Online Appendix for

The Incidental Pundit

Who Talks Politics with Whom, and Why?

William Minozzi
Associate Professor, Department of Political Science, The Ohio State University. minozzi.1@osu.edu. Corresponding author.

Hyunjin Song
Assistant Professor, Department of Communication, University of Vienna. hyunjin.song@univie.ac.at

David M. J. Lazer
Distinguished Professor, Department of Political Science and College of Computer and Information Science, Northeastern University, d.lazer@northeastern.edu

Michael A. Neblo
Associate Professor, Department of Political Science, The Ohio State University. neblo.1@osu.edu

Katherine Ognyanova
Assistant Professor, Department of Communication, School of Communication and Information, Rutgers University. katya@ognyanova.net

Contents

The Friends for Life Study **A1**

Missing Data Imputation **A4**

Permutation-based Simulations Including Leaners **A4**

Identifying Assumptions for Causal Inference **A6**

Goodness-of-fit Assessment **A8**

The Friends for Life Study

The Friends for Life Study is a dormitory-based, multiplex, multi-site, whole network, panel dataset that spans fifteen locations ("chapters"), from 2008 to 2016 (except 2009), for a total of 113 chapter-years. Participants were recipients of a scholarship that required residence in the chapter house throughout the recipient's time in college. One site appears only in 2016, and so we exclude it from the analysis, since we lack longitudinal information about its networks and members. Table A1 reports descriptive statistics on these chapters and their host universities.

Table A1: Chapter Descriptive Statistics

Chapter	Avg. Size	Average Frac. Dem.	Average Frac. Rep.	Average Frac. White	Average Frac. Male	Average Disc. Frac.	University Size (undergrad)
1	44	0.26	0.25	0.81	0.78	0.41	27k
2	112	0.43	0.21	0.61	0.77	0.43	34k
3	56	0.25	0.41	0.96	0.73	0.46	33k
4	62	0.39	0.30	0.68	0.61	0.31	8.4k*
5	48	0.19	0.41	0.92	0.87	0.41	17k
6	60	0.26	0.35	0.87	0.84	0.63	30k
7	66	0.27	0.32	0.84	0.79	0.33	39k
8	53	0.26	0.43	0.93	0.87	0.50	32k
9	41	0.24	0.39	0.94	0.78	0.39	24k
10	37	0.30	0.27	0.65	0.99	0.42	15k
11	34	0.34	0.14	0.72	0.71	0.55	8.4k*
12	66	0.19	0.38	0.93	0.88	0.45	47k
13	42	0.25	0.42	0.80	0.69	0.50	33k
14	62	0.32	0.32	0.94	0.79	0.52	30k

Descriptive statistics are averaged over all survey waves and calculated based on non-imputed values. "Average Disc. Frac." refers to the across-wave average of the percentages of political discussants that respondents reported were members of the chapter. * Private. All others are public universities.

The core of the dataset consists of two pieces: an individual-level panel and a dyadic-level panel. At the individual level, we derived information on the *Chapter* and *Cohort* of each respondent based on information provided by the scholarship organization. We used *Cohort* to identify *School Year*, based on the number of years since an individual entered college. To collect other individual-level variables, we fielded a longitudinal survey with waves every August, during the first week of classes, and November, post-election. Table A2 presents summary statistics on the resulting individual-year dataset. For this paper, we used responses to the following items:

Gender

Item: What is your gender?

Responses: Male, Female

Note: We used these responses to code an indicator for *Female*.

Race/Ethnicity

Item: What racial or ethnic group best describes you? (Check all that apply.)

Responses: White, Black, Hispanic, Asian, Native American, Middle Eastern, Mixed, Other

Note: Multiple responses were accepted. Due to low frequencies, we recoded all multiple responses and all responses in the categories "Native American," "Middle Eastern," and "Mixed" as "Other." We then coded indicators for *Asian*, *Black*, *Latino*, and *White*.

Religion

Item: What is your religion? (Check all that apply.)

Responses: Baptist–any denomination, Protestant (e.g. Methodist, Lutheran, Presbyterian, Episcopal), Catholic, Mormon, Jewish, Muslim, Hindu, Buddhist, Pentecostal, Eastern Orthodox, Other Christian, Other non-Christian, None

Note: Multiple responses were accepted. We recoded all multiple responses that included "Baptist," "Pentecostal," and "Other Christian" as "Multiple Responses (Evangelical)," and all other multiple responses as "Multiple Responses." We then coded *Evangelical* as an indicator that was 1 for responses including "Baptist," "Pentecostal," "Other Christian," and "Multiple Responses (Evangelical)," and 0 otherwise.

Ideology

Item: In general, do you think of yourself as...

Responses: Extremely liberal, Liberal, Slightly liberal, Moderate, Slightly conservative, Conservative, Extremely conservative

Note: We recoded this variable to run from 0 to 1, with higher values indicating more Conservative answers.

Party ID

Item: Generally speaking, do you think of yourself as a...

Responses: Republican, Democrat, Independent, Another party

[If Republican:] Would you call yourself a...

Responses: Strong Republican, Not very strong Republican

[If Democrat:] Would you call yourself a...

Responses: Strong Democrat, Not very strong Democrat

[If Independent or Another Party:] Do you think of yourself as closer to the...

Responses: Republican Party, Democratic Party, Neither

Note: Based on these responses, we coded indicators for *Democrat* and *Republican*, excluding leaners, and *Democrat w/ Leaners* and *Republican w/ Leaners*, including them.

Political Interest

Item: In general, how interested are you in politics and public affairs?

Responses: Very interested, Somewhat interested, Slightly interested, Not at all interested

Note: We rescaled this variable to run from 0 to 1, with higher values indicating more interest.

Political Knowledge

Item: Do you happen to know what job or political office is now held by [Dick Cheney/Joe Biden]?

Responses: [open text box]

Item: Whose responsibility is it to determine whether a law is constitutional or not?

Responses: Congress, The President, The Supreme Court, Don't Know

Item: How much of a majority is required for the US House and Senate to override a Presidential Veto?

Responses: 50% + 1 (simple majority), 60% (three fifths), 66.7% (two thirds), 75% (three quarters), 100% (unanimity), Don't Know

Item: Which party currently holds a majority of seats in the House of Representatives?

Responses: Democratic Party, Republican Party, Neither, Don't Know

Item: Which party currently holds a majority of seats in the Senate?

Responses: Democratic Party, Republican Party, Neither, Don't Know

Note: Correct answers depend on when the survey was fielded. We coded correct answers as "1," and incorrect answers and "Don't Knows" as "0." If an observation contained any non-missing answers to these items, we coded the missing answers as "0." For respondents who answered none of these items, we coded all items as missing. Finally, we summed the correct answers to get a score for *Political Knowledge* ranging from 0 to 5, then rescaled to run from 0 to 1.

Conflict Avoidance

Item: When people argue about politics, I often feel uncomfortable.
Item: When I'm in a group, I stand my ground even if everyone else disagrees with me.
Item: I usually find it easy to see political issues from other people's point of view.
Item: If I'm sure I'm right about a political issue, I don't waste time listening to other people's arguments.
Item: I have no problem revealing my political beliefs, even to someone who would disagree with me.
Item: I would rather not justify my political beliefs to someone who disagrees with me.
Item: I do not take it personally when someone disagrees with my political views.
Item: When I'm in a group, I often go along with what the majority decides is best, even if it is not what I want personally.
Responses: All items were five-point, Likert-type scales with responses ranging from "Strongly agree" (5) to "Strongly disagree" (1).
Note: Different waves of the survey included subsets of these items. We fit an Bayesian latent variable measurement model to estimate an underlying scale for all individuals.

Table A2: Summary Statistics

Variable	Mean	SD	Min	Max	% Missing
Female	0.24	0.42	0	1	4%
Asian	0.02	0.13	0	1	4%
Black	0.04	0.21	0	1	4%
Latino	0.07	0.25	0	1	4%
White	0.82	0.39	0	1	4%
School Year $= 1$	0.28	0.45	0	1	0%
School Year $= 2$	0.25	0.43	0	1	0%
School Year $= 3$	0.23	0.42	0	1	0%
School Year ≥ 4	0.23	0.42	0	1	0%
Evangelical	0.15	0.36	0	1	7%
Democrat	0.34	0.47	0	1	16%
Democrat (w/ leaners)	0.44	0.49	0	1	16%
Republican	0.37	0.48	0	1	16%
Republican (w/ leaners)	0.45	0.50	0	1	16%
Ideology	0.50	0.26	0	1	15%
Political Interest	0.56	0.32	0	1	15%
Political Knowledge	0.73	0.31	0	1	35%
Conflict Avoidance	0.01	0.82	-3.01	3.39	11%

n observations $= 6,248$, and n individuals $= 2,521$.

At the dyadic level, we collected information on the networks within each chapter. In addition to the five networks discussed in the paper (Political, Friend, Time, Esteem, Negative), we also asked about Academic networks in a subset of years. We omitted this network from the paper because we lack data on it in 2008, 2010, and 2016. For each network, each respondent was shown a roster of all individuals in the chapter, and asked to select the ones for which each of the following statements was true:

Political
"I frequently discuss politics, social issues, or current events with this person"

Esteem
"I hold this person in especially high esteem."

Friend
"This person is a close friend."

Time
"I spend a lot of time around this person."

Negative
"Sometimes I do not find it easy to get along with this person."

Academic
"This person has assisted me with my academics."

Missing Data Imputation

Nonresponse—at the item, observation, and individual levels–was an issue, with varying response rates across different waves of the panel. List-wise deletion based on nonresponse would result in substantial loss of information (King et al. 2001), which is especially problematic with network data (Kossinets 2006). Therefore, we imputed missing responses to our non-network items using Amelia II (Honaker, King and Blackwell 2011).

Specifically, we identified the set of variables that form the basis for our eventual analysis. For imputation, we used information from both the August surveys and the November surveys, across all available waves of the survey. These included the categorical variables *Gender*, *Race/Ethnicity*, *Religion*, *Chapter*, and *Cohort*; the ordered variables *Political Interest*, *Political Knowledge*, *Party ID*, and *Ideology*; and the continuous variable *Conflict Avoidance*, estimated using the method described in the previous section. Because *Chapter* and *Cohort* were provided by the organization, there is no missingness on these variables. However, once a respondent had provided an answer to the demographic items, they were not asked the same item in every subsequent survey. In particular, *Gender*, *Race/Ethnicity*, and *Religion* would be asked at most once a year, and in some cases less often. Therefore, our first step was to use last observation carried forward and next observation carried back in succession on these three variables. Finally, we used the amelia function from the Amelia package in R to impute ten complete datasets, taking advantage of the panel nature of the dataset by including time- and respondent-specific information.

