Availability of Better Data on Communication Networks can Undermine Community Enforcement

James D. Campbell
Providence College

January 22, 2018

Abstract

I consider a repeated trust game played between a group of insiders, arranged in a network representing their lines of communication, and a single outsider. Insiders follow a local punishment rule, shunning the outsider if it has cheated them or a neighbor. The object of interest is that the outsider may know either summary statistics about the nature of the network, or know its precise non-anonymous structure. For the outsider to have knowledge of the precise structure may, depending on the shape of the network, increase or decrease the volume of honest interaction that can be sustained. In extreme cases, a small ‘local’ vulnerability to outside exploitation can result in a total breakdown of the chance for mutually beneficial trade across the whole network. Strategic ignorance and obfuscation of network structure may therefore be valuable to both sides, complicating the problem for a network operator seeking to monetize data on the network graph. I discuss recent decisions by Twitter and Facebook in this framework.

Keywords: reputation, community enforcement, network data, privacy JEL: D83, D85,
L14

1 Introduction

Consider a network of insiders who interact repeatedly with an outside entity in a trust game. The outsider faces an immediate temptation to exploit the person with whom they are currently interacting, but a long term incentive to maintain good relationships with insiders. Can the threat of punishment by a person’s network neighbors give the outsider an incentive to ‘play nice’ that is strong enough to sustain honest interaction?

A key consideration is how much information the outsider has on the structure of the network. The outsider may have very precise information, able to see the exact structure of the network, including the exact location and neighborhood of individual nodes. Alternatively, it may have only general information about the network, perhaps knowing summary statistics about the level of connectivity, or knowing the exact structure but not knowing which node occupies which position. I will focus here on how the outsider values these two general types of information about the network in the context of the repeated interaction game.

My main finding is that more precise information about the network structure is not necessarily more valuable to the outsider. The reason is that more information reveals to the outsider exactly which nodes in the network are strongly protected by a large number of neighbors, and which nodes are weakly protected by a small number of neighbors. Poorly protected nodes that may have been masked by high average connectivity in the network are revealed to be vulnerable when the precise structure of the network is revealed.

The volume of ‘honest’ interaction that can be sustained in the repeated interaction game may, depending on the network’s structure, be increased or decreased by the outsider possessing more precise information. Under the regime of the outsider having either the summary statistic or anonymized information, either all nodes are vulnerable to exploitation or none are. In extreme cases, networks in which no nodes were vulnerable under these regimes can have *all* honest interaction destroyed by the outsider obtaining precise network information. These cases are characterized (informally) by a lack of redundant paths and a lack of cliques. The converse is not complete: networks in which all nodes were vulnerable under the regimes with less information can have at best *some*, but not all, honest interaction restored by the outsider obtaining precise network information.

The mechanism here is that local sparseness in a network can lead to global vulnerability. Consider a network in which individuals have a high number of connections on average, but with wide variation in the number of connections per node. Irredundancies—informally speaking, insiders with no or few mutual friends—in one area of such a network can cause a recursive chain reaction of susceptibility to exploitation that can ultimately result in no mutually beneficial interaction taking place. In the opposite direction, local sub-networks that are well-connected among themselves

can be salvaged as secure if their area is revealed by precise network data to be more cyclically connected than the whole-network average.

The general idea of this setup could apply to different types of interaction. A few examples:

Example 1: unverifiable service

In many buyer-seller relationships it is difficult for buyers to concretely verify the quality of the work done or the good provided. The implied cost of bad word-of-mouth or negative online reviews may or may not be enough to induce the seller to provide high quality service. If the seller possesses information on the extent of the buyer's influence, her incentives to provide high quality may now vary depending on the identity of the buyer.

Example 2: equitable treatment

An organization provides customer service. If it has information on the position of individuals in information-sharing networks, it may face a new incentive to neglect or prioritize users based on their visibility to others, in order to promote the appearance of responsiveness. The operating rules for the information-sharing network influence the extent to which the dialogue between the organization and a user is transmitted to others.

Example 3: investigative journalism

Investigative consumer advocacy by local news teams is a classic example of the kind of incentives at play in settings like these. Reluctant businesses are forced to confront complaints after being featured on a news program with wide viewership, where localized 'bad press' from an individual was not enough. This example raises the questions of what real-world institutions may arise in response to the type of problem we study here and how effective they may be.

