fewer than five votes, in which case they receive only the payoff associated with x_i, a random number between 1 and 9. Voting for the safe option, Party S, yields a minimum reward of 10. Finally, the highest reward is achieved when the Underdog reaches the threshold (5 out of 6 votes or more), in which case the participant receives $c + x_i$, the constant plus the random number they have been assigned (for a total payoff between 11 and 19).

Half of the 24 participants in each session were randomly assigned to one of two treatment conditions (among the three types, i.e., control, random, or homophilic network). Within each treatment condition, the twelve participants were randomly assigned to sub-groups (or electorates) of six voters for each round of the experiment. Thus, each election comprised a new electorate of six voters.The participants were informed about this procedure on the computer screens before each election. After each round, or election, we informed players about the result of the election and the payoff they received. The payoffs collected through the 20 rounds were converted to monetary rewards at the end of the experiment (100 points in the game correspond to 5 Euros). Participants were rewarded 9 Euros on average.

Upon completion of the twenty rounds, we asked participants to fill a short survey and informed them of their total gains. The questions measured basic socio-demographic variables. The survey also included an item evaluating their predisposition toward risk, using a question proposed by Dohmen et al. (2010). This survey question reads "How willing are you to take risks in general?" and asks respondents to report their willingness on a 0-10 scale.

Table A1 provides descriptive statistics for the survey variables. We also conducted balance checks to verify that randomization into treatment conditions produced covariate balance. We report the results in Table A2. The models in Table A2 are logistic regressions with the treatment assignment as a dependent variable. Except for one covariate (gender in the random network treatment group), the covariates appear unrelated to the treatments. We replicated the main models presented in the paper with demographic covariates as controls, and the results are substantively the same. Finally, we report screenshots of the experimental module in Figures A1-A5 below.

Table A1: Descriptive Statistics

Variable	Category/Statistic	Value
Experimental Group	Control	24
	Random network	24
	Homophilic network	48
Age	18-24 years old	73
	25-34 years old	22
	35-44 years old	1
Gender	Female	63
	Male	33
Education	High school degree	50
	Some higher education	13
	Bachelor degree	26
	Above bachelor	7
Tolerance to Risk	Mean	5.375
	Std. Deviation	1.98
Total Sample		96

The table presents descriptive statistics for the sample of experimental participants, across all sessions.

Table A2: Balance Checks

	Treatment Group	
	Homophilic Network	Random Network
Aged 25 and above	−1.237	0.469
	(0.643)	(0.764)
Bachelor degree	0.477	−0.227
	(0.990)	(1.063)
High school degree	−0.220	1.019
	(1.036)	(1.126)
Some higher education	−0.053	−0.175
	(1.152)	(1.355)
Gender = Male	−0.448	1.647**
	(0.468)	(0.560)
Tolerance to risk	−0.006	−0.181
	(0.110)	(0.130)
Constant	0.467	−1.460
	(1.251)	(1.377)
Observations	96	96
Log Likelihood	−62.882	−47.352
Akaike Inf. Crit.	139.764	108.703

The table reports binary logistic regressions with the treatment group assignment as a dependent variable. Standard errors are in parentheses. *p<0.05; **p<0.01; ***p<0.001

Figure A1: Instructions (Screen 1)

Instructions

The Parties

Imagine that you are voting in an election with two political parties: A and B. Party A has a strong support in the population and is already expected to win seats in a fictional legislature. Party B can still win seats, but only if it receives at least 5 votes from 6 people currently playing in this experiment, including yourself.

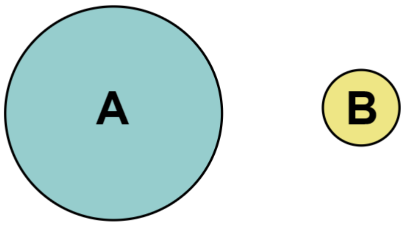

The Points

In each election, you start with 10 points. If you vote for Party A, you receive these 10 points with certainty.

You can accumulate more points if B receives at least 5 votes among the 6 players.

