Rationality and Society 8(1): 111-120 (1996) © Sage Publications

The Forum

The Forum consists of comments on previously published papers and responses by authors. Unlike the Comments section in many other academic journals, the Forum includes solicited as well as unsolicited contributions. We encourage debates by actively seeking points of view contrary to those expressed in articles published in the journal. The Forum is intended to promote an open and critical debate that contributes to the intellectual vitality and further development of rational-choice-based theory and research.

Forum contributions should follow the form used by Rationality and Society *and should not exceed five double-spaced pages.*

COMMENT ON BURT AND KNEZ'S THIRD-PARTY EFFECTS ON TRUST

David Krackhardt
Carnegie-Mellon University

When the editors of this volume asked me to comment on the Burt–Knez piece on trust, I was both pleased and apprehensive about the opportunity. What can one say about a piece that is such an excellent example of dogged and thorough analysis? I enjoyed the paper immensely. In fact, I wish I had written it.

Having said that, I have two comments on the approach used by Burt and Knez. One comment deals with the dependent variable, trust, and the second concerns the independent variable, structure.

Trust is not an easily defined term. Like friendship, trust means different things to different people. Is it reasonable to extend from the results that Burt and Knez present to other realms of trust or to a generalized notion of trust? It is certainly beyond the scope of the current study to be able to answer that question. But, it is one in need of answering.

One of the dominant themes in the paper is that distrust is more affected by various structural factors than trust. But, this conclusion is based on empirical results that do not operationalize trust and distrust as opposite ends of the same dimension. Trust is measured by asking the

respondent to name people they would approach for counsel in looking for a new job. They are allowed to seek the advice of someone inside or outside the organization; and given the inherently sensitive nature of the quest, most people choose (it is assumed) only those whom they would trust with the information that the respondent was considering outside employment. While this may be a conservative measure (trusted people are likely to be left out of the nomination list), it has reasonable face validity.

The measure of distrust is more problematic. The question 'who has made it the most difficult for you to carry out your job responsibilities?' is certainly not an indication of overwhelming support. But is it *distrust*? Burt and Knez justify this measure by citing the various reasons people gave for listing people in this category, many of which do relate to trust-like issues. Some of these are about distrust in that the individual cannot be counted on to carry their own weight ('do not ''follow through on their commitments'' '); some are about selfishness ('egotistical, self-oriented, liar'); some are about abuse of power ('pursue own agenda'; 'had great power and withheld help'). Burt and Knez summarize these comments as suggesting that the respondent was naming individuals who are 'uncooperative'. Perhaps an even better attribution is that the respondents were identifying people they don't like.

Incorporating these stricter interpretations into the results of the paper leads us to a conclusion somewhat different from that of Burt and Knez. For example, one might reasonably state that the presence of exclusive third parties enhances the manager's tendency to dislike someone, and that this effect is stronger than the effect of such third parties on the manager's tendency to trust others. The question remains, then, if we were to operationalize trust and distrust as two ends of the same construct, would distrust still appear more sensitive to these structural factors? I am not sure, but I would like to find out.

Let me turn attention now to the independent variable, structure. The measure of direct ties ($p_{ij} = z_{ij}/\Sigma_q z_{iq}$) as well as the measure of indirect ties ($\Sigma_q p_{iq} p_{qj}$) is standard in Burt's analysis. These constructs have been usefully applied at many levels of analysis: individuals, organizations and industries. That is one of the strengths of the more general theory of Structural Holes (see, for example, the special book review section on Burt's work in *Administrative Science Quarterly*, June 1995).

But little attention has been paid to the z_{ij}, the raw data upon which all these analyses depend. When operating at the individual level of analysis, Burt frequently uses a method that he developed for the

General Social Survey in 1985 (see Burt 1984, for a complete description). In such a method, individuals are sampled from a larger population and asked to recall with whom they engage in certain relations. These relations vary from formal authority relations to informal socializing relations. This 'name generator' reveals a set of N others, in which ego (the respondent) is embedded. To find the structure, ego's account of the strength of these relations with ego and the strength of relations among the set of N actors is taken at face value.