The idea of word-of-mouth as a punishment mechanism has been widely explored in the literature on community enforcement in both many-to-many matching and one-to-many matching repeated games, in which high quality trade is sustained by some strategy for consumers that takes into account previous experience of others. For example, Klein and Leffler (1981) assumes that past play by a firm is public knowledge. Kandori (1992) analyzes cases in which players are randomly matched in each stage and, respectively, know only their own history or know some reputational 'label' attached to each player. Similarly, Okuno-Fujiwara and Postlewaite (1995) allows the 'status' of an individual's matched player to be common knowledge in the game. Ellison (1994) studies a similar random matching setting in which players are anonymous and observe only the play in their own past games.

There is a significant related literature on social punishment in games with repeated bilateral interaction among network insiders. For example, Jackson et al. (2012) analyzes social pressure in a network structure as a means to sustain exchange. Their focus is on exchange structures that are renegotiation-proof due to the protection of mutual friends of parties to a transaction. Their framework fits slightly different applications to the one I will present here, since theirs considers

interactions between members of the network while I consider outsider-insider interactions. This means that the concepts of redundancy and isolation in the network are slightly different, and it permits me to focus on the value of different types of network data to the outsider.

A second example is Ali and Miller (2013), which considers the sustainability of bilateral partnerships among insiders in a network in the presence of two-sided moral hazard. Players can observe the history of play in their own relationships, but cannot observe how their matched partners have played in other relationships. In common with the outcomes in this paper, networks featuring cliques are found to be good for cooperation and payoffs. The recursive diffusion of information is a key consideration in their analysis, whereas I will assume skepticism from insiders so that they never engage in punishment behavior unless and until the outsider cheats a direct neighbor.

Closest to the model I present below is Ahn and Suominen (2001). Their model has an intermediate assumption on information transmission: a single seller transacts with many buyers in an infinitely repeated game, in which one buyer is matched with the seller in each period and some subset of the remaining buyers are selected to be ‘spectators’ to the transaction. Past play and reputations are neither fully public or fully private, but instead information dissipates by successive observation by groups of other players. Their model demonstrates that high quality trade can be supported in every period for a suitably large buyer population, even when the probability of each buyer observing the seller’s choice in a given period is arbitrarily low. This result depends on the assumption that the spectators to each transaction are randomly selected and not determined by the identity of the buyer. This is qualitatively equivalent to the seller not knowing the identity or characteristics of each buyer. An important addition that I make here is to consider selection of spectators according to a pre-existing communication network.

In sum, I will focus here on an outsider-insiders structure and the valuation and strategic implications of the availability of public or tradeable data on the precise structure of networks. In particular I will consider the strategic incentive of the outsider to acquire or ignore precise network data, the strategic incentive of network operators to publicize precise network data, and the social value of precise network data.

2 Repeated game between insiders and outsider

A set of ‘insiders’ interacts with a single outsider over an infinite horizon in discrete time. Denote the number of these insiders by n . The insiders are arranged in a network capturing their lines of communication with each other, in a fashion we will specify shortly. At each date, a single insider is selected at random to interact with the outsider in a stage game. It follows that the probability

of a given insider being selected at some time is $\frac{1}{n}$.

We may interpret the random selection mechanism as a ‘need’ for a product or service arising in the population regularly but by chance, for example a car breaking down. Notice that we therefore do not explicitly match to settings in which needs arise endogenously, and in particular to the setting in which an insider’s propensity to engage the outsider depends on their prior knowledge, who their neighbors are, or the history of the game. Similarly we do not deal with the case in which the outsider strategically chooses who to visit when.¹

The stage game played between insider and outsider has a strategy space and payoffs that are described in the following matrix representation:

		Insider	
		Engage	Don’t engage
Outsider	Honest	1,1	0,0
	Dishonest	1+g,-b	0,0

The outsider discounts future payoffs at the $\delta < 1$, so that their total payoff is given by the discounted sum of their stage payoffs, $\sum_t \delta^t \pi_t$.

