

With Us or Against Us? Networks, Identity and Order in a Virtual World

Magnus Thor Torfason

Working Paper

13-019

August 17, 2012

Copyright © 2012 by Magnus Thor Torfason

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

With Us or Against Us? Networks, Identity and Order in a Virtual World

Magnus Thor Torfason*
Harvard Business School
Entrepreneurial Management Unit

Keywords: Social Norms; Social Networks; Triadic Closure; Social Groups; Group Identity

Running Head: With Us or Against Us

Date: December 2011

*Direct correspondence with Magnus Thor Torfason at mtorfason@hbs.edu. I am grateful to Paul Ingram, James Kitts, Glenn Carroll, Chris Marquis, Mikolaj Piskorski and Eyjolfur Gudmundsson, to Jennifer Rhee for invaluable research assistance, and to participants in seminars at Columbia University, London Business School, University of Michigan, Emory University, the Academy of Management, and the American Sociological Association.

With Us or Against Us? Networks, Identity and Order in a Virtual World

ABSTRACT

Social networks and social groups have both been seen as important to discouraging malfeasance and supporting the global pro-social norms that underlie social order, but have typically been treated either as pure substitutes or as having completely independent effects. In this paper, I propose that interpersonal relationships between individuals with different social identities play a key role in linking local and global norms, and in supporting social order. Specifically, I show that social identity derived from group memberships moderates the effects of social relationships on pro-social norm observance. I test my predictions using a novel empirical setting consisting of a large online virtual environment. I show that the number of within-group relationships increases and the number of an individual's across-group relationships reduces the prevalence of anti-normative behavior. Furthermore, I show that network closure has a qualitatively different effect between within-group ties and across-group ties. The effects of within-group and across-group ties are moderated by both group characteristics and actor experience, providing boundary conditions on the mechanisms presented here. My findings illustrate the need for a more nuanced view of the complex interrelations between institutions, identity, and networks.

INTRODUCTION

“How and why do the members of a group comply with its norms?” This is the central question proposed in Homans’ (1950:282) discussion of social control. Social customs, he observes, are not “natural,” but founded on a complex state of mutual dependence. In a stable social system, this dependence leads to equilibrium with a high degree of conformity to the prevalent norms. I examine in this paper a counterpart to the question posed by Homans. In short, I ask: “How, why, and when do the members of a society deviate from its norms?”

The foundational idea that underlies this study is that the members of a social system are differentially exposed to the influences through which social control is exerted. These differences, which originate in the location of individuals within the social structure, lead to differences in behavior. Deviations lead to sanctions through which the boundaries of what is normative are defined and clarified. In other words, the equilibrium of normative compliance is not a static system in which individuals conform blindly, but a dynamic one in which individuals both explore the boundaries of norms and enforce those boundaries with respect to potential deviants.

I focus on two separate but related mechanisms through which the links between an individual and the social system are formed: *interpersonal relationships* and *group memberships* (Breiger, 1974; Goffman, 1971). The distinction between these mechanisms is sometimes unclear (which is why I provide a detailed discussion of this in the following section), but both have long been understood to be important to social order. Membership in the associational groups that form a part of individuals’ social identity has been studied all the way back to Tocqueville, who noted that “the association gathers the efforts of divergent minds in a cluster and drives them vigorously toward a single goal clearly indicated by it” ([1835] 2000:181).

Interpersonal relationships, which form the basis of social networks, have been implicated by Coleman (1988; 1990), Granovetter (1985), Uzzi (1997) and others as the foundations for social capital and cooperation.

My key argument in this paper is that the effects of interpersonal relationships and group memberships are not independent, but that the way in which these two types of links intersect is fundamental to social order and adherence to norms. The interplay between social networks and social groups is important in at least two ways. First, one can influence the other, for example through individual's tendency to form ties with others members of their social group. Reagans, Zuckerman and McEvily (2004) note that group memberships can influence an individual's relationships, and demonstrate how demographic groups can influence outcomes in this indirect way. Similarly, one could imagine settings where social groups are partially influenced by networks due to interpersonal relationships supporting the formation of a particular social identity. Second, social network structure and group structure can jointly affect how people behave, irrespective of any influence of one on the other.

It is this second type of interplay, namely, how social network structure and group structure combine to support the definition and maintenance of normative boundaries and social order, that I emphasize in this paper. Specifically, I will argue that is only ties that span group boundaries that support adherence to specific norms that support social order. In contrast, I posit that when ties lie within a single social group, they can have either no effect on adherence to such norms, or even the opposite effect. I examine these patterns in a novel empirical context that enables me to measure group memberships and interpersonal relationships independently. Thus, I am able to examine directly how their interrelations shape behavior.

My empirical setting is an immersive, three-dimensional virtual universe that is part of a large online computer game. My data, being derived from the event logs maintained by the game's producer, afford an unusual glimpse into activities that are difficult to monitor in more traditional settings. The game in question is a sophisticated social and economic environment in which individuals, informal networks of friends, and formal groups with explicit identities can interact and accumulate resources through a variety of activities.

The emergence of large-scale electronically mediated environments is a relatively recent phenomenon, and the systematic study of such environments is more recent still. The meanings imputed to social interactions within these environments can be different from what is observed between actors who are physically situated in a particular place, but as Boellstorff (2008) points out, "people find virtual worlds meaningful sites for social action, [so] cultures in virtual worlds exist whether we like it or not." As mediated interaction becomes more prevalent and more important to economic and social life, researchers are increasingly turning their attention to such environments (see e.g. Wimmer and Lewis, 2010; Burt, 2011; Piskorski and Gorbatai, 2011). With this nascent body of research, we are increasingly gaining an understanding of the important question of how and when social processes get transferred from traditional interaction to these new, mediated, interaction environments.

In the next section, I elaborate on the distinction between networks and groups, and lay out the argument for emphasizing specifically the intersection between network and group structure rather than considering each in isolation. I focus on the role that this intersection plays in the link between global societal norms and local group norms, and on the implications for individual behavior. After presenting the hypotheses to be tested, I discuss the particularities of examining social and organizational processes within electronically mediated environments in

general, and within my specific setting in particular, following which I present my data and methodological approach. I then present empirical results on the effect of within-group and across-group ties. I establish boundary conditions on these effects by segmenting the population according to both group characteristics and actor experience. To address issues of causality, I perform longitudinal analysis on a selection of actors who were forced to change their location in the social structure due to exogenous factors. I then conclude with a discussion of the results and their implications.

NETWORKS, GROUPS, AND ORDER

This paper is concerned with how the interaction between *networks* and *groups* affects social order. For such a discussion to be conducted meaningfully, it is essential that the distinction between these concepts be clear. This is particularly important because the social structures to which they refer are in many settings empirically correlated, making the distinction less noticeable. In his studies of public order, Erving Goffman emphasizes that there are two key ways in which an individual is tied to the surrounding social structure. He writes that “the individual is linked to society through two principal social bonds: to collectivities through membership and to other individuals through social relationships” (1971:188). Simmel is making a similar distinction when he notes that the “relations of the group *members* with one another have forms which are diametrically opposed to the interrelations among the *groups*” (1964:33).

Distinguishing between networks and groups

Let us start with the social entities to which an individual is tied through *membership*. In the literature, these have variously been referred to as collectivities, associations, groups, and organizations. Following Simmel, I generally refer to these entities as *groups*. In Simmel’s view, a group’s most basic property is that it provides identity to its members. Simmel writes that “the

groups with which the individual is affiliated constitute a system of coordinates, as it were, such that each new group with which he becomes affiliated circumscribes him more exactly and more unambiguously” (1964:140). Tajfel and Turner refer to these entities as *social groups* and note that such groups “need not depend upon the frequency of intermember interaction, systems of role relationships, or interdependent goals” (1979:40), but on whether the individuals concerned define themselves and are defined by others as members of the group.

