Contagion, Confounding, and Causality: Confronting the Three C's of Observational Political Networks Research

Medha Uppala[1] and Bruce A Desmarais[1,2]

[1]*Center for Social Data Analytics, Pennsylvania State University. Email:* *mvu5040@psu.edu*
[2]*Department of Political Science, Pennsylvania State University*

Abstract

Contagion across various types of connections is a central process in the study of many political phenomena (e.g., democratization, civil conflict, voter turnout). Over the last decade the methodological literature addressing the challenges in causally identifying contagion in networks has exploded. In one of the foundational works in this literature, Shalizi and Thomas (2011) propose a permutation test for contagion in longitudinal network data that is not confounded by selection (e.g., homophily). We illustrate the properties of this test via simulation. We assess its statistical power under various conditions of the data; including the nature of the contagion, the structure of the network through which contagion occurs, and the number of time periods included in the data. We then apply this test to an example domain that is commonly considered in the context of observational research on contagion—the international spread of democracy. We find evidence of the international contagion of democracy. We conclude with a discussion of the practical applicability of the Shalizi & Thomas test to the study of contagion in political networks.

Keywords: social networks, contagion, homophily.

1 Introduction

Contagion has been found to characterize, e.g., individuals' decisions to vote (e.g., Rolfe 2012; Bond *et al.* 2012), the emergence of civil conflicts across countries (e.g., Maves and Braithwaite 2013), and the spread of democracy across countries (e.g., Starr 1991). It is, however, well known that inferences regarding contagion can be confounded by other dynamics that lead connected units to behave in similar ways (Franzese, Hays, and Kachi 2012). Shalizi and Thomas (2011) formalize and analyze the problem of inferring contagion in the presence of homophily. Contagion refers to the influence connected units have on each other, while homophily refers to the tendency for similar units to be connected due to their common traits. The arguments presented by Shalizi and Thomas (2011) apply to any type of dependence of connections on units' traits (e.g., heterophily, whereby dissimilar units tend to form ties)—generally referred to as "selection". We follow their terminology and use "homophily" as synonymous with selection. As a running example, we focus on the spread of democracy across countries. The spread of democracy is a question of contagion vs. homophily; do connected states influence each other to develop democratic institutions or do similarly governed states tend to be connected to each other over time? It is also possible that a state's governing choices are a result of an un-quantified blend of contagion and homophily. The methods we present and illustrate in this note allow the researcher to test for contagion in a way that is not confounded by the presence of homophily.

Given the need to estimate contagion separately from network homophily, it is important to recognize the circumstances where these two effects are confounded. Shalizi and Thomas (2011) explore this idea in detail, specifically considering the problem of identifying contagion in observational longitudinal network data. They analyze the problem within the causal diagram framework (Pearl 1995), and show that in observational social network data, *latent* homophily—tie

formation that is attributable to unmeasured attributes of units—and contagion are 'generically confounded' and cannot be identified without strong parametric assumptions. It is helpful here to build conceptual bridges between what Shalizi and Thomas (2011) refer to as "latent homophliy" and conceptual characterizations of the confounding of influence and homophily in political science, specifically by Franzese, Hays, and Kachi (2012). Franzese, Hays, and Kachi (2012) distinguish between two tie formation mechanisms that confound contagion inferences—common exposure, which occurs when an exogenous variable effects both tie formation and the outcome variable; and endogenous selection (or behavior homophily/heterophily), which occurs when tie values at one time point depend on outcome values from previous time points. Though Shalizi and Thomas (2011) do not explicitly model endogenous selection, the identification problems presented by endogenous selection are equivalent to those presented by latent homophily. Considering our running example, suppose a researcher sought to model the spread of democracy through diplomatic networks (Duque 2018). Latent homophily would confound inferences if, e.g., the researcher failed to measure any important cultural, geographic, economic, or security factors that shaped diplomatic relations between countries and future regime type developments—including countries' histories of regime type developments. More broadly, as political networks research commonly focuses on the factors that explain tie formation (e.g., Minozzi *et al.* 2020), we suspect that the presence of latent homophily in political network data is quite prevalent.

