

Chapter 4

Networks in Their Surrounding Contexts

In Chapter 3 we considered some of the typical structures that characterize social networks, and some of the typical processes that affect the formation of links in the network. Our discussion there focused primarily on the network as an object of study in itself, relatively independent of the broader world in which it exists.

However, the contexts in which a social network is embedded will generally have significant effects on its structure. Each individual in a social network has a distinctive set of personal characteristics, and similarities and compatibilities among two people's characteristics can strongly influence whether a link forms between them. Each individual also engages in a set of behaviors and activities that can shape the formation of links within the network. These considerations suggest what we mean by a network's *surrounding contexts*: factors that exist outside the nodes and edges of a network, but which nonetheless affect how the network's structure evolves.

In this chapter we consider how such effects operate, and what they imply about the structure of social networks. Among other observations, we will find that the surrounding contexts affecting a network's formation can, to some extent, be viewed in network terms as well — and by expanding the network to represent the contexts together with the individuals, we will see in fact that several different processes of network formation can be described in a common framework.

4.1 Homophily

One of the most basic notions governing the structure of social networks is *homophily* — the principle that we tend to be similar to our friends. Typically, your friends don't look like a random sample of the underlying population: viewed collectively, your friends are generally similar to you along racial and ethnic dimensions; they are similar in age; and they are also similar in characteristics that are more or less mutable, including the places they live, their occupations, their levels of affluence, and their interests, beliefs, and opinions. Clearly most of us have specific friendships that cross all these boundaries; but in aggregate, the pervasive fact is that links in a social network tend to connect people who are similar to one another.

This observation has a long history; as McPherson, Smith-Lovin, and Cook note in their extensive review of research on homophily [294], the underlying idea can be found in writings of Plato (“similarity begets friendship”) and Aristotle (people “love those who are like themselves”), as well as in proverbs such as “birds of a feather flock together.” Its role in modern sociological research was catalyzed in large part by influential work of Lazarsfeld and Merton in the 1950s [269].

Homophily provides us with a first, fundamental illustration of how a network's surrounding contexts can drive the formation of its links. Consider the basic contrast between a friendship that forms because two people are introduced through a common friend and a friendship that forms because two people attend the same school or work for the same company. In the first case, a new link is added for reasons that are *intrinsic* to the network itself; we need not look beyond the network to understand where the link came from. In the second case, the new link arises for an equally natural reason, but one that makes sense only when we look at the contextual factors beyond the network — at some of the social environments (in this case schools and companies) to which the nodes belong.

Often, when we look at a network, such contexts capture some of the dominant features of its overall structure. Figure 4.1, for example, depicts the social network within a particular town's middle school and high school (encompassing grades 7-12) [304]; in this image, produced by the study's author James Moody, students of different races are drawn as differently-colored circles. Two dominant divisions within the network are apparent. One division is based on race (from left to right in the figure); the other, based on age and school attendance, separates students in the middle school from those in the high school (from top to bottom in the figure). There are many other structural details in this network, but the effects of these two contexts stand out when the network is viewed at a global level.

Of course, there are strong interactions between intrinsic and contextual effects on the formation of any single link; they are both operating concurrently in the same network. For example, the principle of triadic closure — that triangles in the network tend to “close” as links form between friends of friends — is supported by a range of mechanisms that range from the intrinsic to the contextual. In Chapter 3 we motivated triadic closure by

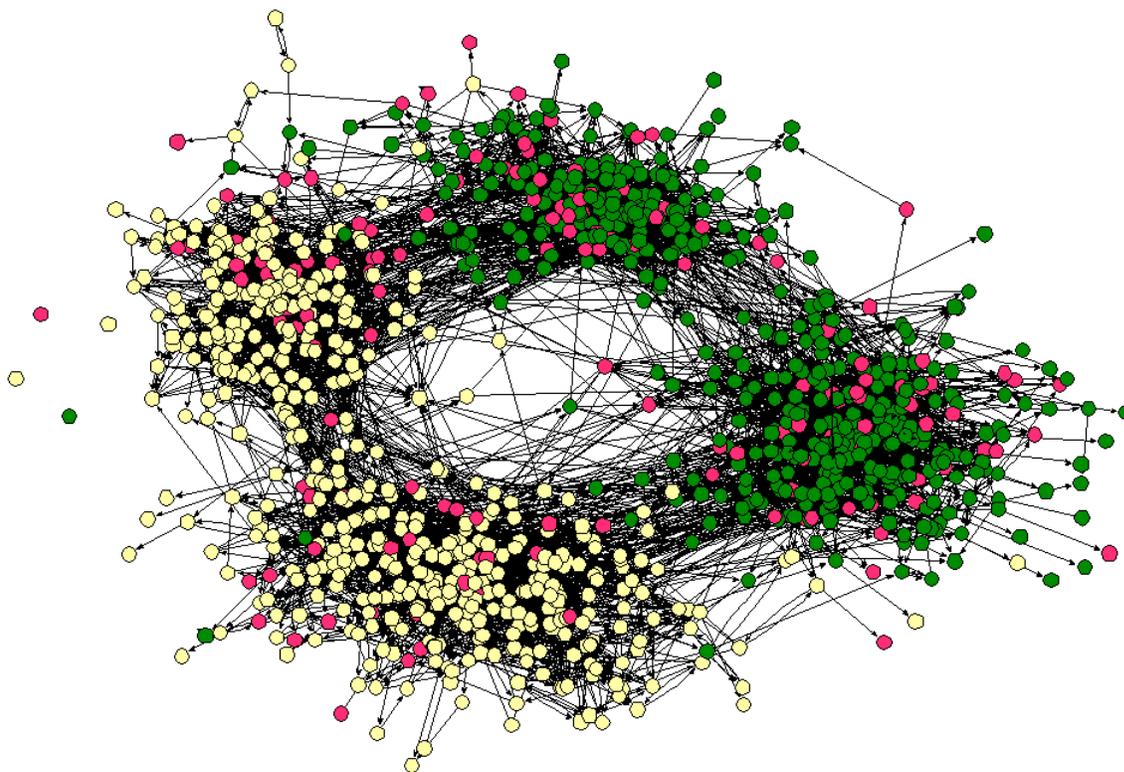


Figure 4.1: Homophily can produce a division of a social network into densely-connected, homogeneous parts that are weakly connected to each other. In this social network from a town’s middle school and high school, two such divisions in the network are apparent: one based on race (with students of different races drawn as differently colored circles), and the other based on friendships in the middle and high schools respectively [304].

hypothesizing intrinsic mechanisms: when individuals B and C have a common friend A , then there are increased opportunities and sources of trust on which to base their interactions, and A will also have incentives to facilitate their friendship. However, social contexts also provide natural bases for triadic closure: since we know that A - B and A - C friendships already exist, the principle of homophily suggests that B and C are each likely to be similar to A in a number of dimensions, and hence quite possibly similar to each other as well. As a result, based purely on this similarity, there is an elevated chance that a B - C friendship will form; and this is true *even if neither of them is aware that the other one knows A*.

The point isn’t that any one basis for triadic closure is the “correct” one. Rather, as we take into account more and more of the factors that drive the formation of links in a social

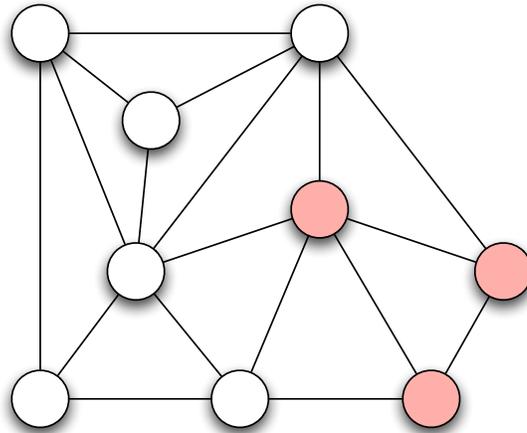


Figure 4.2: Using a numerical measure, one can determine whether small networks such as this one (with nodes divided into two types) exhibit homophily.

network, it inevitably becomes difficult to attribute any individual link to a single factor. And ultimately, one expects most links to in fact arise from a combination of several factors — partly due to the effect of other nodes in the network, and partly due to the surrounding contexts.

Measuring Homophily. When we see striking divisions within a network like the one in Figure 4.1, it is important to ask whether they are “genuinely” present in the network itself, and not simply an artifact of how it is drawn. To make this question concrete, we need to formulate it more precisely: given a particular characteristic of interest (like race, or age), is there a simple test we can apply to a network in order to estimate whether it exhibits homophily according to this characteristic?

Since the example in Figure 4.1 is too large to inspect by hand, let’s consider this question on a smaller example where we can develop some intuition. Let’s suppose in particular that we have the friendship network of an elementary-school classroom, and we suspect that it exhibits homophily by gender: boys tend to be friends with boys, and girls tend to be friends with girls. For example, the graph in Figure 4.2 shows the friendship network of a (small) hypothetical classroom in which the three shaded nodes are girls and the six unshaded nodes are boys. If there were no cross-gender edges at all, then the question of homophily would be easy to resolve: it would be present in an extreme sense. But we expect that homophily should be a more subtle effect that is visible mainly in aggregate — as it is, for example, in the real data from Figure 4.1. Is the picture in Figure 4.2 consistent with homophily?

There is a natural numerical measure of homophily that we can use to address questions

like this [202, 319]. To motivate the measure (using the example of gender as in Figure 4.2), we first ask the following question: what would it mean for a network *not* to exhibit homophily by gender? It would mean that the proportion of male and female friends a person has looks like the background male/female distribution in the full population. Here’s a closely related formulation of this “no-homophily” definition that is a bit easier to analyze: if we were to randomly assign each node a gender according to the gender balance in the real network, then the number of cross-gender edges should not change significantly relative to what we see in the real network. That is, in a network with no homophily, friendships are being formed as though there were random mixing across the given characteristic.

Thus, suppose we have a network in which a p fraction of all individuals are male, and a q fraction of all individuals are female. Consider a given edge in this network. If we independently assign each node the gender *male* with probability p and the gender *female* with probability q , then both ends of the edge will be male with probability p^2 , and both ends will be female with probability q^2 . On the other hand, if the first end of the edge is male and the second end is female, or vice versa, then we have a cross-gender edge, so this happens with probability $2pq$.

So we can summarize the test for homophily according to gender as follows:

Homophily Test: If the fraction of cross-gender edges is significantly less than $2pq$, then there is evidence for homophily.

In Figure 4.2, for example, 5 of the 18 edges in the graph are cross-gender. Since $p = 2/3$ and $q = 1/3$ in this example, we should be comparing the fraction of cross-gender edges to the quantity $2pq = 4/9 = 8/18$. In other words, with no homophily, one should expect to see 8 cross-gender edges rather than than 5, and so this example shows some evidence of homophily.