Permutation-based Simulations Including Leaners

The results in the main paper code leaners (i.e., those who answered our first party ID branching question with either "Independent," "Other Party," etc., but who then responded to the follow-up party lean question by selecting either "Democratic Party" or "Republican Party") as neither Democrats nor Republicans. We chose this coding convention because we are interested purposive homophily based on party, and individuals who attempt to mask their party ID would seem to thwart such intentional selection criteria. Nevertheless, there is ample evidence that these leaners are not independents, but identifiers with major parties (Keith et al. 1986; Petrocik 2009).

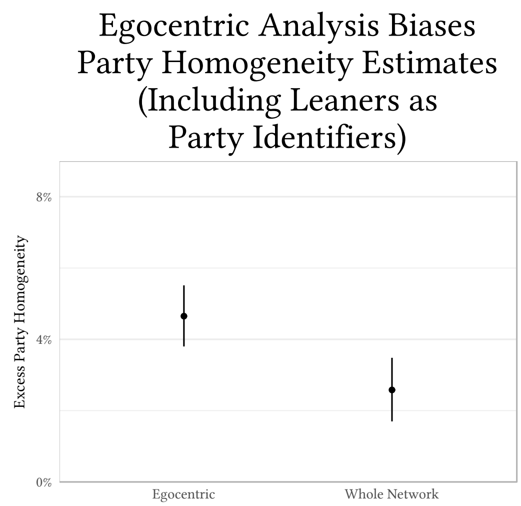

Figure 1: Excess party homogeneity is the difference between observed levels and those that would occur by chance, calculated by permuting party labels. The estimate on the left permutes across all chapters in a year, as with egocentric data alone. Such an approach cannot leverage information about local opportunity structures. The estimate on the right permutes at the chapter-year level, leveraging such whole network information. Ignoring the additional information afforded by whole network data will typically bias estimates. Both estimates include leaners as major party identifiers. Including leaners as major party identifiers does little to change the findings.

Therefore, we re-ran all analysis in the paper coding leaners as major party identifiers. As we note in the paper, very little changes based on this coding decision.

In this section, we present the results of the permutation-based simulations to estimate excess party homogeneity when we include leaners as major party identifiers. Including leaners changes the levels of observed homogeneity, but does little to alter the estimated levels of excess homogeneity, defined as the difference between observed levels and those based on permutations of party labels. For example, whereas 27% of all ties were between explicit major party identifiers, when we include leaners, this proportion rises to 42%. However, the middle 95% ranges of the two are similarly displaced from these observed levels. For explicit major party identifiers, and permuting party labels globally, the middle 95% range is $[24\%, 26\%]$, so the range of estimated excess party homogeneity is 2% to 4%. Including leaners, the middle 95% range is $[39\%, 41\%]$, so the range of estimated excess party homogeneity falls slightly to 1% to 3%.

Figure 1 presents the complete results in an analogous form to that in Figure 3 from the paper. This figure underscores the two important points from the paper. First, if we focus on all ties rather than those between major party identifiers, we find lower rates of excess party homogeneity. Second, if we permute party labels locally rather than globally, we also find lower rates of excess party homogeneity. Both of these findings emerge whether we include leaners as major party identifiers or not. Moreover, there is substantial overlap in the distributions of estimated excess party homogeneity based on both coding decisions.

Identifying Assumptions for Causal Inference

Our focus in the main paper is on prediction, but we can also use our model to yield causal inferences—conditional on appropriate identifying assumptions. Any model identifies causal estimates under the (implausible) assumption that one has correctly specified the data generating process. A more plausible strategy is to assume mean conditional ignorability: for example, the counterfactual expected value of a tie conditional on *Same Party* taking a particular value is equal to our model's estimate. This approach is similar to any observational design that yields causal inferences by adjusting for observed covariates, such as matching or synthetic case control.

This identification strategy requires us to assume both that there are no unobservables correlated with discussion and the causal variable, and that we have not introduced post-treatment bias by including such terms in our model. We have included a large slate of covariates to cope with the first requirement. But the second requirement is likely violated by the inclusion of terms like *Dyadic Stability*. For example, *Same Party* might have caused a discussion tie to form in a previous year, thus changing the value of *Dyadic Stability*. Adjusting for *Dyadic Stability* can therefore introduce bias into our estimates of marginal effects. Here, the goal of predictive accuracy is at odds with that of unbiased causal inference.

Similar issues plague other variables. Structural variables like the closure terms and the reciprocity term are measured contemporaneously with the outcome, and are therefore also susceptible to such bias. Further, the lagged social relationships might also be considered post-treatment, as two people might become friends because of political similarity.

There are no foolproof fixes for post-treatment bias. We can, however, perform a sensitivity analysis with a model that excludes post-treatment variables. Therefore, we re-estimated demographic and political homophily terms using a model that includes only the Demographic and Political terms.

Our key quantity of interest is the average marginal effect. To calculate this, for each dyad we first estimated difference between (1) the probability of discussion if a dyad identifies with the same political group, and (2) that in which the dyad identifies with different groups. Next, we take the mean over all dyads in a chapter to get chapter-level average marginal effects. Finally, we pooled across chapters using the Bayesian model.

We denominate these quantities in percentage point terms, and consider *Same Party* and *Ideological Proximity* in sequence.

First, there are only small changes in the estimates for *Same Party*. As a reminder, according to the Full model, the estimated marginal effect of *Same Party* as reported in the manuscript is 0.3% with 95% interval $[0.2\%, 0.4\%]$. Based on the model that includes only Demographic and Political terms, this value is 0.6% $[0.3\%, 0.8\%]$. Thus, there is little change in the estimates for party-based homophily, which is consistent with at most small levels of post-treatment bias.

The differences across models are a bit larger for *Ideological Proximity*. As a reminder, according to the Full model, the estimated marginal effect of *Ideological Proximity* is 1.0% $[0.8\%, 1.2\%]$. In contrast, the model including only Demographic and Political terms yields an estimate of 2.0% $[1.7\%, 2.4\%]$. This analysis suggests that the estimated marginal effect based on the Full model may have been contaminated by post-treatment bias.

While these auxiliary analyses reveal sensitivity of our Full model, at least for *Ideological Proximity*, they also help reinforce our main claim: there is evidence of limited purposive political homophily, and that it is generally smaller than incidental political talk based on other characteristics. With the Demographic and Political model, we also see increased estimates for *Same Race/Ethnicity* (2.9% $[2.6\%, 3.3\%]$), *Same Gender* (5.2% $[4.9\%, 5.5\%]$), and *Same Cohort* (6.8% $[6.4\%, 7.2\%]$).

Based on this analysis, we draw two conclusions. First, estimates of causal effects based on our predictive model may be underestimates, but generally have the correct sign. Thus, causal interpretations of the quantities presented in the main paper may be conservative, which is why we make no such claims there. Second, the main finding—that political homophily effects are positive and precisely measurable, yet are also generally smaller than demographic homophily effects—is robust to different sets of identifying assumptions.

Goodness-of-fit Assessment

In this section, we present an assessment of goodness-of-fit (GOF) for the *Full* TERGM model from the main paper. To assess GOF, we calculated a set of network statistics for the observed networks and compared them with statistics for simulated networks. For each chapter and imputation, we simulated 100 networks, thus yielding 1000 simulations for each chapter. Networks were simulated with the gof function from the btergm package (Leifeld, Cranmer and Desmarais 2018), using the default values for the number of burn-in iterations (10000) and thinning interval (1000). The statistics we used to assess GOF include the distributions of the seven directed shared partners counts, indegree and outdegree, geodesic distance, walktrap modularity, and directed triads. The figures presented in the next fourteen pages display these results, with one chapter per figure. In each figure, the light grey lines represent the 95% intervals for simulated networks, and the bold black lines represent the statistics for observed networks. This assessment shows that the axillary statistics are within the acceptable range of the simulated distribution in most cases, but with some exceptions. Based on this assessment, we conclude that the *Full* model yields satisfactory goodness-of-fit.

