

Networks and History

PETER BEARMAN,¹ JAMES MOODY,² AND ROBERT FARIS³

¹*Institute for Social and Economic Research and Policy, Columbia University, New York, NY 10027;
e-mail: psb17@columbia.edu.*

²*Department of Sociology, The Ohio State University, Columbus, OH*

³*Department of Sociology, University of North Carolina at Chapel Hill*

Received September 6, 2002; accepted September 6, 2002

Events and event structures compose the constituent elements of history. In order to construct historical accounts of event sequences, historians have to make cases. This article proposes a method for casing historical events. We illustrate the analytic strategy by considering a complex population of interrelated events that make up a narrative of revolution, counter revolution, and revolution in a small village in China. Implications for the methodology of historical social science are discussed. © 2003 Wiley Periodicals, Inc.

Key Words: social networks; social structure; events; event-sequences; casing; historiography; method

In contrast to overly deterministic arguments about fundamental causes, imaginative narratives of the chance catenation of contingent events as revelatory of historical process are the current fashion in historical social science. In the service of such narratives, network imagery is often deployed to describe the seemingly fragile contingent pathways through which complex historical outcomes occur. At first glimpse, networks do appear to provide an appropriate metaphor for chance and contingency, but this is not the case. Instead, consideration of network structures in historical context suggest limited roles for contingency in event dynamics. Consequently, the historical event structures that appear as cases in social science history are much more robust than typically imagined. Nevertheless, some

This is a revised version of work previously published in Social Science History.

events can play more important roles than others in shaping history, and the problem of historical explanation rests on developing a methodology for modeling complex event structures that reveals which events play critical roles in historical outcomes. Such a methodology is the concern of this article in which we propose that application of network models to historical cases can provide answers to such fundamental questions as: when, if ever, do single events change history? What do things mean in historical context, and how do we define cases in historical context.

The argument we propose is simple. The meaning of an event is conditional on its position in a sequence of interrelated events, what historians call a case. Consequently, for those interested in what events mean, casing event sequences is the most fundamental problem that confront historians. In Section 2, we propose a solution, which exploits developments in social network analysis that are relevant for the analysis of complex event structures. We focus

on the similarities between social structures and event structures that point to the applicability of network methods for the analysis of historical data. These similarities also suggest that historical processes may be quite robust to perturbation.

Casing, that is bounding the beginning and end of event sequences, is not dissimilar from a problem in structural analysis: how to specify a boundary on a network. The problem for historical social science involves generating a population of events. Strategies for generating a population of events in historical contexts are briefly described in Section 3. In Section 4, we illustrate the method with respect to a single complex case; revolution, counter-revolution, and revolution in a Chinese village during the period from 1920 to 1950. We exploit modeling techniques for narrative networks [1], to transform narratives into networks. Operations on these networks provide the foundation for our analyses, in which we “test” our casing solution by simulating the future. Finally, we consider whether events can be meaningfully arrayed with respect to their probability of shaping history and describe how such an array could contribute to a new historical method. In the conclusion we indicate how the approach proposed here could alter our thinking about the nature of chance in shaping outcomes, the record of cases that historians consider, and history more generally.

1. THE PROBLEM OF CASING

Casing is necessarily implicated in the simple task of constructing an historical narrative. Likewise, casing is a prerequisite for meaning. Precisely because it is so important, casing is seen as a matter of insight and the judgment that arises from such insight. For most historians, casing is an essentially artistic achievement. Here we propose a strategy for casing historical events that relies less on art and more on method. Not surprisingly, this is not a simple problem. One major complication arises from the future. Because the meaning of an event is conditional on its position in a sequence of inter-related events, it is necessarily impossible to fix forever the meaning of an event, that is, fix forever the end and beginning of a sequence of events, because future events can activate, that is, draw into a new event sequence, past events. We cannot find comfort in the idea that only a certain class of future events could have such a role, because the future occurrence could be as momentous as the storming of the Bastille or as trivial as the discovery of a lost diary. In the latter case, an element of the historians’ craft, discovery (of new events or relations between events) have the capacity to change beginnings and ends, and therefore the specific meaning of events.

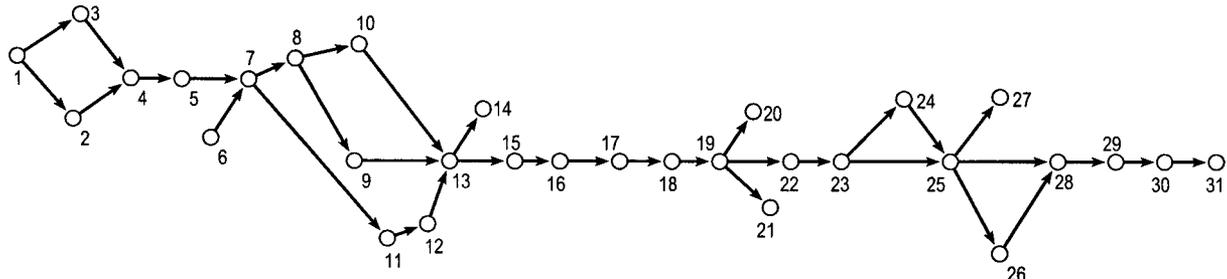
The fact that it is *possible* for the meaning of events to change does not mean that historians should abandon the attempt to develop a strategy for casing event sequences. First, although some events may become activated by discovery or the future, most are never so fortunate. Whatever

meaning most events have is likely fixed completely within a single, specific event sequence, itself fixed within larger, more complex event-sequences. Put another way, neither the discovery of new events nor unknown future occurrences are likely to alter in any way the sequence of events that “dead” events are embedded in, and consequently their meaning is also fixed. Nevertheless, some events have already, and some more may, become embedded in new event sequences following discovery or the occurrence of events in their future. Consequently, we can imagine a distribution of events, defined with respect to their probability of activation, “fluidity of meaning,” or susceptibility to being conditioned by the future. If we can array events with respect to their probability of being conditioned by the future, it follows that event sequences are also characterized by such a distribution, and likewise, congeries of densely interrelated event sequences (what we define as “cases”) are also subject to the same distribution, with some more likely to change than others. This makes intuitive sense and is confirmed by the judgment that historians use. In simple terms, some events, event sequences, and cases are dead. Some events and event sequences are subject to radical revision. We can confidently talk about the meaning of dead events. Our confidence falls with those likely to be alive. The practical problem involves knowing which events, event sequences, and cases are hot potatoes and which are not.