There are a few points to note here. First, the number of cross-gender edges in a random assignment of genders will deviate some amount from its expected value of $2pq$, and so to perform the test in practice one needs a working definition of “significantly less than.” Standard measures of statistical significance (quantifying the significance of a deviation below a mean) can be used for this purpose. Second, it’s also easily possible for a network to have a fraction of cross-gender edges that is significantly *more than* $2pq$. In such a case, we say that the network exhibits *inverse homophily*. The network of romantic relationships in Figure 2.7 from Chapter 2 is a clear example of this; almost all the relationships reported by the high-school students in the study involved opposite-sex partners, rather than same-sex partners, so almost all the edges are cross-gender.

Finally, it’s easy to extend our homophily test to any underlying characteristic (race, ethnicity, age, native language, political orientation, and so forth). When the characteristic can only take two possible values (say, one’s voting preference in a two-candidate election), then we can draw a direct analogy to the case of two genders, and use the same formula

$2pq$. When the characteristic can take on more than two possible values, we still perform a general version of the same calculation. For this, we say that an edge is *heterogeneous* if it connects two nodes that are different according to the characteristic in question. We then ask how the number of heterogeneous edges compares to what we'd see if we were to randomly assign values for the characteristic to all nodes in the network — using the proportions from the real data as probabilities. In this way, even a network in which the nodes are classified into many groups can be tested for homophily using the same underlying comparison to a baseline of random mixing.

4.2 Mechanisms Underlying Homophily: Selection and Social Influence

The fact that people tend to have links to others who are similar to them is a statement about the structure of social networks; on its own, it does not propose an underlying mechanism by which ties among similar people are preferentially formed.

In the case of immutable characteristics such as race or ethnicity, the tendency of people to form friendships with others who are like them is often termed *selection*, in that people are selecting friends with similar characteristics. Selection may operate at several different scales, and with different levels of intentionality. In a small group, when people choose friends who are most similar from among a clearly delineated pool of potential contacts, there is clearly active choice going on. In other cases, and at more global levels, selection can be more implicit. For example, when people live in neighborhoods, attend schools, or work for companies that are relatively homogeneous compared to the population at large, the social environment is already favoring opportunities to form friendships with others like oneself. For this discussion, we will refer to all these effects cumulatively as selection.

When we consider how immutable characteristics interact with network formation, the order of events is clear: a person's attributes are determined at birth, and they play a role in how this person's connections are formed over the course of his or her life. With characteristics that are more mutable, on the other hand — behaviors, activities, interests, beliefs, and opinions — the feedback effects between people's individual characteristics and their links in the social network become significantly more complex. The process of selection still operates, with individual characteristics affecting the connections that are formed. But now another process comes into play as well: people may modify their behaviors to bring them more closely into alignment with the behaviors of their friends. This process has been variously described as *socialization* [233] and *social influence* [170], since the existing social connections in a network are influencing the individual characteristics of the nodes. Social influence can be viewed as the reverse of selection: with selection, the individual characteristics drive the formation of links, while with social influence, the existing links in

the network serve to shape people's (mutable) characteristics.¹

The Interplay of Selection and Social Influence. When we look at a single snapshot of a network and see that people tend to share mutable characteristics with their friends, it can be very hard to sort out the distinct effects and relative contributions of selection and social influence. Have the people in the network adapted their behaviors to become more like their friends, or have they sought out people who were already like them? Such questions can be addressed using *longitudinal* studies of a social network, in which the social connections and the behaviors within a group are both tracked over a period of time. Fundamentally, this makes it possible to see the behavioral changes that occur after changes in an individual's network connections, as opposed to the changes to the network that occur after an individual changes his or her behavior.

This type of methodology has been used, for example, to study the processes that lead pairs of adolescent friends to have similar outcomes in terms of scholastic achievement and delinquent behavior such as drug use [92]. Empirical evidence confirms the intuitive fact that teenage friends are similar to each other in their behaviors, and both selection and social influence have a natural resonance in this setting: teenagers seek out social circles composed of people like them, and peer pressure causes them to conform to behavioral patterns within their social circles. What is much harder to resolve is how these two effects interact, and whether one is more strongly at work than the other. As longitudinal behavior relevant to this question became available, researchers began quantifying the relative impact of these different factors. A line of work beginning with Cohen and Kandel has suggested that while both effects are present in the data, the outsized role that earlier informal arguments had accorded to peer pressure (i.e. social influence) is actually more moderate; the effect of selection here is in fact comparable to (and sometimes greater than) the effect of social influence [114, 233].

Understanding the tension between these different forces can be important not just for identifying underlying causes, but also for reasoning about the effect of possible interventions one might attempt in the system [21, 396]. For example, once we find that illicit drug use displays homophily across a social network — with students showing a greater likelihood to use drugs when their friends do — we can ask about the effects of a program that targets certain high-school students and influences them to stop using drugs. To the extent that the observed homophily is based on some amount of social influence, such a program could have a broad impact across the social network, by causing the friends of these targeted students to stop using drugs as well. But one must be careful; if the observed homophily is arising instead almost entirely from selection effects, then the program may not reduce drug use

¹There are other cognitive effects at work as well; for example, people may systematically misperceive the characteristics of their friends as being more in alignment with their own than they really are [224]. For our discussion here, we will not focus explicitly on such effects.

beyond the students it directly targets: as these students stop using drugs, they change their social circles and form new friendships with students who don't use drugs, but the drug-using behavior of other students is not strongly affected.

Another example of research addressing this subtle interplay of factors is the work of Christakis and Fowler on the effect of social networks on health-related outcomes. In one recent study, using longitudinal data covering roughly 12,000 people, they tracked obesity status and social network structure over a 32-year period [108]. They found that obese and non-obese people clustered in the network in a fashion consistent with homophily, according to the numerical measure described in Section 4.1: people tend to be more similar in obesity status to their network neighbors than in a version of the same network where obesity status is assigned randomly. The problem is then to distinguish among several hypotheses for why this clustering is present: is it

- (i) because of selection effects, in which people are choosing to form friendships with others of similar obesity status?
- (ii) because of the confounding effects of homophily according to other characteristics, in which the network structure indicates existing patterns of similarity in other dimensions that correlate with obesity status? or
- (iii) because changes in the obesity status of a person's friends was exerting a (presumably behavioral) influence that affected his or her future obesity status?

Statistical analysis in Christakis and Fowler's paper argues that, even accounting for effects of types (i) and (ii), there is significant evidence for an effect of type (iii) as well: that obesity is a health condition displaying a form of social influence, with changes in your friends' obesity status in turn having a subsequent effect on you. This suggests the intriguing prospect that obesity (and perhaps other health conditions with a strong behavioral aspect) may exhibit some amount of "contagion" in a social sense: you don't necessarily catch it from your friends the way you catch the flu, but it nonetheless can spread through the underlying social network via the mechanism of social influence.

These examples, and this general style of investigation, show how careful analysis is needed to distinguish among different factors contributing to an aggregate conclusion: even when people tend to be similar to their neighbors in a social network, it may not be clear why. The point is that an observation of homophily is often not an endpoint in itself, but rather the starting point for deeper questions — questions that address why the homophily is present, how its underlying mechanisms will affect the further evolution of the network, and how these mechanisms interact with possible outside attempts to influence the behavior of people in the network.

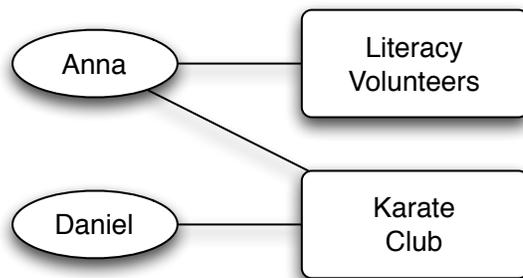


Figure 4.3: An affiliation network is a bipartite graph that shows which individuals are affiliated with which groups or activities. Here, Anna participates in both of the social foci on the right, while Daniel participates in only one.

4.3 Affiliation

Thus far, we have been discussing contextual factors that affect the formation of links in a network — based on similarities in characteristics of the nodes, and based on behaviors and activities that the nodes engage in. These surrounding contexts have been viewed, appropriately, as existing “outside” the network. But in fact, it’s possible to put these contexts into the network itself, by working with a larger network that contains both people and contexts as nodes. Through such a network formulation, we will get additional insight into some broad aspects of homophily, and see how the simultaneous evolution of contexts and friendships can be put on a common network footing with the notion of triadic closure from Chapter 3.

In principle we could represent any context this way, but for the sake of concreteness we’ll focus on how to represent the set of activities a person takes part in, and how these affect the formation of links. We will take a very general view of the notion of an “activity” here. Being part of a particular company, organization, or neighborhood; frequenting a particular place; pursuing a particular hobby or interest — these are all activities that, when shared between two people, tend to increase the likelihood that they will interact and hence form a link in the social network [78, 161]. Adopting terminology due to Scott Feld, we’ll refer to such activities as *foci* — that is, “focal points” of social interaction — constituting “social, psychological, legal, or physical entit[ies] around which joint activities are organized (e.g. workplaces, voluntary organizations, hangouts, etc.)” [161].

Affiliation Networks. As a first step, we can represent the participation of a set of people in a set of foci using a graph as follows. We will have a node for each person, and a node for each focus, and we will connect person A to focus X by an edge if A participates in X .

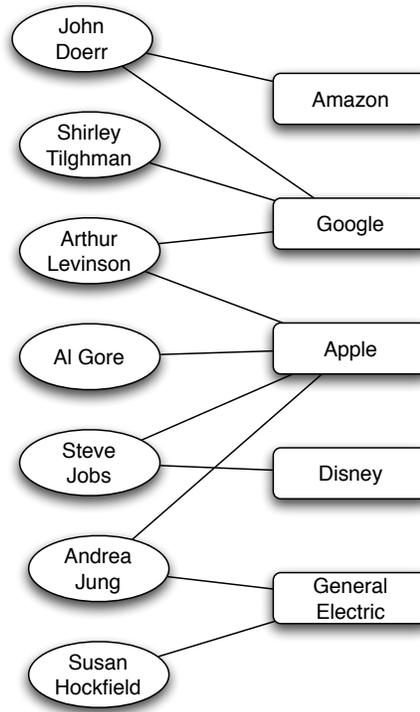


Figure 4.4: One type of affiliation network that has been widely studied is the memberships of people on corporate boards of directors [301]. A very small portion of this network (as of mid-2009) is shown here. The structural pattern of memberships can reveal subtleties in the interactions among both the board members and the companies.