Your Personal Bonus

In each election, all players will be given a personal number of bonus points between 1 and 9, chosen at random. You will not know the bonus points allocated to all the other players.

If party B gets 5 votes or more, you gain the initial 10 points **plus** your personal bonus. However, if you choose B and that party does not get at least 5 votes, you only get the bonus and **lose** the initial 10 points. (See the example below.)

Your Task

Your task is to end up with as many points as possible across 20 elections. These points will be converted in real money after the experiment, using the conversion 100 game points = 5.0 EUR.

Example

Suppose that your personal bonus is: 2

The following table shows you the points you would win based on your decision and the outcome of the election:

Choice	Payoff
You vote for A and B receives 4 votes or less	10
You vote for B and B receives 4 votes or less	2
You vote for A or B and B receives 5 votes or more	12

Next

Figure A2: Homophily Treatment (Screen 2)

Election # 1

You are playing with five other participants in this room. Your voting group is modified randomly at each election.

To help you make a decision, the personal bonuses of two other players in your election are shown in the blue circles.

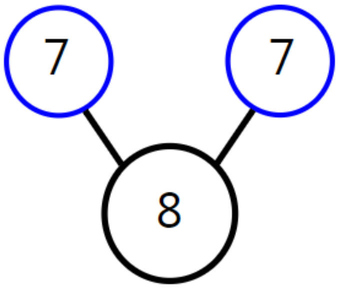

Your personal bonus

Vote Decision

Who do you vote for?

Vote for A   Vote for B

(Note: Clicking one of the two buttons will register your choice. You will not be able to return.)

Points you could win in this round:

Choice	Payoff
You vote for A and B receives 4 votes or less	10
You vote for B and B receives 4 votes or less	8
You vote for A or B and B receives 5 votes or more	18

Click here to display the instructions for the game.

Figure A3: Random Treatment (Screen 2)

Election # 1

You are playing with five other participants in this room. Your voting group is modified randomly at each election.

To help you make a decision, the personal bonuses of two other players in your election are shown in the blue circles.

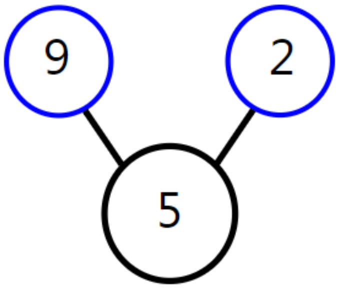

Your personal bonus

Vote Decision

Who do you vote for?

Vote for A   Vote for B

(Note: Clicking one of the two buttons will register your choice. You will not be able to return.)

Points you could win in this round:

Choice	Payoff
You vote for A and B receives 4 votes or less	10
You vote for B and B receives 4 votes or less	5
You vote for A or B and B receives 5 votes or more	15

Click here to display the instructions for the game.

Figure A4: Control Treatment (Screen 2)

Election # 1

You are playing with five other participants in this room. Your voting group is modified randomly at each election.

1

Your personal bonus

Vote Decision

Who do you vote for?

Vote for A   Vote for B

(Note: Clicking one of the two buttons will register your choice. You will not be able to return.)

Points you could win in this round:

Choice	Payoff
You vote for A and B receives 4 votes or less	10
You vote for B and B receives 4 votes or less	1
You vote for A or B and B receives 5 votes or more	11

Click here to display the instructions for the game.

Figure A5: Results Example (Screen 3)

