Collective Decisions on Networks: information biased and voting coordination

ALEXANDER J. STEWART, University of Houston

MOHSEN MOSLEH, Massachusetts Institute of Technology

MARINA DIAKONOVA, University of Oxford

ANTONIO ARECHAR, Massachusetts Institute of Technology

DAVID G. RAND, Massachusetts Institute of Technology

JOSHUA B. PLOTKIN, University of Pennsylvania

Distorted and false information threaten to disrupt public discourse and democratic decision making [1-3]. Social media platforms are particularly vulnerable, because they allow users to shut out dissenting voices [4,5] while providing adversarial actors with anonymity and opportunity to target messages for maximal effect [6]. Two distinct but intertwined threats have received considerable attention: information campaigns using fake news [1,7] and automated bots [8], and the growth of polarized political debate [9]. These issues pose a profound social problem. Progress requires that we develop basic science to understand how networks that constrain the flow of information influence group decision-making.

Here we develop a voter game to study collective decisions under incomplete information. The game is simple enough to analyze mathematically and to use in controlled human-subject experiments, yet it retains the salient features of real-world collective decisions. Players are split into two parties of equal size and allowed to change their voting intention over time in response to continuously updated polling data. The aggregate polling data seen by a player is determined by her placement on a directed graph, called the influence network. Players are aware that polls represent a subset sampled from the entire population. At the end of the game, players receive the maximum payoff B if the final vote share for their assigned party exceeds a super-majority threshold V, where $V>1/2$; they receive a lower payoff $b<B$ if the opposing party's vote share exceeds V; and they receive no payoff if both parties fail to reach the threshold V, which we call "deadlock" (Figure 1). The possibility of deadlock forces players to consider both their personal preferences and the voting intentions of others in their decision-making. We initially assume the payoff to the losing team is positive, $b>0$, which reflects a "compromise worldview": it is preferable to reach some decision than to end in deadlock (There is ample evidence that large majorities of Americans, for example, adopt a compromise worldview in their attitudes towards political decisions [10]). This payoff scheme captures the common practical value of broad consensus in collective decisions, even as individuals pursue partisan preferences.

The structure of the influence network has profound, and surprising, effects on vote outcomes. Even when parties are equally matched in influence and representation, bias can arise when the parties differ in how their influence is assorted across the network. The phenomenon of information gerrymandering arises when one party punches above its weight by distributing its influence on a network so as to flip a disproportionate number of persuadable voters. To study information gerrymandering, we first define the "influence assortment" of an influence network. Positive influence assortment means players are predominately exposed to the voting intentions of members from their own party; and negative influence assortment means players are predominately exposed to the opposing party. Information gerrymandering arises when parties have asymmetric influence assortment. We quantify information gerrymandering as the difference in assortment between a party P and its opposition. Our model predicts that a party with a positive influence gap will benefit from information gerrymandering.

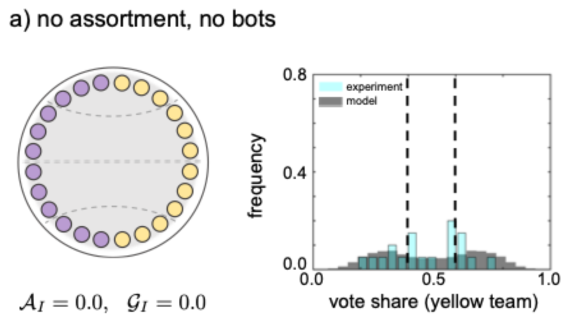

f) mean vote share and consensus frequency

	vote share model	vote share experiment	consensus model	consensus experiment
no assortment	0.5 (0.41, 0.59)	0.48 (0.42,0.54)	0.63 (0.35, 0.85)	0.55 (0.25, 0.7)
asymmetric assortment	0.6 (0.56, 0.63)	0.57 (0.53,0.6)	0.47 (0.2, 0.7)	0.45 (0.2, 0.65)
asymmetric bots	0.56 (0.54, 0.59)	0.53 (0.51,0.55)	0.19 (0.05,0.33)	0.18 (0.03, 0.28)
symmetric assortment	0.5 (0.46, 0.55)	0.51 (0.47,0.55)	0.18 (0.0, 0.35)	0.15 (0.0, 0.3)
symmetric bots	0.5 (0.49, 0.51)	0.5 (0.49, 0.51)	0.01 (0, 0.05)	0 (0.0)