A very simple example of such a graph is depicted in Figure 4.3, showing two people (Anna and Daniel) and two foci (working for a literacy tutoring organization, and belonging to a karate club). The graph indicates that Anna participates in both of the foci, while Daniel participates in only one.

We will refer to such a graph as an *affiliation network*, since it represents the affiliation of people (drawn on the left) with foci (drawn on the right) [78, 323]. More generally, affiliation networks are examples of a class of graphs called *bipartite graphs*. We say that a graph is bipartite if its nodes can be divided into two sets in such a way that every edge connects a node in one set to a node in the other set. (In other words, there are no edges joining a pair of nodes that belong to the same set; all edges go between the two sets.) Bipartite graphs are very useful for representing data in which the items under study come in two categories, and we want to understand how the items in one category are associated with the items in the other. In the case of affiliation networks, the two categories are the people and the foci, with each edge connecting a person to a focus that he or she participates in. Bipartite

graphs are often drawn as in Figure 4.3, with the two different sets of nodes drawn as two parallel vertical columns, and the edges crossing between the two columns.

Affiliation networks are studied in a range of settings where researchers want to understand the patterns of participation in structured activities. As one example, they have received considerable attention in studying the composition of boards of directors of major corporations [301]. Boards of directors are relatively small advisory groups populated by high-status individuals; and since many people serve on multiple boards, the overlaps in their participation have a complex structure. These overlaps can be naturally represented by an affiliation network; as the example in Figure 4.4 shows, there is a node for each person and a node for each board, and each edge connects a person to a board that they belong to.

Affiliation networks defined by boards of directors have the potential to reveal interesting relationships on both sides of the graph. Two companies are implicitly linked by having the same person sit on both their boards; we can thus learn about possible conduits for information and influence to flow between different companies. Two people, on the other hand, are implicitly linked by serving together on a board, and so we learn about particular patterns of social interaction among some of the most powerful members of society. Of course, even the complete affiliation network of people and boards (of which Figure 4.4 is only a small piece) still misses other important contexts that these people inhabit; for example, the seven people in Figure 4.4 include the presidents of two major universities and a former Vice-President of the United States.²

Co-Evolution of Social and Affiliation Networks. It’s clear that both social networks and affiliation networks change over time: new friendship links are formed, and people become associated with new foci. Moreover, these changes represent a kind of *co-evolution* that reflects the interplay between selection and social influence: if two people participate in a shared focus, this provides them with an opportunity to become friends; and if two people are friends, they can influence each other’s choice of foci.

There is a natural network perspective on these ideas, which begins from a network representation that slightly extends the notion of an affiliation network. As before, we’ll have nodes for people and nodes for foci, but we now introduce two distinct kinds of edges as well. The first kind of edge functions as an edge in a social network: it connects two

²The structure of this network changes over time as well, and sometimes in ways that reinforce the points in our present discussion. For example, the board memberships shown in Figure 4.4 are taken from the middle of 2009; by the end of 2009, Arthur Levinson had resigned from the board of directors of Google (thus removing one edge from the graph). As part of the news coverage of this resignation, the chair of the U.S. Federal Trade Commission, Jon Leibowitz, explicitly invoked the notion of overlaps in board membership, saying, “Google, Apple and Mr. Levinson should be commended for recognizing that overlapping board members between competing companies raise serious antitrust issues, and for their willingness to resolve our concerns without the need for litigation. Beyond this matter, we will continue to monitor companies that share board members and take enforcement actions where appropriate” [219].

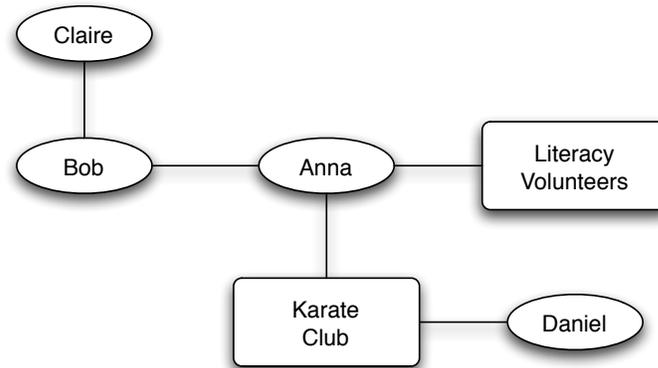


Figure 4.5: A social-affiliation network shows both the friendships between people and their affiliation with different social foci.

people, and indicates friendship (or alternatively some other social relation, like professional collaboration). The second kind of edge functions as an edge in an affiliation network: it connects a person to a focus, and indicates the participation of the person in the focus. We will call such a network a *social-affiliation network*, reflecting the fact that it simultaneously contains a social network on the people and an affiliation network on the people and foci. Figure 4.5 depicts a simple social-affiliation network.

Once we have social-affiliation networks as our representation, we can appreciate that a range of different mechanisms for link formation can all be viewed as types of *closure processes*, in that they involve “closing” the third edge of a triangle in the network. In particular, suppose we have two nodes B and C with a common neighbor A in the network, and suppose that an edge forms between B and C . There are several interpretations for what this corresponds to, depending on whether A , B , and C are people or foci.

- (i) If A , B , and C each represent a person, then the formation of the link between B and C is triadic closure, just as in Chapter 3. (See Figure 4.6(a).)
- (ii) If B and C represent people, but A represents a focus, then this is something different: it is the tendency of two people to form a link when they have a focus in common. (See Figure 4.6(b).) This is an aspect of the more general principle of selection, forming links to others who share characteristics with you. To emphasize the analogy with triadic closure, this process has been called *focal closure* [259].
- (iii) If A and B are people, and C is a focus, then we have the formation of a new affiliation: B takes part in a focus that her friend A is already involved in. (See Figure 4.6(c).) This is a kind of social influence, in which B 's behavior comes into closer alignment

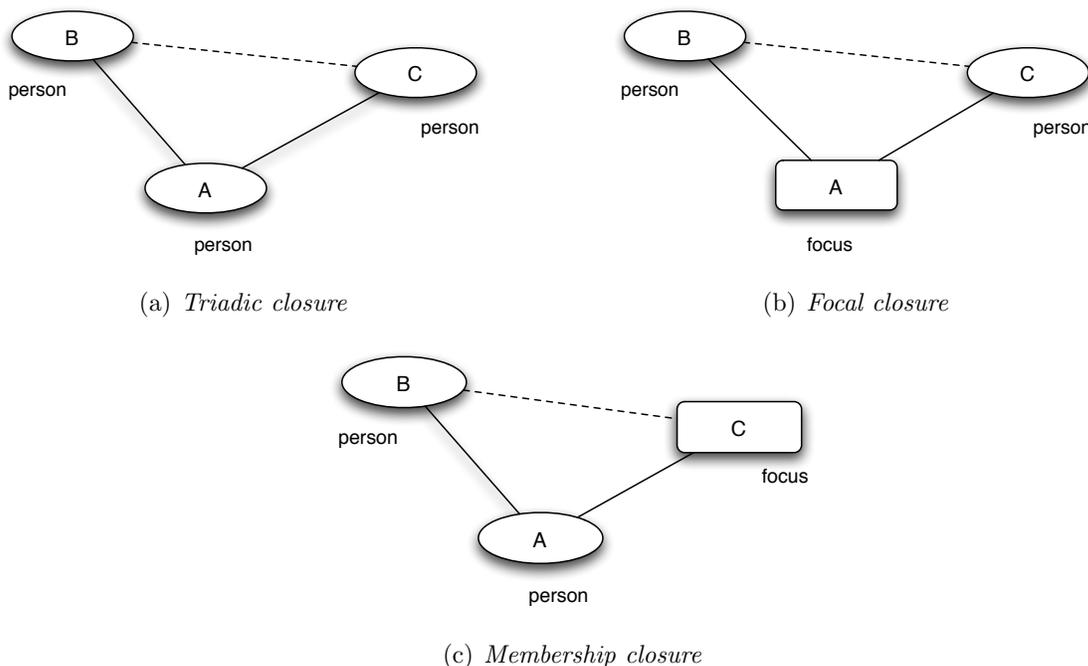


Figure 4.6: Each of triadic closure, focal closure, and membership closure corresponds to the closing of a triangle in a social-affiliation network.

with that of her friend *A*. Continuing the analogy with triadic closure, we will refer to this kind of link formation as *membership closure*.

Thus, three very different underlying mechanisms — reflecting triadic closure and aspects of selection and social influence — can be unified in this type of network as kinds of closure: the formation of a link in cases where the two endpoints already have a neighbor in common. Figure 4.7 shows all three kinds of closure processes at work: triadic closure leads to a new link between Anna and Claire; focal closure leads to a new link between Anna and Daniel; and membership closure leads to Bob’s affiliation with the karate club. Oversimplifying the mechanisms at work, they can be summarized in the following succinct way:

- (i) *Bob introduces Anna to Claire.*
- (ii) *Karate introduces Anna to Daniel.*
- (iii) *Anna introduces Bob to Karate.*

4.4 Tracking Link Formation in On-Line Data

In this chapter and the previous one, we have identified a set of different mechanisms that lead to the formation of links in social networks. These mechanisms are good examples

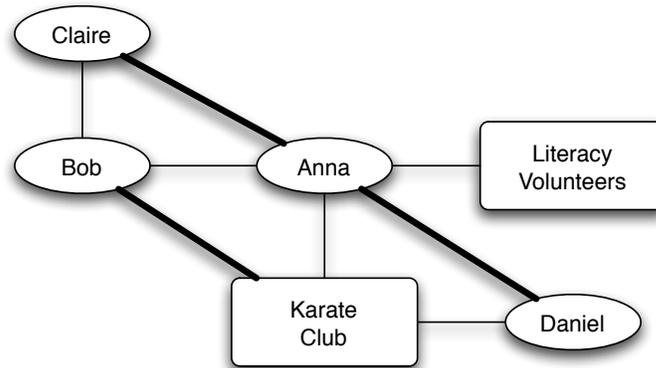


Figure 4.7: In a social-affiliation network containing both people and foci, edges can form under the effect of several different kinds of closure processes: two people with a friend in common, two people with a focus in common, or a person joining a focus that a friend is already involved in.

of social phenomena which are clearly at work in small-group settings, but which have traditionally been very hard to measure quantitatively. A natural research strategy is to try tracking these mechanisms as they operate in large populations, where an accumulation of many small effects can produce something observable in the aggregate. However, given that most of the forces responsible for link formation go largely unrecorded in everyday life, it is a challenge to select a large, clearly delineated group of people (and social foci), and accurately quantify the relative contributions that these different mechanisms make to the formation of real network links.