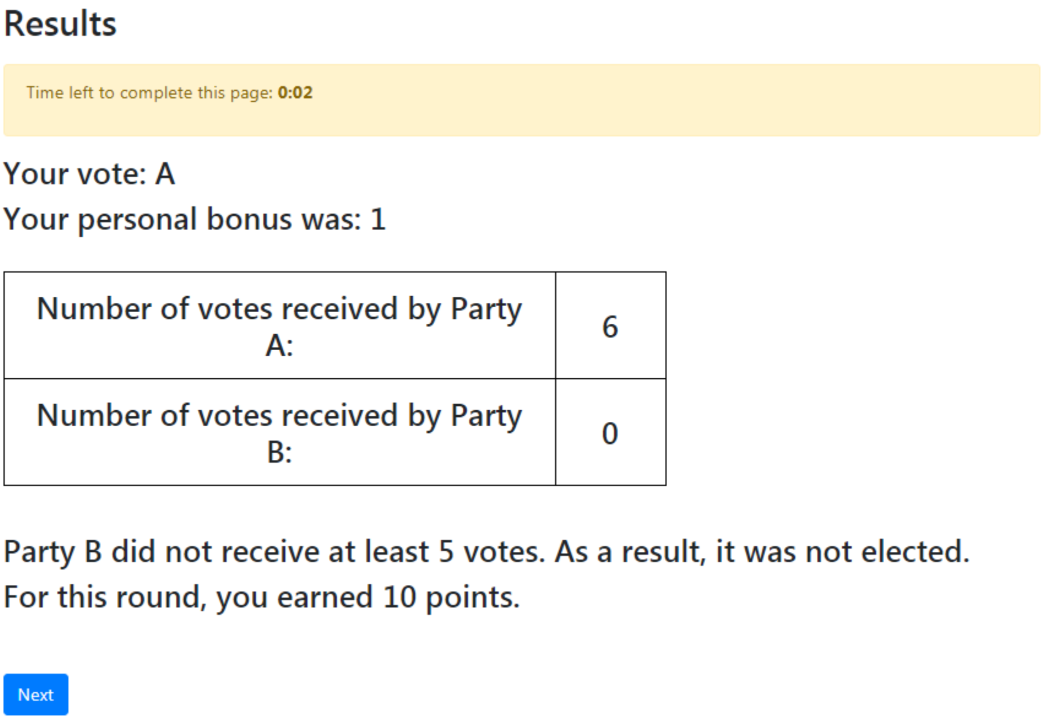

Additional Results

This section provides additional results in support of the empirical findings presented in the main text. The two subgroups in the homophilic treatment are defined in terms of the private payoffs (the x_i variable), to reproduce the shared affinities of voters in reinforcing networks. A rigorous test of our hypotheses would consist of estimating treatment effects for a constant x_i, as we did in the multivariate models presented in the main text. Put another way, the conditional average treatment effect corresponds to

$$\mathbb{E}[y_i(1) - y_i(0)|x_i]$$

where $y_i(1)$ is the binary vote choice in the treatment group and $y_i(0)$ in the control group. This quantity isolates the effect of network information from the effect of the voter's strength of preference for the underdog. Table A3 below reports cross-tabulations based on subsamples of participants: we compute the proportion of underdog votes only for participants with $x_i > 5$ for the High Signal treatment, and only for participants with $x_i < 5$ in the Low Signal treatment.

Table A3: Cross-tabulation of the vote for the underdog, for restricted subsamples

	Vote Choice	
	Safe Option (S)	Underdog (U)
Subsample: $x_i < 5$		
Control	94.47%	5.53%
Homophilic Treatment (Low Signal)	94.92%	5.08%
Observations	593	
Subsample: $x_i > 5$		
Control	63.60%	36.40%
Homophilic Treatment (High Signal)	49.31%	50.69%
Observations	591	

Once the individual payoffs are taken into account, the difference in proportions is statistically significant only for the subgroup receiving high signals from the homophilic network. In other words, voters are equally likely to choose the safe option when their own sincere preference for the underdog is weak, whether or not they observe reinforcing preferences in their network ($p \approx 0.85$; bootstrapped cluster-robust p-value). As mentioned in the main text, this finding is consistent with the observed tendency of players to choose the safe option in previous stag hunt game experiments (Skyrms, 2013). On the other hand, when the network brings together voters with a strong preference for the underdog, network signals have a significant impact on the decision to coordinate on the underdog. In that case, the (conditional) average treatment effect is roughly 14.3 percentage points ($p \approx 0.01$). The fact that homophilic networks induce stronger coordination effects in the latter group explains why, when considering the homophilic treatment group as a whole (i.e., both those receiving high and low signals), the support for the underdog is higher overall than in the other two comparison groups (Table 2 of the main text).