The clear advantage of this method over more traditional network data-collection techniques is that one can draw a reasonably sized sample from a large population. This is a serious advantage. In traditional methods, which require the researcher to identify a population and then gather data on the entire set of relations among all pairs of actors in the population, studies are limited in scope to relatively small populations, often comprised of fewer than 50 actors. The GSS approach allows data to be gathered on a much larger group of people (in the Burt–Knez case, several hundred), and these people can be sampled from an even larger population (in their case, more than 3000).

But this advantage comes with a cost. It is assumed that people can and do report accurately on the ties that exist between themselves and others, and that they can and do report accurately on the ties that exist between the others that they have named. We know that people are often not particularly accurate at recalling event-specific relations they have with others (Bernard et al. 1984), although the correlations between such recollections and actual behavior is not zero (Burt and Bittner 1981). On the other hand, it has been shown that people are remarkably good at recalling recurring relations they engage in, such as frequently having lunch with someone (Freeman 1987) or co-attending a seminar series with someone (Freeman, Romney and Freeman 1987). But we know precious little about how accurately ego can assess whether the others are connected to each other. This latter point is critical, since it is through such assessments that the 'third-party' effects in the Burt–Knez paper are identified.

To illustrate the scope of this problem, I draw on a data set that allows us to assess these relations from many points of view. In a study of managers in a small manufacturing firm (Krackhardt 1987), I asked the 21 managers to tell me who they typically go to for help and advice about work-related matters. In addition, I asked them to whom they thought everyone else went for help and advice. Formally, these data comprise a three-dimensional cognitive social structure, $\mathbf{A}_{ijk}$, where i is

the actor who potentially seeks advice, j is the actor from whom the advice is potentially sought, and k is the perceiver of this $i \rightarrow j$ relation. $A_{ijk} = 1$ if k thinks that i goes to j for advice; $A_{ijk} = 0$ otherwise. In a sense, these data are quite similar to the Burt–Knez network data, except that I was able to get the perceptions of the entire population of managers (since there were only 21 of them).

What these data allow us to do is examine the question of whether people agree on the existence of relations among some defined set of peers. For example, i may claim that she goes to j for advice, but j does not confirm this claim ($A_{iji} = 1$ and $A_{ijj} = 0$). We do not know who is correct in this claim, but a reasonable argument could be made that if both i and j agree that i goes to j for advice then it is more likely to be an accurate statement than if they do not agree on this fact. Using this logic, we can construct a conservative matrix of agreed-upon advice relations, A*, where $A^*_{ij} = 1$ if $A_{iji} = 1$ and $A_{ijj} = 1$; otherwise $A^*_{ij} = 0$. Such a matrix is called a 'locally aggregated matrix', or LAS (Krackhardt 1987). In other contexts, such an aggregation has been used to estimate the 'actual' underlying structure, as distinct from what the actors think it is (Krackhardt 1990).

By looking at the subset of nodes that person k claims to be tied to, I can approximate what the Burt–Knez method of assessing the data would look like. That is, for each respondent k, I can take the nodes that k claims to be tied to (either k goes to them for advice or they go to k for advice). I also have their perceptions of what ties k thinks exist among those alters, just as Burt and Knez have done to assess the indirect ties. The interesting question becomes, then, to what extent are k's perceptions matched by the conservative LAS matrix?

There are two types of inconsistencies that we can identify in answering this question. First, there are the ties that exist in the LAS matrix that do not exist in k's map (errors of omission); second, there are ties that k claims exist among alters that do not appear in the LAS matrix (errors of commission). Taking the LAS structure as a better estimate of the 'actual' network, I find that on the average 30 percent of the ties that exist in the LAS network were incorrectly omitted by the respondent k. And, on the average, respondents incorrectly 'created' a tie 26 percent of the time when, according to the LAS structure, none existed. Given that the density in the LAS structure was about one-third, this translates to about a .42 average correlation between the perceived and the 'actual' (LAS) structure.