The availability of data from large on-line settings with clear social structure has made it possible to attempt some preliminary research along these lines. As we emphasized in Chapter 2, any analysis of social processes based on such on-line datasets must come with a number of caveats. In particular, it is never *a priori* clear how much one can extrapolate from digital interactions to interactions that are not computer-mediated, or even from one computer-mediated setting to another. Of course, this problem of extrapolation is present whenever one studies phenomena in a model system, on-line or not, and the kinds of measurements these large datasets enable represent interesting first steps toward a deeper quantitative understanding of how mechanisms of link formation operate in real life. Exploring these questions in a broader range of large datasets is an important problem, and one that will become easier as large-scale data becomes increasingly abundant.

Triadic closure. With this background in mind, let's start with some questions about triadic closure. Here's a first, basic numerical question: how much more likely is a link to

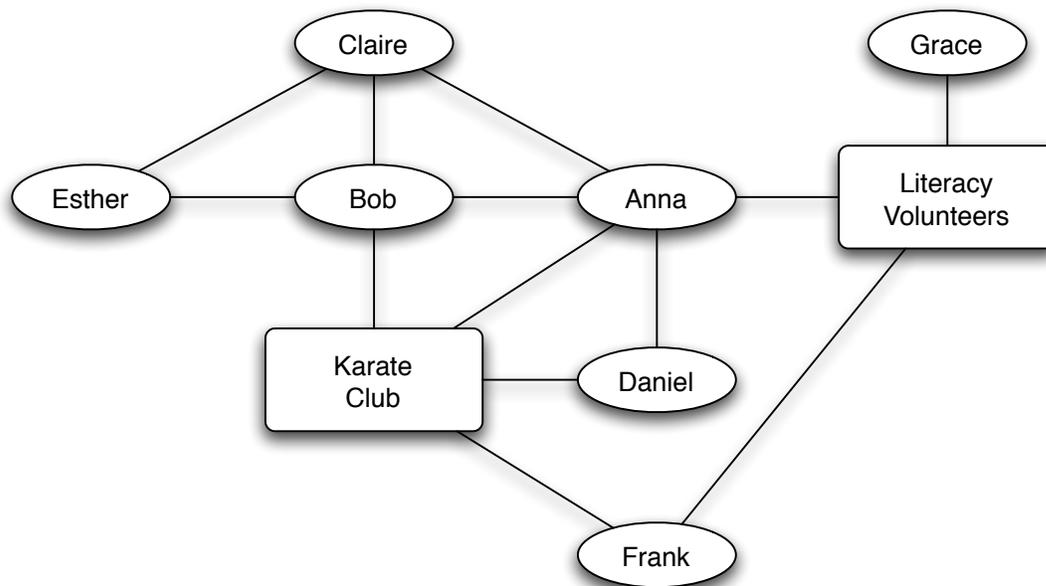


Figure 4.8: A larger network that contains the example from Figure 4.7. Pairs of people can have more than one friend (or more than one focus) in common; how does this increase the likelihood that an edge will form between them?

form between two people in a social network if they already have a friend in common? (In other words, how much more likely is a link to form if it has the effect of closing a triangle?)

Here's a second question, along the same lines as the first: How much more likely is an edge to form between two people if they have *multiple* friends in common? For example, in Figure 4.8, Anna and Esther have two friends in common, while Claire and Daniel only have one friend in common. How much more likely is the formation of a link in the first of these two cases? If we go back to the arguments for why triadic closure operates in social networks, we see that they all are qualitatively strengthened as two people have more friends in common: there are more sources of opportunity and trust for the interaction, there are more people with an incentive to bring them together, and the evidence for homophily is arguably stronger.

We can address these questions empirically using network data as follows.

- (i) We take two snapshots of the network at different times.
- (ii) For each k , we identify all pairs of nodes who have exactly k friends in common in the first snapshot, but who are not directly connected by an edge.
- (iii) We define $T(k)$ to be the fraction of these pairs that have formed an edge by the time

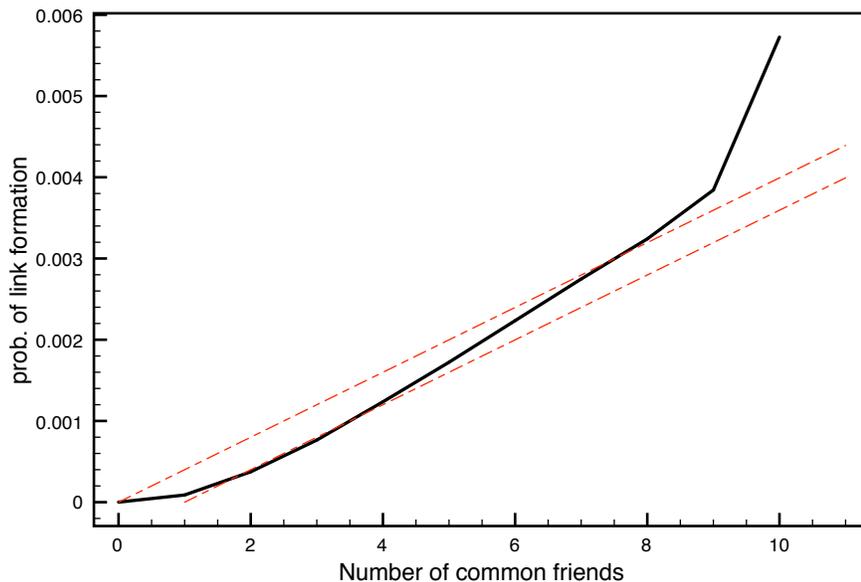


Figure 4.9: Quantifying the effects of triadic closure in an e-mail dataset [259]. The curve determined from the data is shown in the solid black line; the dotted curves show a comparison to probabilities computed according to two simple baseline models in which common friends provide independent probabilities of link formation.

of the second snapshot. This is our empirical estimate for the probability that a link will form between two people with k friends in common.

- (iv) We plot $T(k)$ as a function of k to illustrate the effect of common friends on the formation of links.

Note that $T(0)$ is the rate at which link formation happens when it does not close a triangle, while the values of $T(k)$ for larger k determine the rate at which link formation happens when it does close a triangle. Thus, the comparison between $T(0)$ and these other values addresses the most basic question about the power of triadic closure.

Kossinets and Watts computed this function $T(k)$ using a dataset encoding the full history of e-mail communication among roughly 22,000 undergraduate and graduate students over a one-year period at a large U.S. university [259]. This is a “who-talks-to-whom” type of dataset, as we discussed in Chapter 2; from the communication traces, Kossinets and Watts constructed a network that evolved over time, joining two people by a link at a given instant if they had exchanged e-mail in each direction at some point in the past 60 days. They then determined an “average” version of $T(k)$ by taking multiple pairs of snapshots: they built a curve for $T(k)$ on each pair of snapshots using the procedure described above, and then

averaged all the curves they obtained. In particular, the observations in each snapshot were one day apart, so their computation gives the average probability that two people form a link per day, as a function of the number of common friends they have.

Figure 4.9 shows a plot of this curve (in the solid black line). The first thing one notices is the clear evidence for triadic closure: $T(0)$ is very close to 0, after which the probability of link formation increases steadily as the number of common friends increases. Moreover, for much of the plot, this probability increases in a roughly linear fashion as a function of the number of common friends, with an upward bend away from a straight-line shape. The curve turns upward in a particularly pronounced way from 0 to 1 to 2 friends: having two common friends produces significantly more than twice the effect on link formation compared to having a single common friend. (The upward effect from 8 to 9 to 10 friends is also significant, but it occurs on a much smaller sub-population, since many fewer people in the data have this many friends in common without having already formed a link.)

To interpret this plot more deeply, it helps to compare it to an intentionally simplified baseline model, describing what one might have expected the data to look like in the presence of triadic closure. Suppose that for some small probability p , each common friend that two people have gives them an independent probability p of forming a link each day. So if two people have k friends in common, the probability they fail to form a link on any given day is $(1 - p)^k$: this is because each common friend fails to cause the link to form with probability $1 - p$, and these k trials are independent. Since $(1 - p)^k$ is the probability the link fails to form on a given day, the probability that it does form, according to our simple baseline model, is

$$T_{\text{baseline}}(k) = 1 - (1 - p)^k.$$

We plot this curve in Figure 4.9 as the upper dotted line. Given the small absolute effect of the first common friend in the data, we also show a comparison to the curve $1 - (1 - p)^{k-1}$, which just shifts the simple baseline curve one unit to the right. Again, the point is not to propose this baseline as an explanatory mechanism for triadic closure, but rather to look at how the real data compares to it. Both the real curve and the baseline curve are close to linear, and hence qualitatively similar; but the fact that the real data turns upward while the baseline curve turns slightly downward indicates that the assumption of independent effects from common friends is too simple to be fully supported by the data.

A still larger and more detailed study of these effects was conducted by Leskovec et al. [272], who analyzed properties of triadic closure in the on-line social networks of LinkedIn, Flickr, Del.icio.us, and Yahoo! Answers. It remains an interesting question to try understanding the similarities and variations in triadic closure effects across social interaction in a range of different settings.

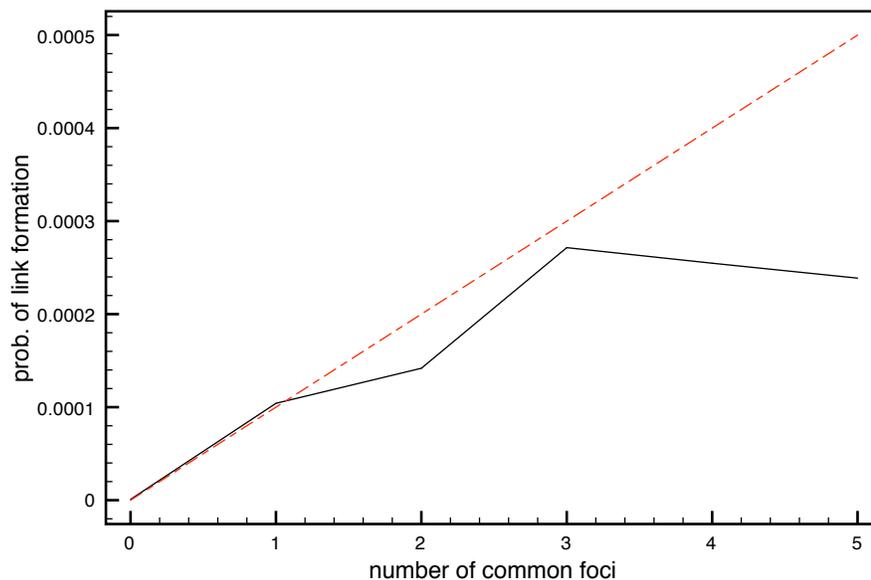


Figure 4.10: Quantifying the effects of focal closure in an e-mail dataset [259]. Again, the curve determined from the data is shown in the solid black line, while the dotted curve provides a comparison to a simple baseline.