An important implication of the above is that the concept of groups as applied in this paper is not simply shorthand for cliques or densely connected clusters.¹ Although group members often associate frequently with one another, this is not always the case. A case in point is the set of Boston Red Sox fans. Most participants are clear about whether they are a part of this group, and some members bear prominent insignia to advertise their membership. But any given Red Sox fan can have relationships with only a small fraction of the other members and is sure to also have relationships with non-members. Yet, membership in this group nevertheless creates a bond that moderates behavior toward other members, even when interaction between group members is minimal or non-existent. Other examples of groups include Republicans, Hispanic and Latino Americans, red-heads, teenagers, and academics. Of course, groups do not have to correspond to nationally recognized categories for their members to perceive them as meaningful. Individual families, corporate divisions, and native tribes are also groups in the sense considered here. Groups may have a greater or lesser degree of formal organization, but this is not the key criterion. Rather, the key determinant is whether the question of membership

¹ It is worth noting that in social psychology, the existence of groups (often referred to as “small groups”) is sometimes taken to imply exactly the sort of interaction among members that characterizes cliques or densely connected clusters; the Simmelian conception of groups therefore bears a greater resemblance to “minimal groups” (Tajfel et al., 1979) than to “small groups”.

in a group is clear in the minds of its members (and typically in the minds of others with whom they associate), and whether group membership contributes to its members' social identity.

On the other hand, the set of people with whom a particular individual is friends is *not* a group in the above sense because they are not conscious of, let alone defined by, their commonality in that case. A celebrity entourage could perhaps constitute an illustrative exception to this, the relationship to the celebrity possibly being salient enough to induce a conscious idea of membership. But leaving such curious hybrids aside, friendship is first and foremost a link between individuals rather than a link between an individual and a group. Social relationships, including not only friendship but also trust, love, and other types of ties, can exist between individual members of the same group, and shared group membership can strengthen such ties, but social relationships can also span group boundaries. Thus, a given set of people and the associated *relationships* among them does not constitute a group (at least not in the sense used here), but rather a *network*. We can summarize the distinction between the aforementioned concepts as follows: *memberships* are properties of *groups*, whereas *relationships* are properties of *networks*.

Views on the structural origins of order

The importance of groups to social order has been highlighted by numerous researchers. Hechter and Kanazawa, who focus on group structures in their study of social order in Japan, define order as “the extent to which citizens comply with important norms – written laws and other mores that the state attempts to enforce” (1993:458), and suggest that the high level of order observed in Japan is due to the success of local groups, such as firms, families, and neighborhood organizations, in facilitating normative compliance. According to their solidaristic

theory of order, “global order emerges from mechanisms producing group solidarity – that is, from spontaneous local order” (1993:458).

Much earlier, the role of groups in supporting norms and local order was emphasized by Tocqueville, who paid a great deal of attention to groups and associations as part of the distinctive social thread of the United States. In particular, he noted how, when a particular viewpoint or opinion is “represented by an association, it is obliged to *take a clearer and more precise form*” ([1835] 2000:181, emphasis added). The goals of the associations referenced in Tocqueville’s analysis were at least partly political in nature, but the idea that rules and norms are strengthened when they are associated with a group rather than with an ephemeral network of individuals has been upheld in numerous works. Putnam (2000), for example, emphasizes the importance of groups even when they have no explicit social goals, associating the collapse of American community with the decline of activities as mundane as league bowling.

A number of more recent accounts stress the importance of networks in promoting social order and discouraging anti-normative behavior. Perhaps the most prominent among those accounts is that of Granovetter, whose idea of “embeddedness” has underpinned a large stream of literature within economic sociology on the impact of network structure. Granovetter advances embeddedness as a solution to the problem of order, noting that “the embeddedness argument stresses [...] the role of concrete personal relations and structures (or ‘networks’) of such relations in generating trust and discouraging malfeasance” (1985:490). According to this view, interpersonal relationships at the dyadic level provide a tie of both trust and information transfer, and the extent and structure of such relationships provides the context in which economic action takes place. The idea that interpersonal relationship structures are at the heart of human social interaction has enjoyed immense success and support in recent years, leading Nee

and Ingram to observe that “network ties” have become “the bedrock concept of modern structural sociology” (1998:22).

Social networks analysis lends itself well to quantitative analysis and allows a relatively tight link between theory and prediction, and the relationship between network structure and order is a clear example of this advantage of structural analysis. Granovetter (1985; 2005) emphasizes how order is supported by densely knit networks that offer thick channels of information and understanding. These networks support the development and crystallization as well as monitoring and enforcement of norms. Coleman’s (1988; 1990) exploration of social capital, which strengthens the emphasis on cohesion and closure by outlining how third parties enable stronger monitoring and enforcement of norms, is generally consistent with the embeddedness account of the importance of dense networks in general. The embeddedness arguments of Granovetter and Coleman have also been supported by empirical studies (See e.g. Uzzi, 1997; Fernandez, Castilla and Moore, 2000). But others have pointed out that closed cliques can sometimes have a highly negative effect on social order. Burt (1999; 2000; 2001), for example, has emphasized that social capital depends crucially on bridging ties that span structural holes in a network. Stressing the importance of such ties in providing a wide range of different perspectives, he argues, in contrast to Granovetter and Coleman, that closure is a double-edged sword, and that actually, “third parties amplify probability of distrust in weak relationships” (Burt, 1999:224). This is supported by empirical analyses in which Burt demonstrates that “closure’s association with distrust and character assassination is as strong as its association with trust” (Burt, 2000:352).

The commonality of such approaches is in their emphasis on network structural analysis as a means to address the problem and in downplaying the role of identity-based groups with

respect to the problem of order. Emirbayer and Goodwin pointed to this “priority of relations over categories” (1994:1414), as a general characteristic of social network analysis, whereby groups tend to be treated as constructs that can, for the most part, be derived from the network structure rather than existing alongside it and moderating its influence on social order and social capital. In practice, the network approach can thus lead to a view where bilateral ties completely replace groups and identity as a basis for analysis, rather than complementing them.

Emirbayer and Goodwin are of course not the only ones to observe that identity and relationships both play a part in social processes. For example, although Moody and White (2003) focus on the relational aspect of social structure in their account of social cohesion and solidarity, they explicitly note that “solidarity can be partitioned into an ideational component, referring to members' identification with a collectivity, and a relational component” (2003:104). Cook and Hardin also acknowledge that behavior can be guided either by (network-based) reciprocity or by (group-based) norms, even if they then proceed on the premise that when one is enacted the other has “little additional role to play” (2001:336).

A number of researchers have investigated empirically the influence of group membership on the formation of interpersonal relationships (See e.g. Mehra, Kilduff and Brass, 1998; Reagans, Zuckerman et al., 2004; Entwisle, Faust et al., 2007). Fewer studies have treated the two as having a joint influence on other outcomes, although the study of demographic minorities on corporate boards by Westphal and Milton (2000) is a notable exception. Westphal and Milton examine the relative influence that different board members have on corporate decision making, and show that network ties have a disproportionate effect on individual board members' influence when the individuals are members of demographic minority groups.

Similarly, it seems clear that to understand the origins of social order, we should not limit ourselves to either groups or networks. Rather, we should incorporate both of these in the analysis whenever possible, and pay special attention to the intersection of these two aspects of social structure. This is especially important when considering the relationship between local order and global order, and in the following section I discuss how groups and networks jointly influence the way in which local norms get aggregated and whether they contribute to order at the societal level or undermine it.

THE LINK BETWEEN LOCAL AND GLOBAL ORDER

In this paper, I follow Hechter and Kanazawa (1993) in treating social order as the extent to which citizens comply with important norms. Although their decision to focus on norms endorsed by states is useful in some contexts, I go beyond the identity of the endorser, treating social order instead as deriving from the extent to which the members of a society adhere to *pro-social* norms that operate at the *global*, or societal, level.² This definition thus excludes both aberrational norms destructive to everyone affected by them and local, discriminatory norms, the benefit of which to a particular group comes at the expense of their effect on another. These types of norms, although certainly present in many societies, do not contribute to social order as understood by either Hechter and Kanazawa or the structural sociologists mentioned above, who treat social order as being related to social capital and cohesion.

The question of if and how local behavior and norms aggregate to the level of society is thus a key part of the study of social order, but one to which neither pure network accounts nor pure group accounts provide a satisfactory answer. The idea that an interpersonal relationship

² It is worth emphasizing that in the discussion below, ‘global’ is generally treated as synonymous with ‘society-level’ and ‘local’ as synonymous with ‘group-level’. Although the argument could in many cases be treated as hierarchical and referring to links between global and local structures in the abstract (including, for example, those between a large corporation and its departments), such a generalization is not pursued in this paper.

between two individuals increases the chances that they act in ways that are mutually beneficial them is, of course, quite plausible. It is also reasonable to assume that this effect would be especially strong when the two are embedded in a cohesive network of counterparts who can monitor compliance. But in applying this approach to the linking together of the norms of an entire society, one is faced with the issue that key linkages must occur over bridging ties where monitoring opportunities are weak or nonexistent.