Such numbers are not encouraging, if our required assumption is that the perceived networks as recounted by Burt and Knez are reasonable

representations of the actual network structure one is embedded in. But more critical than actual hit rates is the question, how does this error rate affect the Burt–Knez measure of direct and indirect tie strength, which are central to the paper's results?

To answer this question, I calculated p_{ij}, the direct measure of tie strength from ego (i) to alter (i) based on both the ego's (k's) perception and on the LAS structure.[1] There were 283 i–j pairs, where i was the ego (perceiver k) and j was one of the set of connected alters. The correlation between the p_{ij} scores based on the perceived and the LAS matrices was .60. I conducted a similar analysis on indirect ties, calculated as $\Sigma_q p_{iq} p_{qj}$. The correlations between the measures based on the perceived vs LAS matrices was .55. One might easily construe these correlations, .60 for direct ties and .55 for indirect ties, to be indicators of construct validity.

Should we interpret these numbers as reducing our confidence in the results presented by Burt and Knez? That is, especially given the importance of the effect of 11 exclusive third parties (Table 2 and Figures 4 and 5), which relies on the respondent's ability to accurately record ties among alters, how much faith can we proclaim in their conclusions?

I have two thoughts on this. First, if the data-collection methods used by Burt and Knez were truly faulty, such that people's accounts were only marginally better than random measures, then it would be difficult to obtain their substantial results. Construct validity scores of .55 and .6, while modest, may be sufficiently better than random to capture enough information to allow their results to emerge. If this is the case, just think how much stronger their results would be if they had access to better indicators of the 'actual' structure.

A second perspective, though, suggests a different interpretation. Perhaps what is critical is what people perceive the network to be around them, not what it actually is. Burt himself (1982, Ch. 5) recognized the importance of perceived networks in predicting social comparison outcomes. It has been suggested elsewhere that such perceptions can occasionally be more crucial than real networks (Kilduff and Krackhardt 1994; Krackhardt 1987). This realization can lead us to radically different interpretations of what would otherwise be regarded as straightforward results—for instance:

> From this perspective, the BKS findings [that people's perceptions of their interactions do not match their actual interaction behavior] constitute evidence that one should not bother collecting behavioral data, since they do such a poor job of capturing the cognitions that live in people's heads. (Krackhardt 1987,110)

Following this line of reasoning, one might expect perceived networks to do a better job of predicting trust and distrust than 'actual' networks would. One problem with this interpretation is that the theoretical explanation behind the Burt–Knez result is based on 'actual' networks, not on subjective perceptions of them. Burt and Knez envision their perceptual data as reflecting 'true' networks, and their theoretical predictions are firmly based on this assumption.

Therein lies the dilemma. Do the results based on perceived ties predict trust better than those that would be drawn from a behavioral sample? This is an empirical question. If the answer is yes, then their theory needs to be modified to explain why. If the answer is no, then I would argue that the Burt–Knez results need to be grounded in better representations of the 'actual' network to assess its true power.

NOTE

1. I used Burt's standard formula for calculating $p_{ij} = (z_{ij} + z_{ji}) / (\Sigma_q z_{iq} + z_{qi})$ as given in *Structural Holes*, p. 51. There is no mention in note 5 of the Burt–Knez paper about whether the data were 'symmetrized' as part of this calculation, but since Burt ordinarily does use both z_{ij} and z_{ji}, I kept with that tradition.

REFERENCES

Bernard, H. Russell, Peter Killworth, David Kronenfeld and Lee Sailer. 1984. 'The Problem of Informant Accuracy: The Validity of Retrospective Data.' *Annual Review of Anthropology* 13: 495–517.