Focal and Membership Closure. Using the same approach, we can compute probabilities for the other kinds of closure discussed earlier — specifically,

- *focal closure*: what is the probability that two people form a link as a function of the number of foci they are jointly affiliated with?
- *membership closure*: what is the probability that a person becomes involved with a particular focus as a function of the number of friends who are already involved in it?

As an example of the first of these kinds of closure, using Figure 4.8, Anna and Grace have one activity in common while Anna and Frank have two in common. As an example of the second, Esther has one friend who belongs to the karate club while Claire has two. How do these distinctions affect the formation of new links?

For focal closure, Kossinets and Watts supplemented their university e-mail dataset with information about the class schedules for each student. In this way, each class became a focus, and two students shared a focus if they had taken a class together. They could then compute the probability of focal closure by direct analogy with their computation for triadic closure, determining the probability of link formation per day as a function of the number of shared foci. Figure 4.10 shows a plot of this function. A single shared class turns out to have roughly the same absolute effect on link formation as a single shared friend, but after this the

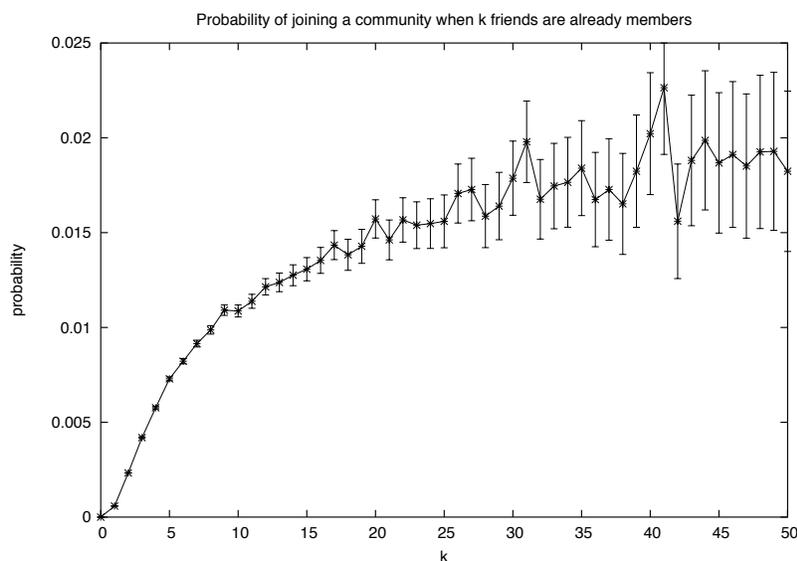


Figure 4.11: Quantifying the effects of membership closure in a large online dataset: The plot shows the probability of joining a LiveJournal community as a function of the number of friends who are already members [32].

curve for focal closure behaves quite differently from the curve for triadic closure: it turns downward and appears to approximately level off, rather than turning slightly upward. Thus, subsequent shared classes after the first produce a “diminishing returns” effect. Comparing to the same kind of baseline, in which the probability of link formation with k shared classes is $1 - (1 - p)^k$ (shown as the dotted curve in Figure 4.10), we see that the real data turns downward more significantly than this independent model. Again, it is an interesting open question to understand how this effect generalizes to other types of shared foci, and to other domains.

For membership closure, the analogous quantities have been measured in other on-line domains that possess both person-to-person interactions and person-to-focus affiliations. Figure 4.11 is based on the blogging site LiveJournal, where friendships are designated by users in their profiles, and where foci correspond to membership in user-defined communities [32]; thus the plot shows the probability of joining a community as a function of the number of friends who have already done so. Figure 4.12 shows a similar analysis for Wikipedia [122]. Here, the social-affiliation network contains a node for each Wikipedia editor who maintains a user account and *user talk page* on the system; and there is an edge joining two such editors if they have communicated, with one editor writing on the user talk page of the other. Each

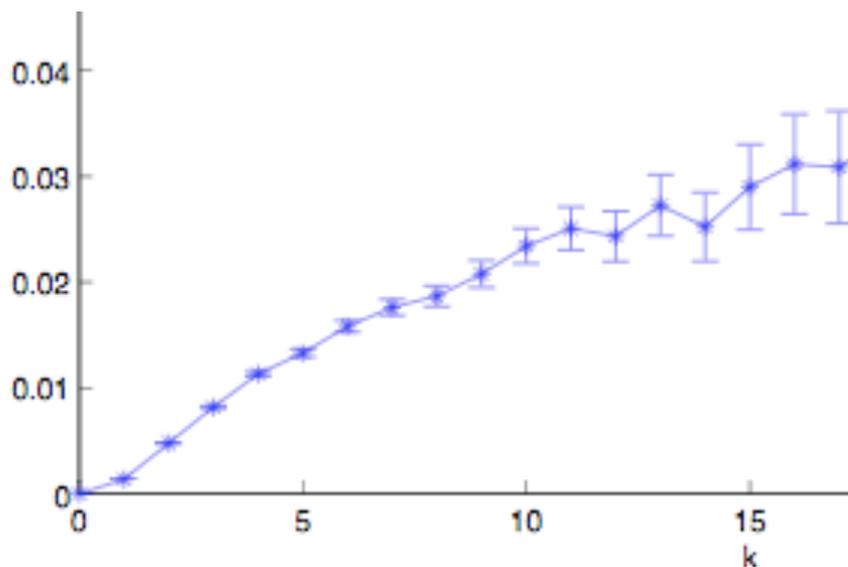


Figure 4.12: Quantifying the effects of membership closure in a large online dataset: The plot shows the probability of editing a Wikipedia articles as a function of the number of friends who have already done so [122].

Wikipedia article defines a focus — an editor is associated with a focus corresponding to a particular article if he or she has edited the article. Thus, the plot in Figure 4.12 shows the probability a person edits a Wikipedia article as a function of the number of prior editors with whom he or she has communicated.

As with triadic and focal closure, the probabilities in both Figure 4.11 and 4.12 increase with the number k of common neighbors — representing friends associated with the foci. The marginal effect diminishes as the number of friends increases, but the effect of subsequent friends remains significant. Moreover, in both sources of data, there is an initial *increasing* effect similar to what we saw with triadic closure: in this case, the probability of joining a LiveJournal community or editing a Wikipedia article is more than twice as great when you have two connections into the focus rather than one. In other words, the connection to a second person in the focus has a particularly pronounced effect, and after this the diminishing marginal effect of connections to further people takes over.

Of course, multiple effects can operate simultaneously on the formation of a single link. For example, if we consider the example in Figure 4.8, triadic closure makes a link between Bob and Daniel more likely due to their shared friendship with Anna; and focal closure also makes this link more likely due to the shared membership of Bob and Daniel in the karate club. If a link does form between them, it will not necessarily be *a priori* clear how to attribute it to these two distinct effects. This is also a reflection of an issue we discussed

in Section 4.1, when describing some of the mechanisms behind triadic closure: since the principle of homophily suggests that friends tend to have many characteristics in common, the existence of a shared friend between two people is often indicative of other, possibly unobserved, sources of similarity (such as shared foci in this case) that by themselves may also make link formation more likely.

Quantifying the Interplay Between Selection and Social Influence. As a final illustration of how we can use large-scale on-line data to track processes of link formation, let's return to the question of how selection and social influence work together to produce homophily, considered in Section 4.2. We'll make use of the Wikipedia data discussed earlier in this section, asking: how do similarities in behavior between two Wikipedia editors relate to their pattern of social interaction over time? [122]

To make this question precise, we need to define both the social network and an underlying measure of behavioral similarity. As before, the social network will consist of all Wikipedia editors who maintain talk pages, and there is an edge connecting two editors if they have communicated, with one writing on the talk page of the other. An editor's behavior will correspond to the set of articles she has edited. There are a number of natural ways to define numerical measures of similarity between two editors based on their actions; a simple one is to declare their similarity to be the value of the ratio

$$\frac{\text{number of articles edited by both } A \text{ and } B}{\text{number of articles edited by at least one of } A \text{ or } B}, \quad (4.1)$$

For example, if editor A has edited the Wikipedia articles on *Ithaca NY* and *Cornell University*, and editor B has edited the articles on *Cornell University* and *Stanford University*, then their similarity under this measure is $1/3$, since they have jointly edited one article (*Cornell*) out of three that they have edited in total (*Cornell*, *Ithaca*, and *Stanford*). Note the close similarity to the definition of *neighborhood overlap* used in Section 3.3; indeed, the measure in Equation (4.1) is precisely the neighborhood overlap of two editors in the bipartite affiliation network of editors and articles, consisting only of edges from editors to the articles they've edited.³

Pairs of Wikipedia editors who have communicated are significantly more similar in their behavior than pairs of Wikipedia editors who have not communicated, so we have a case where homophily is clearly present. Therefore, we are set up to address the question of selection and social influence: is the homophily arising because editors are forming connections with those who have edited the same articles they have (selection), or is it because editors are led to the articles of those they talk to (social influence)?

³For technical reasons, a minor variation on this simple similarity measure is used for the results that follow. However, since this variation is more complicated to describe, and the differences are not significant for our purposes, we can think of similarity as consisting of the numerical measure just defined.

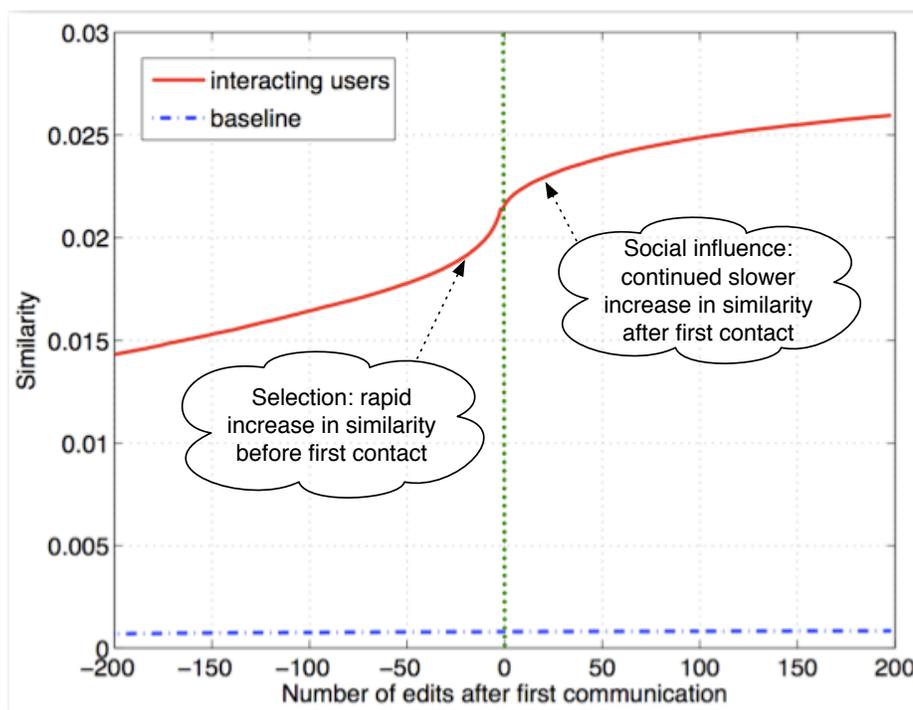


Figure 4.13: The average similarity of two editors on Wikipedia, relative to the time (0) at which they first communicated [122]. Time, on the x -axis, is measured in discrete units, where each unit corresponds to a single Wikipedia action taken by either of the two editors. The curve increases both before and after the first contact at time 0, indicating that both selection and social influence play a role; the increase in similarity is steepest just before time 0.