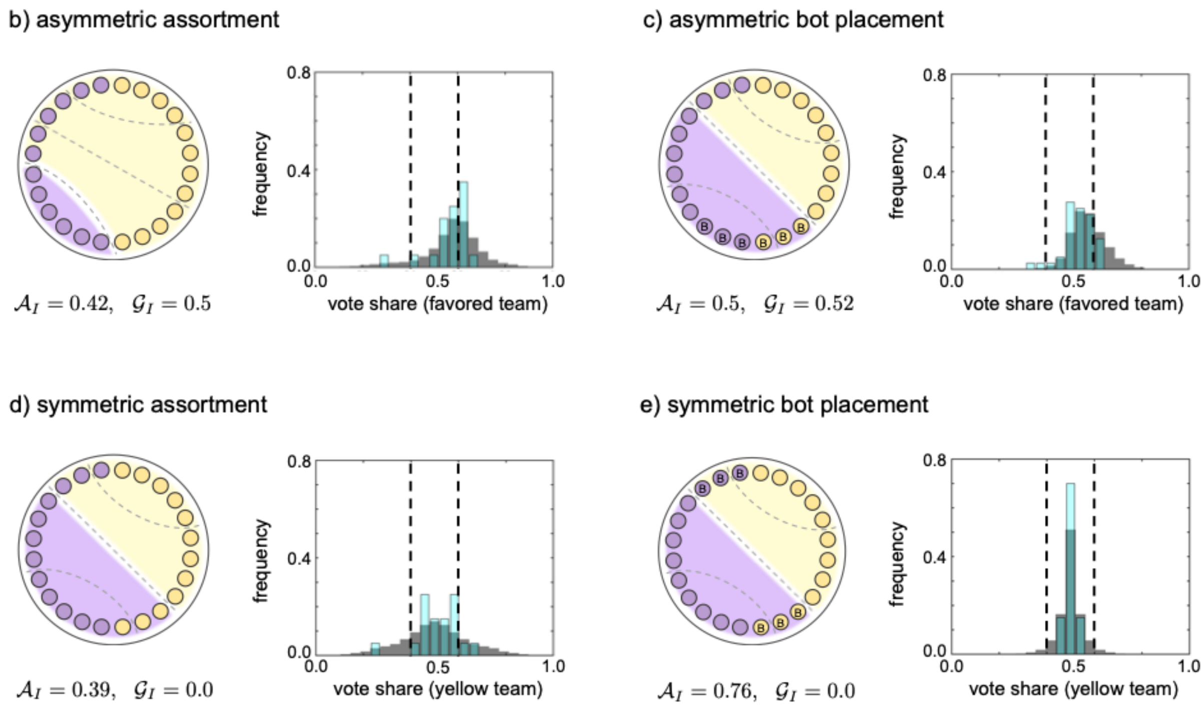

Figure 1 ***Undemocratic outcomes and polarization in the voter game.*** *(a) We inferred the parameters of our behavioral model from experiments in a baseline condition: networks with no influence assortment and no influence gap (A_I=0 and G_I=0). We used the model to predict the distribution of voting outcomes in four additional treatments each with $\geq$ 20 replicates (panels b-e, model predictions in gray and experimental results overlayed in light blue). (b) Information gerrymandering (G_I=0.5) produced vote shares as large as 67% for the more assorted party, which received a mean vote share of 57% across experimental replicates, consistent with the model prediction. (c) Asymmetric placement of six zealot bots also favored the party with a positive influence gap, resulting in vote shares as large as 63% and a mean vote share of 53%. (d) Symmetric influence assortment gave neither party an advantage, and the frequency of a consensus (15%) was dramatically reduced compared to networks without assortment (55%). (e) Symmetrically placed bots gave neither party an advantage and resulted in deadlock for all replicates. Dashed lines indicate thresholds V=60% and 1-V=40%.*

We confirm the predicted effects of information gerrymandering in pre-registered social network experiments varying only the structure of the influence network (Figure 1). To perform these experiments, we recruited at a total of 2,520 unique subjects from Amazon Mechanical Turk, over the course of 120 experimental sessions in total. We ran 20 replicate sessions for each of five experimental conditions. Once joining a session, subjects were randomly assigned to one of the two teams (purple or yellow), and were asked to take a tutorial, followed by four multiple-choice questions about the game rules and the payment structure. Each session of the voter game was conducted with exactly 24 players equally assigned to each team. Based on the experimental condition in the session, a directed network with 24 nodes was constructed and populated with players according to their teams. Each session of the voter game lasted for 240 seconds. Each subject was provided with real-time aggregated polling information of the voting intentions of those subjects assigned to their polling group; their own current vote; and their assigned team. Subjects could change their vote at any time before the timer ran out. We stored all updates of voting intentions of each player, the polling information seen by each player, and the aggregate votes of the whole group.