Groups that provide their members with a clear social identity can alleviate the need for direct relational monitoring because membership in the group is preconditioned on the internalization of particular norms and values negotiated within the group. In strongly solidaristic groups, such norms can be quite effective in constraining behavior. Thus, the idea that solidarity and identity can be the foundations for local order within the group is unproblematic. But the strength of a particular group identity is directly related to the clarity of separation between those who conform to the identity and those who do not. As Tilly notes: “collective identities [...] *center on boundaries* separating us from them” (2005:7, emphasis added). This implies a fundamental tradeoff between the strength of local group identities and the potential for those identities to overlap in ways that link multiple local orders to form a single global order, and it is far from clear that one implies the other in all settings. In fact, Lyons even suggests that “social cohesion and informal social control may facilitate hate crime, especially in communities whose identities are rooted in racial homogeneity” (2007:824). According to Hechter and Kanazawa (1993), respect for order and norms in the abstract is forged through memberships in highly solidaristic groups. But as the example presented by Lyons demonstrates, the particular content of the norms thus forged can diverge considerably between groups, significantly undermining order at the societal level. Although divergence in content would be reduced by blurring the

identity-based boundaries on which solidarity is founded, this would at the same time be expected to diminish respect for order and normative adherence in general.

[Figure 1 about here]

This dilemma can be resolved by departing from attempts to arrive at a social structure capable of maintaining high social order by focusing only networks or only on groups. Rather, I propose that high levels of global order and adherence to societal-level norms are most likely to be found when solidaristic groups that support strong identities at the local level are linked together by boundary-spanning interpersonal relationships. Across societies, this would imply that a greater prevalence of such relationships was associated with greater societal order. Within a single society there is likely to be variation in both social structure and behavior, and at the individual level, the people who are most likely to comply with global societal norms are precisely those who are engaged in interpersonal relationships that span group boundaries. This intersection between groups and networks is illustrated in figure 1, which highlights how one's view of social structure changes when one considers groups and networks jointly, rather than focusing only on network structure. There are three key reasons for expecting boundary spanning relationships (denoted with jagged lines in the figure) to engender greater adherence to societal norms.

First, relationships that span group boundaries provide greater exposure to and appreciation of societal norms. When two individuals in a dyad do not share a social identity at the group level, the shared understanding on which their interaction is based must be rooted in the global rather than local culture. As Blumer (1969) noted, when individuals are engaged in joint action together, they bring to that action their respective symbolic objects, meanings and interpretations. The more issues that individuals from different social groups address together,

the greater the attention to those meanings and interpretations that they share – namely those that originate in the global culture. This renders commonality more salient than distinction, increases attention to the global society rather than the respective local groups, and increases the probability that the norms that support global order are transferred through these links. In contrast, interaction between individuals who share a narrow social identity will not tend to be framed in such terms. In fact, such interactions may even be framed in terms of the distinction between the local context to which the interlocutors belong and the global context. It is therefore clear that a tie within a group boundary is not likely to have the same effect on adherence to global norms as a boundary spanning tie, and when all is taken together it seems unclear that such ties should have an effect at all.

Second, interaction increases empathy with others. When those others are part of the focal individual's outgroup, such empathy has been shown to generalize to the entire outgroup (Pettigrew and Tropp, 2006; see also Pettigrew, 1998). A greater generalized empathy towards those people is therefore likely significantly to enhance one's appreciation of norms that shape individuals' behaviors in ways that are likely to benefit outgroup members with whom they do not have a direct personal relationship. Although empathy should not be expected to play a role in supporting all conceivable norms, the norms that it would encourage are precisely the global, pro-social, norms that are generally posited to support societal order.

Third, potential violators of global societal norms (as opposed to local group norms) are exposed to monitoring and sanctions to a greater extent by boundary-spanning relationships than by relationships within a particular group. Group affiliation is known to affect sanctioning behavior in a variety of ways (see e.g. Bernhard, Fehr and Fischbacher, 2006; Goette, Huffman and Meier, 2011). In the context of societal norms, the content of which may differ from the

norms of an individual group, sanctioning within a group is likely to be triggered by behavior that threatens, rather than strengthens, the identity of the group. Thus, when a potential for divergence between local and global norms arises, within-group sanctions are likely to focus exclusively on the local aspect.

These three mechanisms form the basis for a set of hypotheses about individual variation in adherence to global, pro-social norms. Hypotheses 1 and 2 derive directly from these mechanisms, while hypotheses 3 and 4 concern boundary conditions in terms of how the effects vary for specific subsets of the population. My starting point is with the intuitive assertion that the mere existence of social relationships supports social order by discouraging the violation of societal norms. This view is consistent with Granovetter's idea that "norms – shared ideas about the proper way to behave – are clearer, more firmly held and easier to enforce the more dense a social network" (2005:34). The idea that social relationships serve to integrate an individual into society is widely held among social networks researchers and leads to the straight-forward prediction that the more interpersonal relationships an individual has, the less likely he is to violate societal norms. However, in emphasizing the interrelation between social networks and social groups I establish an important criterion regarding which types of ties are likely to support which types of norms. On the basis of the forces discussed above – primarily salience of commonality and increased empathy – I predict that adherence to global societal norms is governed by ties that span the boundaries of such groups rather than ties between individuals who share a single local identity:

Hypothesis 1: Interpersonal relationships that span local group boundaries will reduce an individual's propensity to violate global norms, but relationships that lie within a single group will not.

The above hypothesis concerns the effect of direct relationships, but the wider structure of a network is also generally posited to influence individuals' behavior. In particular, triadic closure in an individual's immediate network neighborhood is often considered important to adherence to social norms because it facilitates monitoring and sanctions. Social network theorists, however, reach somewhat different conclusions about the effects of this mechanism. The role of coordinated monitoring and social influence in support of norms is emphasized by Granovetter (2005) and Coleman (1988). In contrast, Burt (1999) emphasizes closure's potentially negative effect on social capital, noting that high levels of triadic closure are accompanied with a lower diversity of information to which individuals are exposed. At a systemic level, this would be expected to lead to balkanization of the social structure, divergence in local norms, and weak and ineffective societal norms at the global level. We thus have two opposing predictions related to the effect of closure on adherence to global societal norms.

An appreciation of the interplay between networks and groups reconciles to some degree these opposing predictions. Because the positive effects of closure on social order depends on effective monitoring and sanctions, and the negative effects have to do mainly with balkanization of the network, I expect closure to play the most constructive role in the presence of the former, and absence of the latter, effect. Monitoring and sanctioning of societal norms, as noted above, operate differently within and across groups. In particular, sanctions are more likely to facilitate compliance with global norms when the triads involved span group boundaries. Such triads are also less likely to lead to a balkanization of networks because, although the individual relationships do not necessarily bridge structural holes, their bridging of identity boundaries mitigates against balkanization. I thus predict the effect of closure to be conditional on the intersection of network and group structure as follows:

Hypothesis 2a: Closure in an individual's immediate network will *reduce* the individual's propensity to violate global norms when the underlying triads span group boundaries.

Hypothesis 2b: Closure in an individual's immediate network will *increase* the individual's propensity to violate societal norms when the underlying triads lie within a single group.

Both hypotheses 1 and 2 relate to aggregate levels of anti-normative behavior in society, and regardless of the mechanism by which structural patterns lead to norm violation, such behavior would be expected undermine the social order of a society. It is, however, also important to consider when these mechanisms are likely to have a greater or lesser role. I propose above that a key role of boundary-spanning ties is that they link the content of local norms in different groups, without undermining solidarity at the group level and respect for norms in general. This has implications for how group and network structure would affect behavior in groups whose norms have for some reason diverged from societal norms.