Because every action on Wikipedia is recorded and time-stamped, it is not hard to get an initial picture of this interplay, using the following method. For each pair of editors A and B who have ever communicated, record their similarity over time, where “time” here moves in discrete units, advancing by one “tick” whenever either A or B performs an action on Wikipedia (editing an article or communicating with another editor). Next, declare *time* 0 for the pair A - B to be the point at which they first communicated. This results in many curves showing similarity as a function of time — one for each pair of editors who ever communicated, and each curve shifted so that time is measured for each one relative to the moment of first communication. Averaging all these curves yields the single plot in Figure 4.13 — it shows the average level of similarity relative to the time of first interaction, over all pairs of editors who have ever interacted on Wikipedia [122].

There are a number of things to notice about this plot. First, similarity is clearly increasing both before and after the moment of first interaction, indicating that both selection and

social influence are at work. However, the the curve is not symmetric around time 0; the period of fastest increase in similarity is clearly occurring before 0, indicating a particular role for selection: there is an especially rapid rise in similarity, on average, just before two editors meet.⁴ Also note that the levels of similarity depicted in the plot are much higher than for pairs of editors who have not interacted: the dashed blue line at the bottom of the plot shows similarity over time for a random sample of non-interacting pairs; it is both far lower and also essentially constant as time moves forward.

At a higher level, the plot in Figure 4.13 once again illustrates the trade-offs involved in working with large-scale on-line data. On the one hand, the curve is remarkably smooth, because so many pairs are being averaged, and so differences between selection and social influence show up that are genuine, but too subtle to be noticeable at smaller scales. On the other hand, the effect being observed is an aggregate one: it is the average of the interaction histories of many different pairs of individuals, and it does not provide more detailed insight into the experience of any one particular pair.⁵ A goal for further research is clearly to find ways of formulating more complex, nuanced questions that can still be meaningfully addressed on large datasets.

Overall, then, these analyses represent early attempts to quantify some of the basic mechanisms of link formation at a very large scale, using on-line data. While they are promising in revealing that the basic patterns indeed show up strongly in the data, they raise many further questions. In particular, it natural to ask whether the general shapes of the curves in Figures 4.9–4.13 are similar across different domains — including domains that are less technologically mediated — and whether these curve shapes can be explained at a simpler level by more basic underlying social mechanisms.

4.5 A Spatial Model of Segregation

One of the most readily perceived effects of homophily is in the formation of ethnically and racially homogeneous neighborhoods in cities. Traveling through a metropolitan area, one finds that homophily produces a natural spatial signature; people live near others like them, and as a consequence they open shops, restaurants, and other businesses oriented toward the populations of their respective neighborhoods. The effect is also striking when superimposed on a map, as Figure 4.14 by Möbius and Rosenblat [302] illustrates. Their images depict the

⁴To make sure that these are editors with significant histories on Wikipedia, this plot is constructed using only pairs of editors who each had at least 100 actions both before and after their first interaction with each other.

⁵Because the individual histories being averaged took place at many distinct points in Wikipedia's history, it is also natural to ask whether the aggregate effects operated differently in different phases of this history. This is a natural question for further investigation, but initial tests — based on studying these types of properties on Wikipedia datasets built from different periods — show that the main effects have remained relatively stable over time.

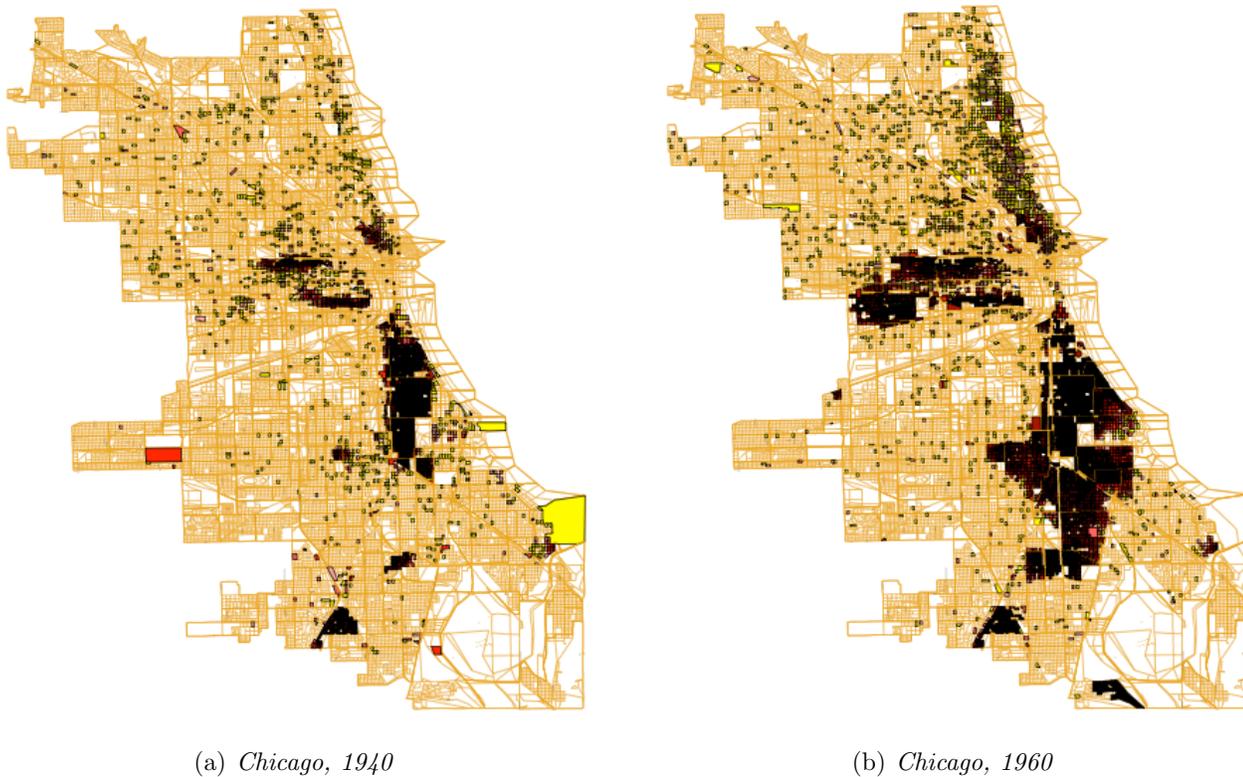


Figure 4.14: The tendency of people to live in racially homogeneous neighborhoods produces spatial patterns of segregation that are apparent both in everyday life and when superimposed on a map — as here, in these maps of Chicago from 1940 and 1960 [302]. In blocks colored yellow and orange the percentage of African-Americans is below 25, while in blocks colored brown and black the percentage is above 75.

percentage of African-Americans per city block in Chicago for the years 1940 and 1960; in blocks colored yellow and orange the percentage is below 25, while in blocks colored brown and black the percentage is above 75.

This pair of figures also shows how concentrations of different groups can intensify over time, emphasizing that this is a process with a dynamic aspect. Using the principles we’ve been considering, we now discuss how simple mechanisms based on similarity and selection can provide insight into the observed patterns and their dynamics.

The Schelling Model. A famous model due to Thomas Schelling [365, 366] shows how global patterns of spatial segregation can arise from the effect of homophily operating at a local level. There are many factors that contribute to segregation in real life, but Schelling’s model focuses on an intentionally simplified mechanism to illustrate how the forces leading to segregation are remarkably robust — they can operate even when no one individual explicitly wants a segregated outcome.

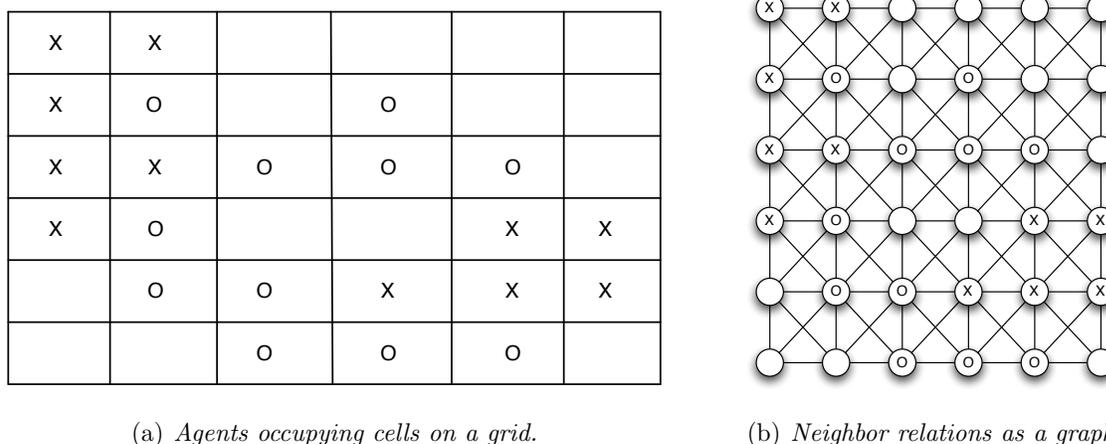


Figure 4.15: In Schelling’s segregation model, agents of two different types (X and O) occupy cells on a grid. The neighbor relationships among the cells can be represented very simply as a graph. Agents care about whether they have at least some neighbors of the same type.