Shalizi and Thomas (2011) present a few ideas regarding how to make inferences on contagion in observational dynamic network data despite the presence of latent homophily. One of these is a permutation test that requires minimal assumptions regarding the structure of contagion and no assumptions regarding the structure of homophily. Specifically, since the test relies on associations across time-lagged data, it must be assumed that contagion does not completely manifest and then dissipate within a single time period—that the contagion effects persist for more than at least one time period. The test does not condition at all on any network structure, and does not rely on any assumptions regarding the structure of homophily/selection. In this paper, we implement this permutation test and show, through simulation, that it provides a sensible first step to uncover the presence of contagion in longitudinal social network data. We illustrate the use of the test on the dynamics of contagion of democracy, for which we find evidence.

2   Shalizi & Thomas Test

Shalizi and Thomas (2011) present a test for contagion that that does not condition on the ties between nodes (units). The process is to randomly permute the nodes in a social network into two groups (J_1 and J_2), and estimate the relationship between the outcome variable in one group and the time-lagged counterpart of the other group, while controlling for the time-lagged outcome of the current group. By iterating over all possible (or a large number of) partitions, and averaging over all iterations, "there will be a non-zero predictive ability if and only if there is actual contagion," in the social network. While the power of this test is low when the time series is short, the random partition of nodes into bins assures that the analysis is not confounded by conditioning on ties (i.e., two-node groups) that are themselves potentially formed according to homophily.

The steps of the test are as follows:

1. Given longitudinal network data, randomly partition the nodes into two bins, J_1 and J_2.
2. Aggregate, by, e.g., taking the mean of the outcome variable Y over all the bin nodes, at each time step, resulting in an aggregated time series of $\bar{Y}_{J_1}(t), \bar{Y}_{J_1}(t-1), \dots, \bar{Y}_{J_1}(1)$ for bin J_1 and time series $\bar{Y}_{J_2}(t), \bar{Y}_{J_2}(t-1), \dots, \bar{Y}_{J_2}(1)$ for bin J_2.
3. Estimate the relationship between $\bar{Y}_{J_i}(t)$ and $\bar{Y}_{J_k}(t-1)$, adjusting for $\bar{Y}_{J_i}(t-1)$, where $(i, k) \in \{(1, 2), (2, 1)\}$. We use ordinary least squares regression, but other estimators could be used.
4. Repeat steps 1 to 3. The total number of partitions possible for equal bin sizes is $\binom{n}{n/2}$.
5. The test for contagion is conducted by calculating empirical *p*-values with respect to the distri-

bution of estimated relationships between $\bar{Y}_{J_i}(t)$ and $\bar{Y}_{J_k}(t-1)$. A left(right)-tailed p-value is given by the proportion of estimated relationships that are less(greater) than zero.

Intuitively, this test is designed to detect a diffuse contagion signal whereby, due to the presence of contagion between some of the nodes in the two randomly partitioned groups, the aggregated values across the two groups are not independent. This indirect form of signal detection is necessary to avoid conditioning the contagion estimate on the network structure, which activates the confounding presented by latent homophily.

The Shalizi & Thomas test is a valuable tool in the study of contagion through political networks. It does, however, have a few limitations that are important to note. First, it a hypothesis test only, allowing one to evaluate the sharp null hypothesis of no contagion. It does not offer estimates of contagion parameters, or even the capacity to test for contagion through specific networks. Second, the contagion signal that the test relies on is the association of outcome values in $J_1(J_2)$ with recent values of $J_2(J_1)$, controlling for recent values of $J_1(J_2)$. The presence of this signal requires that the system embeds memory of the contagion effect. If contagion manifests and then dissipates completely within one time period—something that could happen if the time units are too aggregated—the test will fail to detect a signal. Third, the test can fail in the presence of a form of quasi-contagion that behaves like "interference" as discussed in the experimental literature (Bowers, Fredrickson, and Panagopoulos 2013). If one unit's covariate value affects the outcome value of another unit, this can look, to the Shalizi & Thomas test, like contagion through the outcome variables, but it is actually a more subtle form of cross-unit dependence. To give an example of this dynamic, major policy decisions (e.g., business or trade shut downs due to the COVID-19 pademic) made in a country may effect the economy of the country making the decision, as well as the economies of other countries (Cronert 2022). This dynamic would look like economic contagion to the Shalizi & Thomas test, but it is actually a form of cross-border economic dependence based on policymaking effects.