Strong Theory and Thin History

The stronger the theory the thinner the history, a truism that is revealed most clearly when one sets out to represent history as a network of events connected by flows of causation. Historical accounts of events, especially those proffered by social science historians tend to have an uniform appearance. They start with a relatively dense cluster of inter-related events. These, typically macro-level events (e.g., fiscal crisis, agrarian crisis, crisis in confidence/legitimacy) flow into a narrow stream of specific micro-level events. A thin pathway (sparsely connected with very little redundancy, few cycles, etc) moves through time, ultimately inducing a pivotal event that is characterized by high out-degree, impacting multiple event sequences and providing (typically) the boundary of the “case.” Figure 1 is a network graph of a standard historical narrative, in this case, the story of revolution and counter revolution in a Chinese Village.

In Figure 1, nodes are specific events that took place, and edges are links between events (causal or logical) implicit or explicit in the narrative. Time moves, in general, from left to right. The pivotal event(s) are those in the center of the graph, bounded by the beginning and end of the narrative, which composes the case. The narrative is very thin as a network, just sparsely connected. This implies that the theory that gave rise to the specific story is strong, because theory involves denying data. Thin narrative accounts are

FIGURE 1

Traditional historical narratives.

the product of specific theories that direct the historian to identify some events as salient and to deny other events as not salient. History involves selection of events to interconnect into a narrative. To have a theory requires that we know the end of the story to direct the selection of events. But this is a problem. How are we to know the beginning and end if they alone tell us what the events mean?

Rather than focus on the selection of events, we now consider the implicit theory of history as characterized by thin lines without independent pathways connecting causes and events. An irony is that with strong theory we are soon driven to contemplation of butterfly effects as driving history. In the Figure 1 narrative, there are many critical points through which only one-path flows. Butterfly effects would be pronounced if a small perturbation has the consequence of deleting (or adding) a node or line between events. If the event or link were absent, could we really imagine that the revolution would not occur? The problem is not parsimony of explanation *per se*. Many parsimonious accounts traversing the same field from different end points can generate populations of dense event structures. The problem is too few sets of eyes. The core methodological trick is to integrate views from multiple perspectives.

2. SOCIAL NETWORKS AND HISTORICAL SOCIAL SCIENCE

Over the past decade, influential articles that rely on network analysis—on substantively important historical topics—from the organization of the Medici to Ottoman state building and beyond to the Paris Commune have been published [2–7]. Network imagery and methods provide insight into specific mechanisms and processes by focusing on the middle-range, above isolated individuals yet below whole social formations. These studies have provided new insight into the role that social relations play in structuring, and blocking, action, and more abstractly, they have provided a new language for describing the dense, often knotted and cyclical, interrelated levels of social relations, sym-

bolic constructions, and practices (seen as flows in a network) that compose tangible social structures in historical and contemporary settings.

These notable achievements have not come without costs. The detailed reconstruction of social structure, defined with respect to pattern across multiple relations, necessary for network analysis has often led to a heightened commitment to highly particular explanations, and a reluctance to abstract structure *per se* away from specific contexts. Consequently, much of the work in historical social science that uses networks looks prosopographical—an approach to relational data which is limited because it is unable to provide an analytic scaffolding for meaningful comparison across cases with respect to interpretable structural parameters. On the other hand, the emphasis on context has been a useful palliative to counter a more disturbing trend in social science history, the idea that rational choice models can serve an explanatory, as versus heuristic, function. It is ironic that a method (structural network analysis) designed for comparison across contexts celebrates particularity as the principal barrier to a theory that denies the salience of all contexts (despite protestation to the contrary). [Rational choice modelers would deny this by pointing to how their models embed context (as values, goods, costs, etc) into actors' decision frameworks. But the fact that all contexts are equally easy to embed into the model gives the ghost away.]

Equally ironic is the strange marriage between relational and contingency theorists. As with networks, contingency has been an important “discovery” for historical social scientists and currently serves as the principal challenge to older models in historical social science that focus on the macro-level determinants of social change without sufficient attention to (social, relational, symbolic, etc.) mechanisms. The principal metaphors are drawn from the fact that social network observations, like historical observations are tied and interdependent. In social networks and in history there is the sense that the fact of interdependence

means that subtle change can concatenate wildly through a system and cumulate into unanticipated historical and/or structural change [8]. The idea is attractive, but wrong. Tangible social structures *build on* and depend on local fluidity and disruption for stability [9,10]. (We can only observe social structures that are robust. Nonrobust social structures do not last long enough to observe. A popular idiom explains what makes structures robust. Love, like a tree, can weather storms better if it bends.)

Robust structures absorb fluidity at the micro-level by virtue of specific structural features that “exploit” interdependence. Network data on a population are locally dense, yet globally sparse, often cyclic, knotted, and characterized by a redundancy of ties. (There are many more similarities. One similarity, which we exploit subsequently, is that the characteristics of global social networks can be meaningfully ascertained by sampling local networks, an argument that is often implicit in historical narratives.) Social structures share these features with historical structures. Aside from radical revisionists, most historians would also agree that historical data exhibit tie redundancy, e.g., the idea that there are multiple independent pathways through which causal effects flow. Cycles in historical data appear when future events condition past events, drawing out of the past new relations to other events. In social networks, local density, knottiness, redundancy, and cyclicity give rise to the complex social structures that organize the relational world. Although analytically separable, they entail each other. Cyclicity gives rise to redundancy, redundancy gives rise to local density, and density gives rise to knots, generating macro-level cohesive properties from a host of independent micro-processes. Our interest here is to show that it is the same with event structures. We demonstrate that actual event structures arising from historical data have a similar structure, where order appears at the aggregate level, a product of micro-level fluidity. Consequently, representations of event structures as thin narratives, and consequently subject to “butterfly effects,” are largely mistaken.