The general formulation of the model is as follows. We assume that there is a population of individuals, whom we’ll call *agents*; each agent is of type X or type O. We think of the two types as representing some (immutable) characteristic that can serve as the basis for homophily — for example, race, ethnicity, country of origin, or native language. The agents reside in the cells of a grid, intended as a stylized model of the two-dimensional geography of a city. As illustrated in Figure 4.15(a), we will assume that some cells of the grid contain agents while others are unpopulated. A cell’s *neighbors* are the cells that touch it, including diagonal contact; thus, a cell that is not on the boundary of the grid has eight neighbors. We can equivalently think of the neighbor relationships as defining a graph: the cells are the nodes, and we put an edge between two cells that are neighbors on the grid. In this view, the agents thus occupy the nodes of a graph that are arranged in this grid-like pattern, as shown in Figure 4.15(b). For ease of visualization, however, we will continue to draw things using a geometric grid, rather than a graph.

The fundamental constraint driving the model is that each agent wants to have at least some other agents of its own type as neighbors. We will assume that there is a *threshold* t common to all agents: if an agent discovers that fewer than t of its neighbors are of the same type as itself, then it has an interest in moving to a new cell. We will call such an agent *unsatisfied* with its current location. For example, in Figure 4.16(a), we indicate with an asterisk all the agents that are unsatisfied in the arrangement from Figure 4.15(a), when the threshold t is equal to 3. (In Figure 4.16(a) we have also added a number after each agent. This is simply to provide each with a unique name; the key distinction is still whether each agent is of type X or type O.)

X1*	X2*				
X3	O1*		O2		
X4	X5	O3	O4	O5*	
X6*	O6			X7	X8
	O7	O8	X9*	X10	X11
		O9	O10	O11*	

(a) *An initial configuration.*

X3	X6	O1	O2		
X4	X5	O3	O4		
	O6	X2	X1	X7	X8
O11	O7	O8	X9	X10	X11
	O5	O9	O10*		

(b) *After one round of movement.*

Figure 4.16: After arranging agents in cells of the grid, we first determine which agents are *unsatisfied*, with fewer than t other agents of the same type as neighbors. In one round, each of these agents moves to a cell where they will be satisfied; this may cause other agents to become unsatisfied, in which case a new round of movement begins.

The Dynamics of Movement. Thus far, we have simply specified a set of agents that want to move, given an underlying threshold; we now discuss how this gives the model its dynamic aspect. Agents move in a sequence of *rounds*: in each round, we consider the unsatisfied agents in some order, and for each one in turn, we have it move to an unoccupied cell where it will be satisfied. After this, the round of movement has come to an end, representing a fixed time period during which unsatisfied agents have changed where they live. These new locations may cause different agents to be unsatisfied, and this leads to a new round of movement.

In the literature on this model, there are numerous variations in the specific details of how the movement of agents within a round is handled. For example, the agents can be scheduled to move in a random order, or in an order that sweeps downward along rows of the grid; they can move to the nearest location that will make them satisfied or to a random one. There also needs to be a way of handling situations in which an agent is scheduled to move, and there is *no* cell that will make it satisfied. In such a case, the agent can be left where it is, or moved to a completely random cell. Research has found that the qualitative results of the model tend to be quite similar however these issues are resolved, and different investigations of the model have tended to resolve them differently.

For example, Figure 4.16(b) shows the results of one round of movement, starting from the arrangement in Figure 4.16(a), when the threshold t is 3. Unsatisfied agents are scheduled to move by considering them one row at a time working downward through the grid, and each agent moves to the nearest cell that will make it satisfied. (The unique name of each agent in the figure allows us to see where it has moved in Figure 4.16(b), relative to the initial state in Figure 4.16(a).) Notice that in some concrete respects, the pattern of agents has become more “segregated” after this round of movement. For example, in Figure 4.16(a), there is only a single agent with no neighbors of the opposite type. After this first round of movement, however, there are six agents in Figure 4.16(b) with no neighbors of the opposite type. As we will see, this increasing level of segregation is the key behavior to emerge from the model.

Larger examples. Small examples of the type in Figures 4.15 and 4.16 are helpful in working through the details of the model by hand; but at such small scales it is difficult to see the kinds of typical patterns that arise. For this, computer simulation is very useful.

There are many on-line computer programs that make it possible to simulate the Schelling model; as with the published literature on the model, they all tend to differ slightly from each other in their specifics. Here we discuss some examples from a simulation written by Sean Luke [282], which is like the version of the model we have discussed thus far except that unsatisfied agents move to a random location.

In Figure 4.17, we show the results of simulating the model on a grid with 150 rows and

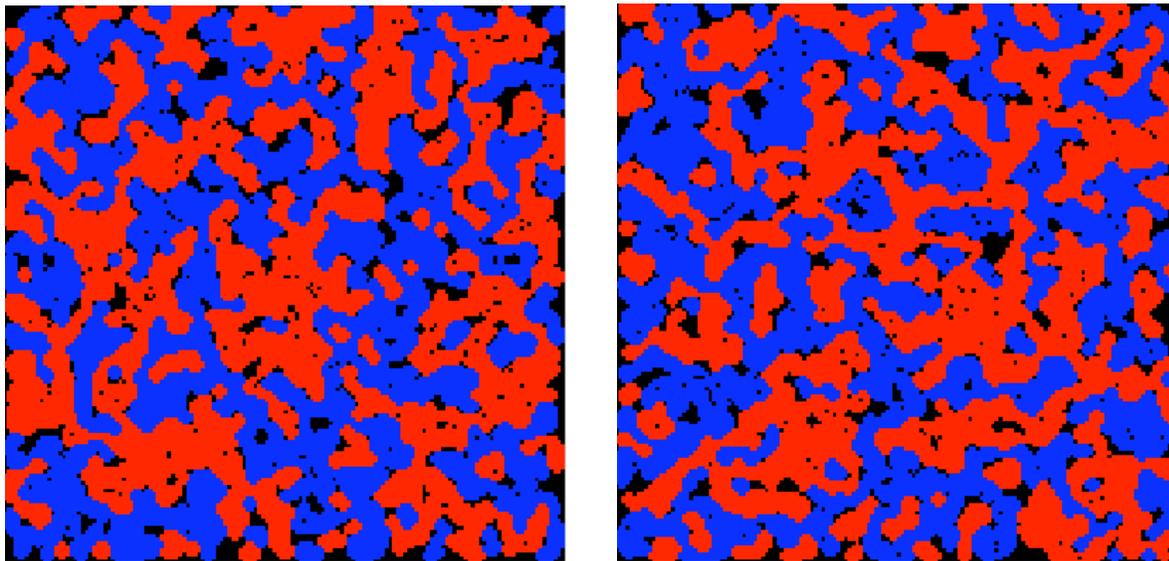
(a) *A simulation with threshold 3.*(b) *Another simulation with threshold 3.*

Figure 4.17: Two runs of a simulation of the Schelling model with a threshold t of 3, on a 150-by-150 grid with 10,000 agents of each type. Each cell of the grid is colored red if it is occupied by an agent of the first type, blue if it is occupied by an agent of the second type, and black if it is empty (not occupied by any agent).

150 columns, 10,000 agents of each type, and 2500 empty cells. The threshold t is equal to 3, as in our earlier examples. The two images depict the results of two different runs of the simulation, with different random starting patterns of agents. In each case, the simulation reached a point (shown in the figures) at which all agents were satisfied, after roughly 50 rounds of movement.

Because of the different random starts, the final arrangement of agents is different in the two cases, but the qualitative similarities reflect the fundamental consequences of the model. By seeking out locations near other agents of the same type, the model produces large homogeneous regions, interlocking with each other as they stretch across the grid. In the midst of these regions are large numbers of agents who are surrounded on all sides by other agents of the same type — and in fact at some distance from the nearest agent of the opposite type. The geometric pattern has become segregated, much as in the maps of Chicago from Figure 4.14 with which we began the section.

Interpretations of the Model. We've now seen how the model works, what it looks like at relatively large scales, and how it produces spatially segregated outcomes. But what broader insights into homophily and segregation does it suggest?

The first and most basic one is that spatial segregation is taking place even though no

X	X	O	O	X	X
X	X	O	O	X	X
O	O	X	X	O	O
O	O	X	X	O	O
X	X	O	O	X	X
X	X	O	O	X	X

Figure 4.18: With a threshold of 3, it is possible to arrange agents in an integrated pattern: all agents are satisfied, and everyone who is not on the boundary on the grid has an equal number of neighbors of each type.

individual agent is actively seeking it. Sticking to our focus on a threshold t of 3, we see that although agents want to be near others like them, their requirements are not particularly draconian. For example, an agent would be perfectly happy to be in the minority among its neighbors, with five neighbors of the opposite type and three of its own type. Nor are the requirements globally incompatible with complete integration of the population. By arranging agents in a checkerboard pattern as shown in Figure 4.18, we can make each agent satisfied, and all agents not on the boundary of the grid have exactly four neighbors of each type. This is a pattern that we can continue on as large a grid as we want.

Thus, segregation is not happening because we have subtly built it into the model — agents are willing to be in the minority, and they could all be satisfied if we were only able to carefully arrange them in an integrated pattern. The problem is that from a random start, it is very hard for the collection of agents to find such integrated patterns. Much more typically, agents will attach themselves to clusters of others like themselves, and these clusters will grow as other agents follow suit. Moreover, there is a compounding effect as the rounds of movement unfold, in which agents who fall below their threshold depart for more homogeneous parts of the grid, causing previously satisfied agents to fall below their thresholds and move as well — an effect that Schelling describes as the progressive “unraveling” of more integrated regions [366]. In the long run, this process will tend to cause segregated regions to grow at the expense of more integrated ones. The overall effect is one in which the local preferences of individual agents have produced a global pattern that none of them necessarily intended.