Figure 2: Goodness-of-fit for Chapter 1

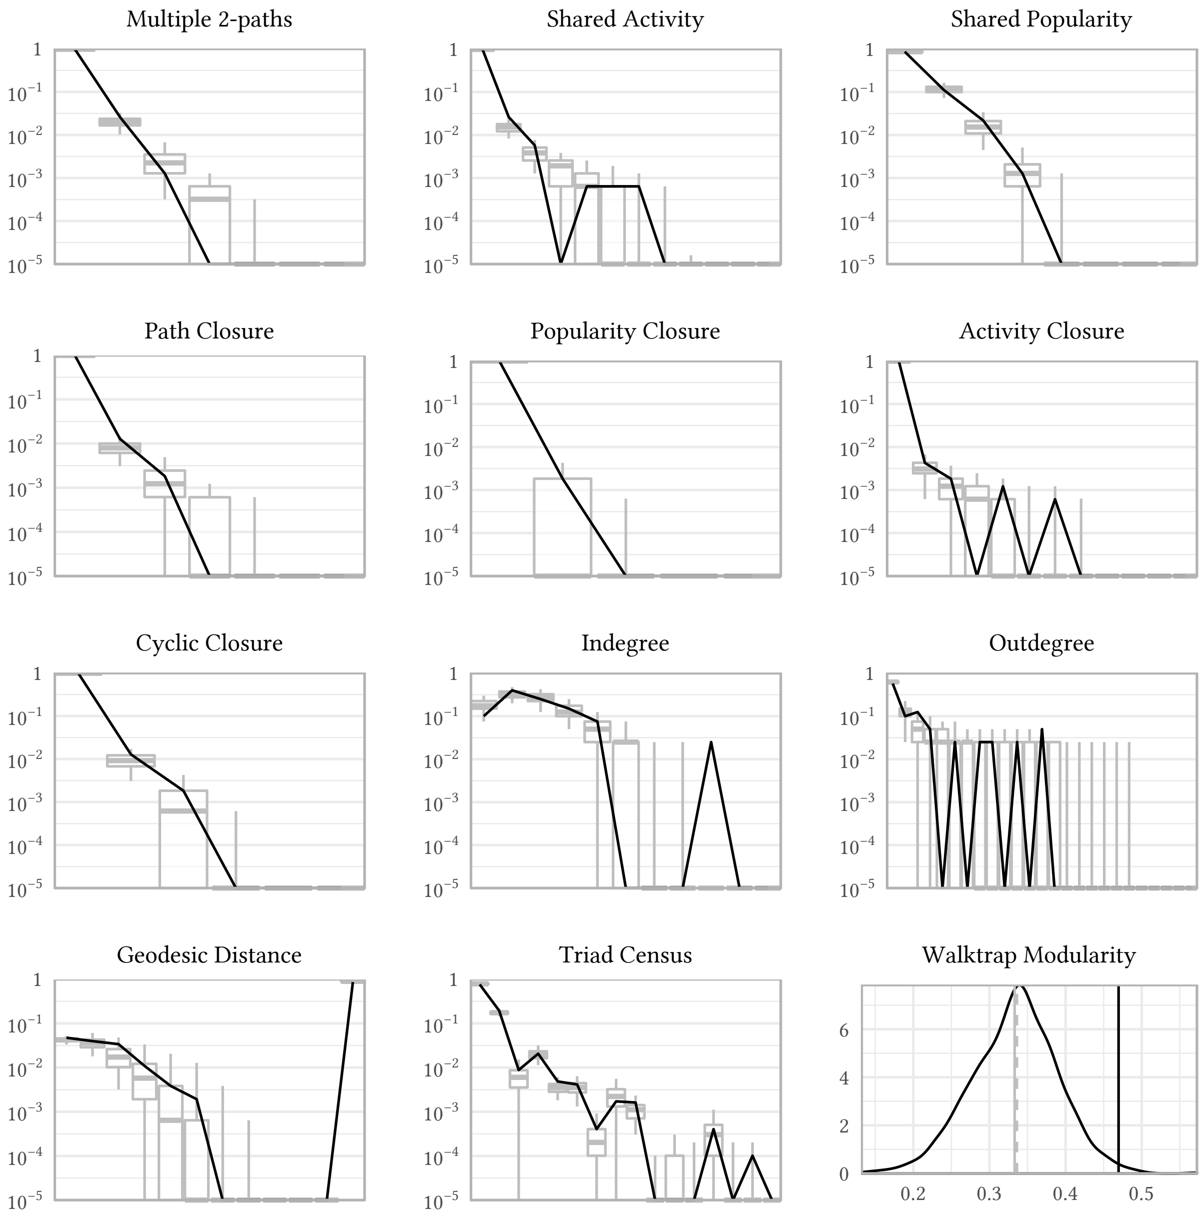

Figure 3: Goodness-of-fit for Chapter 2

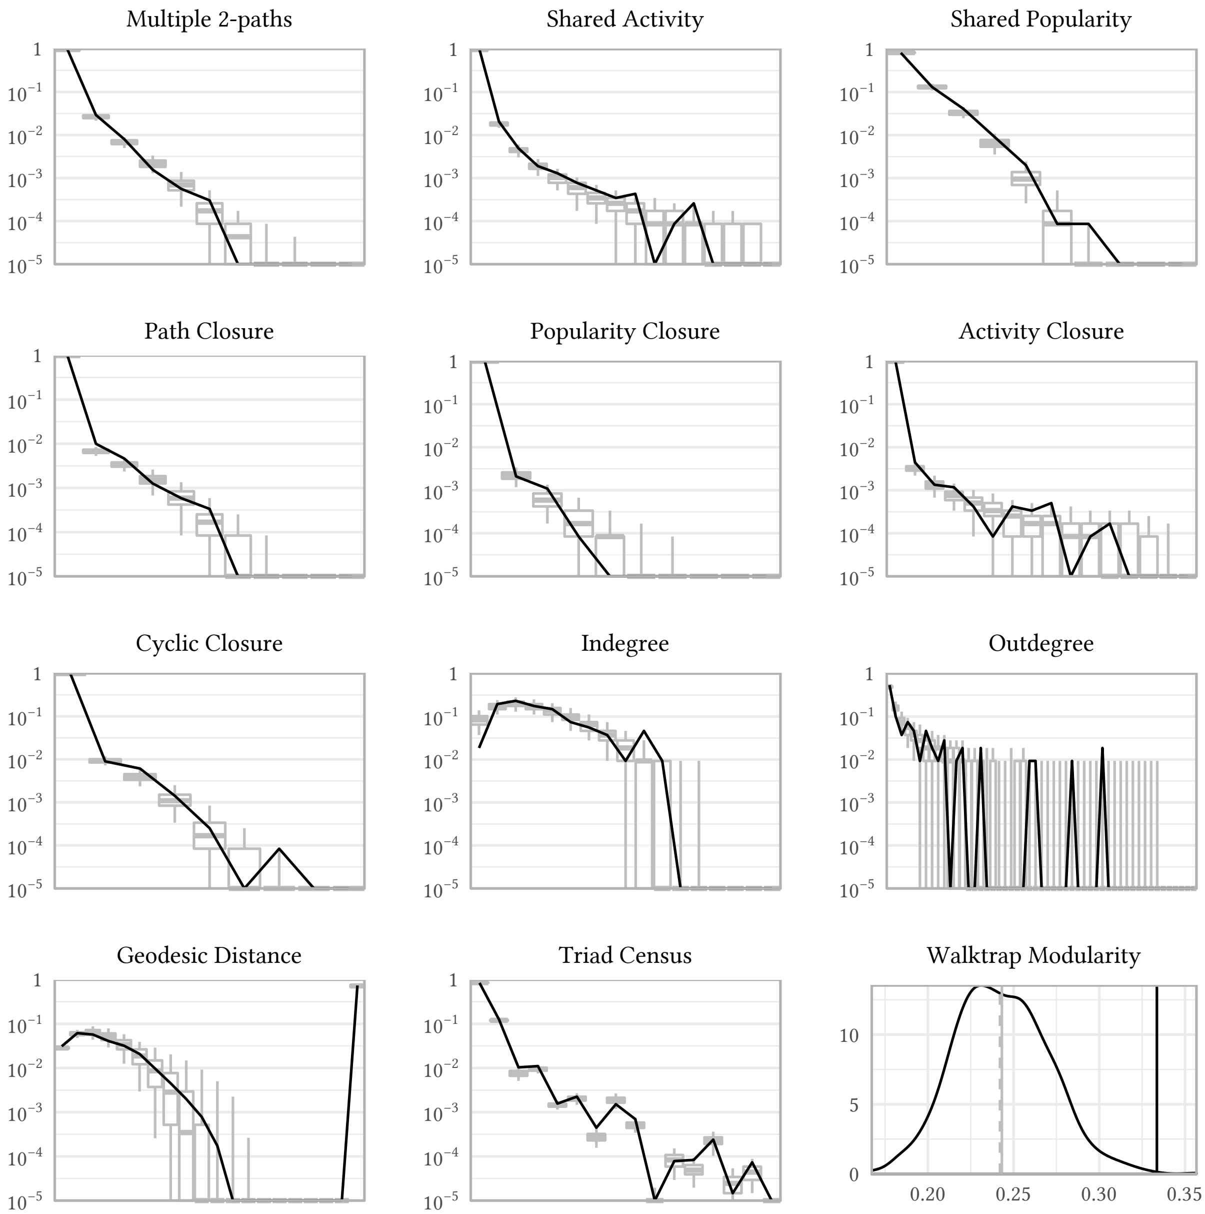

Figure 4: Goodness-of-fit for Chapter 3

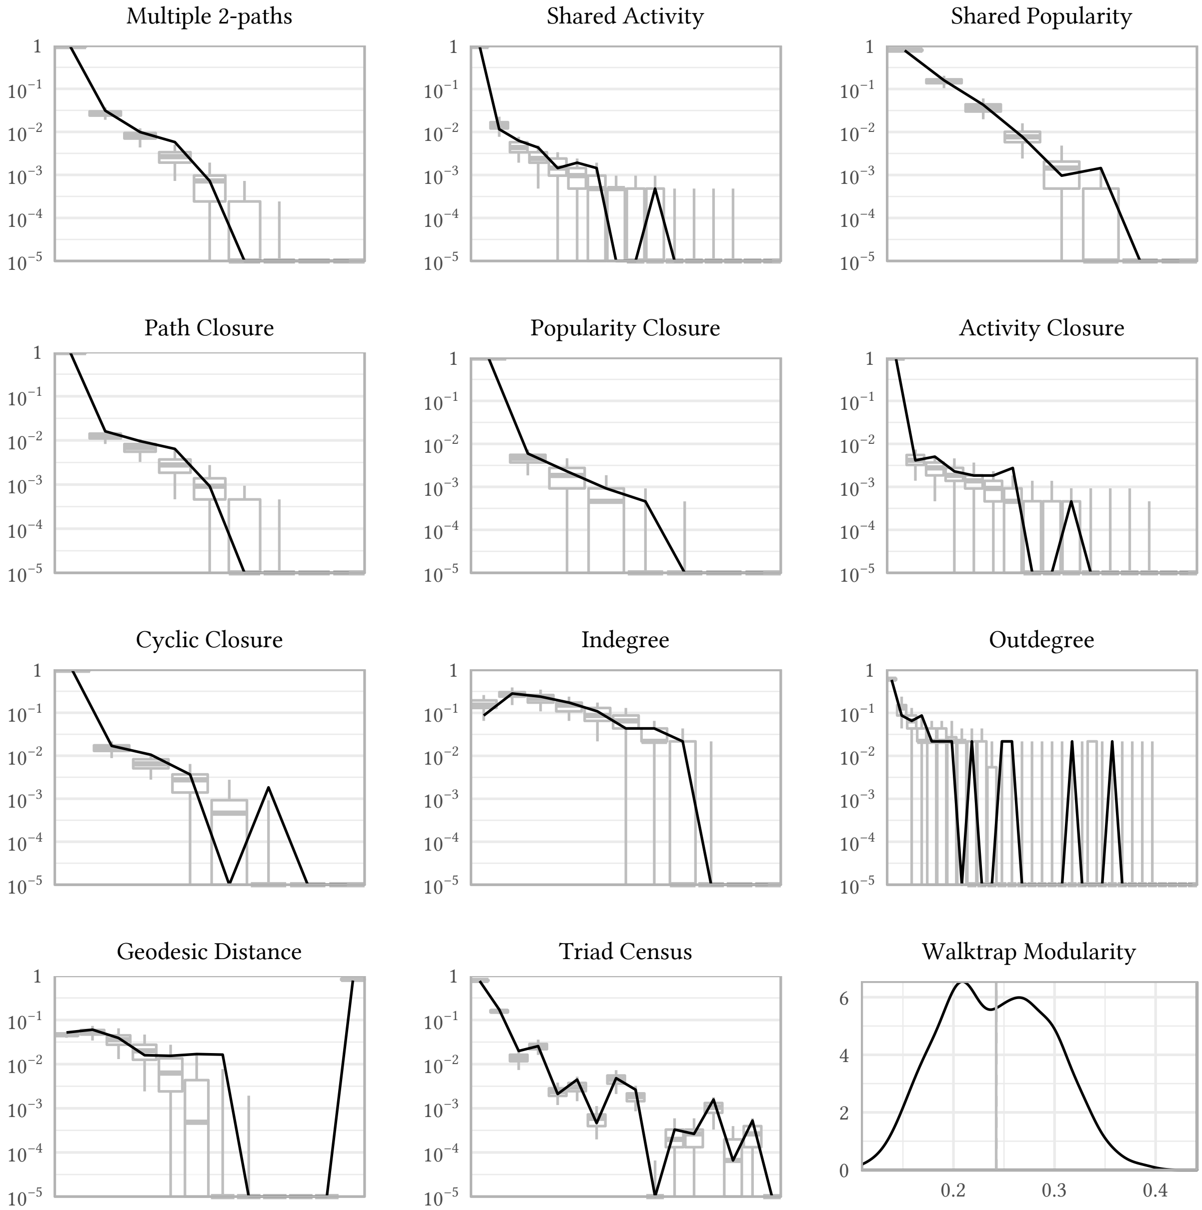

Figure 5: Goodness-of-fit for Chapter 4

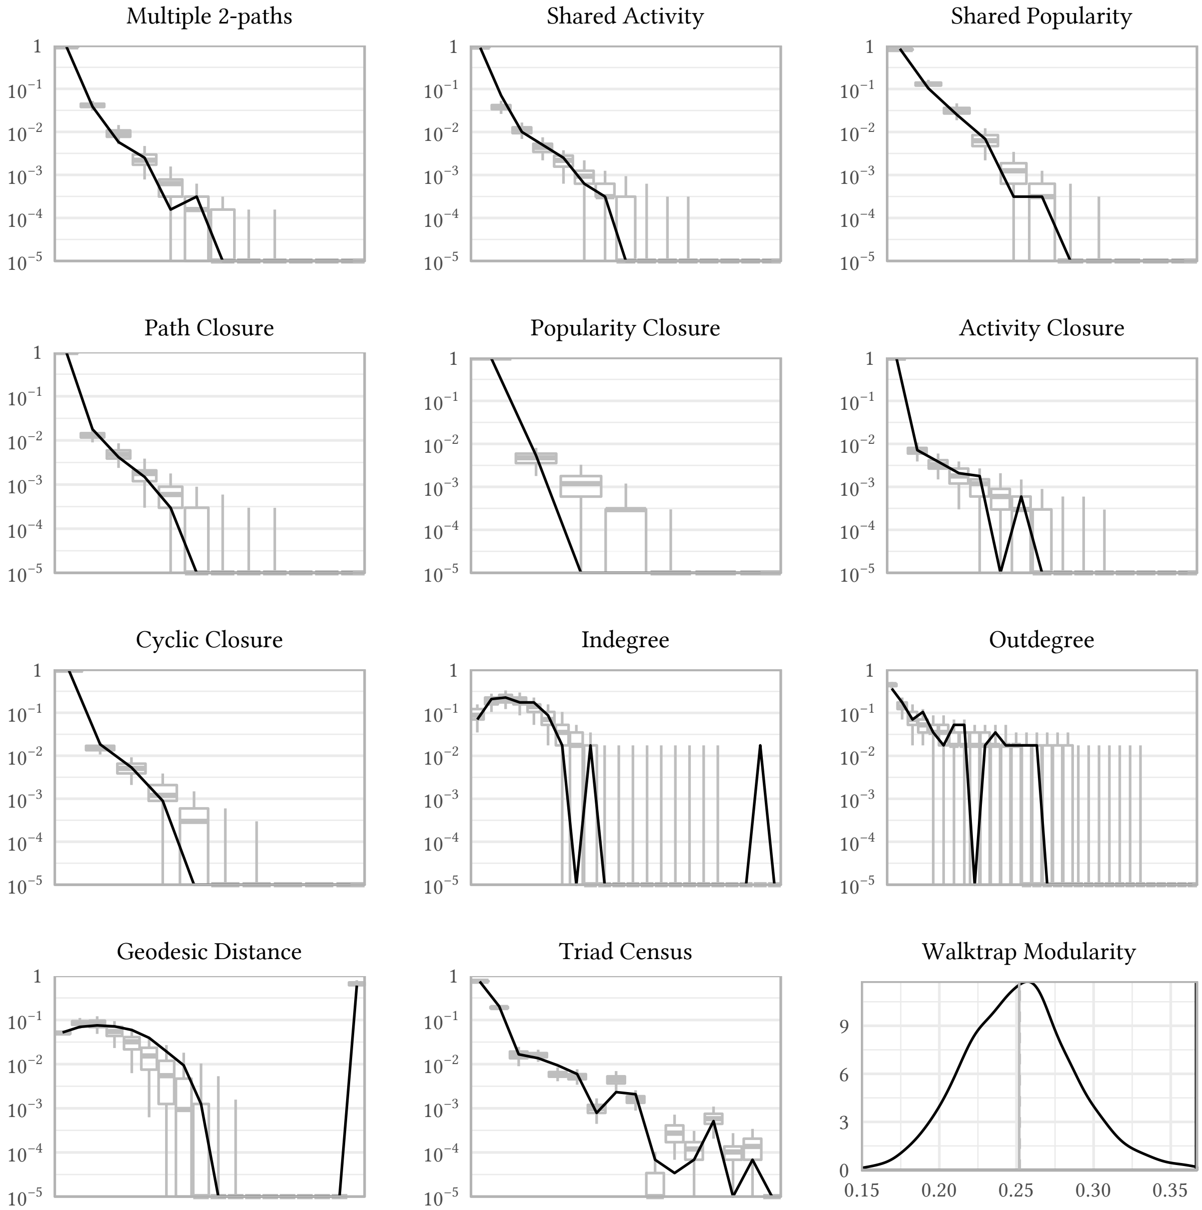

Figure 6: Goodness-of-fit for Chapter 5