If the content of local norms in a group is for some reason markedly different from that of the global norms that prevail in the wider society, interaction that is framed in terms of the local culture is then likely to be not only neutral towards the global culture, but actually hostile to it. And on the flipside, when members of a deviant group interact with others who do not share that local identity, and frame that interaction in terms of global meanings, this is likely to be especially important to enhancing appreciation of the global norm rather than the local, divergent norm. In the case of such deviant groups, we would therefore expect within-group ties to be even less effective in reducing anti-normative behavior, but across-group ties to be even more effective, compared to non-deviant groups (groups whose local norms are consistent with global norms):

Hypothesis 3a: Across-group ties will reduce anti-normative behavior *more* in deviant groups than in non-deviant groups.

Hypothesis 3b: Within-group ties will reduce anti-normative behavior *less* in deviant groups than in non-deviant groups.

The predictions I have presented until now are derived from what we know about social behavior and structural constraint in general settings, and apply to a social system in equilibrium. As I noted in the introduction, however, the fact that patterns of norm violations are in equilibrium does not imply that they are static. Individuals both explore the boundaries of norms and enforce those boundaries with respect to others, and it is through these joint actions that those norms get clarified and reinforced. In particular, individuals who enter a particular social system need to learn about the prevalent norms, how adherence to them is construed, and about the social consequences of violating them. In a stable social system, global norms pertaining to social order are typically well understood, but local variations require greater familiarity with the environment.

Consequently, individuals who are new to a social system can gain an understanding of institutionalized society-wide norms relatively quickly, but as tenure within the environment increases, we would expect greater salience of local norms and lesser salience of global norms. Anecdotal evidence of such gradual evolution in normative adherence is ample. One example is that of the speed at which travelers along the interstate choose to travel. In this case the institutionalized global norm – the state-wide speed limit – is very clear to anyone regardless of experience. However, this limit is enforced very differently at different stretches, and locals who are familiar with the environment learn that in some places it is normative to exceed the limit by significant amounts, whereas in other places the norm calls for maintaining a speed very close to the limit. Other examples may not be as graphic, but are just as real and important. For example,

new employees are introduced to company-wide rules quickly, through orientation sessions and handbooks, but as they gain experience they learn that different departments interpret corporate rules differently, and local norms about how to navigate the rules gradually become more accessible to the employee than the literal content of the rules themselves.

As individuals gain a better understanding of the variation of norms in an environment and the context in which norms are interpreted shifts from the global to the local, the relative importance of within-group and across-group ties also shifts. A rookie is inclined to treat the global norm as the default, and it is only through interaction within the local group, and joint interpretation based on the local identity, that attention shifts to the local. In contrast, a veteran is intimately familiar with local variation in norms. Thus, for the veteran, the decision to emphasize global norms is likely to be dependent on interaction rooted in the global culture – that is, interaction across group boundaries. This results in the following prediction:

Hypothesis 4: As an individual's tenure in the (global) social environment increases, the relative impact of across-group ties will increase, and the relative impact of within-group ties will decrease

Below, I will test these hypotheses empirically in the context of EVE Online, an online computer game. When applying theoretical predictions to specific empirical settings, one must always pay attention to how more general hypothesis map on to specific circumstances. In a novel empirical setting, this is even more pertinent, and in the next section I therefore describe in some detail the online environment from which my data derive, before going on to discuss the specific variables and methods I use for my quantitative analysis.

THE SOCIAL REALITY OF A VIRTUAL WORLD

I test my hypotheses in the context of an online game, named “EVE Online,” in which large numbers of people interact within a virtual environment. The state of this environment is

persistent, meaning that participants' actions can have permanent effects on their surroundings and that those effects are potentially visible to all other participants. It also means that each of the 300,000 participants maintains a permanent in-game persona, or avatar, that must be retired to safe location before the participant logs off so that the avatar is not vulnerable while the participant pursues life outside the online world.

Participants must construct their avatars upon joining the game and endow them with various traits including gender, race, and physical and mental abilities. An avatar's customizable physical appearance is the face others see in their interactions with the participant, both within the virtual environment and on the forums maintained by the game developer. In addition to the avatar's appearance, participants choose a vehicle – a space ship – which is used for travel through the virtual space of the game.

The spatial representation of the virtual environment is a collection of solar systems, each containing a number of planets and other natural bodies as well as artificial installations. Movement within each solar system is accomplished using realistic three-dimensional navigation, movement between systems with the aid of “jump-gates,” each of which is connected to a single neighboring system. To travel longer distances, participants must utilize a series of such gates. The universe that faces them thus resembles an interconnected network of solar systems.

Perhaps more important than the spatial structure of the environment is its social structure. One way in which this gets established is through formal groups, which any participant can create and which other participants can join (with the consent of existing members). In my data, I observe around 6,000 such groups; each group typically consists of between 5 and 100 members, although some larger and smaller groups exist as well. When characters are first

created they automatically become members of one of several default groups, but usually move on relatively quickly, to groups established by other participants (or they establish their own groups). It is possible to move freely between these participant-managed groups as well, although many participants spend significant periods in the same group.

The groups are referred to as “corporations,” and their formal governance structures would be familiar to most organizational researchers, since they have been intentionally designed to resemble those of real-world corporations. Each of these is headed by an individual with the title of CEO and includes one or more directors. A number of other titles and roles can also be assigned to members. The goals of these groups are quite varied, but a common theme is an emphasis on camaraderie within the group, support to new members, and ambitions of establishing the group as a holder of territory or as a significant actor within an alliance of multiple such groups (EVE Forum, 2009).

The ways in which participants pursue these goals are similarly diverse. Participants can engage in trade, either as merchants or as speculators, and some even offer financial services. Others collect revenue by providing transportation services such as shipping or couriering on a contract basis. There are also a number of ways of creating resources directly from the virtual environment, through mining, construction and similar things; raw material or processed items can then be sold on markets. And finally, participants have the option of taking resources from others by force. Such force can be small-scale, in the form of robbery, piracy, hostage taking and other methods; it can also be on a larger scale, through organized military action.

Much of players’ participation in the game is associated with the groups they enter. Some groups maintain rules that govern how much time is spent online, and participants are often expected to support group activities such as trade or conflict missions. Each group has a name

and a logo, and some also have customized insignia denoting particular titles. Participants' identities are thereby tied to the identity of their group and participants' actions are shaped by group membership as well as by their own social network.

The specific behavior examined in this paper is the practice, in a conflict situation, of continuing to inflict damage upon an opponent who has been defeated. I can observe this because of a feature that enables a participant's avatar to escape after his or her ship destroyed in combat. When a ship is destroyed, its pilot can eject in an emergency vehicle, called an "escape pod," a weaponless emergency vehicle the only function of which is to get the pilot to a safe location. The institutional structure of the game clearly denotes destruction of escape vehicles as anti-normative, the clearest signal of which is that although the game's institutional agents (avatars controlled not by people but by the programming of the environment) will sometimes destroy participants' regular vehicles, they will never destroy escape pods, even when their occupants have violated the rules of the game. The practice is also reminiscent of World War II when fighter pilots were sometimes killed while attempting to parachute to safety after their planes had been shot down (Blackburn and Hammel, 2006), and it would be well described by the saying about "kicking people while they're down." These negative connotations, coupled with the institutional infrastructure, provide a global baseline in the form of a global societal norm against the destruction of escape pods.

However, in accordance with the designer's goal of making the environment more realistic by allowing a wide variety of behaviors as "part of the game," it is nevertheless possible for players to destroy such pods. The destruction of opponents' vehicles in the game is referred to as "killing", and the specific case of destroying escape pods is typically referred to as "pod killing," or simply "podding" someone. I will follow this convention and refer to this behavior as

“pod killing,” and specific instances of it as “pod kills.” Because the norm against pod killing operates at the level of the global society, and is pro-social in its content, it is a good choice for examining the determinants of social order within the game environment – although, as discussed below, it does exhibit some local variation. Consistent with Hechter and Kanazawa (1993), I treat the patterns of compliance and deviance with respect to this norm as an important aspect of order.

DATA AND METHODOLOGY

Virtual worlds generate enormous amounts of information, but just as in the physical worlds, the vast majority of this information is ethereal and disappears before anyone has the chance to observe or notice it. In virtual environments such as EVE Online, those elements of the instantaneous state that do get recorded are saved in one or more databases residing on a server cluster managed by those responsible the environment (in this case, the producer of the game). Much of even this information is relatively short-lived, but some information is usually preserved in log tables, either for operational reasons or with the explicit intention of analyzing it. Such log tables afford an opportunity to explore a number of actions that are difficult to record in a non-intrusive manner in the real world. I received access to such logs directly from CCP Games, the company that develops and maintains the online environment of EVE Online.