These payoffs are constructed to have two key features. First, the insider can effectively ‘shun’ the outsider by choosing not to engage the outsider, in which case both parties earn a payoff of zero. Second, if engaged, the outsider faces a short-term temptation to behave dishonestly that we shall contrast with their long-term incentives. There are two unspecified parameters. $g > 0$ is the ‘extra’ payoff to the outsider for dishonesty as compared to honesty in the one-shot game. $b > 0$ is the loss to a swindled insider as compared to never having engaged the outsider in the first place. However, since our analysis is qualitative, the relative magnitude of these parameters does not play a big role.

Dishonesty is a weakly dominant strategy for the outsider in the stage game. The Nash equilibria of the stage game have no engagement: the insider chooses not to engage, and the outsider chooses dishonesty with sufficiently high probability. The object of our interest is the possibility of subgame perfect Nash equilibria in the repeated game in which all insiders—or, as we shall see, as many as possible—choose to engage the outsider whenever they are selected and are treated honestly by the outsider.

An insider’s *neighbors* are those other insiders who are precisely one degree away in the network. A concept that will be important for our analysis is a *neighborhood* for some insider i , which we define as the insider themselves plus their neighbors. Let us denote the number of insiders in the neighborhood of i by x_i . Following naturally from this, we will denote the average neighborhood size in the network of insiders by $\bar{x}$.

¹This gives rise to the possibility that the order of visitation could signal something to insiders, an idea discussed in, for example, Campbell (2015).

Assume that insiders use a *local punishment rule* that restricts them to the so-called ‘grim strategy’ in the case in which they are aware that a network neighbor was treated dishonestly in the past:

Definition 1. Local punishment rule: *if in any prior period*

- (i) *someone in an insider’s neighborhood was selected to play the stage game,*
- (ii) *that person engaged the outsider, and*
- (iii) *the outsider behaved dishonestly,*

then the insider shall play ‘don’t engage’ whenever they are selected to play the stage game.

This rule has several noteworthy characteristics. First, it implies that either knowledge of deviations does not travel beyond neighbors, or that insiders do not punish dishonest treatment of those who are not their neighbors. Punishment behavior is not contagious (this contrasts with the focus of the majority of the prior literature discussed earlier, which concerns the percolation of information and inference through the network). Second, it is not necessarily sequentially rational, in the sense that there is no consideration by the insider of any possible continuation game in which they may receive a higher payoff by *not* following the local punishment rule. Our conception of punishment can therefore be interpreted as both pro-social (or, if you prefer, altruistic) and also conservative in its treatment of hearsay.

Our local punishment rule, when combined with a strategy profile in which all insiders always choose engage and the outsider always chooses to be honest, comprises the ‘unforgiving strategy profile’ in Ahn and Suominen (2001). The question that we will focus on is: to what extent does the local punishment rule insulate insiders against being treated dishonestly by the outsider?

We may briefly note that there are surely other modeling approaches that would capture similar forces to the one at play here. One alternative, for example, could be to view network position as reflecting one’s outside option in a bargaining situation, with implications for the disagreement payoff. The outcome of a bargaining process may then change depending on whether this network data is known to the counter-party.

2.1 Continuation payoffs and putting a bound on honest trade

The problem we analyze here is familiar from generic reputation games: cheat now versus future cooperation. In checking whether strategies to form an equilibrium in this game, we must check the outsider’s *local* continuation payoff in the event that they behave honestly against the one-shot

game from deviating. The ‘local’ here refers to the restriction on insiders’ strategies from Definition 1: the repercussions from dishonesty are felt only in the neighborhood of the victim.

As in all repeated games, the continuation payoff and individual strategies could in general be heavily path dependent in an equilibrium. We are taking a typical approach to this problem by seeking to describe an upper bound on the amount of honest interaction. In this we are assisted by the observation that, by forward induction, in equilibrium no continuation payoff can include a payoff to the outsider of more than 1 in any future period t . That is, the outsider’s strategy for the surrounding area cannot include being honest now while ‘waiting’ to defect against another insider in the neighborhood, since in equilibrium that subsequent insider will not buy. This means that the maximal local continuation payoff to the firm is that associated with perpetually honest transactions with the selected insider and all of their neighbors.

2.2 The outsider knows average connectivity

First consider a situation in which the outsider knows a summary statistic about the network: the average number of neighbors across all insiders. A situation in which the outsider knows broadly similar summary statistics, for example density, would be qualitatively similar to this one. So would a situation in which the outsider knew the precise structure of the network but could not tell which insider was which when they are matched in the stage game. The crucial point is that the outsider cannot meaningfully distinguish insiders from each other when they are randomly selected to play the stage game.