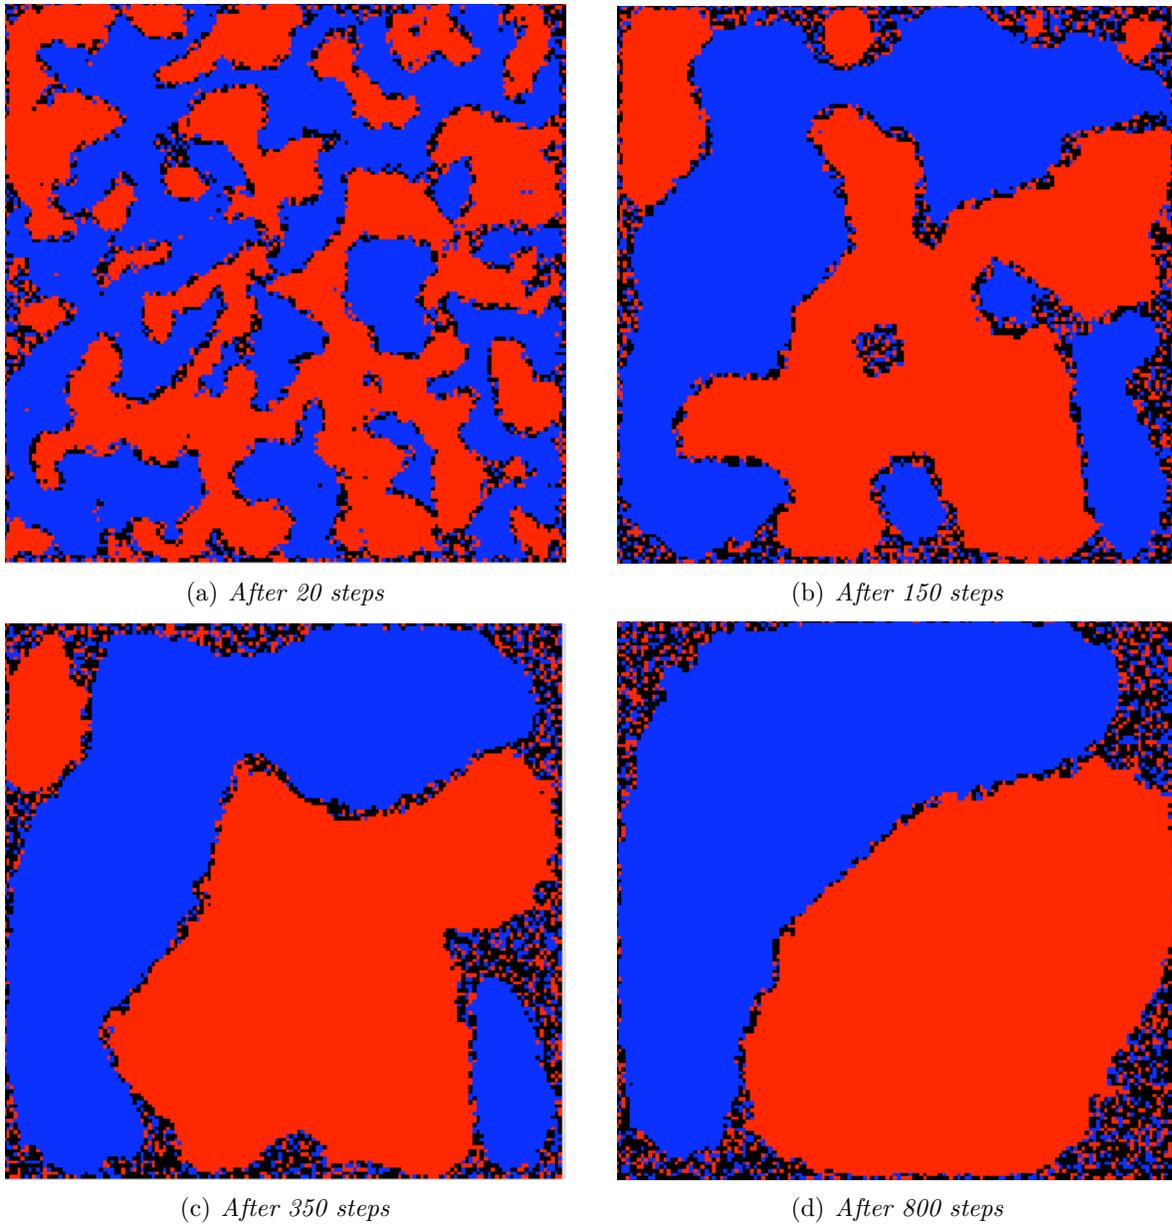


Figure 4.19: Four intermediate points in a simulation of the Schelling model with a threshold t of 4, on a 150-by-150 grid with 10,000 agents of each type. As the rounds of movement progress, large homogeneous regions on the grid grow at the expense of smaller, narrower regions.

This point is ultimately at the heart of the model: although segregation in real life is amplified by a genuine desire within some fraction of the population to belong to large clusters of similar people — either to avoid people who belong to other groups, or to acquire a critical mass of members from one’s own group — we see here that such factors are not necessary for segregation to occur. The underpinnings of segregation are already present in a system where individuals simply want to avoid being in too extreme a minority in their own local area.

The process operates even more powerfully when we raise the threshold t in our examples from 3 to 4. Even with a threshold of 4, nodes are willing to have an equal number of neighbors of each type; and a slightly more elaborate checkerboard example in the spirit of Figure 4.18 shows that with careful placement, the agents can be arranged so that all are satisfied and most still have a significant number of neighbors of the opposite type. But now, not only is an integrated pattern very hard to reach from a random starting arrangement — any vestiges of integration among the two types tends to collapse completely over time. As one example of this, Figure 4.19 shows four intermediate points in one run of a simulation with threshold 4 and other properties the same as before (a 150-by-150 grid with 10,000 agents of each type and random movement by unsatisfied agents) [282]. Figure 4.19(a) shows that after 20 rounds of movement, we have an arrangement of agents that roughly resembles what we saw with a lower threshold of 3. However, this does not last long: crucially, the long tendrils where one type interlocks with the other quickly wither and retract, leaving the more homogeneous regions shown after 150 rounds in Figure 4.19(b). This pulling-back continues, passing through a phase with a large and small region of each type after 350 rounds (Figure 4.19(c)) eventually to a point where there is only a single significant region of each type, after roughly 800 rounds (Figure 4.19(d)). Note that this is not the end of the process, since there remain agents around the edges still looking for places to move, but by this point the overall two-region layout has become very stable. Finally, we stress that this figure corresponds to just a single run of the simulation — but computational experiments show that the sequence of events it depicts, leading to almost complete separation of the two types, is very robust when the threshold is this high.

Viewed at a still more general level, the Schelling model is an example of how characteristics that are fixed and unchanging (such as race or ethnicity) can become highly correlated with other characteristics that are mutable. In this case, the mutable characteristic is the decision about where to live, which over time conforms to similarities in the agents’ (immutable) types, producing segregation. But there are other, non-spatial manifestation of the same effect, in which beliefs and opinions become correlated across racial or ethnic lines, and for similar underlying reasons: as homophily draws people together along immutable characteristics, there is a natural tendency for mutable characteristics to change in accordance with the network structure.

As a final point, we note that while the model is mathematically precise and self-contained, the discussion has been carried out in terms of simulations and qualitative observations. This is because rigorous mathematical analysis of the Schelling model appears to be quite difficult, and is largely an open research question. For partial progress on analyzing properties of the Schelling model, see the work of Young [420], who compares properties of different arrangements in which all agents are satisfied; Möbius and Rosenblat [302], who perform a probabilistic analysis; and Vinković and Kirman [401], who develop analogies to models for the mixing of two liquids and other physical phenomena.

4.6 Exercises

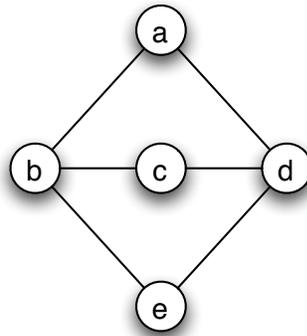


Figure 4.20: A social network where triadic closure may occur.

1. Consider the social network represented in Figure 4.20. Suppose that this social network was obtained by observing a group of people at a particular point in time and recording all their friendship relations. Now suppose that we come back at some point in the future and observe it again. According to the theories based on empirical studies of triadic closure in networks, which new edge is most likely to be present? (I.e. which pair of nodes, who do not currently have an edge connecting them, are most likely to be linked by an edge when we return to take the second observation?) Also, give a brief explanation for your answer.
2. Given a bipartite affiliation graph, showing the membership of people in different social foci, researchers sometimes create a *projected graph* on just the people, in which we join two people when they have a focus in common.
 - (a) Draw what such a projected graph would look like for the example of memberships on corporate boards of directors from Figure 4.4. Here the nodes would be the

seven people in the figure, and there would be an edge joining any two who serve on a board of directors together.

- (b) Give an example of two different affiliation networks — on the same set of people, but with different foci — so that the projected graphs from these two different affiliation networks are the same. This shows how information can be “lost” when moving from the full affiliation network to just the projected graph on the set of people.

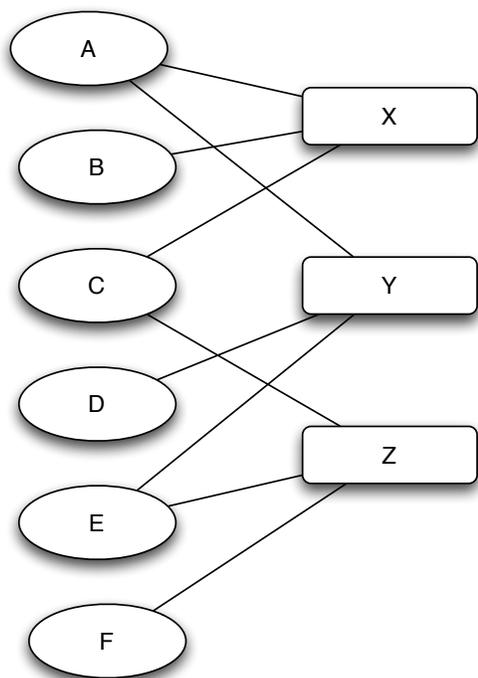


Figure 4.21: An affiliation network on six people labeled A – F , and three foci labeled X , Y , and Z .

3. Consider the affiliation network in Figure 4.21, with six people labeled A – F , and three foci labeled X , Y , and Z .
 - (a) Draw the derived network on just the six people as in Exercise 2, joining two people when they share a focus.
 - (b) In the resulting network on people, can you identify a sense in which the triangle on the nodes A , C , and E has a qualitatively different meaning than the other triangles that appear in the network? Explain.

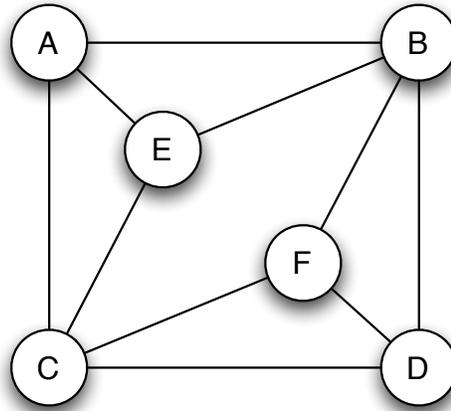


Figure 4.22: A graph on people arising from an (unobserved) affiliation network.

4. Given a network showing pairs of people who share activities, we can try to reconstruct an affiliation network consistent with this data.

For example, suppose that you are trying to infer the structure of a bipartite affiliation network, and by indirect observation you've obtained the projected network on just the set of people, constructed as in Exercise 2: there is an edge joining each pair of people who share a focus. This projected network is shown in Figure 4.22.

- (a) Draw an affiliation network involving these six people, together with four foci that you should define, whose projected network is the graph shown in Figure 4.22.
- (b) Explain why any affiliation network capable of producing the projected network in Figure 4.22 must have at least four foci.