3   Simulation

To evaluate the performance of the Shalizi & Thomas test, we conduct a simulation study. We vary the time lengths, homophily and contagion conditions, and network structure, in order to understand the performance of the Shalizi & Thomas test. The test was implemented on four separate simulation conditions—one that includes contagion only, one that includes endogenous homophily only, one that includes both contagion and endogenous homophily, and one that includes endogenous homophily plus a time shock to outcome values. The last condition, the time shock, is included to evaluate the test's performance with time-based common exposure. The time shock is tuned to create, on average, a correlation of 0.5 in the Y values of units at the same time. The data generation models we use are similar to the ones used by Shalizi and Thomas (2011). We set the parameter values to assure that all of the data generating processes result in stationary outcome data.[1] The data generation models are outlined below.

Contagion only data:

1. Begin with n nodes in a network and each node i is assigned a scalar latent variable $X_i \sim \mathcal{U}(0, 1)$.
2. Generate directed ties between every pair of nodes (i, j) with probability $logit^{-1}(-3|x_i - x_j|)$. The smaller the different between x_i and x_j, the higher the probability of a tie between i and j. This produces an $n \times n$ adjacency matrix A where $A_{i,j} = 1$ represents a directed tie[2].
3. Initiate a starting value for the time series data. We use $Y_i(0) = 0.25x_i + N(0, 0.06^2)$
4. Given the adjacency matrix from step 2, contagion incorporated time series data is simulated as

---

1. We thank anonymous reviewers for pointing out that our initial simulation design created non-stationary data. We use the pCADFtest function in the R package punitroots to test for stationarity.
2. Technically, this represents generating data according to homophily, but Y does not cause A in this simulation, and thus cannot induce confounding.

$Y_i(t) = 0.25x_i + 0.3Y_i(t-1) + 0.7\overline{Y_k(t-1)} + N(0, 1^2)$ where $k = 1, \ldots, n$ and $A_{i,k} = 1$

Homophily Only data:

1. Begin with n nodes in a network and each node i is assigned a scalar latent variable $X_i \sim \mathcal{U}(0, 1)$.
2. Initiate a starting value for the time series data. We use $Y_i(0) = 0.25x_i + N(0, 0.06^2)$
3. Generate directed ties between every pair of nodes (i, j) with probability $logit^{-1}(-3|Y_i - Y_j|)$. This produces an $n \times n$ adjacency matrix A where $A_{i,j} = 1$ represents a directed tie.
4. Produce the values for the next time step as $Y_i(t) = 0.25x_i + 0.3Y_i(t-1) + N(0, 1^2)$.
5. Update the adjacency matrix while keeping the current number of network ties constant and the probability of the tie is proportional to the similarity of $Y_i(t)$ and $Y_j(t)$ [3]. This results in a dynamic network where the ties are correlated with the outcome variable.
6. Repeat steps 4 and 5 for each time point.

Contagion and Homophily data:

1. Begin with n nodes in a network and each node i is assigned a scalar latent variable $X_i \sim \mathcal{U}(0, 1)$.
2. Initiate a starting value for the time series data. We use $Y_i(0) = 0.25x_i + N(0, 0.06^2)$
3. Generate directed ties between every pair of nodes (i, j) with probability $logit^{-1}(-3|Y_i - Y_j|)$. This produces an $n \times n$ adjacency matrix A where $A_{i,j} = 1$ represents a directed tie.
4. Given the ties in the adjacency matrix, contagion incorporated time series data is simulated as $Y_i(t) = 0.25x_i + 0.3Y_i(t-1) + 0.7\overline{Y_k(t-1)} + N(0, 1^2)$ where $k = 1, \ldots, n$ and $A_{i,k} = 1$.
5. Update the adjacency matrix while keeping the current number of network ties constant and the probability of the tie is proportional to the similarity of $Y_i(t)$ and $Y_j(t)$[4].
6. With the updated adjacency matrix, repeat steps 4 and 5 until the desired length of time series data is achieved.