All data were anonymized before being transferred. Although the in-world characters do operate under pseudonyms, this was not deemed to amount to sufficient anonymization, so a second round of anonymization was performed, with the goal of detaching the data not only from the real-life participants, but also from their in-game characters. The reason for this was two-fold. First, it seemed likely that it would in some cases be possible to link characters with the people behind them and defeat the anonymization. Second, the characters themselves often

represent a significant investment on behalf of their owners, so any harm to a character or a character's reputation could well constitute material harm to its owner. All variables that were used for the quantitative analysis (and are described below) derive from these anonymized activity logs.

Dependent variable

My key dependent variable is individuals' propensity to engage in a particular type of antinormative behavior – pod killing. I examine the prevalence of pod killing using logs that detail the ship type, time and location of destruction for all ships lost in combat, as well as the identities of both the perpetrator and the victim. By examining the ship type variable, I can identify all instances where escape pods are destroyed and classify events as either “regular kills” or pod kills. Since the opportunity to pod kill depends on having first engaged in a regular kill, but it is only the specific tendency to engage in pod killing that is antinormative rather than the general tendency towards conflict, the variable of interest is the propensity to pod kill *conditional on* having first engaged in a regular kill. I therefore pair each pod kill with the regular kill that preceded it. The resulting variable, a dummy variable named *Pod Kill*, thus consists of one observation for each regular kill that takes place. If a regular kill is immediately followed by a pod kill (with the same individuals involved), this variable takes the value of one. If no pod kill follows the regular kill, the variable takes the value of zero.

Network structure

A key aspect of my study is the extraction of actors' interaction networks within the game. EVE Online is a virtual world that offers a sense of space through which participants can move just as people do in the physical world. Consequently, people can “hang out together” while they are engaged in the trade, shipping, conflict or other activities that they perform in

order to accumulate money and resources, and I operationalize a network tie as existing if two actors repeatedly move together in close proximity through this virtual space within the game environment.

The sense of space and place is the defining feature that sets virtual environments apart from other types of mediated interaction, and is their main source of attraction. Markham (1998) referred to this characteristic as “presence”, and as Boellstorff (2008) notes, the idea of presence is inherently bound to the concepts of place and time. The way in which such “presence” supports social relationships and intimacy is well captured in one of Boellstorff’s anecdotes from Second Life, another virtual environment, where he notes that: “One day I was im-ing [instant messaging] a friend when they suddenly teleported to my [virtual] location, saying ‘I just wanted to say hi in person – it feels so rude to only talk through ims.’” The fact that people are seeking exactly this type of togetherness when they participate in virtual environments makes the use of virtual “co-presence” an attractive method for measuring social relationships. This method also compares favorably to methods of measuring direct communication due to technical issues, most importantly because the individuals who have the strongest ties are actually likely to communicate directly through IP telephony programs such as Skype, leaving no trace of their communication in the available data. Finally, because of technical reasons related to load distribution on the server infrastructure, information about location and movement has been collected in much greater detail and with greater reliability than other potential sources of network data, despite the significant amount of space required to do so.

Similarly, the decision to use co-movement rather than static measures of co-location was taken for both theoretical and practical reasons. Theoretically, the goal orientation associated with co-movement is a stronger signal of affiliation than mere co-location. To take a real-world

example, two people may spend an entire workday in the same building without engaging in any social interaction. If they are repeatedly observed leaving the building together for lunch at a nearby restaurant, however, one can be more confident that they have a social tie. In the virtual world, individuals similarly spend considerable periods of time at particular places because of the availability of various resources. Being at a particular place at the same time as others is therefore not as good an indicator of a social relationship as is entering and leaving those places at the same time. Practically, the number of potential dyads formed by looking at raw movement data over a given period of time is very high and places significant demands on the computing power available for processing them; the number of potential dyads among all individuals who merely happen to be in the same place at any point during the same period would be greater by orders of magnitude.

I obtained information about co-movement from logs³ denoting when individual characters move (or “jump”) between solar systems. This is not as uncommon an occurrence as someone bound to the same planet for their entire life might presume: In the course of a typical day, the approximately 100,000 players active in the game make some three million jumps, or about 30 jumps per character on average. Using the logs, I extracted pairs of characters that made a series of such jumps between the same places at the same time, and classified those who did so repeatedly as having a positive relationship between each other.

Although this approach does offer a powerful way to discern network structure, it is important to ensure that the behavior used to construct the network ties really is indicative of a social relationship. The first and most basic concern is of course that this approach must not simply pick up random occasions of individuals moving in the same direction at the same time.

³ These logs consist of the exact time of each jump, the destination, and character making the jump. Each batch contains data spanning two weeks, which totals about 40 million jumps.

A greater threat to inference would be if the way that the construction of the network lead to systematic errors. For example, if mass-transit and commuting were elements of the game, this would occasion the risk of spurious ties, as unrelated individuals would be making the same commute at the same time. Another risk would be that observed co-movement might indicate not a positive, but a negative relationship, because of enemies repeatedly pursuing each other in conflict situations. To guard against these and other threats to the validity of the measures, I therefore performed several rounds of processing to the data specifically intended to reduce the chance of spurious ties, and performed a number of tests (described below) on the end result.

In the first round of processing, I extract information about every instance in which two individuals jump to the same destination within five minutes of each other. I refer to this as “joint jumping.” Since people’s paths often pass without any specific intention on their behalf, many or most of these joint jumps are therefore likely to be random occurrences: It is not merely passing through the same place that is meaningful in this context, but actually going somewhere *together* – and doing so repeatedly. I therefore perform two additional rounds of processing to isolate those instances that are truly indicative of a social relationship.

In the second round of processing, I extract the instances in which two individuals make a series of joint jumps within a short time frame. Based on an examination of the patterns of travel, I chose to use the occurrence of at least five joint jumps within a 30-minute interval, which I refer to as a “joint trip.” Due to computational limitations, I process this information in hourly batches.⁴

⁴ To ensure that every instance that satisfies the criteria for a joint trip is counted, I include a 30-minute overlap with the next batch, but retain only instances that start in the first 60 minutes of each batch. Using this method, it is possible that a long joint trip that spans more than one hourly interval is counted as two shorter trips instead. However, since the third round of processing requires at least two trips to have occurred with more than a 24 hour interval between them, this slight inaccuracy cannot lead to a tie being inferred where one would not have been otherwise.

In the third and final round of processing, I retain only pairs that have made multiple trips together *and* made their joint trips over an interval of *at least* 24 hours. By applying this minimum interval, I can be reasonably certain that these instances of repeated interaction are truly multiple intentional interactions, and that I am not simply misdiagnosing one long interaction as two separate ones. This also ensures that only players who engage repeatedly in joint movement over more than one session of play are included, thus assuring that all but the most serendipitous instances of joint travel by strangers have been purged from the data. The remaining pairs constitute the ties in the social network used in my analysis.

Independent variables

Having constructed the social network for all 300,000 characters in the game, I calculated individual-level network measures to include in the regression analysis using igraph (Csárdi and Nepusz, 2006), a software library designed specifically for the analysis of large-scale social networks. The key network measures used in the regression estimates are node degree and node constraint. Node degree is simply the number of ties that an actor has with others and is captured in the variable *Total number of actor's ties*. Node constraint is the measure developed by Burt (2004) to capture the extent to which an actor's network spans structural holes. This measure ranges from zero to one, with higher values for individuals with greater closure in their network neighborhood (and thus fewer structural holes). This measure is captured in the variable *Actor constraint score*.

In addition to the basic network measures, I am interested in the combined structure of the network and the social groups that individuals are members of. Group membership can be captured in a straight-forward manner, since the data on each character in the game includes an

identifier denoting which “corporation” he or she is a member of.⁵ The group identifiers are combined with the basic network data to segregate the network into within-group ties and across-group ties. A tie between two individuals who are both members of the same such group is classified as a within-group tie, but a tie between two individuals who are in different groups is classified as an across-group tie. The segregated network is then used to construct network measures that parallel the ones described above, but which reflect whether they are based on ties that lie completely within a single group, or whether they span group boundaries. These network measures are captured in the variables *Ties within group boundaries*, *Ties across group boundaries*, *Actor constraint within-group only*, and *Actor constraint across-group only*.