In this situation the *expected* cost of dishonesty to the outsider is losing the possibility of any positive payoff in the future in the neighborhood of the currently selected insider. The probability of someone in that neighborhood being selected at a given future date is $\frac{\bar{x}}{n}$, the size of the neighborhood relative to the total population.

An equilibrium in which all insiders engage and the outsider behaves honestly in each period exists only if

$$\frac{1}{1-\delta} \left(\frac{\bar{x}}{n} \right) > 1 + g \quad (1)$$

$$\delta > 1 - \left(\frac{\bar{x}}{n} \right) \frac{1}{1+g} \quad (2)$$

This is in precisely the same spirit as the result from Ahn and Suominen (2001), in which high quality trade can be sustained indefinitely as long as there are a sufficiently large number of spectators per transaction. Since the outsider is conditioning on the *average* neighborhood size here, there is no distinction between spectators chosen at random each period, as in Ahn and Suominen (2001), or fixed but unknown to the outsider, as in the present framework.

The intuition for this result is the well-known ‘shadow of the future’ (Dal Bó, 2005). If the expected loss from ‘cheating’ an insider today exceeds the one-shot gain from doing so, then if all insiders choose to engage the outsider in each period, the outsider’s best response is to cooperate throughout.

2.3 The outsider knows insiders’ position in the network

Next consider a situation in which the outsider has fully personalized data on the structure of the network and can identify each insider by their position in it. The outsider can now condition their behavior on the location of the matched insider selected to play the stage game, and different neighborhood sizes mean different intensity of punishment under the local enforcement rule.

Denote by x_1 the size of the neighborhood of the insider with the smallest number of neighbors. Since the discrepancies in continuation payoffs across neighborhoods are now known to the outsider, we can observe that there exists an equilibrium with perpetual full engagement and honest trade if and only if

$$\frac{1}{1-\delta} \left(\frac{x_1}{n} \right) > 1 + g \quad (3)$$

$$\delta \geq 1 - \left(\frac{x_1}{n} \right) \frac{1}{1+g}. \quad (4)$$

This is a ‘weakest link’ argument. Given the local enforcement rule, it is not credible that the outsider would choose to play honestly if engaged by the insider with neighborhood x_1 . Their social protection is not sufficient to outweigh the outsider’s one-shot gain from dishonesty. In a network with a nondegenerate degree distribution, perpetual engagement and honest interaction requires a strictly more patient outsider when the outsider has full information on the network position of each insider.

We can develop the reasoning from Section 2.1 to identify the full set of insiders who are *vulnerable* in this sense. The maximal continuation payoff in the neighborhood of i is given by the discounted stream of future payoffs of 1 in that neighborhood, $\frac{x_i}{n} \frac{1}{1-\delta}$. It is therefore unambiguously better for the outsider to be dishonest with insider i if

$$1 + g > \left(\frac{x_i}{n} \right) \frac{1}{1-\delta} \quad (5)$$

This is just an inversion of condition 4 to identify insiders who will not engage the outsider in equilibrium, since they will never be treated honestly.

But this is recursive. Consider the algorithm to construct a *vulnerable set* V of insiders from the original network graph g :

1. Consider the insider j with the smallest neighborhood, x_1 .
2. Check condition 5. If it is not satisfied, end, else go to 3.
3. Put j in V . Update g to $g' = g - j$. Return to 1.

The resulting vulnerable set are those insiders who will never receive honest treatment from the outsider in equilibrium.

A very simple example of recursive vulnerability is shown in Figure 1. The three connections of the node at the center of the four-person star do not provide social protection if the neighbors are themselves vulnerable.

Figure 1: Vulnerability can be recursive

If the threat of punishment in a neighborhood of size 2 cannot incentivize the outsider to behave honestly with the peripheral insiders, then we can prune the periphery from consideration in the continuation game, leaving the central insider on an island with no meaningful social protection. This recursion has a flavor of contagion arguments, but in the outsider-insiders framework the intuition is not cutting off ties with someone who wronged you, but rather losing the protection of someone who cannot help to protect you.