Our model predicts that information gerrymandering will skew the final vote towards the party with a positive influence gap ($G_I>0$); and that party will achieve a winning consensus more often than its opposition. Both of these predictions were validated experimentally (signed rank test, one-sided $p=0.003$ and binomial test, one-sided $p=0.02$, Figure 1b), demonstrating that a party does indeed gain a significant advantage by information gerrymandering. If one party benefits from information gerrymandering then, understanding this, the opposing party will naturally seek to do the same. The party at a disadvantage, with $G_I<0$, can redress the imbalance only by increasing the influence assortment of its members. But when both parties have equally high levels of influence assortment ($A_I>0$ and $G_I=0$) neither party will have an advantage. In fact, our model predicts that both parties will suffer from their self-constructed echo-chambers, resulting in deadlock more often than in the case of no influence assortment (Figure 1f). This prediction was also validated experimentally ($t=2.5$, one-sided $p=0.006$, Figure 1), demonstrating that increasing your party's influence assortment, while a rational response to information gerrymandering by your opponent, decreases the rate of consensus and therefore decreases payoffs for both parties. Another way to achieve the same advantage is to encourage players to adopt a zero-sum worldview and play as zealots. In the context of the voter game zealot bots always project the intention to vote for their party regardless of the polls. When one party's zealot bots are so deployed ($G_I>0$) our model predicts the vote will be skewed in its favor and the party will win a consensus more often than its opposition. The first of these predictions was validated experimentally (sign rank test, one-sided $p=0.002$, Figure 1) and the second was not statistically significant (binomial test, one-sided $p=0.2$, Figure 1). Thus, a party receives some advantage from information gerrymandering by zealot bots. But if both parties seek to use bots in the same way then overall influence assortment increases, neither party receives an advantage, and deadlock occurs in all experimental replicates (Figure 1e).

Furthermore, we identify extensive information gerrymandering in real-world influence networks, including online political discussions in the lead up to United States federal elections, and also in historical patterns of bill co-sponsorship in the United States Congress and European legislatures. Our analysis provides an account of the vulnerabilities of collective decision-making to systematic distortion by restricted information flow. Our analysis also highlights a group-level social dilemma: information gerrymandering can enable one party to sway decisions in its favor, but when multiple parties engage in gerrymandering the group loses its ability to reach consensus and remains trapped in deadlock.

Political polarization and echo chambers are the focus of intense research and public discussion. Unraveling the psychological mechanisms at play when people interact with different identity groups, opposing viewpoints, hot-button topics like climate change, fake or misleading news, trusted versus distrusted sources, and bots is vital for understanding decision making in real-world settings. Furthermore, affective polarization – negative attitudes to members of the other party, rather than to specific policies – is of great importance as it may cause people to adopt a zero-sum worldview. Nevertheless, our study on the voter game highlights how sensitive collective decisions are to information gerrymandering on an influence network, how easily gerrymandering can arise in realistic networks, and how widespread it is in real-world networks of political discourse and legislative process. Our analysis provides a new perspective and a quantitative measure to study public discourse and collective decisions across diverse contexts.

ACKNOWLEDGEMENT

This abstract is based on the following paper:
Stewart, A. J., Mosleh, M., Diakonova, M., Arechar, A. A., Rand, D. G., & Plotkin, J. B. (2019). Information gerrymandering and undemocratic decisions. *Nature*, *573*(7772), 117-121.

REFERENCES

1 Lazer, D. M. *et al.* The science of fake news. *Science* **359**, 1094-1096 (2018).
2 Vosoughi, S., Roy, D. & Aral, S. The spread of true and false news online. *Science* **359**, 1146-1151 (2018).
3 Stewart, A. J. *et al.* Information gerrymandering and undemocratic decisions. *Nature* **573**, 117-121 (2019).
4 Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A. & Bonneau, R. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science* **26**, 1531-1542 (2015).
5 Bakshy, E., Messing, S. & Adamic, L. A. Exposure to ideologically diverse news and opinion on Facebook. *Science* **348**, 1130-1132 (2015).
6 Matz, S. C., Kosinski, M., Nave, G. & Stillwell, D. J. Psychological targeting as an effective approach to digital mass persuasion. *Proceedings of the national academy of sciences* **114**, 12714-12719 (2017).
7 Pennycook, G. & Rand, D. G. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition* **188**, 39-50 (2019).
8 Woolley, S. C. Automating power: Social bot interference in global politics. *First Monday* **21** (2016).
9 Iyengar, S., Sood, G. & Lelkes, Y. Affect, not ideologya social identity perspective on polarization. *Public opinion quarterly* **76**, 405-431 (2012).
10 Gallup. *Americans favor compromise to get things done in washington*, <https://news.gallup.com/poll/220265/americans-favor-compromise-things-done-washington.aspx> (