In examining the differences between groups, I include a dummy variable to capture whether an individual is a member of a group that has high levels of deviance or one that has low levels of deviance. This measure is constructed by calculating the average propensity that the actor’s fellow group members have for engaging in pod kills. Since the simple group-level tendency to engage in pod kills is a function of the focal actor’s own propensity for pod kills, it is structurally correlated with the dependent variable (in other words, a group could be seen to be deviant simply because the focal actor is deviant). I therefore always exclude the focal actor in the construction of this measure and rely *only* on the behavior of other members of his or her group. If the level of deviance among *other group members* is greater than the median level in society, I classify the group as deviant; if the level is lower than the median level, I classify the group as non-deviant. This measure is captured in the variable *Deviant group (dummy)*. To separate the effect of network ties in more and less deviant groups, I include the interaction

⁵ As was mentioned in the previous section, there do exist two types of groups, the “default groups” that characters are initially assigned to and the “participant-managed groups” that they subsequently join. However, the individuals who are engaged in conflict are overwhelmingly members of participant-managed groups (only 3% of the *Pod Kill* variable observations stem from members of default groups). The analytical results are essentially the same regardless of whether those observations are included or excluded.

between this variable and the number of ties within and across group boundaries. The two interactions are captured in the variables *Ties within boundaries of deviant group (interaction)*, and *Ties across boundaries of deviant group (interaction)*.

I also include data on actors' experience, or tenure, within the game. Experience is calculated on the basis of points that characters accumulate over time and are needed to endow characters with new skills. Since these points are accumulated solely on the basis of the time an account is active, they provide a good measure of how long a character has been active in the game, and I include this as the variable *Actor experience*. I also account for the size of the group of which an individual is a member. I operationalize this simply as the number of members in the group that a focal actor is a member of, and include this as the variable as *Size of actor's group (thousands)*. Since all data was anonymized before it was exported, it was unfortunately not possible to combine it with textual analysis, such as mission statements of individual corporations or other public statements made by characters. The anonymization method also precluded the inclusion of demographic information on the actual people who control the character.

[Table 1 about here]

Before performing the main analysis, I calculated correlations and summary statistics of the key variables; these are reported in table 1. I also examined the data in detail to ensure that the variables captured the constructs that they were intended to capture. In particular, I examined the degree to which the social network structure was consistent with what would be expected for a network or interpersonal relationships with positive valence. The possibility that the pattern was random could be ruled out relatively quickly, since the number of repeated instances of joint movement was vastly higher than it would be if dyads engaged in joint trips at random (of the 23

million total trips observed among the 300.000 participants, only 15 dyads would be expected to experience repeated trips if this were a random process, but the actual number is 3 million). The intersection between the network structure and group structure is also as expected: A tie is much more likely to exist between two participants if they are members of the same group than if they are not. The pattern of aggressive behavior further supports the interpretation of ties as measuring positive relationships, rather than incorrectly picking up interactions between adversaries, since instances of killing where the victim and perpetrator have a tie are essentially non-existent. The same applies to group memberships: Essentially no instances of an individual killing a fellow group member were observed. Finally, I examined the data on joint movement to see if there was any evidence that patterns of repeated joint movement were caused by commuting, and found no indication that this was the case. The movement data does indicate high levels of activity in general between around 8pm and 2am (in the US time zones), as would be expected given that the participants, who are for the most part located in the US, tend to go online in the evening, after work. Within that pattern, however, there is great variance between timing of the individual trips that are used to establish a tie between two actors, so repeated joint movement does not seem to be caused by unrelated individuals simply moving around the virtual environment according to the same external schedule.

Statistical methodology

The statistical methodology I use to test the four hypotheses complements the way in which the dependent variable and network measures were constructed. In my main analysis I examine behavior at the level of the event. I then break down behavior by different subpopulations in three ways: First I separate groups according to whether they are deviant or not. Second I incorporate group-level fixed effects to distinguish between influences on

individuals and groups. Third, I perform analysis separately for subpopulations based on experience. I also include two sets of robustness checks. To address the important issue of causality, I perform longitudinal analysis on a subset of the population which I identified as being forced to change their location in the social structure due to exogenous circumstances. Second, I replicate the main analysis using individuals as the unit of observation, to ensure robustness to the choice of level of analysis.

In examining pod kills, it is important to distinguish between the tendency to destroy an escape pod, specifically, and more general tendency to be involved in conflict, the latter being a regular part of the game and, by its nature, correlated with the amount of time that players spend on the game each day. As I describe above, the dependent variable in the main analysis includes one observation for each instance where a regular vehicle is destroyed. This variable takes the value of one if this is then followed by the destruction of an escape pod, and zero otherwise. The dependent variable therefore explicitly captures an aggressor's propensity to engage in a pod kill conditional on having first destroyed the victim's non-pod ship. I use standard logit regression for the statistical analysis and report White standard errors, which are robust to heteroscedasticity. To allow for non-independent observations within each group, the standard errors are clustered by group identifier.

To examine differences between groups whose local norms are congruent with global norms and groups whose local norms are incongruent with them, I use the *Deviant Group (dummy)* variable and its interaction with within-group and across-group ties. To examine the extent to which particular variables act on an individual versus a group, I add fixed effects for each group. Group fixed effects allow the interpretation of significant coefficients as acting on an individual. To a lesser extent, the observation that coefficients that are significant in the initial

analysis cease to be so when fixed effects are added, can also be taken as a tentative indication that a particular variable influences the group as a whole.

A key consideration in this context is the direction of causality. There are clear theoretical reasons to expect social structure to have an impact on behavior. However, if particular locations facilitate normative violation, or if actors believe that they might, we would also expect those who are interested in violating norms to choose such locations, potentially leading to an observed effect due to reverse causality. To guard against this threat, I performed longitudinal analysis, where I examined whether individual changes in social structure are correlated with changes in behavior. Even a longitudinal result could arguably result from a reverse causal process if individuals, desiring to change their behavior, were to systematically change their location in the social structure to match the behavior they desire to start. Rather than performing the longitudinal analysis using every individual, I attempted to identify individuals who changed their social structure not because they chose to, but because they had to.

To locate individuals who were forced to make changes in their social structure, I ran a set of analyses on the data intended to find groups that were dissolved by their leaders. Although no explicit logs are kept of group dissolutions, I obtained data on participants' historical transfers between groups. By going through this data I was able to indirectly identify groups whose membership suddenly went to zero, and which did not experience any subsequent activity. I was interested in selecting groups that did not obviously close simply because of ongoing attrition, while also ensuring that the sample was not unduly influenced by one or a few large events. I therefore included medium-sized groups that initially had between 5 and 100 members, but whose size then went to zero in five days or less. I chose to focus on such closures because according to my discussions with the game developers and others involved with the game, such

closures tend to be associated with group leaders suddenly having less time to devote to the game because of demands external to the game. The reasons cited for this included changes in employment situation, requests by spouses or family, and in some cases a general feeling that the game was taking too much time for the player. Due to data anonymization, I was unfortunately not able to confirm the circumstances of the specific closures that I identified, but the general pattern of the data was consistent with this interpretation.

Although a number of those who experienced the closure of the group did not remain active after the event, I identified a set of individuals who experienced closures, and whom I could confirm as being active both for at least a month before the closure event, and for at least a month after the event. By calculating the network and group structure for each of these individuals both before and after the closure event, I was able to relate changes in the number of within-group and across-group ties with changes in whether they engaged in pod kills at those points in time.

To assess the robustness of the result with respect to the unit of analysis, I performed comparable analysis at the level of the individual, in which the independent variable is the fraction of that player's kills that are pod kills (rather than destruction of non-pod ships). In the individual-level analysis, I weight the observations by the frequency with which the focal participant engages in the destruction of other vessels.