We may briefly mention a feature of the vulnerable set that flows from the assumptions of the local enforcement rule. Let us say that a network has a vulnerable set composed of some insiders but not all. At the margin of the vulnerable set—assuming that the network graph is connected overall—we will have neighbors who straddle the boundary of the set. That is, one insider may not engage the outsider, knowing that they will not receive honest treatment, while their neighbor does engage the outsider, knowing that they will receive honest treatment.

The interesting aspect here is that the local enforcement rule has prescribed that if the vulnerable insider *did* ever engage—perhaps a ‘tremble’ or a misunderstanding—then the non-vulnerable insider is bound to shun the outsider for the rest of time. This is the sense in which the local enforcement rule is not sequentially rational. It therefore captures motivation outside the model, perhaps pro-social incentives to maintain friendships by sacrificing payoffs to stick up for a friend.

Absent such motivation, we may circumvent the issue by modifying the punishment rule such that non-vulnerable insiders do not punish the cheating of vulnerable insiders.

A second sense in which the punishment rule is not sequentially rational is that if any neighbor is cheated, an insider may prefer not to punish the outsider if they have sufficiently many remaining uncheated neighbors to support local cooperation. Since this does not depend on whether the cheated neighbor was vulnerable, it cannot be circumvented by the same modification of the punishment rule.

Another way to say all of this is that the equilibrium with maximal honest engagement is built on punishment threats that never come to pass. Although it is outside the scope of this model, in a different setting we may explore the implications of the insiders having to conjecture and ‘feel out’ the propensity of the outsider to cheat, for example because they don’t know the firm’s level of time preference.

For a formal analysis of these issues, in the related literature on social punishment to sustain bilateral interactions between network insiders, Bloch et al. (2008), Lippert and Spagnolo (2011), and Joshi and Mahmud (2016) all characterize stable network structures with explicit concern for sequential rationality.

2.4 How network location data affects the vulnerable set

The outsiders’s knowledge of the network structure can increase, decrease or leave unchanged the upper bound on honest engagement in the repeated game. Which of these happens depends on the exact structure of the network, and we may draw some general conclusions here. In this section we will stick to the interpretation of ‘knowledge’ of the network as capturing whether the data available to the outsider is anonymous with respect to the insider’s location or not.

We have already seen from equation 4 that it is not possible for the outsider’s knowledge of non-anonymous location data to admit a full honest engagement equilibrium if one did not exist without that data. A simple case in which the outsider knowing the location of insiders in the network matters is a three-node line (Figure 2).

Figure 2: Three-node line

In the three-node line the average neighborhood size is $\frac{7}{3}$ but the smallest neighborhood size is 2. The anonymous data thus imposes more ‘discipline’ on the outsider. There is some level of

patience such that fully honest engagement can be perpetually sustained with anonymous data but not without.

This is an example that generalizes: the network structures least conducive to fully honest engagement are trees. These structures are such that there is precisely one path between any two nodes in the network. For example, Figure 3 shows a single parent node with two children, each of which has two children of its own, and so on. If a neighborhood of size 2 is not sufficient to protect an insider, recursion will erode this tree entirely from bottom to top. In sum: if the structure of connections among insiders is a tree, the vulnerable set is either empty or includes every insider in the network. Either all are vulnerable, or none are. The same logic extends to tree segments of larger network structures. For example, if the parent node of Figure 3 was a bridge between the pictured segment and the rest of a larger network structure, then this segment would either be wholly contained in the vulnerable set or wholly outside it.

Figure 3: Trees are the structure most fragile to recursive vulnerability

Conversely, cycles and cliques—which generate local redundancy in paths between two nodes—are helpful. Figure 4 adds an edge to the four-person star.

Figure 4: Adding a cycle

In the non-anonymous world, ‘mutual friends’ are helpful to an insider. When insiders have mutual friends, the risk of a recursive vulnerability emerging in their neighborhood is reduced.

Notice too that there is an externality effect operating on an insider’s incentive to create or maintain links. With anonymous data, an insider can ‘hide’ behind average connectivity in the network. Some of the benefit to an individual of forming new connections or fostering mutual friendships is dispersed. By contrast, with non-anonymous data, the insider’s individual connectivity is exposed. Their incentive to form new connections or foster mutual friendships is therefore sharper, since the benefit flows to their own protection.