Homophily and Time Shock data:

1. Begin with n nodes in a network and each node i is assigned a scalar latent variable $X_i \sim \mathcal{U}(0, 1)$.
2. Initiate a starting value for the time series data. We use $Y_i(0) = 0.25x_i + N(0, 0.06^2)$
3. Generate directed ties between every pair of nodes (i, j) with probability $logit^{-1}(-3|Y_i - Y_j|)$. This produces an $n \times n$ adjacency matrix A where $A_{i,j} = 1$ represents a directed tie.
4. Produce the values for the next time step as $Y_i(t) = 0.25x_i + 0.3Y_i(t-1) + N(0, 1^2)$.
5. Add a temporal shock at each timestep as $Y(t) = Y(t) + N(0, 1^2)$
6. Repeat steps 3,4 and 5 until the desired length of time series data is achieved.

These simulation parameters result in networks with density of 0.3-0.4, tie reciprocity levels of 0.5-0.6, and first-order autocorrelation in Y of 0.3-0.8. Descriptive visualizations are presented in the online appendix. The contagion vs. homophily-only data generation models are outlined as directed acyclic graphs in Figure 1.

3.1   Results

All of the results presented in this paper reflect a Shalizi & Thomas test run comprising of 10,000 partition iterations. The results of our simulation study are presented in Figures 2 and 3. The implementation of the test on the contagion-only data and the contagion-plus-homophily data shows that the test identifies a positive contagion signal. On the homophily-only data and the data with homophily-plus-time shock (presented in the online appendix), the test does not identify a consistent contagion signal, i.e. it estimates a signal centered on zero, and actually some negative bias with short time series—a result that is consistent with the negative "Hurwicz" bias that arises with dynamic models fit to short time series data (Nickell 1981; Franzese, Hays, and Cook 2016). In the homophily-only case, the standard deviation of the signal begins to stabilize after around 30 time steps of data.

---

3. We use the `simulate` function from the R package `ergm` to update the adjacency matrix while keeping the tie count and in turn, the network density constant.

4. ibid.
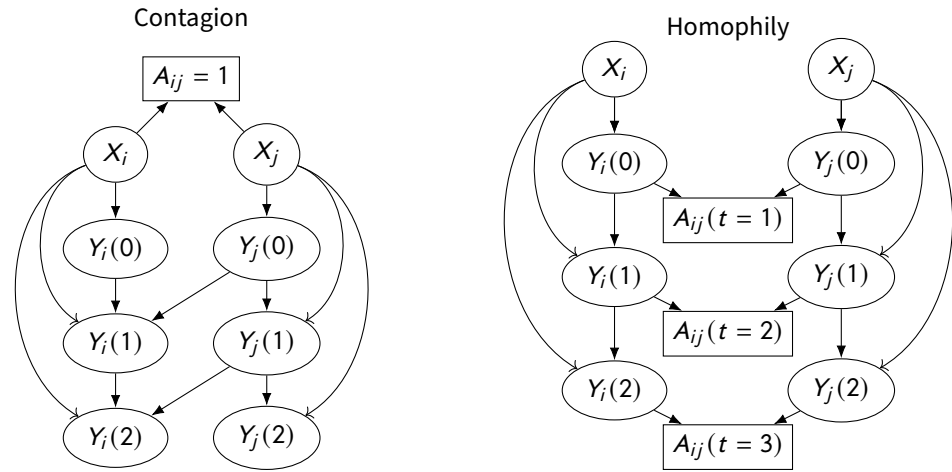
**Figure 1.** DAGs for the contagion vs. homophily only data generation models used for the simulation data.

When there is contagion in the data, the test is more variable the lower the time series lengths. For time series lengths greater than 20 steps, the signal converges to approximately 0.35, with or without endogenous homophliy in the data. In the first plot in Figure 3 we present summary estimates of the performance of the Shalizi & Thomas test. We summarize the test's performance at the 0.05 and 0.10 (two-tailed) significance levels, and consider both power and Type-1 error. With fewer than ten time steps, Type-1 error is high, and power is quite low, suggesting that this test should simply not be used with a relatively short time series. As a point of comparison, we estimated the correctly specified regression model on the simulated data using ordinary least squares, and found the power to exceed 0.90 even with one time point. With more than ten time steps, Type-1 error is slightly above the nominal significance levels and converges to the nominal levels with a longer time series. Statistical power converges to 1.0 with a long time series.

In addition to the true strength of the contagion effect, we expect the performance of the the Shalizi & Thomas test to improve as the network becomes more dense, as density determines the degree to which nodes are subject to contagion effects. The results from the contagion-only simulation at varying levels of network density, are presented in the online appendix. With relatively low density (0.05) the signal is weaker, but the signal strength levels out with densities of 0.30 or greater.