Burt, Ronald S. 1982. *Toward a Structural Theory of Action*. New York: Academic Press.

Burt, Ronald S. 1984. 'Network Items and the General Social Survey.' *Social Networks* 6: 293–339.

Burt, Ronald S. and William M. Bittner. 1981. 'A Note on Inferences Regarding Network Subgroups.' *Social Networks* 3: 71–88.

Burt, Ronald S. and Marc Knez. 1995. 'Kinds of Third-party Effects on Trust.' *Rationality and Society* 7: 255–92.

Freeman, Sue. 1987. Verbal Response Data as Predictors of Social Structure. PhD Dissertation, University of California, Irvine.

Freeman, Linton C., A. Kimball Romney and Sue Freeman. 1987. 'Cognitive Structure and Informant Accuracy.' *American Anthropologist* 89: 310–25.

Kilduff, Martin and David Krackhardt. 1994. 'Bringing the Individual Back In: A Structural Analysis of the Internal Market for Reputation in Organizations.' *Academy of Management Journal* 37: 87–107.

Krackhardt, David. 1987. 'Cognitive Social Structures.' *Social Networks* 9: 109–34.

Krackhardt, David. 1990. 'Assessing the Political Landscape: Structure, Cognition, and Power in Organizations.' *Administrative Science Quarterly* 35: 342–69.

A FURTHER NOTE ON THE NETWORK STRUCTURE OF TRUST

REPLY TO KRACKHARDT

Ronald S. Burt and Marc Knez
University of Chicago

To extract maximum value from David Krackhardt's comments, consider them in broad perspective. Beginning with prisoner's dilemma (PD) deliberations as a baseline for analyzing trust, there are two ways to bring context into the analysis; one cultural, the other structural.

The cultural approach involves studying context in terms of player beliefs. Players engage the game according to the beliefs of the culture in which they live, or the culture they believe relevant to the game. Kinds of cultures produce kinds of players. This is the path taken by Aaron Wildavsky (1992) in his provocative treatment of PD games as context dependent. Drawing on his work with Michael Thompson and Richard Ellis distinguishing kinds of cultures (Thompson, Ellis and Wildavsky 1990), Wildavsky (1992) distinguishes kinds of rationality within a PD game. The theoretical and empirical gains from the cultural analysis can be elusive because it is difficult to get a formal handle on culture, but Wildavsky's cultural analysis is undeniably a prism that reveals the spectrum of opinion that reads *Rationality and Society* (see the debate from critics in the special issue containing Wildavsky's analysis and subsequent Forum discussion).

We took the structural approach in Burt and Knez (1995), which involves going one step behind culture to study context in terms of the network structure through which players internalize culture. Context in the structural approach is a pattern of variably strong connections between PD players and third parties outside the game. Players enter the game variably exposed to gossip about one another communicated through third parties. The structural approach differs from the cultural in that predictions are based on contrasts between forms of interpersonal connection rather than contrasts between modes of personal belief. The touted virtues of the structural approach are that context is more readily measured for empirical research, and more productively captured by parameters in formal theory. Both virtues are relevant to Krackhardt's (1995) comments on our analysis.

We were after an easily observed, structural primitive that could be used in formal models of trust and empirical studies of organizations. Our method was generic—go into the field, observe dyadic situations that sometimes lead to trust and other times lead to distrust, then measure the effects of how players are connected indirectly through third parties witnessing and gossiping about the game.

We present argument and evidence from senior managers to support three conclusions: (a) far and away the most important trust factor is the past relation between two people as in the usual PD game; (2) dense indirect ties through third parties are a significant secondary trust factor—they amplify the dyadic condition. Where two people are inclined to trust one another, dense indirect ties through third parties amplify that inclination to certainty. Where two people are inclined to distrust one another, dense indirect ties through third parties amplify that suspicion into certainty; (3) the form of the indirect ties through third parties predisposes density amplification to the positive or negative. The positive effect of network density is associated with mutual third parties. The negative effect of network density is associated with exclusive third parties.