3. GENERATING A POPULATION OF EVENTS FROM INTERCALATING NARRATIVES

In conventional historical accounts the end determines the beginning and hence the elements to be arrayed in the narrative. To case an event, which may be in multiple inter-related subsequences, we need a population of events around which we can draw a beginning and an end. Two distinct strategies for building a population of events are possible, short-path snowball sampling and intercalating narratives. The idea of short-path snowball sampling is to start with a large sample of events and use snowball sampling techniques to generate a population of events. A variety of sampling strategies for networks [see 11,12 for first-steps], can be deployed to build populations of historical events. Here, we illustrate the second strategy, intercalating

narratives, to demonstrate our method for casing. The data we use are life stories. Like historical accounts, life stories presume an end (a standpoint). Telling stories involves arraying elements selected from a rich and inexhaustible plate of cultural goods—people, places, things, events, ideas, and so on—into narrative sequences that are oriented toward a particular end, in such a way as to be a plot. The end allows the author to select from an endless sea of events just those events he or she sees as important (on the basis of a theory) for the story to be revealed. In contrast to formal histories, life-stories have features that make them ideal for our goal, the most important of which is a weak theoretical structure.

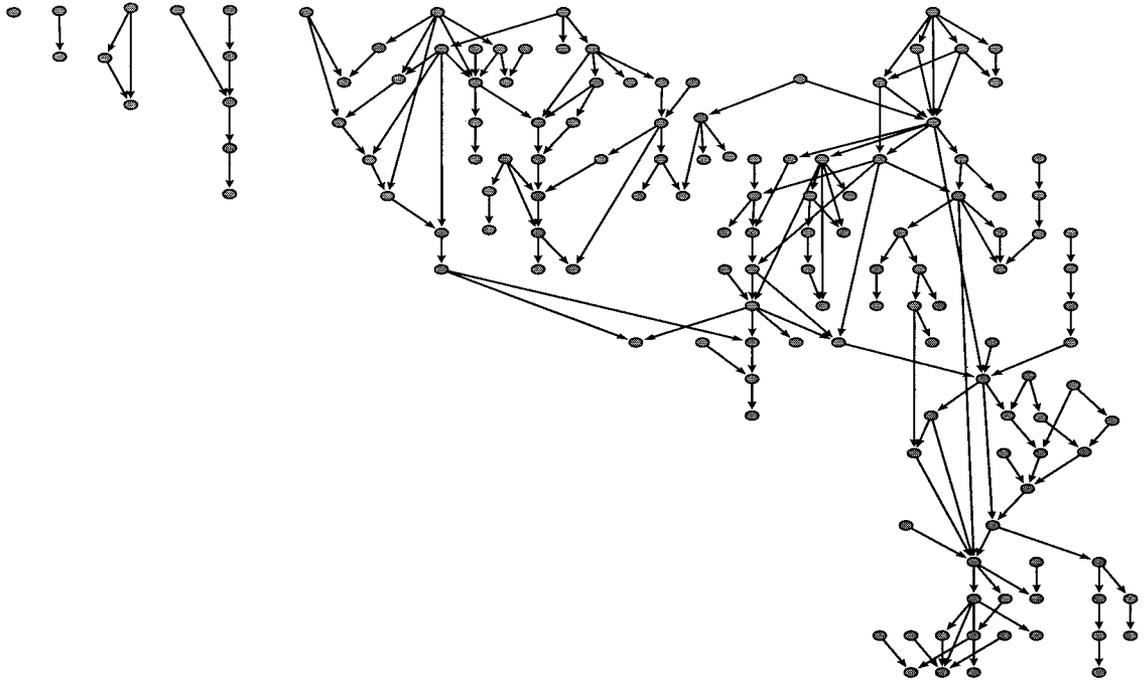
To illustrate we use 14 life stories from Chinese villagers whose experiences encompassed agrarian revolt in the countryside, counter-revolution, a revolution, and then the encoding of a revolutionary regime into an institutional framework. The context is a small village in Northern China. The stories are taken from *Report from a Chinese Village* [13]. The book contains a collection of life stories of the villagers of Liu Ling village, in Northern China near Yanan. Myrdal conducted interviews there in 1961. Figure 2 provides a graph representation of two of the life stories we use. By treating events as nodes and relations between events as arcs, narrative sequences of elements are transformed into networks. By representing complex event sequences as networks, we are able to observe and measure structural features of narratives that may otherwise be difficult to see.

In these graphs, elements of the narrative life-story are treated as nodes that are connected by narrative clauses, represented by arcs. A narrative clause is a clause that is temporally ordered in such a way as moving it involves changing the meaning of the subsequence in which it is embedded. Free clauses, by contrast, can be moved without changing the meaning of a subsequence or the narrative as a whole [1,14,15]. We code only narrative clauses as arcs, linking one event (or element) to another over time. The elements (nodes) of the narratives are heterogeneous in scope and range, ranging from greeting conquering troops with tea, to a staged battle between the KMT and the Communists. The former event tied the landowners’ sons to the KMT; the latter resulted in an imaginary defeat of the Communists. The idea behind this mirage was to trick the KMT leadership into thinking the Communists had been crushed by local KMT forces so that both forces could resist the Japanese.