It is not always the case, though, that non-anonymous data reduces the upper bound on honest engagement.

Figure 5: Adding a cycle

The network in Figure 5 has an average neighborhood size of $\frac{11}{3}$. This is the level of discipline on the outsider from anonymous social protection. With non-anonymous data, however, the two peripheral ‘spokes’ are revealed to have neighborhoods of size 2, while each of the nodes in the completely connected subgraph in the center have neighborhoods of size 4 in the event that the peripheral spokes are removed. There is therefore some discount factor δ such that all are vulnerable

in the anonymous data case, but the central four nodes are not vulnerable in the non-anonymous data case.

The question at hand then becomes the minimum neighborhood size within subgraphs. The recursive process of identifying the vulnerable set can be stopped by a firewall of connectivity among a subset of insiders. In that sense, the problem of vulnerability here is a local one, but an inescapably collective one. This is because while protection in the sense of this game does not require a network-wide, top-down solution, we have also seen that for an individual to increase their number of neighbors is not alone enough to increase protection. Rather protection can be achieved piecemeal by actions to increase connectivity within a sufficiently populous neighborhood by forging mutual connections among a set of insiders.

How many steps away an insider has to reach to forge those mutual connections depends on how high the propensity to cheat is for the outsider. Is it enough to link friends with each other, or does the insider have to work harder to link friends of friends? The answer depends on the strength of incentives to cheat and cooperate in the stage game, and on the degree of patience by the outsider. These parameters are potentially informed by the cultural and legal setting, as well as by inherent preferences.

A general implication of the analysis of this game is that the outsider may prefer to commit *ex ante* to ignorance of the precise structure of relationship networks among insiders. By remaining in the dark, it avoids the risk that it will be cursed by the knowledge of who is vulnerable and who is not, leading to less honest engagement than before. Another way to say the same thing is that it is possible in this game for the value of non-anonymous data on insider connections to be negative for the outsider.

3 Network operators and examples

Consider online social networks as an example of an application of the model. As of 2015, 78% of Fortune 500 companies have active Twitter accounts, and 74% have Facebook pages (Barnes et al., 2016). No doubt there is a promotional aspect to this kind of activity, for example to spread information about products (Campbell, 2012) or target launch promotions (Campbell, 2015).

However, there is evidence that there is a customer service motive at play too, such that companies use Twitter to respond to public, visible complaints and concerns. In Figure 6 we see telecommunications company Comcast quickly respond to a customer complaint, in sharp contrast to the stereotype of being stuck on hold when calling for support on the telephone.

Figure 6: Comcast responding to a tweet in 9 minutes

The model in this paper speaks to possibly conflicting implications of this type of social media interaction between a company and a customer. On one hand, a public online forum may be thought of as providing a larger network neighborhood for a customer than their offline circle of friends, increasing the discipline imposed on the outsider. On the other hand, the same publicness may reveal differences in the extent of each individual's network connections, decreasing the discipline imposed on the outsider when engaging with sparsely connected customers. A business may feel a stronger incentive to placate a complaint on Twitter from a person with one million followers than from a person with one hundred.

As the model showed, these are not independent issues: the recursive nature of the second could unravel any benefit from the first. Together they mean that the outsider faces a commitment problem. Ex ante it would prefer to tie its hands to eschew information about insiders' network positions to avoid unraveling, but ex post it would prefer to learn this information in order to treat different insiders differently.

In this environment, the network operator has discretion over the rules of the game. It can influence who sees interactions between insider and outsider, which we may view as influencing the connectivity of the communication network. It can also choose the extent to which the outsider can observe or purchase individualized data on the network neighborhood of each insider. These related and conflicting concerns are characteristic of the structure of popular online communication platforms in general. In some sense all such platforms present hybrid cases that do not initially seem to fit neatly into either anonymous or non-anonymous designations, but which taken as a

whole can be interpreted through the lens of the model of community enforcement.

3.1 Twitter

On Twitter, communication is varied and layered: all communication is in theory public unless a user has opted out, but in practice swamped by vast volume unless it is sought. Direct messages are a channel for one-to-one private communication, and ‘mentions’ are something of a hybrid.