RESULTS

To illustrate the social structure of EVE Online as a whole, I begin with a descriptive examination of the structure of the "group-level network." That is, I examine the social network as seen when each group is considered as a node, although the underlying structure is based on the ties established between individual players through co-movement. Figure 2, which illustrates

this structure graphically, shows the network structure of the 100 largest player groups (“corporations”) in EVE Online laid out in a three-dimensional space according to the ties each group’s members have with those of all other groups. The volume of each node is proportional to the number of members in each group. The shade of each node represents the density of ties among group members⁶, darker nodes representing groups with denser within-group networks. An interesting feature of this figure is that the groups at the core of the network, where we observe extensive across-group ties, are also the ones whose within-group network is less dense.

[Figure 2 about here]

Having examined the network structure descriptively, I proceed to regression analysis in order to test the hypotheses presented above. Table 2, which demonstrates the relationship between network structure and behavior at the event-level, shows the results of logit regressions that examine the relationship between actors’ individual and network characteristics and the probability that they will destroy an overpowered opponent’s escape capsule.

[Table 2 about here]

The baseline predictions implied by a pure network approach (ignoring the intersection of network and group structure) are presented in model 1, which examines the effects of both *Total number of actor’s ties* and *Actor constraint score*. If a greater number of network relationships led to a reduction in anti-normative behavior, we would expect a negative coefficient for *Total number of actor’s ties*. However, this coefficient is in fact slightly positive, although not significantly so. The coefficient for *Actor constraint score* is also not statistically significant. In

⁶ When measuring density at the group level I use “average number of ties per member” because, in contrast to some alternative density measures, this measure can be meaningfully compared for different groups even when they are different in size.

other words, modeling this process using only network structure without accounting for group structure uncovers no meaningful patterns of behavior.

In model 2, I separate the network measures according to whether the underlying ties span group boundaries or lie within a single group. This separation allows me to test hypotheses 1, as well as hypotheses 2a and 2b. The results of this model differ starkly from those of model 1. The model shows a strong effect of *Ties across group boundaries*, and the negative sign of this coefficient suggests that such ties significantly reduce the propensity for deviation: Individuals who have many relationships with people who do not share the narrow social identity implied by their group membership are less likely to violate the global societal norm. As predicted, *Ties within group boundaries* do not reduce the propensity to violate the norm. In fact, not only is the coefficient not negative – it is significantly *positive*. In other words, the model suggests that within-group ties are not simply neutral with respect to anti-normative behavior; they actually seem to have the effect of *increasing* the propensity for deviation. This model thus both provides support for hypothesis 1 and explains why a model that does not account for group boundaries shows no effect: the positive effect of ties between members of the same group cancels out the negative effect of ties that span group boundaries, resulting in no average effect. A model that accounts only for the impact of network structure without considering how this interacts with group structure thus fails to capture an important aspect of this pattern.

Model 2 also shows the divergent effects of network closure within a single group and when the underlying triads span different groups. The effect of *Actor constraint within-group only* is positive, suggesting that closed networks within groups are conducive to the violation of societal norms, as predicted by hypothesis 2b. This is consistent with the idea that closed networks that reinforce identity boundaries lead to the balkanization of the social structure,

undermining social order at the global level. Notably, this effect exists even after controlling for the simple number of ties that an individual has within the group, suggesting that a closed structure has an impact independent of mere number of network ties that an individual has. The effect of *Actor constraint across-group only* is negative, as predicted by hypothesis 2a, but it falls slightly short of statistical significance.

Model 3, which incorporates additional covariates in the analysis, shows no effect for *Size of actor's group (thousands)*, but the effect of *Actor experience* is positive, suggesting that experienced participants are more likely to engage in pod kills. At first glance, the positive relationship between pod kills and experience might seem surprising because individuals who have participated in the game for a long time might be expected to have internalized the norms more fully. But because the norm in question is a global, institutionally supported norm that is visible even to beginners in the game, this pattern is not inconsistent with the idea presented in hypothesis 4 – that as participants gain experience, their attention to norms turns from the global baseline to local variation. Players might thus be experiencing influence from divergent local subcultures within the game that encourage different behavior. The inclusion of these covariates does not affect the estimates for the number of ties and constraint scores, with the exception that the negative effect of *Actor constraint across-group only* now reaches statistical significance, providing support for hypothesis 2a.

[Figure 3 about here]

The models presented above demonstrate how patterns of within- and across-group ties affect anti-normative behavior overall. They do not, however, explain whether this happens because participants in certain structural locations tend to deviate individually and independently of others, or because particular structural locations lend themselves to the formation of deviant

subcultures. Although both mechanisms lead to reduction in global order and increased aggregate levels of deviance, the specific patterns of deviance implied by each mechanism are different. Figure 3 examines the average levels of deviation in different groups, and demonstrates that groups differ considerably in this respect, suggesting the possibility that the effects observed in models 2 and 3 are mediated by the formation of deviant subcultures in groups with particular relationship patterns. I examine this possibility in models 4 and 5.

Model 4 incorporates a measure of whether an individual is a member of a deviant group or not. Not surprisingly, the coefficient for *Deviant group (dummy)* is positive and significant: If those you identify with tend to violate the global norm, you are more likely to do so as well. More interesting are the estimates for the interactions between this variable and the number of within-group and across-group ties. The coefficient for *Ties within boundaries of deviant group (interaction)* is positive, indicating that in deviant groups, the individuals who are better connected within the group are more likely to engage in antinormative behavior. The coefficient for *Ties across boundaries of deviant group (interaction)* is negative, indicating that for members of deviant groups, having ties that span group boundaries reduces antinormative behavior even more than for members of non-deviant groups. These coefficients thus provide support for both hypotheses 3a and 3b.

The patterns for the non-interacted variables in model 4 are also consistent with expectations. The main effect of *Ties across group boundaries* remains negative and significant, indicating that such ties restrain anti-normative behavior regardless of whether the group that the individual identifies with is deviant or not. However, the results now show no significant main effect of *Ties within group boundaries*. Since the omitted category in this regression is non-deviant groups, this indicates that conditional on a group being non-deviant, within-group ties do

not lead to deviance, although there is no evidence that they hinder it either. A test of the relative values of the within-group and across-group coefficients confirms that the difference between *Ties across group boundaries* and *Ties within group boundaries* remains statistically significant ($p < .01$), confirming that even in non-deviant social groups, social relationships have a very different impact on behavior based on whether they span group boundaries or not.

Model 5 addresses the question of individual or group-level deviance in a different manner, by incorporating fixed effects for each group. This means that any influences that affect the whole group equally are not reflected in the estimated coefficients. When the fixed effects are included, the across-group measures remain significant and negative, suggesting that across-group ties support norms at the individual level, even controlling for any group effects. The within-group measures are now insignificant, suggesting that the positive relationship between within-group ties and antinormative behavior may be driven primarily by the fact that high levels of within-group connectivity facilitate the evolution of divergent norms within the social group as a whole. As before, the model confirms the material difference in the effect of within-group and across-group ties, since the difference between the *Ties across group boundaries* and *Ties within group boundaries* is statistically significant ($p < .01$).

[Figure 4 about here]

In addition to variation between groups, one might expect that socialization patterns would lead to differences between individual behaviors depending on their experience with a given social environment. As I discussed above, model 3 provides initial evidence that this is indeed the case. As stated in hypothesis 4, this might occur if individuals primarily pay attention to global norms upon entering the environment (because they are clearer and easier to observe),

but come to appreciate local nuances as they gain experience. In figure 4, I explore this possibility in more detail.

Figure 4 was constructed by separating players into quintiles according to experience and then estimating the effect of within-group and across-group ties for each quintile. The figure shows that for inexperienced members, *Ties within group boundaries* increase the propensity to engage in anti-normative behavior, whereas *Ties across group boundaries* have a limited effect. For experienced members, however, the pattern is reversed: *Ties across group boundaries* have a significant effect, but the effect of *Ties within group boundaries* is not significant.⁷ This pattern is consistent with hypothesis 4: For newcomers, the default level of attention is the global, and the key determinant of whether they engage in deviant behavior is therefore whether their attention is drawn to local norms because of exposure to a strong subculture through a rich set of internal ties. In contrast, individuals who have spent a long time in the environment are well aware of local nuances and normative variation. For them, the link to the global, society-level, norms that is provided by ties to the broader population outside their own group is therefore especially important.

Robustness checks

The above results paint a consistent picture along multiple dimensions, but below I present additional analyses intended to examine the robustness of the results over different units of analysis, as well as using longitudinal analysis to address issues of causality.