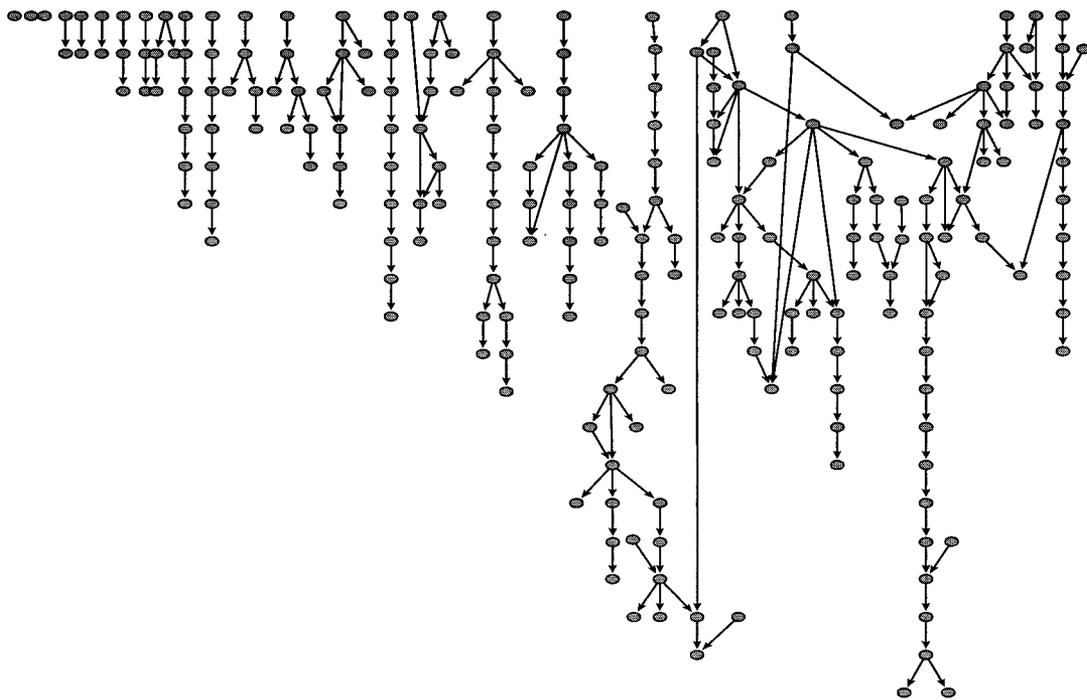
In Figure 2, narrative time moves from the top to the bottom of the page. The left-right axis is not substantively interpretable. Narrative depth is represented by the number of arcs connecting events. In this instance, for example, the two events at the bottom of Figure 2B have a narrative depth of 17, that is, there are 17 steps from the bottom to a starting event at the top of the graph. A characteristic of these stories is that they are structurally very different from the stories of professional historians. They have many disconnected ele-

FIGURE 2

Panel A: Li Hsiu-tang



Panel B: Pai Yu-teh



Narrative networks.

ments. Events are mentioned but are not necessarily tied. Across subsequences, it is impossible to walk from the early events to later events without a break. Not surprisingly, life-stories are denser and more complex than conventional historical narratives. They tend to have deep narrative flow. They are more complex because ordinary people are not trained as theorists. Therefore, they have trouble denying data. The life stories we work with exhibit heterogeneity. Some accounts are thin (Figure 2A), whereas others are convoluted (Figure 2B). Each of these stories has a different end point. The narrators are standing in different places. The end of the stories thus involves different outcomes.

The fact that they are standing in different places directs the selection of the elements that they choose to account for their end. By analogy, one might consider a set of professional accounts of the same sequence of events, each standing in a different position. All of the stories cover the same village and village events over the same time, and consequently, the field which they traverse, and the events which they refer to, overlap considerably. We exploit this overlap by intercalating stories to generate a population of inter-related events, which provides a new data structure, and consequently, points to new strategies for analysis. These new directions are taken up below in Section 4.

4. MAKING AND TESTING A CASE

Between 1920 and 1950, China was transformed. Reform, revolution, and warfare wracked the countryside. Our data arise from one of thousands of villages in Northern China. They are about events in this village and their connection to distant events occurring in other villages and cities and countries, the character and context of which was likely unimaginable to the villagers who lived in Liu Ling. Our problem, is to develop a method to case inter-related event sequences. In order to make a case, we first need a population of events and we need information about their relation. The second step is to draw a boundary on the nodes in the graph. The problem (and solution) is known as the boundary-specification problem [16]. Drawing on an old tradition in the social network literature, we can isolate cases by defining a partition on the population of events. Standard clustering techniques are, however, not appropriate for our problem, because arcs connecting dense regions of a graph (bridge-nodes) might well play an important role in the narrative sequence we are trying to capture. Instead, we adopt a new strategy, which is to identify all bicomponents on the population [17]. A component of a graph is a maximal connected subgraph. A maximal sub-graph is one which cannot be made larger and still retain the property that there is a path between all pairs of nodes in the sub-graph and that there is no path between a node in the component and a node not in the component. A bicomponent is a component that has the property that all nodes are connected by at least two different independent paths and

that the addition of a node requires that it is connected to two nodes in the subgraph. The central idea is that a case, seen as a set of interconnected events produced by multiple intercalated narratives must have the property of at least a bicomponent. A bicomponent is not necessarily a case. *It is a candidate for a case.* We define cases as bicomponents that are robust to discovery or future activation.

Figure 3 reports all of the events mentioned in the 14 histories of the Chinese villagers we work with, intercalated to form a single graph. Almost 2000 unique events are mentioned, each event is represented by a circle. Events that are in more than one narrative are shaded. Narrative time moves from the top to the bottom of the page. In some regions of the graph, where events and their relations are especially dense, arcs are invisible. Events that are tied to one another by arcs in these dense regions appear to overlap in the graph. Events to the left side of the figure are embedded in event sequences that are not tied to events on the right side of the figure.