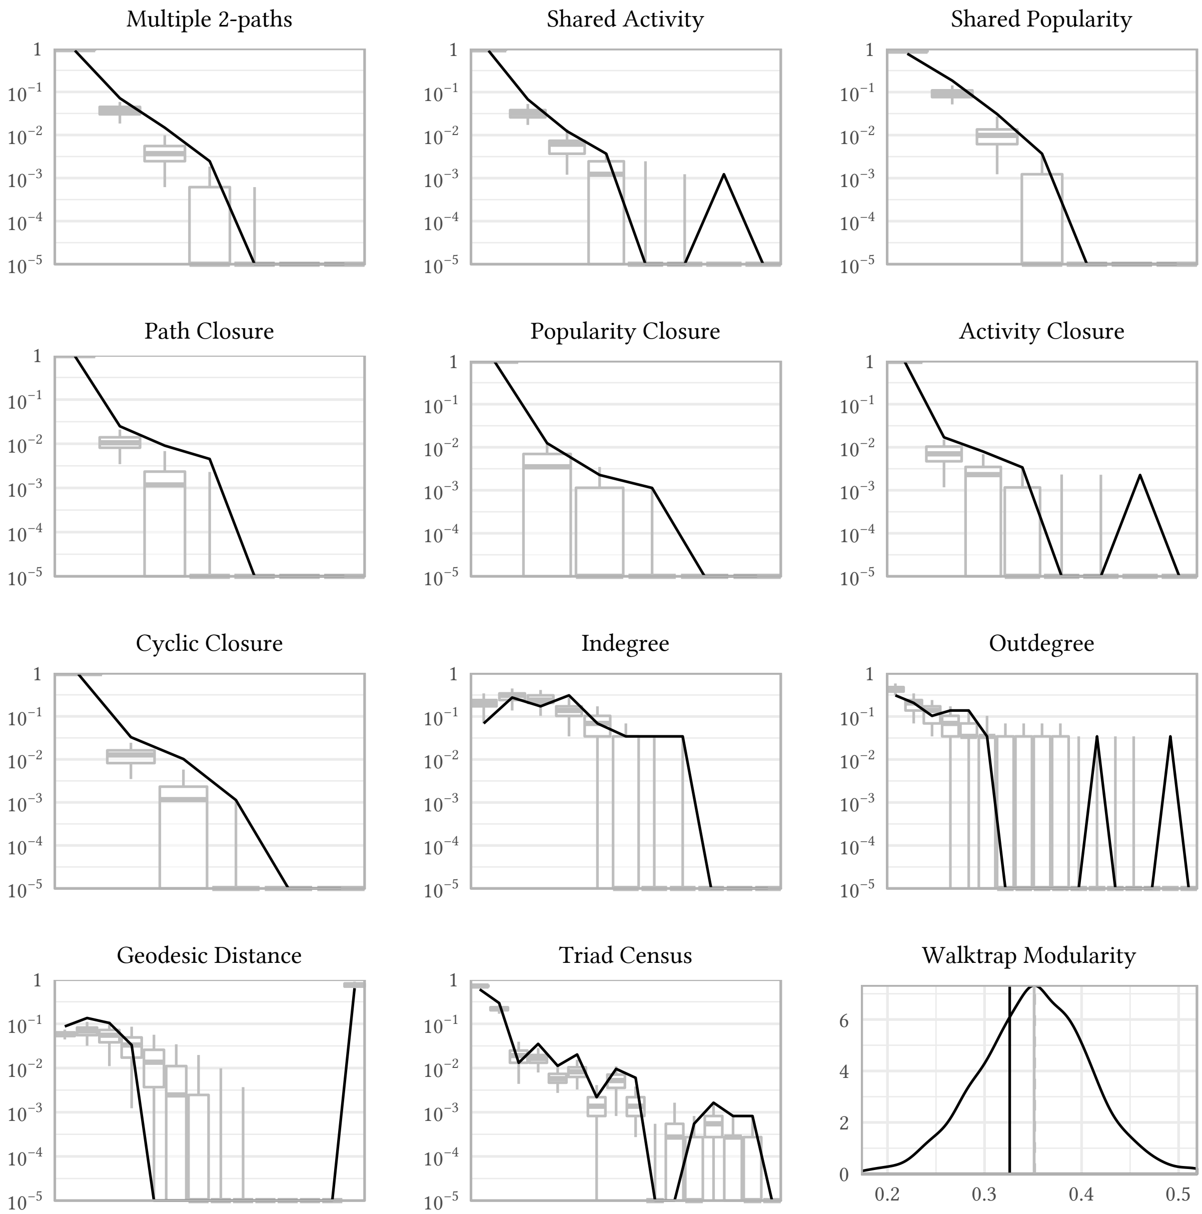

Figure 7: Goodness-of-fit for Chapter 6

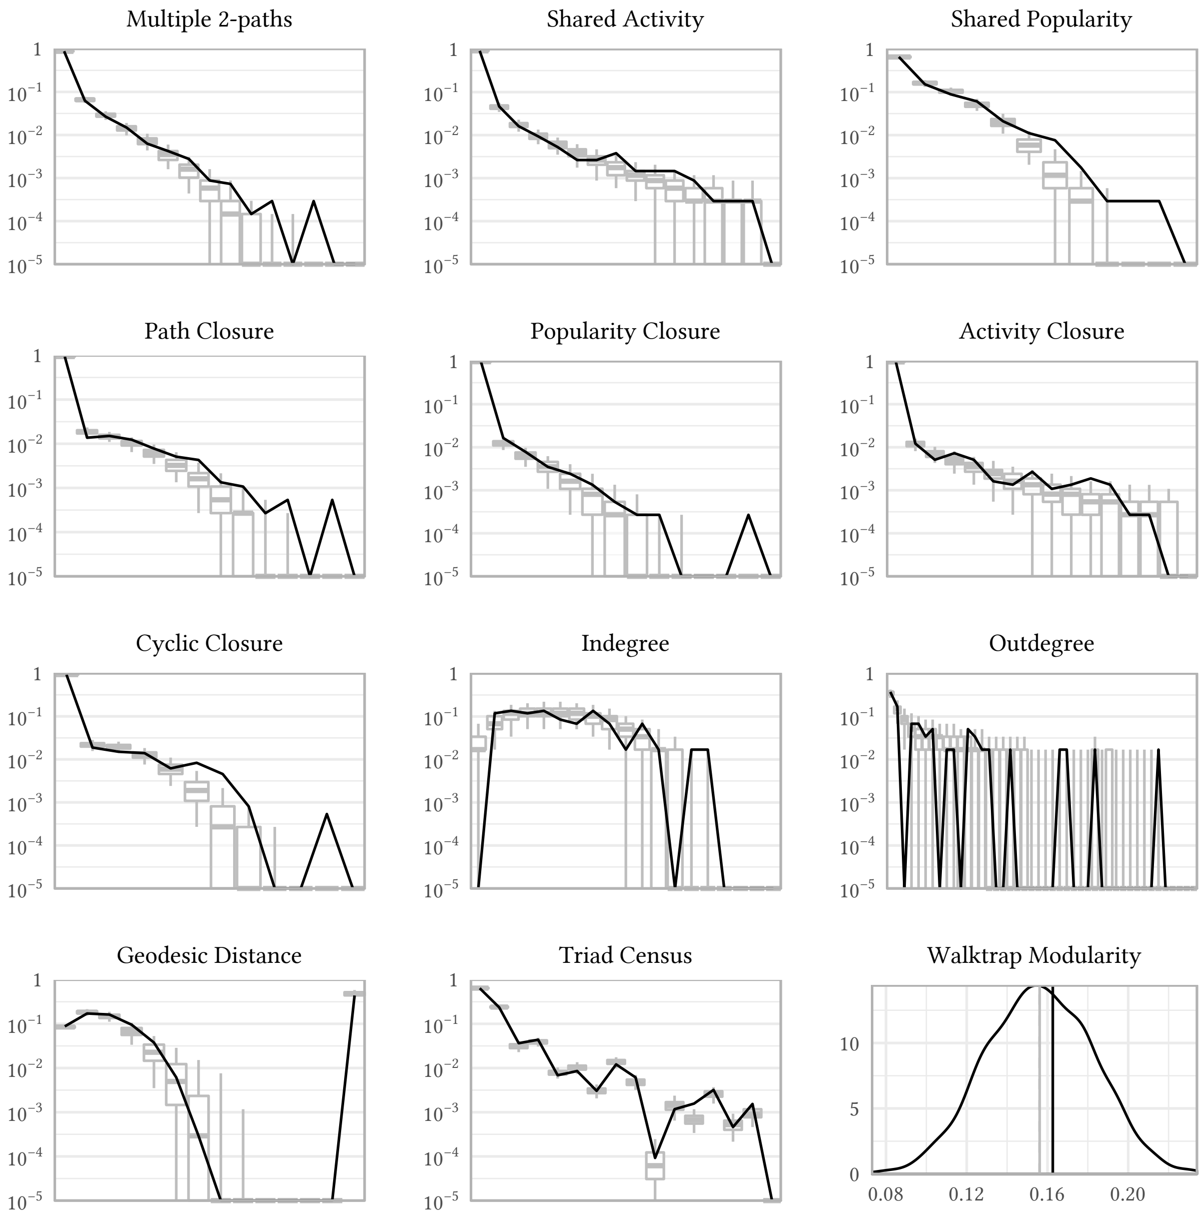

Figure 8: Goodness-of-fit for Chapter 7

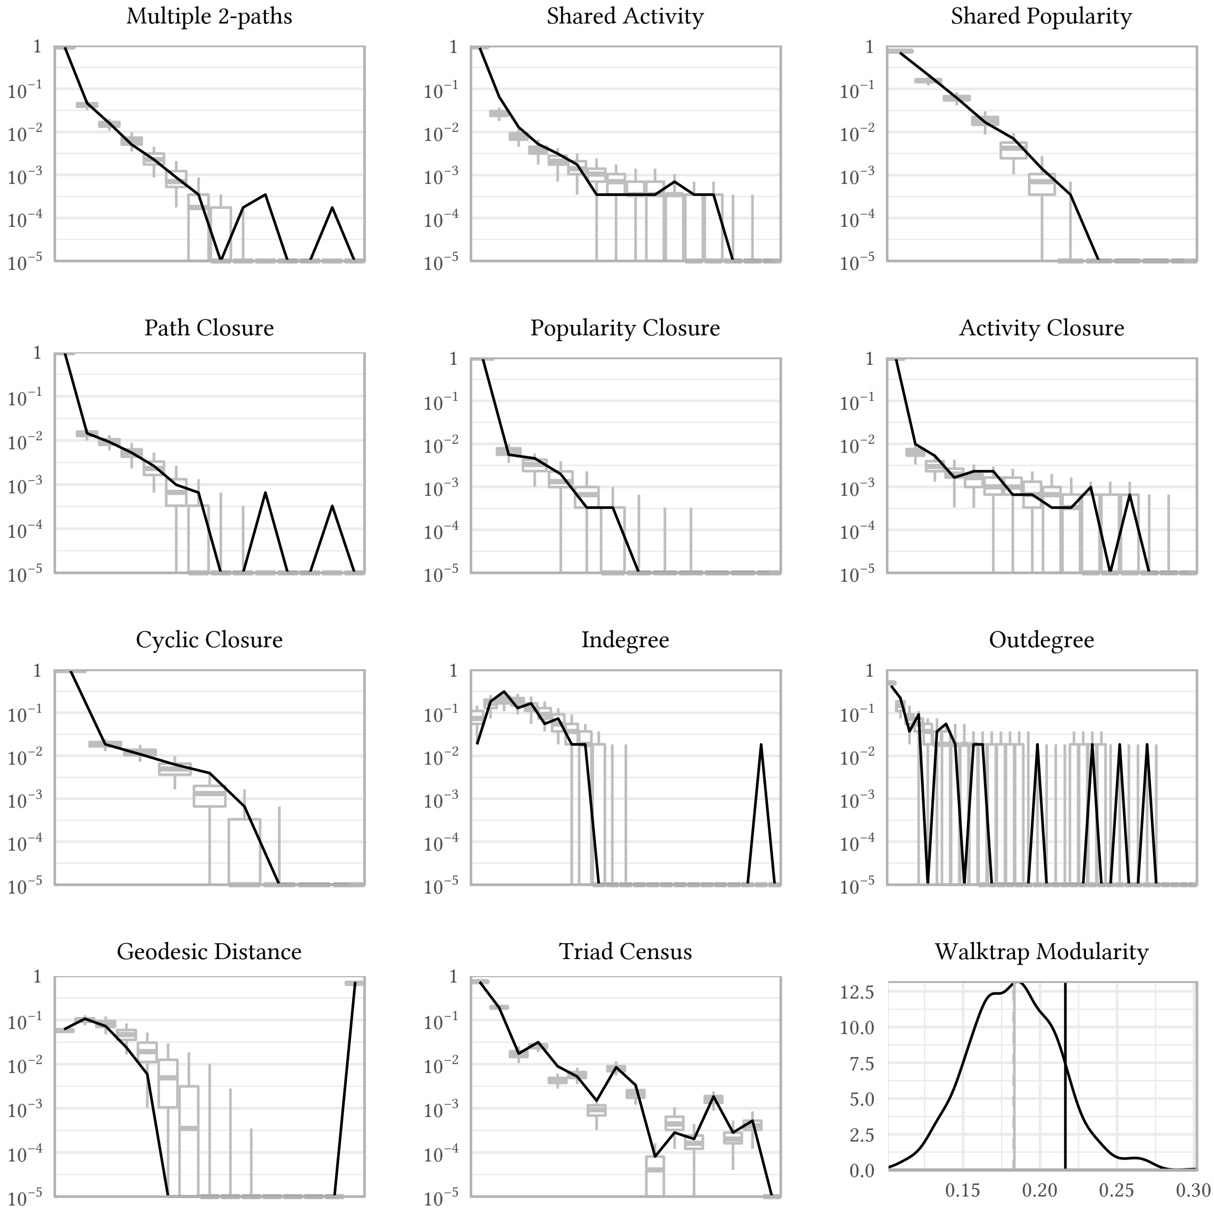

Figure 9: Goodness-of-fit for Chapter 8

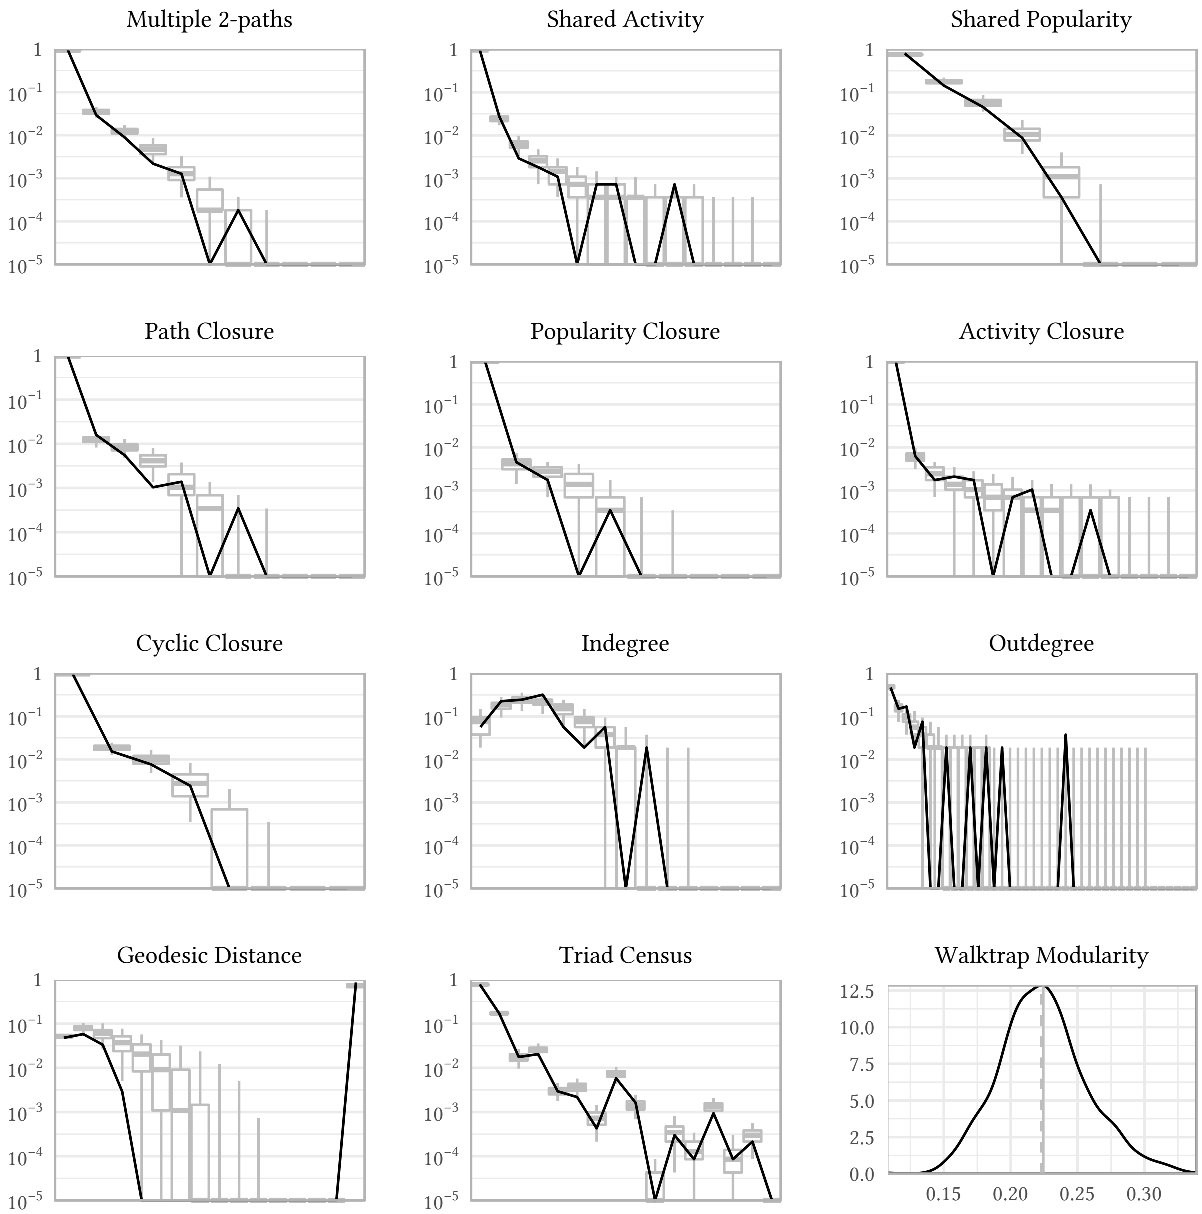

Figure 10: Goodness-of-fit for Chapter 9

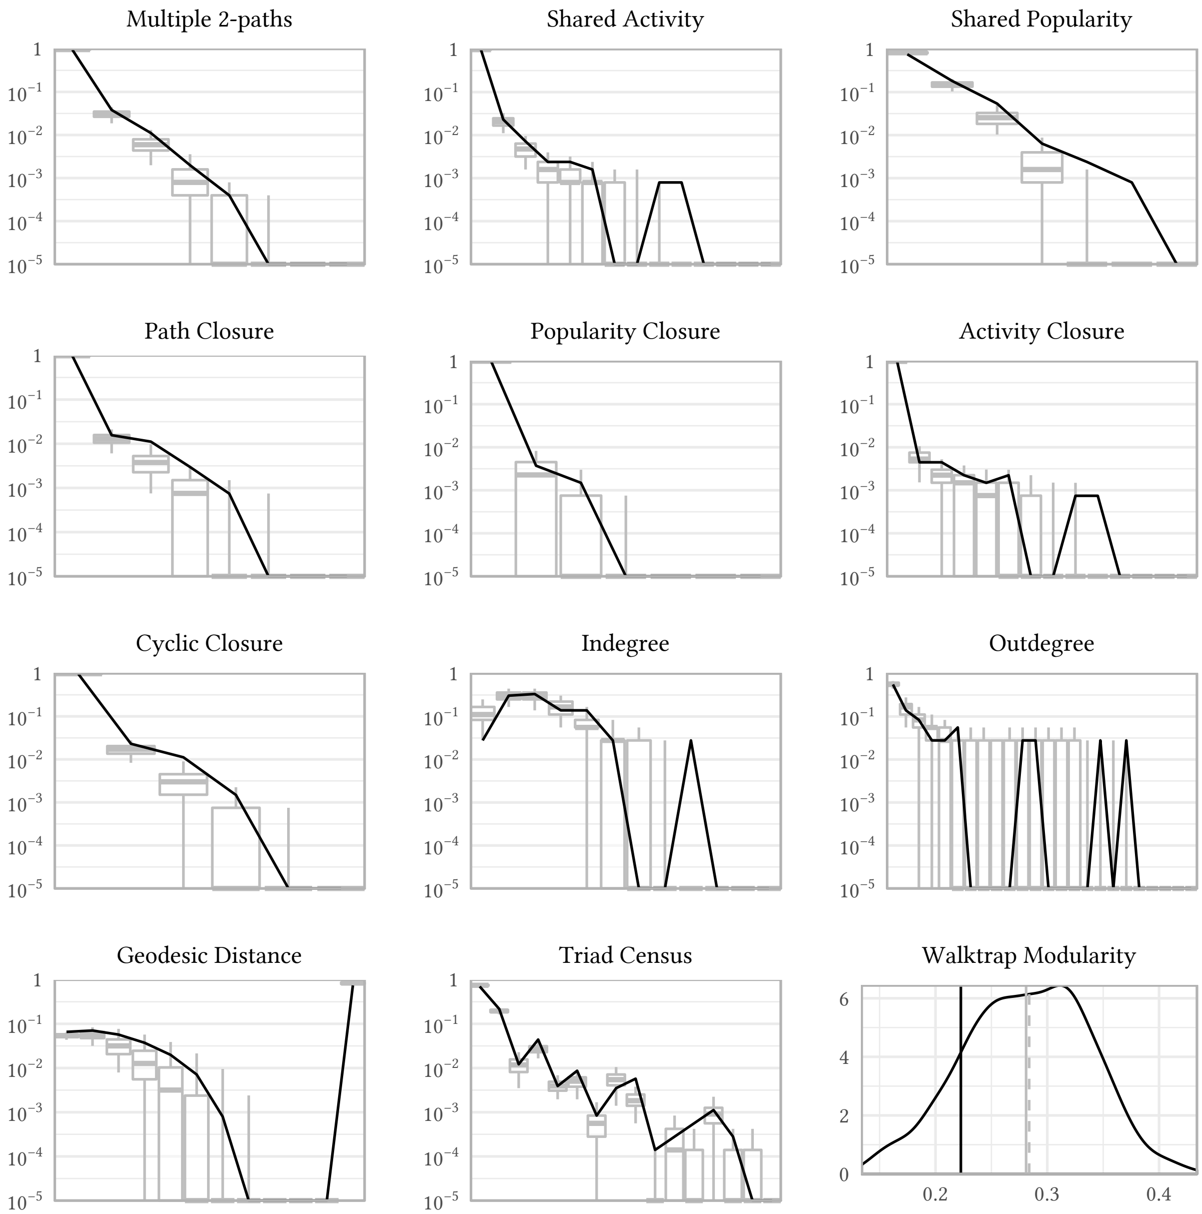