4   Application: The Spread of Democracy

When it comes to contagion dynamics, one of the domains of political science that would benefit most from the use of an observational hypothesis test that is not confounded by selection is the study of the contagion of country-level outcomes across international networks. Examples include the spread of specific policies (Towns 2012), civil conflict (Forsberg 2014), and democracy (Epstein 2005). Due to their national scales and substantial human effects, it is difficult and unethical to design randomized experiments that would provide design-based tests for dynamics such as conflict or democracy contagion. The necessity of working with observational data in this context underpins the importance of the Shalizi & Thomas test.

We focus on the international contagion of democracy through the analysis of Polity scores. The Polity IV Annual Time Series, 1800-2018 (Center for Systemic Peace) is a database that tracks and compiles regime changes and regime authority in countries with a total population greater than 500,000 in 2018. This database is extensively used in political science to study regime changes and effects of regime authority in countries over time. This is also a classic database used to study the contagion of democracy among networked countries over time and hence the implementation
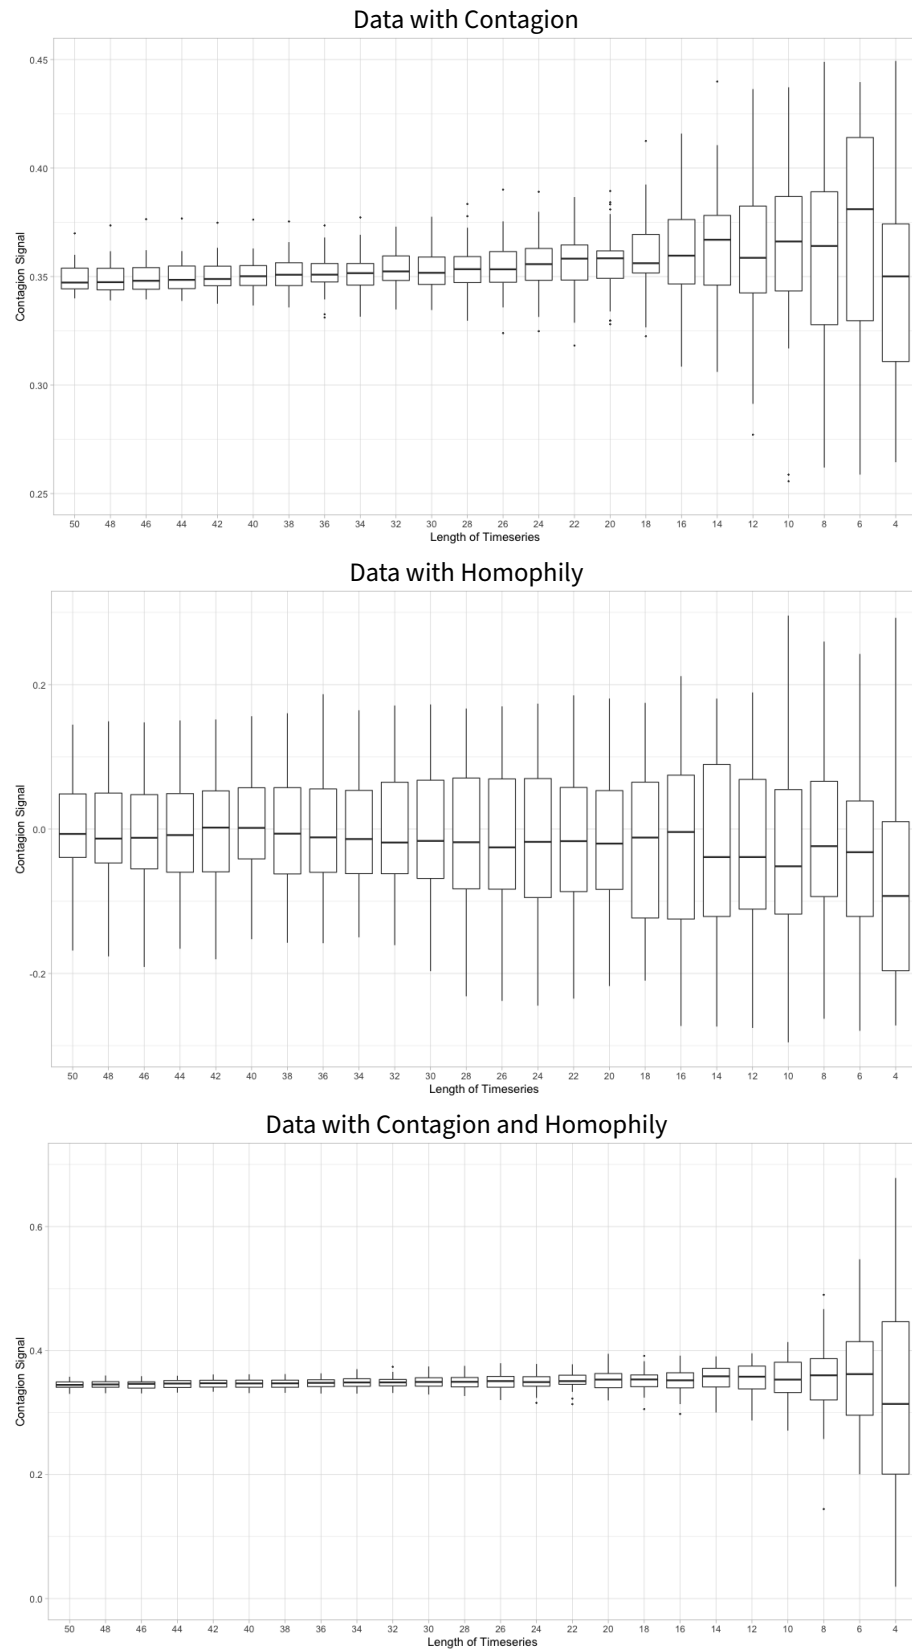
Figure 2. Results from the Monte Carlo simulation study of the Shalizi & Thomas test.
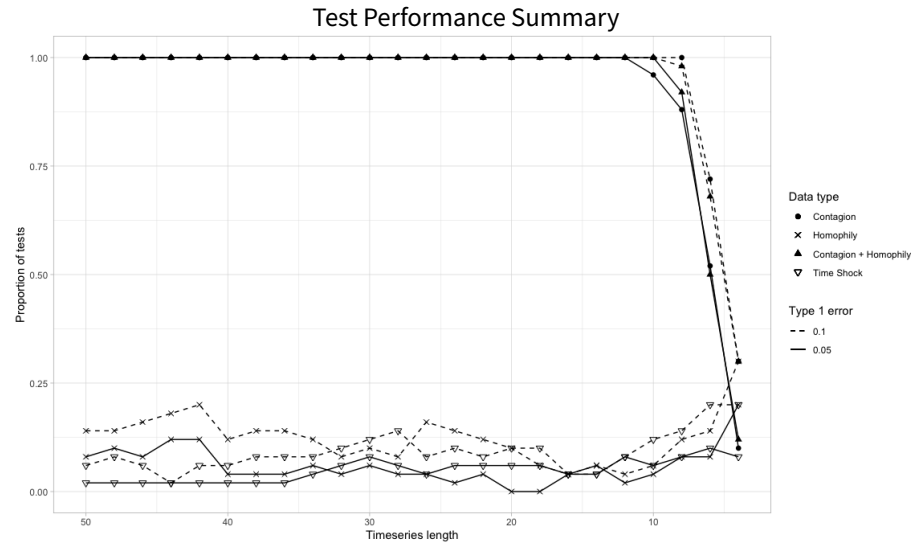
Figure 3. The plot summarizes power and type-1 error of the Shalizi and Thomas test in our simulation study.