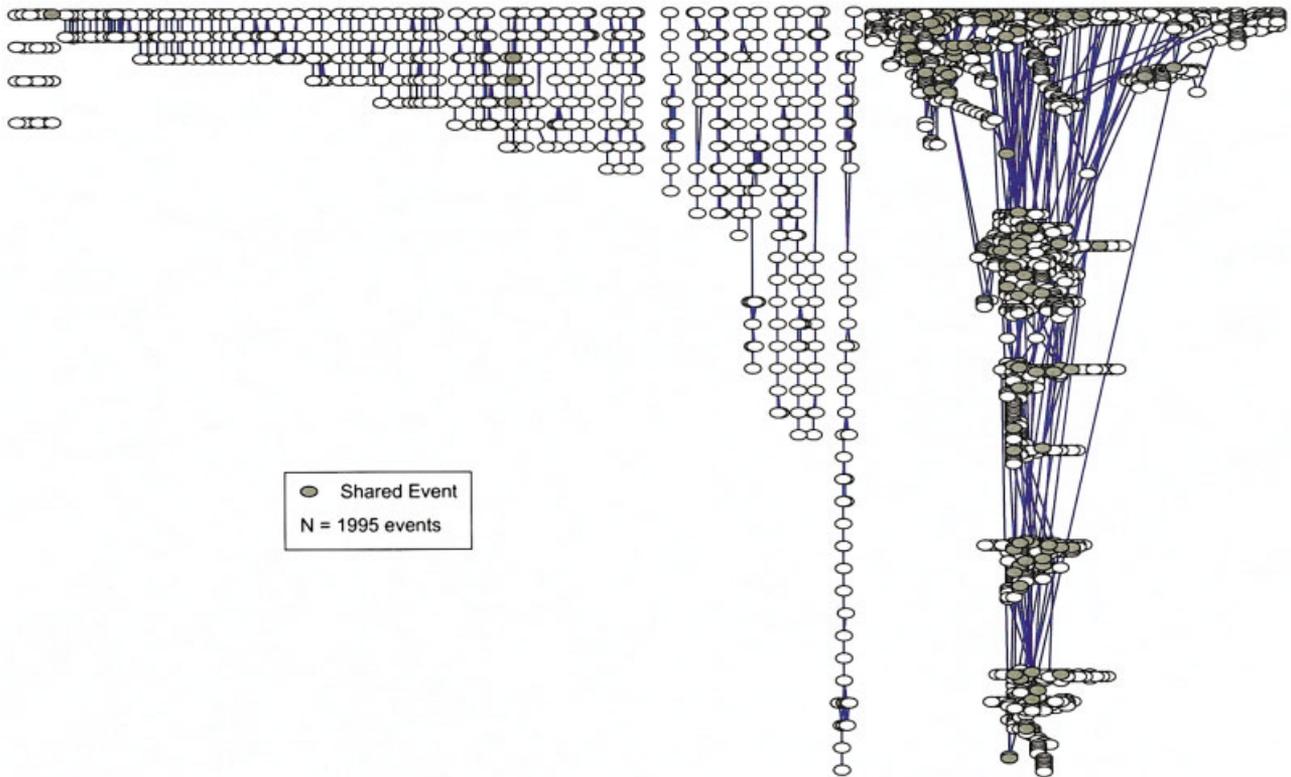
This is our population of events. Of course, there a millions of events not present. They might belong to some other history but not this history. But some of the events that are present look like they do not belong to this history either; e.g., no pathway connects them to other events. Happenings without relations are just happenings. The relations they have with other events not in our population may make them part of history, but not the history of the case we are working on. Figure 4 identifies and represents the major component. Note that we have moved from 1995 events, many of which were not connected with any other events, to a smaller set of roughly 1476 events, all of which were clustered together on the right hand side of Figure 4.

As before, narrative time moves from the top to the bottom of the page, overlapping events are connected by invisible arcs, and events shared across multiple narratives are shaded. One could consider a component a case. The substantive problem is that it is too fragile. The deletion of any number of single arcs or nodes (causal relations or events) would result in a partition of the component into multiple discreet subgraphs. Our strategy is to define a candidate case as a bicomponent, insisting that all events be connected by at least two independent pathways and to test its robustness to the future. The largest bicomponent contains 493 events. Figure 5 represents the structure of this bicomponent, following the template used in earlier figures. Figure 5 highlights events shared across multiple narratives. *This is the candidate case.*

5. TESTING CASES

In order to know what an event means one has to embed it in a sequence of inter-related events, which are in turn embedded in larger sequences that compose a case. Some cases are more robust than others. Robust cases are composed of elements that even if activated by the future (or by

FIGURE 3



All events from Liu Ling.

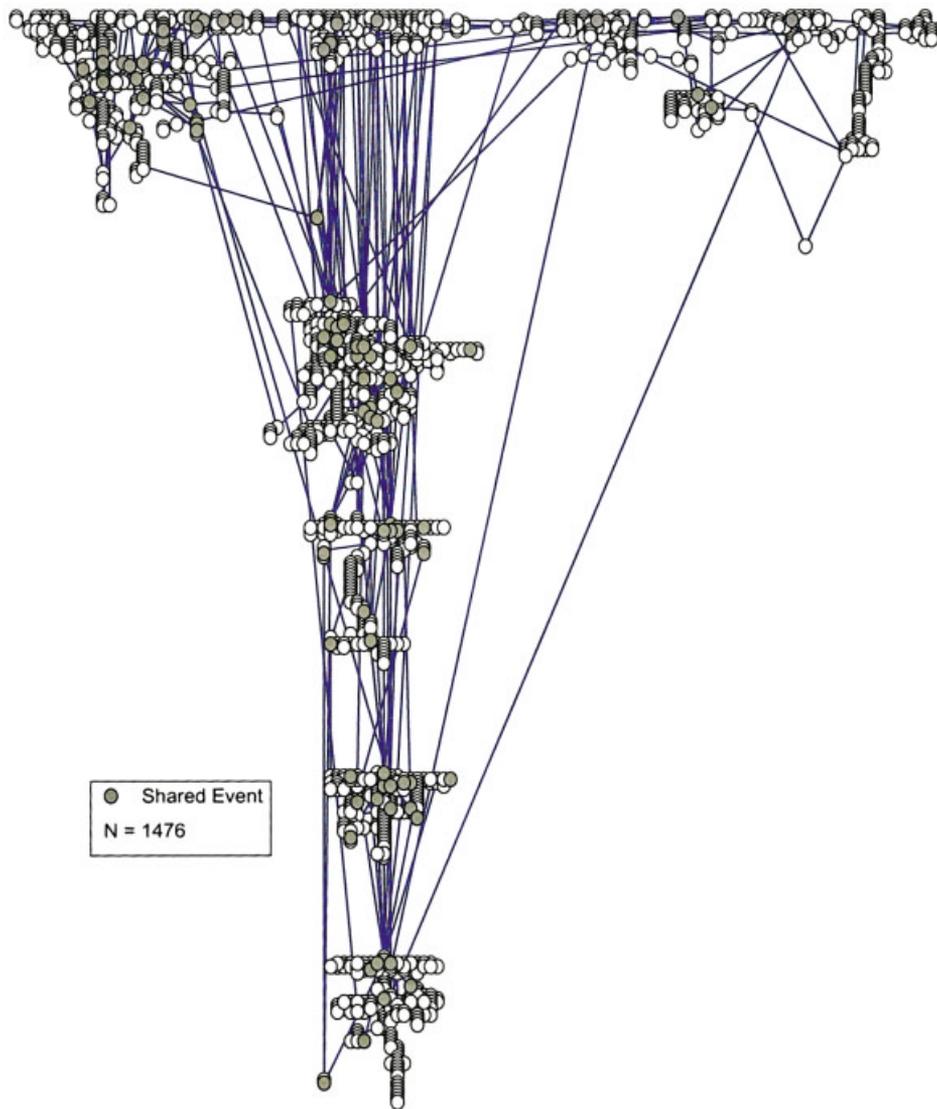
discovery) do not change the case. It is possible to assess case robustness by simulating the effect of the future. The byproducts are an assessment of case robustness and an inventory of events arrayed with respect to the probability that they will be case breakers. Figure 6 reports the robustness of our candidate case, its resilience to both minor and major perturbation. The criteria we use is the RAND statistic, which reports the extent of classification agreement when a randomly selected pair of elements (here, events) are classified in the same way (either belonging to the same cluster, or belonging to different clusters) across two partitions of a matrix. The adjusted statistic corrects for chance overlap [18, Eq. 9] and reports the agreement between two subgraphs beyond chance expectation.