In September 2016, TechCrunch reported on new features launched by Twitter to deliver customer service functionality (Perez, 2016). Interestingly, though, it is not clear whether the changes are designed to increase or decrease visibility of complaints. On one hand, the new features include a publicly visible ‘responsiveness’ metric, to indicate how quickly the company responds to customers. On the other hand, it admits new direct messaging functionality that in principle nudges customers and companies to take their conversations out of the public network and into private messaging.

The tension here is suggestive of the type of conflicting incentives present in our repeated game model: publicness is a mixed blessing. In Figure 6, Comcast may have been disciplined to respond to the customer complaint in the public forum of Twitter, but it would perhaps prefer such interactions to be conducted in private. Twitter’s new features appeared to favor taking interactions from a large-neighborhood public setting to an atomized-neighborhood private setting.

By taking the content of the interaction between insider and outsider out of the public view, the platform may feel that it is protecting the outsider from contagious bad publicity. However, as we have seen, exposing and isolating weakly-connected insiders makes the structure of the communication network less conducive to community enforcement. A policy like this could therefore have the unintended consequence of reducing the amount of mutually beneficial engagement between insiders and outsiders.

3.2 Facebook

We see a similar ambiguity in Facebook’s Pages, launched in 2007 as a tool for companies to interact with customers, but with an important difference. Pages permitted companies to see the network ‘reach’ of customers who chose to engage, but also makes communication between brand and customer visible to all other engaged consumers.

One way to view this quasi-public communication is, as discussed in the analysis of the model, to bind the hands of the ‘outsider’ in the face of differentiated network data. Rather than attempt to anonymize the insider’s position, Pages can instead render it irrelevant by essentially creating

communication links among all insiders who have ‘liked’ the outsider’s Page. The communication network for community enforcement is no longer only the revealed reach of a customer’s prior social neighborhood, but the more richly connected neighborhood that incorporates the Page.

This effect is quite different from that of Twitter’s tools. Both platforms permit various types of communication between insider and outsider on a spectrum from private to public. But the direction that each platform’s policies push is toward opposite ends of that spectrum. The approach that enlarges the insider’s neighborhood and publicizes communication within it is better justified by the lessons of the community enforcement model, since this fosters network structures with less vulnerability.

3.3 Privacy concerns

Facebook, Twitter, LinkedIn and Instagram have all increasingly offered fine controls to users that allow them to tailor the visibility of individual communications, rather than tie them to the same communication network for all purposes. This discretion is designed in part to address privacy concerns, and the agency of insiders to choose their network on the fly is one direction in which the model we have analyzed here could be extended.

As we have seen, a novel concern raised by the ambiguous value of non-anonymous network data is what the effect could be on the contracting problem between network operators and outsiders. It is tempting to view personalized data on insiders as valuable to the outsider but harmful to insiders who would prefer their data to be kept private. How can a network operator monetize its detailed information on people’s connections if there is a possibility that the knowledge might be toxic?

Prior literature has studied the interplay between consumers’ control over data on their characteristics and relationships on one hand, and commercial use of that data on the other hand. A common theme in previous work in this area is the tension between the ability of organizations to improve their product offerings through the use of consumer data and consumer distaste for feeling violated by the use of this data (Miller and Tucker, 2009, Goldfarb and Tucker, 2011, Tucker, 2014, Campbell et al., 2015).

The ability of social networks to provide community protection is an additional factor that may contribute to regulatory concerns in this area. It suggests one sense in which the availability and trade of data on the relationships and lines of communication among consumers across various platforms could be bad for *all* parties. Local and aggregate trust in long-term relationships may be eroded by the publicness of network data in unexpected ways. The evolving tools employed and offered by modern communication platforms reflects a searching in the dark for the appropriate balance among many competing concerns. But, in this community enforcement application, it can

be that anonymizing data may be beneficial to all—good for the privacy of network insiders, and good for the incentives of outsiders who would engage them.

3.4 Lessons for platform owners

In sum, the model in this paper suggests several implications for platform operators. First, individualized data on insiders' positions in the network graph may have negative value to outsiders. This is likelier when connectivity is on average high but varies widely. Data on insiders' characteristics would then be more valuable to the outsider and engagement would be higher with anonymized data on connectivity. Conversely, if connectivity is on average low but varies widely, network data would be more valuable to the outsider and engagement would be higher with non-anonymized data on connectivity.