Figure 11: Goodness-of-fit for Chapter 10

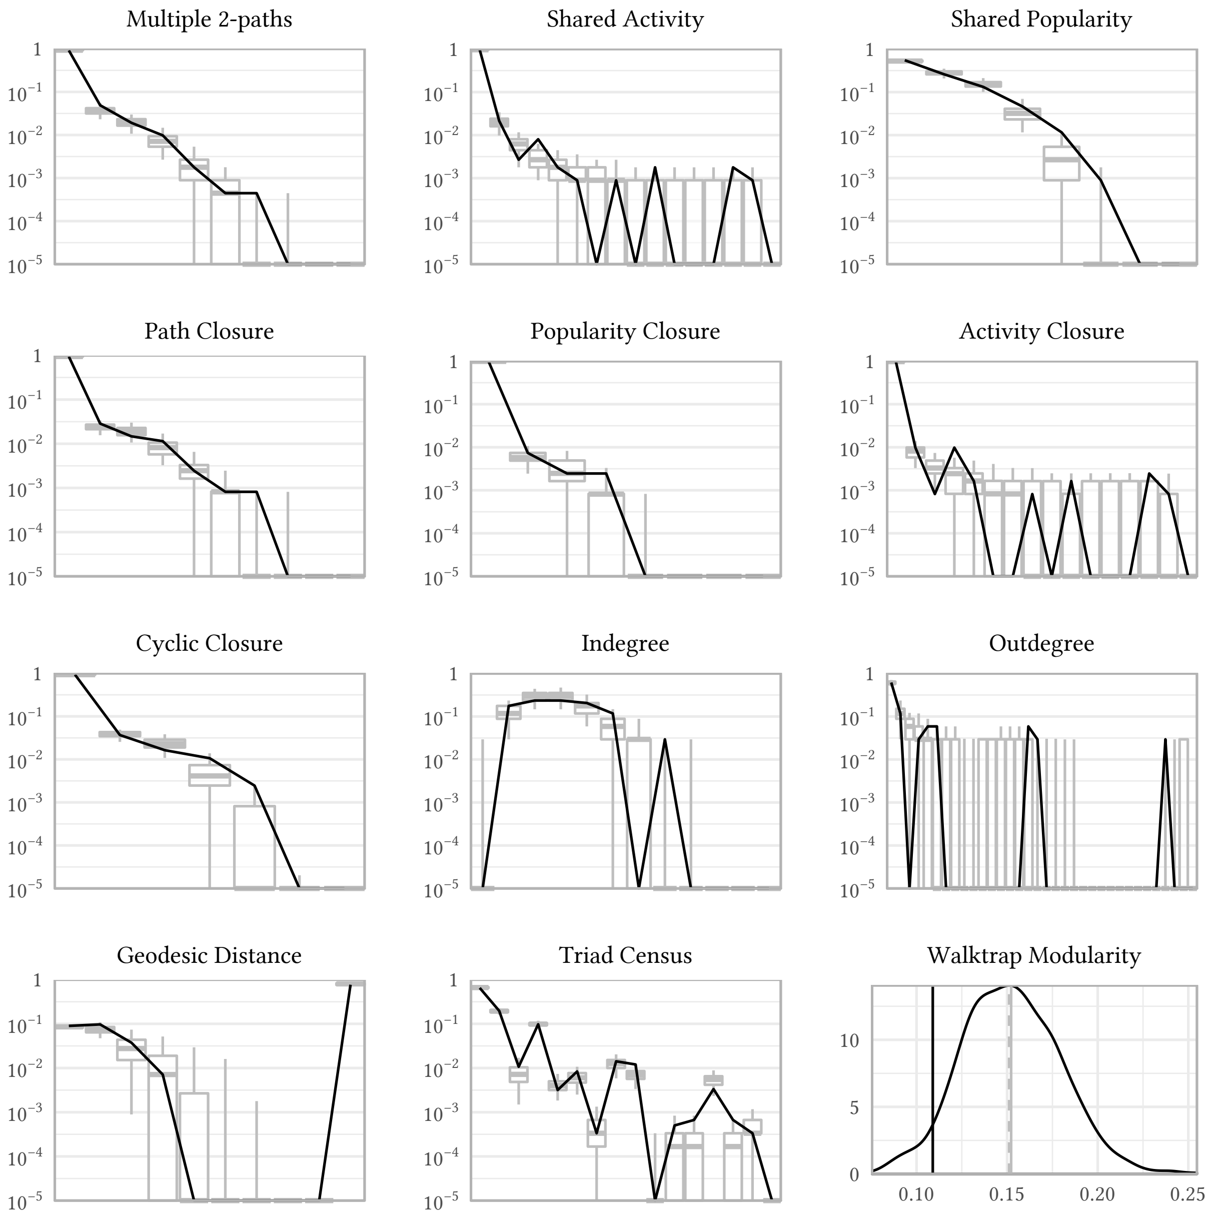

Figure 12: Goodness-of-fit for Chapter 11

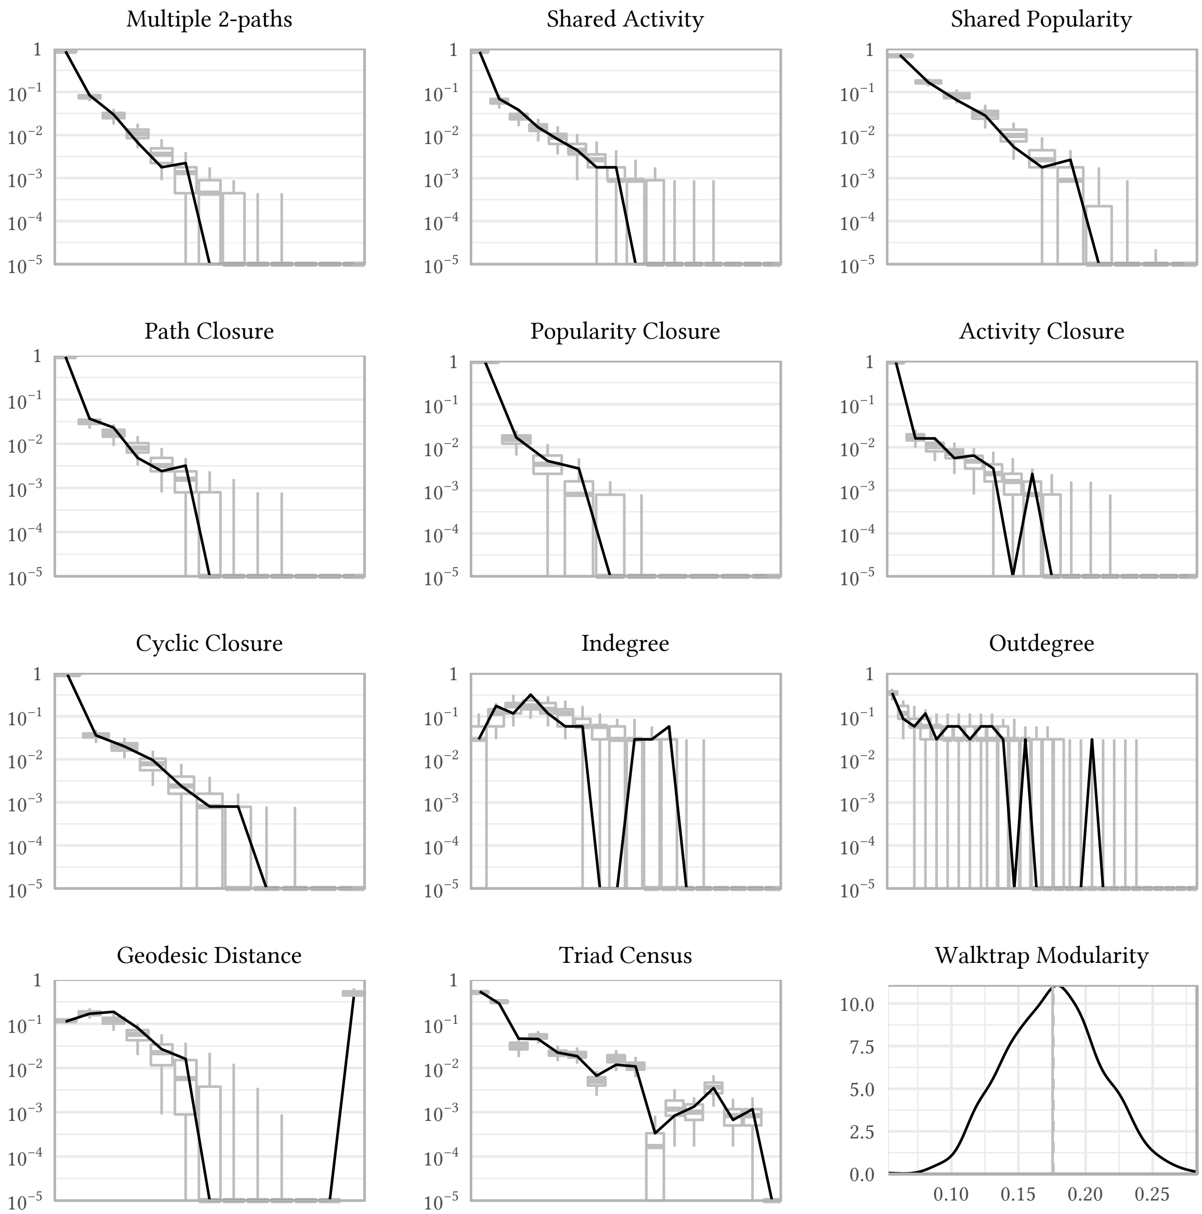

Figure 13: Goodness-of-fit for Chapter 12

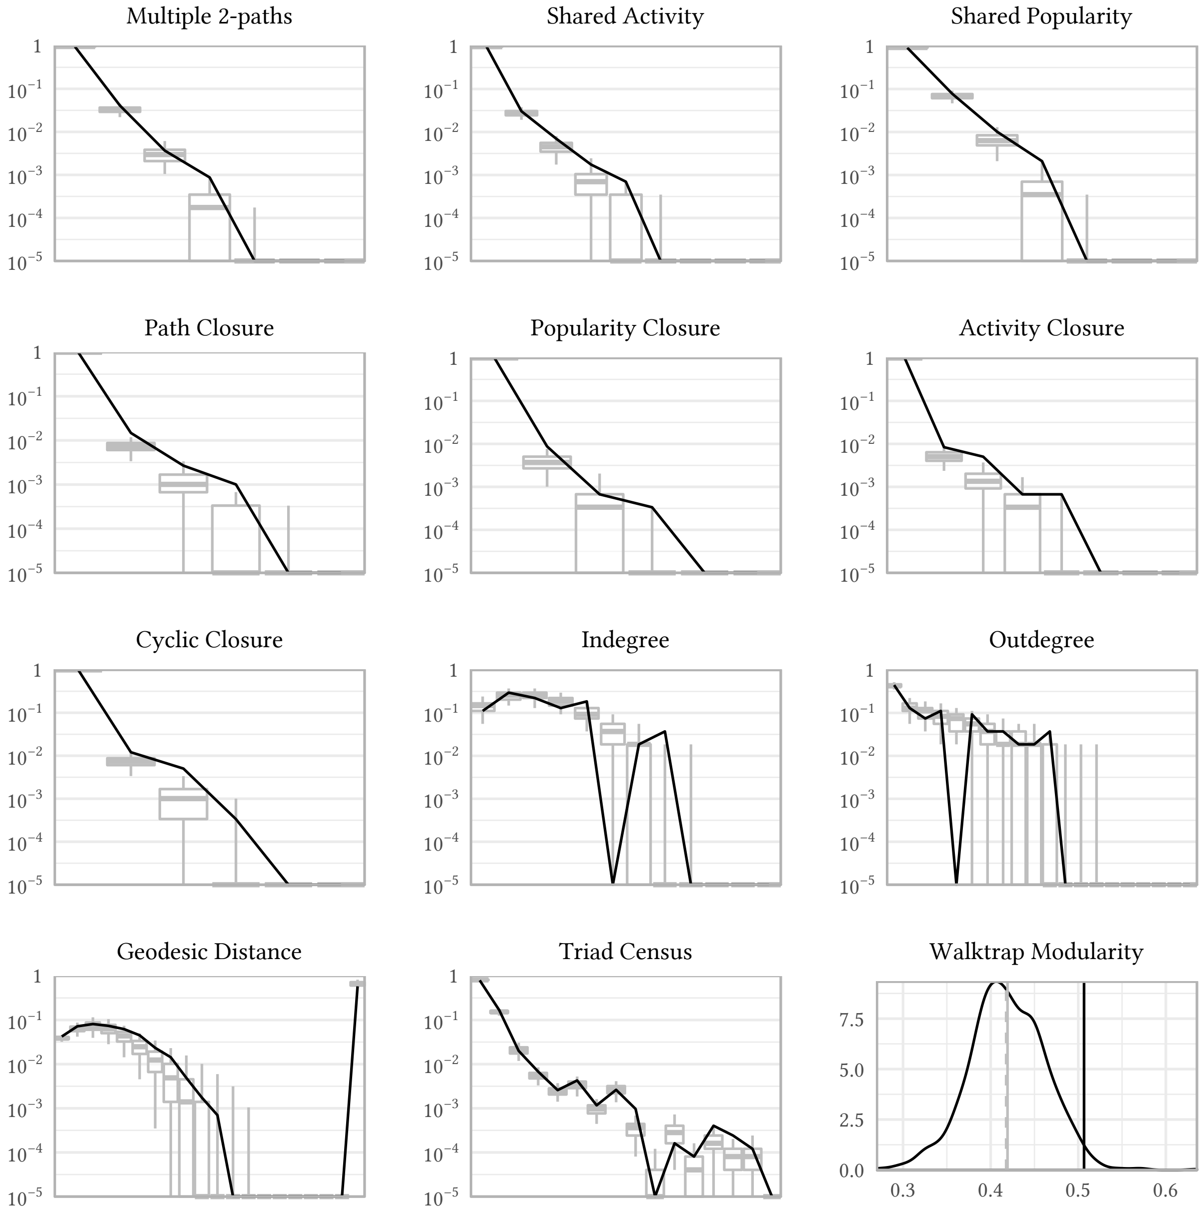

Figure 14: Goodness-of-fit for Chapter 13

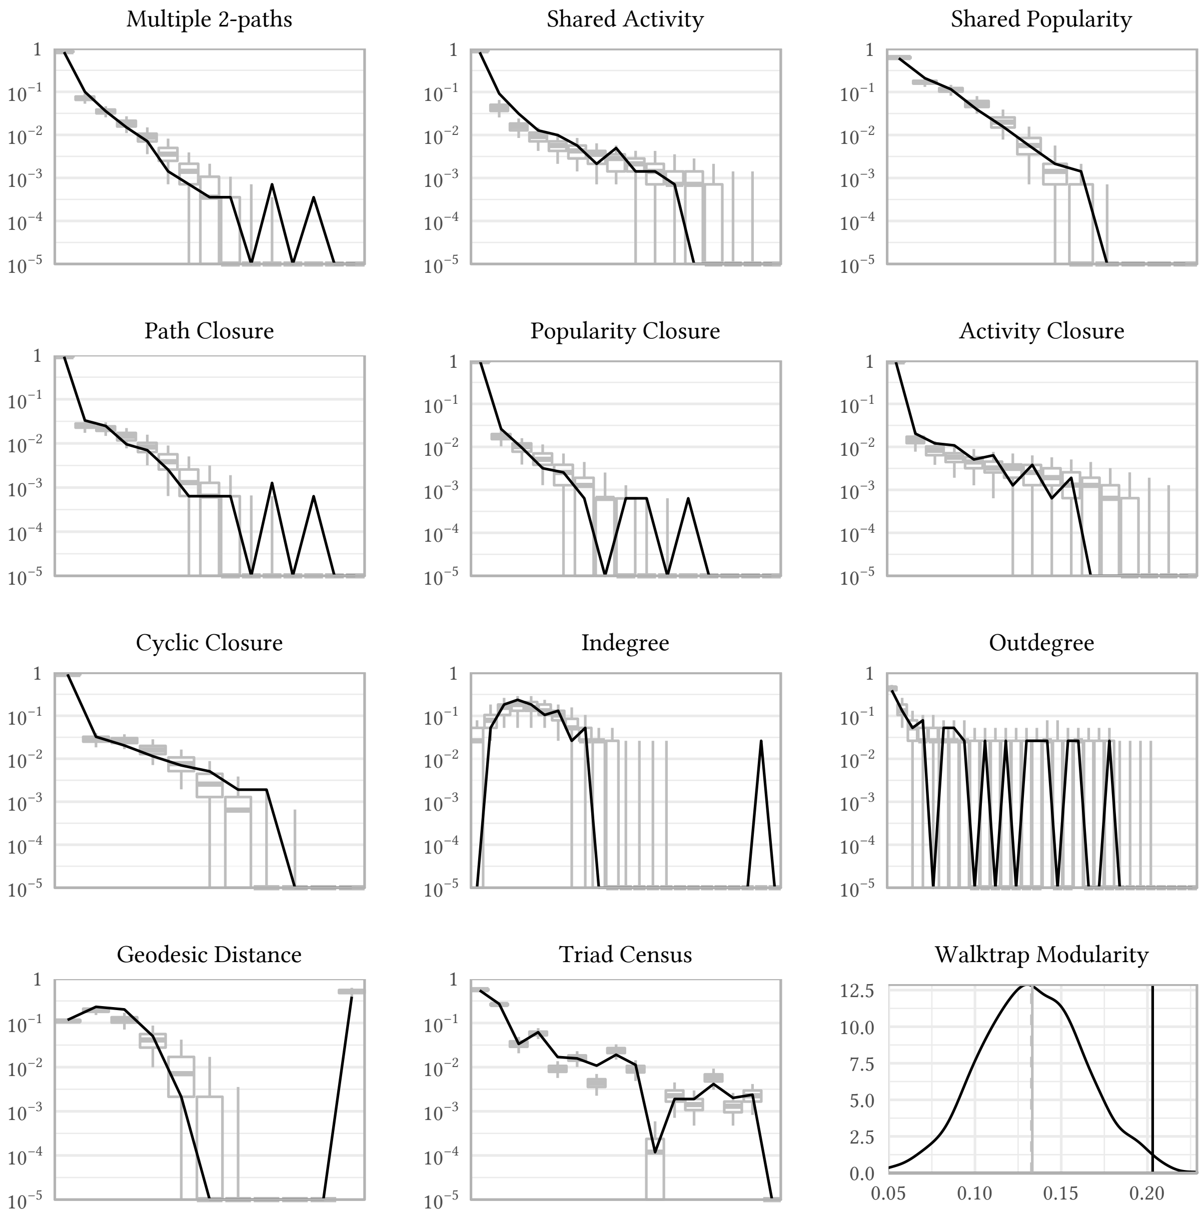