The left side of Figure 6 reports the extent of agreement between the initial events that compose the initial bicomponent ($n = 479$) and the events that compose a second bicomponent potentially altered by the random addition of from 1 to 10 new edges to one or more of the 1995 events that compose the event universe of Liu Ling. In other words, we add some number of random lines to connect previously disconnected events in Liu Ling. Adding edges changes the structure of the original graph (much like the discovery of a

new “fact” might connect two events previously thought disconnected). We then reduce the new graph to its largest bicomponent and compare the bicomponent from the original graph to the new bicomponent. For each case, we run the same simulation 500 times, assessing the effect of adding 1, 2, 3,...10 edges. The dark horizontal line reports the median effect; the shaded cross-hatch reports the interquartile range. Tailing away from the shaded areas are dots that report the extreme effects of adding edges.

It should be obvious that the case is robust to the impact of adding one edge. In the average instance, there is *no* change. In the worst-case scenario, adding a single line results in agreement between the two candidate cases, which is 93% greater than expected by chance. Butterfly effects are possible, but exceedingly rare. A similar pattern is observed for the addition of two or three new relations. Structure breaks down a bit with more and more radical alterations of the original graph. By the time 10 new lines are added, the overlap between the two candidate cases falls to 90% greater than expected by chance. The scope of change is significant, much like the discovery of a new archive, multiple additions would lead to (re)connecting elements of the underlying data structure, thereby poten-

FIGURE 4

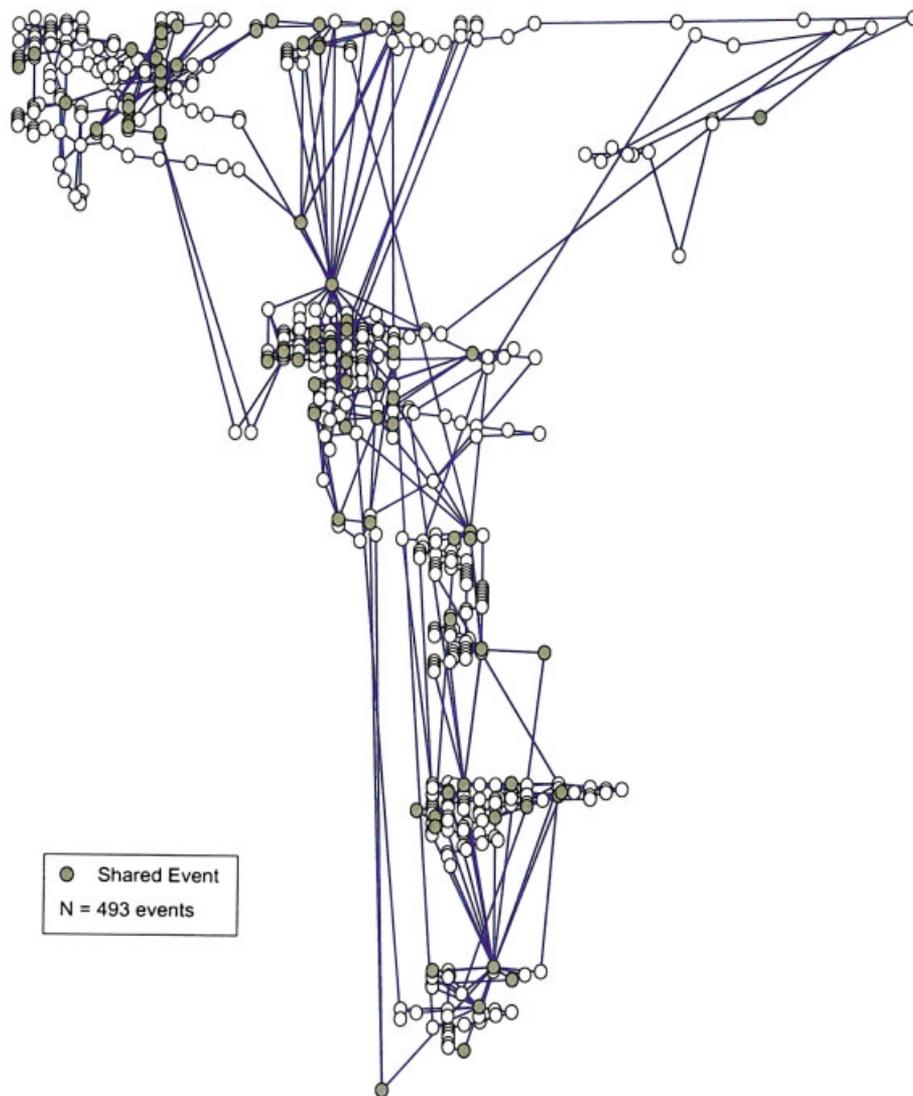


Largest component in Liu Ling.

tially changing their meaning by changing the case in which they are embedded. The simultaneous alteration of multiple causal relations can have a deep multiplier effect. Case instability results from specific combinations (conjunctions) of multiple, simultaneous, changes to the underlying data.