Second, if the platform operator wishes to take actions to enhance the value of network data to outsiders, it matters whether the data is anonymized or not. Average connectivity is important for honest and trustworthy engagement if the data is anonymized, but the operator must pay attention to the whole distribution of connectivity, not just the average, if data is personalized. Local sparseness and vulnerability in some area of the network cannot be ignored due to the problem of recursivity.

Third, actions by the platform operator that foster network closure, cliques, and visibility of communication between insider and outsider are good for well-functioning community enforcement. It may seem that outsiders are well served by the platform operator taking communication between them and insiders out of the public part of the network, but this is counterproductive. A commitment to communication visible to an enlarged neighborhood binds the outsider to better long-run outcomes in the face of short-run temptation to shirk.

4 Concluding comments

In general, the value data on communication networks has for interested parties and for society is a difficult question to answer. Here we have explored one sense in which recursion and non-linearities can complicate valuation, but there are certainly others. What we can say for sure is that the problem is very different in nature from placing a value on other kinds of demographic data that allow for a better match of product to tastes. The importance of information and communication in an array of economic models and settings makes data on communication networks a unique case rich with potential unintended consequences.

References

- Ahn, I. and Suominen, M. (2001). Word-of-Mouth Communication and Community Enforcement. *International Economic Review*, 42(2):339–415.
- Ali, S. N. and Miller, D. A. (2013). Enforcing Cooperation in Networked Societies. Working paper. <https://dl.dropboxusercontent.com/u/1258389/Website/cliques.pdf>.
- Barnes, N. G., Lescault, A. M., and Holmes, G. (2016). The 2015 Fortune 500 and Social Media: Instagram Gains, Blogs Lose. *UMass Dartmouth Center for Marketing Research*.
- Bloch, F., Genicot, G., and Ray, D. (2008). Informal insurance in social networks. *Journal of Economic Theory*, 143(1):36–58.
- Campbell, J. D. (2012). Targeting informative messages to a network of consumers. *Review of Network Economics*, 11(3):Article 5.
- Campbell, J. D. (2015). Localized price promotions as a quality signal in a publicly observable network. *Quantitative Marketing and Economics*, 13(1):27–57.
- Campbell, J. D., Goldfarb, A., and Tucker, C. (2015). Privacy Regulation and Market Structure. *Journal of Economics & Management Strategy*, 24(1):47–73.
- Dal Bó, P. (2005). Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games. *American Economic Review*, 95(5):1591–1604.
- Ellison, G. (1994). Cooperation in the Prisoner’s Dilemma with Anonymous Random Matching. *The Review of Economic Studies*, 61(3):567–588.
- Goldfarb, A. and Tucker, C. E. (2011). Privacy regulation and online advertising. *Management Science*, 57(1):57–71.
- Jackson, M. O., Rodriguez-Barraquer, T., and Tan, X. (2012). Social Capital and Social Quilts: Network Patterns of Favor Exchange. *The American Economic Review*, 102(5):1857–1897.
- Joshi, S. and Mahmud, A. S. (2016). Sanctions in networks: “the most unkindest cut of all”. *Games and Economic Behavior*, 97(1):44–53.
- Kandori, M. (1992). Social Norms and Community Enforcement. *The Review of Economic Studies*, 59(1):63–80.

- Klein, B. and Leffler, K. B. (1981). The role of market forces in assuring contractual performance. *The Journal of Political Economy*, 89(4):615.
- Lippert, S. and Spagnolo, G. (2011). Networks of relations and word-of-mouth communication. *Games and Economic Behavior*, 72(1):202–217.
- Miller, A. R. and Tucker, C. (2009). Privacy protection and technology adoption: The case of electronic medical records. *Management Science*, 55(7):1077–1093.
- Okuno-Fujiwara, M. and Postlewaite, A. (1995). Social Norms and Random Matching Games. *Games and Economic Behavior*, 9(1):79–109.
- Perez, S. (2016). Twitter rolls out new features for businesses running customer service accounts. TechCrunch. <https://techcrunch.com/2016/09/15/twitter-rolls-out-new-features-for-businesses-running-customer-service-accounts/>. Accessed 18 October 2016.
- Tucker, C. (2014). Social Networks, Personalized Advertising, and Privacy Controls. *Journal of Marketing Research*, 51(5):546–562.