The steps forward from here are three: (1) build third-party effects into rational choice models of trust to strengthen their descriptive power; (2) study third-party effects in the field for corroborating evidence and a better understanding of how gossip animates third-party effects; and (3) run experiments to test our field results and see how third-party effects change previous PD results obtained without third-party effects held constant.

Krackhardt moves in a fourth direction. In the best spirit of referee critics, he asks for a pause to reconsider the evidence and alternative images of third-party effects. His first comment is that our sociometric indicator of distrust might indicate dislike. No doubt. People rarely like the people they distrust. At the other extreme, people often have positive sentiments for the people they trust.

Now what? It is difficult to see how abstracting trust and distrust into positive and negative sentiment improves our grasp of context dependence: our empirical predictions are unaffected (network density still has its positive and negative amplifying effects via network balance theory), but we lose the theoretical connection with rational actors deliberating PD moves, and we descend into a semantic debate over the true meaning of the sociometric indicators. Our preference is (a) to work within the simple theoretical framework of PD deliberations (rather than generalizing to feelings of positive and negative sentiment), and (b) to

focus on the construct validity of our indicators (do they have the correlations with other variables that they ought to have if they are measuring trust and distrust?) backed up with respondent quotes explaining why they cited the people we discuss as distrusted. The quotes describing distrust relations in our paper are about cooperation denied, which makes the sociometric indicator eminently suitable for a PD analogy.

Second, Krackhardt raises the possibility of discrepancy between cognitive and behavioral social structure. What a manager reports to us about relations among the third parties around him or her could be different in the reports of another manager. Quite possibly.

Now what? How does this improve our grasp of context dependence? It doesn't, but why it doesn't is a useful consideration. Here are three ways to respond to discrepancy between cognitive and behavioral social structure in our analysis:

(a) 'Exactly.' Manager perceptions of the social structure around them is what affects their trust of other people. Therefore, their perceptions are what we want to measure to test network effects on trust, and our estimates of network effects are correct as reported.

(b) 'Thank you.' Manager reports on social structure are distorted by their personal biases such that our measures of third-party network structure are stochastic regressors. Therefore, our estimates of network effects are conservative. Estimates corrected for the errors in our predictor network variables would be stronger than the estimates reported in our paper.

(c) 'Oy.' Manager reports on social structure are distorted by their personal biases such that evidence of third-party effects is systematically distorted. The estimates reported in our paper are spurious or suppressed, and we have a problem.

But Krackhardt doesn't argue that manager cognitions deviate from behavioral data such that our evidence of third-party effects is systematically distorted. His comment is keyed to option (a), which means that our analysis is correct as it stands, and option (b), which means that our evidence of third-party effects is an understatement (cf. the two thoughts with which Krackhardt closes his comment). Professor Krackhardt is a leading figure among the several people working to understand systematic differences between behavioral and cognitive social structures. Option (c) remains the goal that will make this methodological work (like response effects in survey research more generally) theoretically significant. We look forward to future research in which an understanding of discrepancies between behavioral and cognitive social

structures is used to more accurately model context effects on individual behavior and opinion—such as third-party effects on trust.

REFERENCES

Burt, Ronald S. and Marc Knez. 1995 'Kinds of Third-party Effects on Trust.' *Rationality and Society* 7: 255–92.

Krackhardt, David. 1996. 'Comment on Burt and Knez's Third-party Effects on Trust.' *Rationality and Society* 8: 113–18.

Thompson, Michael, Richard Ellis and Aaron Wildavsky. 1990. *Cultural Theory*. Boulder, CO: Westview.

Wildavsky, Aaron. 1992. 'Indispensible Framework or just Another Ideology? Prisoner's dilemma as an antihierarchical game.' *Rationality and Society* 4: 8–23.