The effect of deleting relationships is much less pronounced. Even in extreme cases, deleting 10 edges, and thus potentially up to 20 nodes, the two candidate cases remain remarkably similar. Here, the contrast between our case and traditional historical narratives (or even the component we

identify earlier) is marked. These findings are not artifactual, and they provide insight into the structure of a case. If one were to delete an edge from a minimally connected bicomponent, the result would be a partition of the component into subgraphs and hence significantly lower classification agreement than we observe. The robustness of the case to deletion implies that the bicomponent is composed of multiple dense clusters and that events which compose each cluster are linked by more than two independent pathways. This structure is closer to that of social structure writ large. The local density of real event structures protects

FIGURE 5

Largest bicomponent, with shared events.

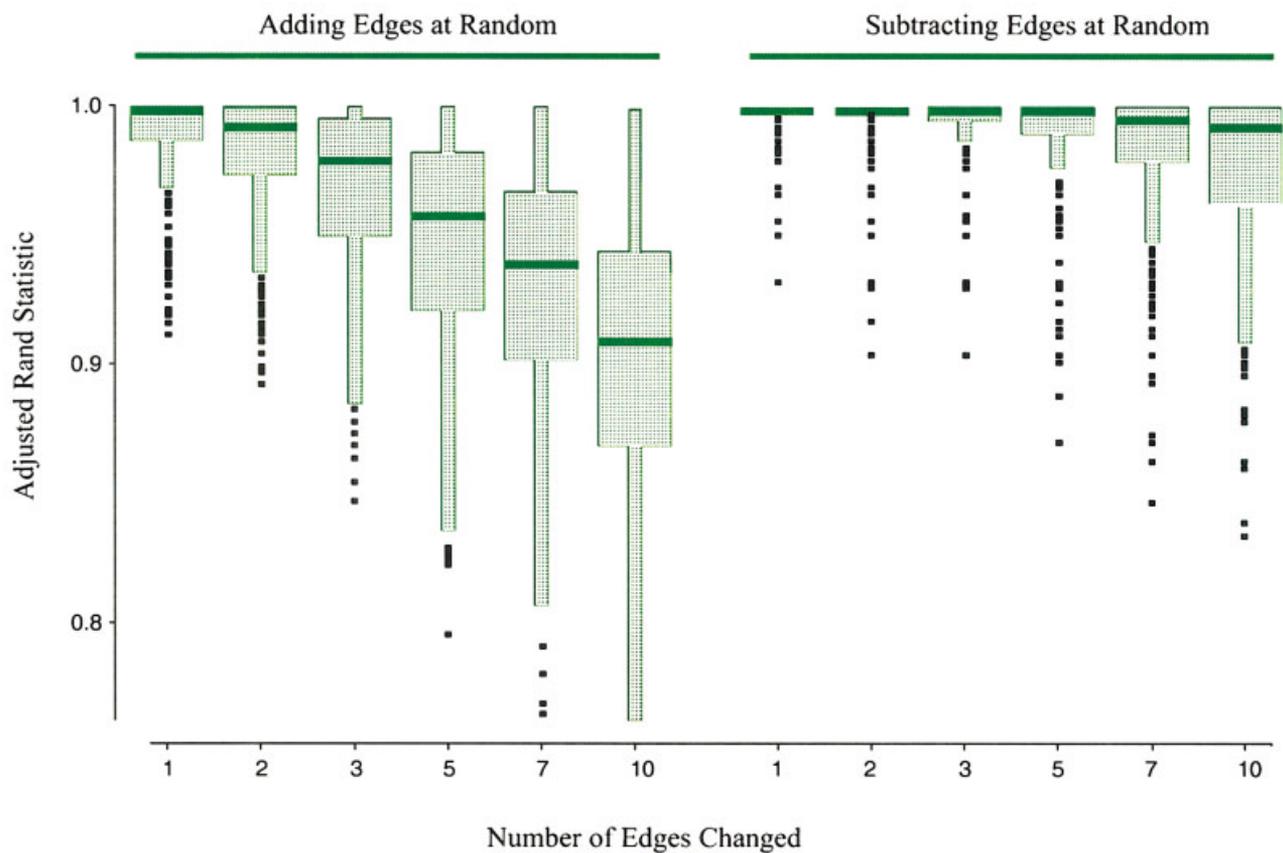
cases from collapsing from perturbations that have the effect of deleting causal relationships between historical events.

Case-breakers

For cases that collapse under subtle pressure (by adding or deleting one or a few lines), one could have little confidence in the meanings ascribed to an event. With cases that are robust to the future, the meaning of the events that compose the case are fixed. It follows that if others followed the same research strategy, they would reveal the same case. Consequently, they would agree on the meaning of the event. As useful is an inventory of events arrayed with

respect to their probability of breaking the case. This array would allow historical social scientists to learn about the structural characteristics of events that have the potential to touch off case-breaking effects. From the tails in both panels of Figure 6, it is clear that in some instances, adding or subtracting one edge can break the case. These are pivotal events. Pivotal events may be induced in ways not already implied by the proximal cohesion of initial event clusters. One mechanism (differentiation) is that an early event cluster connects multiple subsequent event clusters, in each case through multiple independent paths. A second mechanism (convergence) is that separate early event clusters connect to the same subsequent event clusters, in each case

FIGURE 6



Case resilience to discovery.

through multiple independent paths. Various combinations of differentiation may also be visible. In the first case (differentiation), what looks like an unitary event cluster splits into multiple event clusters. In the second case (convergence) we observe the reverse kind of structure.