The logistic regressions in Table 3 of the main text report a similar finding, while also controlling for risk tolerance and the round of the experiment. Holding constant the private signal x_i to 5 and the round to 5, the difference in the predicted probability of voting for the underdog, contrasting the homophily (high signal) with the control group, is 14.1 percentage points (see Table A4). The difference varies from 5.4 to 18.2 percentage points when changing the value of the private signal of the participant from 2.5 to 7.5, respectively.

Table A4: Difference in Predicted Probabilities (Table 3, Model 3)

x_i	Homophily (High)	Control	Difference
5.0	0.331	0.190	+0.141
2.5	0.109	0.055	+0.054
7.5	0.668	0.486	+0.182

The table reports predicted probabilities of voting for the Underdog under two treatment conditions, as well as the difference in predicted probabilities between groups, computed from Model 3, Table 3 in the main text. The probabilities are calculated after setting the level of risk tolerance to the middle of the scale (the value of 5), the round number to 5, and by varying the value of the private signal x_i and the treatment condition.

Note that our design ensures that the distribution of preferences is the same across all treatment groups. Even when participants received a signal that two peers also have a high payoff for selecting the underdog, the ex-ante distribution of payoffs remained exactly the same as that used in the other treatment groups. In short, the homophilic treatment should have

Figure A6: Vote for underdog by experimental group and round

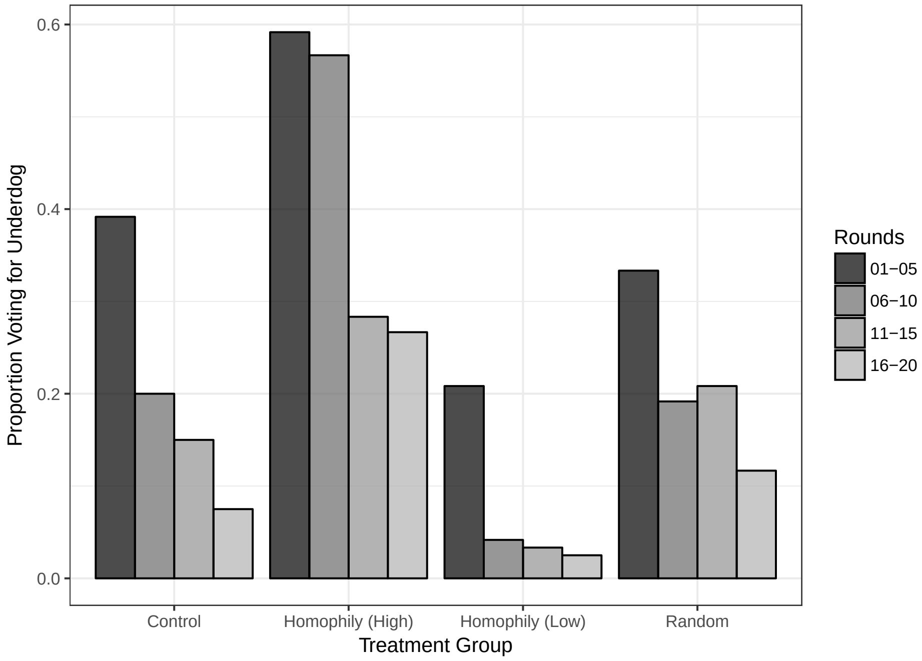

little impact on purely rational grounds, as it changes nothing about the baseline calculations. Moreover, we randomized individual payoffs at every single election, such that participants observed first hand that the assignment of a high or low payoff was equally likely. Nonetheless, we still observe a clear difference in behavior between treatment conditions across rounds. Figure A6 plots the distribution of support for the underdog over time. Although learning effects are noticeable, the tendency to use homophilic network signals for equilibrium selection appears to last for the duration of each session.

References

Chen, Daniel L., Martin Schonger, and Chris Wickens (2016). oTree: An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance 9*, 88–97.