of the Shalizi & Thomas test on this database is highly relevant.

We implemented the test on a 50-year subset of the democracy 'democ' score in the Polity IV data from 1969 to 2018 comprising of 118 countries. This democracy index (on a 11-point scale) scores how 'institutionalized' democracy is within a nation. While there are different indices that characterize regime type, the democ score is one of the common indices used to study the degree of democracy in countries' governments (Marshall *et al.* 2002). Due to substantial evidence that the panel of democracy scores is non-stationary, we apply the test to the panel of first differences in democracy scores. The distribution of 10,000 contagion estimates is visualized in the online appendix. The contagion signal was 0.169, and the proportion of estimates under zero (the one-tailed p-value) was 0.005. We find reliable evidence of positive contagion of democracy. This is an important finding, as it compliments and replicates the result from several model-based observational studies that democracy spreads through international networks, but we do not rely on a methodology that requires us to (a) identify the network through which it spreads, or (b) select the other, potentially confounding, factors for which to adjust our estimates.

5   Discussion

The presence of contagion dynamics in political processes can have substantial and important substantive implications. For example, the adoption of an innovative policy solution in one state/city/country would lead to innovation elsewhere, democratic reforms in one country could eventually lead to a more democratic future beyond that country's borders. Contagion in political turnout means that get-out-the-vote efforts have effects beyond those voters who are directly engaged by activists. As researchers are often limited in their material or ethical capacities to answer questions about contagion experimentally, we are forced to make inferences with observational data. In this paper we study and apply a hypothesis test for contagion proposed by Shalizi and Thomas that is not confounded by homophily. The test has limitations, but it is, to our knowledge, the only testing framework that can reliably differentiate contagion and homophily in observational data. We see this test as a compliment to the use of methods that rely on structural parametric assumptions to model contagion (e.g., Snijders 2017), offering researchers a non-parametric robustness check in testing for contagion. In an application to the international contagion of democracy, we reject the null hypothesis of no contagion, and find evidence for the international spread of democracy.

Supplementary Material

For supplementary material accompanying this paper, please visit https://doi.org/10.1017/pan.xxxx.xx.

References

Bond, R. M., C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489 (7415): 295–298.

Bowers, J., M. M. Fredrickson, and C. Panagopoulos. 2013. "Reasoning about interference between units: A general framework." *Political Analysis* 21 (1): 97–124.

Center for Systemic Peace. *Polity IV Annual Time Series, 1800-2018.* https://www.systemicpeace.org/inscrdata.html.

Cronert, A. 2022. "Towards a Swiss Army Knife State? The changing face of economic interventionism in advanced democracies, 1980–2015." *Review of International Political Economy* 29 (2): 477–501.

Duque, M. G. 2018. "Recognizing international status: A relational approach." *International Studies Quarterly* 62 (3): 577–592.

Epstein, R. A. 2005. "NATO enlargement and the spread of democracy: evidence and expectations." *Security Studies* 14 (1): 63–105.

Forsberg, E. 2014. "Transnational transmitters: Ethnic kinship ties and conflict contagion 1946–2009." *International Interactions* 40 (2): 143–165.

Franzese, R. J., J. C. Hays, and S. J. Cook. 2016. "Spatial-and spatiotemporal-autoregressive probit models of interdependent binary outcomes." *Political Science Research and Methods* 4 (1): 151–173.

Franzese, R. J., J. C. Hays, and A. Kachi. 2012. "Modeling history dependence in network-behavior coevolution." *Political analysis* 20 (2): 175–190.

Marshall, M. G., T. R. Gurr, C. Davenport, and K. Jaggers. 2002. "Polity IV, 1800-1999: comments on Munck and Verkuilen." *Comparative Political Studies* 35 (1): 40–45.

Maves, J., and A. Braithwaite. 2013. "Autocratic institutions and civil conflict contagion." *The Journal of Politics* 75 (2): 478–490.

Minozzi, W., H. Song, D. M. Lazer, M. A. Neblo, and K. Ognyanova. 2020. "The incidental pundit: Who talks politics with whom, and why?" *American Journal of Political Science* 64 (1): 135–151.

Nickell, S. 1981. "Biases in dynamic models with fixed effects." *Econometrica: Journal of the econometric society,* 1417–1426.

Pearl, J. 1995. "Causal diagrams for empirical research." *Biometrika* 82 (4): 669–688.

Rolfe, M. 2012. *Voter turnout: A social theory of political participation.* Cambridge University Press.

Shalizi, C. R., and A. C. Thomas. 2011. "Homophily and contagion are generically confounded in observational social network studies." *Sociological Methods and Research* 40 (2): 211–239.

Snijders, T. A. 2017. "Stochastic actor-oriented models for network dynamics." *Annual review of statistics and its application* 4:343–363.

Starr, H. 1991. "Democratic dominoes: Diffusion approaches to the spread of democracy in the international system." *Journal of Conflict Resolution* 35 (2): 356–381.

Towns, A. E. 2012. "Norms and social hierarchies: Understanding international policy diffusion "from below"." *International organization* 66 (2): 179–209.