One simple strategy for identifying high-impact edges/nodes is to loop over each edge (or pair of nodes) one at a time, delete or add it, and calculate an adjusted RAND statistic for the resulting bicomponents. This generates a systematic potential impact score for each edge, under the assumption that it could be deleted (or added between nodes) by some future event. At the boundaries of our case lie smaller, relatively dense event clusters. Whether or not events which lie on the boundary of cases are pivotal depends on the structure of the smaller event clusters that, like moons, are suspended on the periphery of the focal case. In this instance, pivotal events are exclusively located within the semi-dense regions of the bicomponent.

6. DISCUSSION

This article exploits network methods for doing history. By focusing on networks as useful for the method of historical social science, new solutions to old problems have appeared. The deepest problem is what do events mean. The central idea of this article is that the meaning of events is conditional on their position in a sequence of events, and hence, the central problem for historical social science is casing event sequences, in order to induce beginnings and ends. Old solutions to casing are all around. They rest on knowing the end, having a theory to guide the selection of events back toward some beginning. The structure of history appears as a sand clock. All of the tangible causal energy is locked into thin behavioral streams that appear subject to all sorts of contingency. It takes little vision to see that, like nested Russian dolls, the inside of one history provides the outside skin for another. At each remove, what appears globally sparse is revealed to be locally dense, and vice versa.

Network methods provide a way to exploit this fractal characteristic of event structures, if we can reveal them. We illustrate a simple strategy for generating and revealing dense event structures, as a new unit of analysis. The strategy we illustrate is to intercalate multiple stories. The historical event structures that our method produces are characterized by cyclicity, redundancy, and local density. Because they are structures they have meaningful parameters. They conform to our intuitive understanding of a case, as something that envelopes events within a boundary, either by virtue of similar structural principals organizing relations between elements, or deep structuration through memory or cultural encoding. They also conform to our intuitive understanding of how history unfolds as the result of multiple sources operating through multiple pathways at multiple levels of observation.

If history has this structure, it follows that contingency, while possible, is constrained by deeply complex, fractal event structures that absorb events of the present and future. It could hardly be otherwise. How then can it be that contingency and chance play such large roles in historical understanding? One conservative answer is suggested above. Some events break cases. Because this is the case, a central contribution of the methodology we propose is to yield a consistent array of events with respect to their probability of serving as case-breakers. Such an array will, at the least, help historians empirically demonstrate which events are critical for their case. In the long run, better understanding of case breakers within cases, should provide footing for abstraction across cases—the principle aim of historical sociology.

There are more radical possibilities. One could of course, with only some irony, simply assert that the emphasis on chance and contingency results from disciplinary pressures. There is some truth to this assertion, although perhaps not where one would first look. The truth lies in the commitment that the discipline(s) have to the old cases. If new work must remain within the boundaries of accepted cases, the most convenient new arguments will gravitate toward contingency as explanation. It could be otherwise, though. If judgment has produced the real cases to consider, the method proposed in this article will induce those cases, and only those cases.

At the same time, the method we describe allows for the induction of cases—dense event structures robust to slight permutation—for which we have no words. And here, perhaps, lies the avenue for new insight into the past. Our guess

is that historians' commitment to the structure of known cases has significantly limited our understanding of events, event sequences, and the nature of the past, in much the same way that sociologists' commitment to the reality of categorical descriptions of the present limited understanding of the social structures within which the stuff of life is organized, expressed, and enacted. To know for sure, of course, one has to await subsequent applications of network methods—more sophisticated than those utilized here—to history.

REFERENCES

1. Bearman, P.; Stovel, K. *Becoming a Nazi: A model for narrative networks*. Poetics, Forthcoming, 1999.
2. Barkey, K.; Van Rossen, R. Networks of contention: Villages and regional structure in the seventeenth century Ottoman empire. *Am J Sociol* 1997, 102(5), 1345–1382.
3. Bearman, P. *Relations into Rhetorics*; Rose Monograph Series; ASA: New Brunswick, NJ, 1993.
4. Brudner, L.; White, D. Class, property and structural endogamy: Visualizing networked histories. *Theory and Society* 1997, 25, 161–208.
5. Gould, R.V. *Insurgent Identities: Class, Community, and Protest in Paris from 1848 to the Commune*; University of Chicago Press: Chicago, IL, 1995.
6. Gould, R.V. Patron-client ties, state centralization, and the whiskey rebellion. *Am J Sociol* 1996, 102(5), 400–429.
7. Padgett, J.; Ansell, C. Robust action and the rise of the Medici, 1400–1434. *Am J Sociol* 1993, 98, 1259–1319.
8. Emirbayer, M.; Goodwin, J. Network analysis, culture, and the problem of agency. *Am J Sociol* 1994, 99, 1411–1454.
9. Tilly, C. *Durable Inequality*; University of California Press: Berkeley, 1999.
10. White, H. *Identity and Control*; Princeton University Press: Princeton, NJ, 1992.
11. Frank, O. Sampling and estimation in large social networks. *Social Networks* 1978, 1, 91–101.
12. Granovetter, M. Network sampling: Some first steps. *Am J Sociol* 1977, 81(6), 1287–1303.
13. Myrdal, J. *Report From a Chinese Village*; Pantheon Books: New York, NY, 1965.
14. Labov, W. *Language in the Inner City*; University of Pennsylvania Press: Philadelphia, 1972.
15. Franzosi, R. *From Words to Numbers*. Unpublished Manuscript. Oxford University, England, 1999.
16. Wasserman, S.; Faust, K. *Social Network Analysis. Methods and Applications*; Cambridge University Press: Cambridge, MA, 1994.
17. White, D.; Schnegg, M.; Brudner, L.; Nutini, H. *Status Groups and Structural Endogamy: Compadrazgo in Rural Tlaxcala, Mexico*. Unpublished manuscript, 1999.
18. Morey, R.; Agresti, A. The measurement of classification agreement: An adjustment to the RAND statistic for chance agreement. *Educational Psychological Measurement* 1984, 44, 33–37.