Figure 15: Goodness-of-fit for Chapter 14

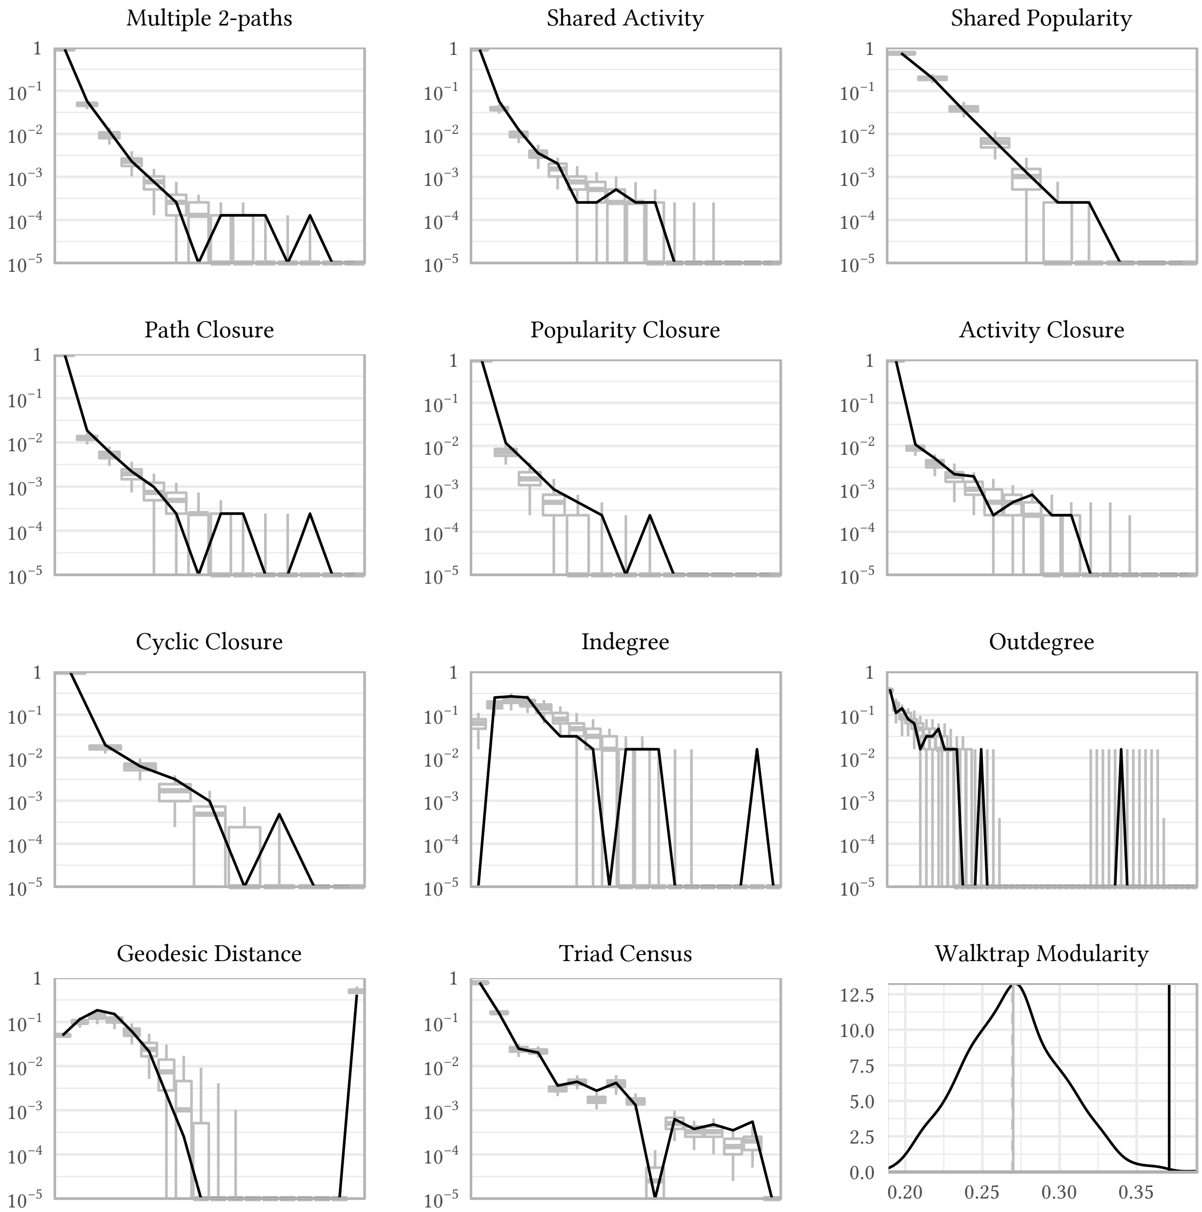

References

Honaker, James, Gary King and Matthew Blackwell. 2011. "Amelia II: A Program for Missing Data." *Journal of Statistical Software* 45:1–47.

Keith, Bruce E., David B. Magleby, Candice J. Nelson, Elizabeth Orr, Mark C. Westlye and Raymond E. Wolfinger. 1986. "The Partisan Affinities of Independent 'Leaners'." *British Journal of Political Science* 16(2):155–85.

King, Gary, James Honaker, Anne Joseph and Kenneth Scheve. 2001. "Analyzing Incomplete Political Science Data: An Alternative Algorithm for Multiple Imputation." *American Political Science Review* 95:49–69.

Kossinets, Gueorgi. 2006. "Effects of Missing Data in Social Networks." *Social Networks* 28:247–68.

Leifeld, Philip, Skyler J. Cranmer and Bruce A. Desmarais. 2018. *Temporal Exponential Random Graph Models with btergm: Estimation and Bootstrap Confidence Intervals.* R package version 1.9.1.
URL: *https://CRAN.R-project.org/package=btergm*

Petrocik, John R. 2009. "Measuring Party Support: Leaners Are Not Independents." *Electoral Studies* 28(4):562–572.