[Table 3 about here]

⁷ The moderating effect of experience on the relationship between network structure and anti-normative behavior is shown by separating the population into subgroups to facilitate interpretation. An alternative method, constructing interaction variables from network variables and experience, yields the same result.

Table 3 shows the results of analysis at the level of the individual rather than of the event. In these models, the dependent variable is the fraction of times that an individual engages in pod kills after destroying a vehicle, and the regressions are estimated using fractional dependent variable GLM. Model 6 includes the basic network variables; model 7 incorporates additional covariates. The results are generally consistent with the results in table 2: The coefficient for *Ties within group boundaries* is positive and the coefficient for *Ties across group boundaries* is negative; both are statistically significant. The results for other variables, including within-group and across-group closure as well as experience, are very similar to those in models 2 and 3.

[Table 4 about here]

In table 4, I report the results of my longitudinal analysis. These models focus on those individuals who I had identified as being forced to change their location in the social structure because the groups they were members of ceased to exist. The outcome of interest is whether those individuals engaged in pod kills before and after the change in their social structure. The models are first difference regressions based on two periods for each individual. The timing of each period is not chosen according to an absolute time, but relative to when the focal individual's group was closed. Thus, for each individual I included one period between 4 and 6 weeks before his or her group was closed, and one period between 4 and 6 weeks after. The independent variables are the change in the number of within-group and across-group ties between the periods.

Model 8 is the baseline model, and the dependent variable is the change in a dummy variable indicating whether an individual engaged in pod kills in each period. The coefficient for *Change in within-group ties* is positive and significant, indicating that when individuals who are forced to find a new group end up with more within-group ties after the move than before it, they

are more likely to start engaging in antinormative behavior and less likely to stop doing so. The coefficient for *Change in across-group ties* is not significant in this specification, which is somewhat surprising given the consistent effects of across-group ties in other models (although the difference between the two coefficients remains significant, $p < .01$). One potential reason for why an effect might not be observed could be that the effect of across-group ties simply takes longer to materialize than the effect of within-group ties. High numbers of ties within a group are likely to be associated with a relatively speedy socialization process, but the effect of changes in across-group ties may occur in a slower and more nuanced manner. Another reason, however, may be that the effects of ties are not symmetric when it comes to starting to engage in anti-normative behavior and when it comes to stopping it.

To examine the possibility of asymmetric effects on the transition into, and out of, anti-normative behavior, model 9 therefore focuses exclusively on changes in the dependent variable that are larger than zero. In this specification, the dependent variable takes the value of one only for individuals who do not engage in pod kills in the first period, but who do engage in them in the second period (all other cases are coded as zero). In this model, the coefficient for *Change in across-group ties* is negative and significant. That is, the more across-group ties that an individual has in the second period relative to the first period, the less likely he or she is to start engaging in anti-normative behavior. The coefficient for *Change in within-group ties* remains positive and significant.

Model 10 focuses on changes in the dependent variable that are smaller than zero: Cases where individuals engage in anti-normative behavior in the first period but not in the second are assigned a value of negative one, but all other cases are assigned a value of zero. Since transitions out of anti-normative are assigned a negative value, the direction of estimated effects

is comparable with models 8 and 9. In this model, the coefficient for *Change in within-group ties* remains positive, but *Change in across-group ties* is not significant. Thus, an increase in the number of within-group ties is associated with a smaller chance of transitioning out of anti-normative behavior, but changes in across-group ties do not seem to cause individuals who have already started engaging in antinormative behavior to stop it. Taken together, these results show a consistent pattern of across-group ties being associated with less anti-normative behavior, and with a reduced tendency to initiate such behavior. In contrast, within-group ties are associated with more anti-normative behavior, an increased tendency to initiate it, and a decreased tendency to cease it.

DISCUSSION

Over the years, scholars of social constraints on behavior have emphasized different aspects of the social structure. Polanyi (1944; 1957) based his analysis on a *matrix of institutions* as a foundation of social structure. In contrast, Simmel (1964) took the *web of group affiliation* as a starting point, and a third view of social structure was proposed by Granovetter (1985) in considering the *network of interpersonal relations* to be at the heart of social influence. In this paper, I present evidence that interpersonal ties and affiliations with identity-supplying groups do not act independently to constrain behavior. Rather, the effect of each depends crucially on the other. In fact, when group membership is ignored, the analysis shows no effect of network ties on anti-normative behavior. When group boundaries are considered, however, the conditional effects of the network structure are clearly exposed.

In accordance with established network theory, I find that network ties do have an important effect on compliance with the norms that are supportive of social order within the environment I examine. But it is only ties that span group boundaries that reduce the tendency to

violate these norms. When ties lie entirely within the same group, I find either the opposite relationship or no relationship at all, and a comparison of the estimated coefficients shows a significant and consistent difference between the effect of within-group and across-group ties. The effect seems to be even stronger when the pattern of ties is characterized by network closure: even after controlling for the effect of direct ties, within-group and across-group patterns of network closure have independent effects that parallel those for direct ties. The results suggest that dense networks within a group with a shared social identity, coupled with the absence of ties that span group boundaries, may be the social structure most conducive to local isolation from global norms of conduct.

The relationship between local and global norms is further illuminated by an examination of how the effects of social relationships vary between groups. I find the divergent effects of within-group and across-group ties to be even more pronounced for deviant groups than for groups whose local norms are congruent with global norms. The analysis also suggests that although across-group ties reduce antinormative behavior by directly influencing individuals, the positive relationship between within-group ties and deviance may result from the fact that such ties facilitate the development of divergent local norms at the level of the group. Taken together, these results are broadly consistent with the idea that the linkage between local and global norms is most effective when identity-supplying social groups are combined with interpersonal relationships that span group boundaries.

The evolution of behavior over time provides further support for this interpretation. I examine this evolution in two different ways, by comparing cohorts with different levels of experience with the social environment, and by performing longitudinal analysis on a select set of individuals who have been forced to change their social structure due to external factors. The

cohort analysis shows that for individuals who have recently joined the game, ending up in dense within-group networks is especially predictive of norm violation, whereas for more experienced participants it is the absence of across-group ties that leads to antinormative behavior. This is the pattern we would expect if, when participants first join the game, they tend mainly to be attentive to the global norms of this new environment; it is only when such novices find themselves in the presence of a strong and consistent influence from a divergent subculture that they are likely to pay attention to local norms of deviance. As participants gain more skill and confidence, the default scope of attention shifts from global norms to the local norms of the particular part of the environment in which they find themselves, and the key determinant therefore becomes the presence or absence of ongoing social influence from the broader environment.

The longitudinal analysis suggests that when individuals must find a new place in the social structure because the groups to which they belong are eliminated, they are more likely to start engaging in antinormative behavior if their change in circumstances leads to an increase in within-group ties, or if it leads to a decrease in across group ties. Within-group ties also seem to influence the probability that individuals stop engaging in antinormative behavior, and the divergent effects of within-group and across-group ties are evident in these models as well. These longitudinal results are consistent with both the main analysis and the comparison of different cohorts. Taken together, the dynamic patterns of behavior paint a coherent picture of the joint influence of social relationships and social groups. These analyses also support the causal interpretation of the results, since alternative causal mechanisms would be unlikely to give rise to the persistent patterns I observe.

This examination of the joint influence of networks and groups is conducted in a particular social context, and naturally, the question arises to what extent the results are

applicable to different environments. As Wiesenfeld, Raghuram and Garud (1999) have noted, virtual ways of organizing are increasingly prevalent, and contexts where interpersonal interaction is electronically mediated in one way or another are becoming more important for the material outcomes of firms and individuals. Knorr Cetina and Bruegger (2002) demonstrated that even minimally interactive environments, such as the discussion boards of foreign exchange traders, facilitate strong norms and effective sanctioning mechanisms, and an emerging body of research is beginning to shed light on social processes in other types of mediated environments, including Facebook (Wimmer and Lewis, 2010), Second Life (Burt, 2010), and Wikipedia (Piskorski and Gorbatai, 2011). Ideally, research on mediated interaction should proceed on two fronts. On one hand, such research can advance the theoretical understanding of social processes in a general manner. On the other hand, it is important to understand how and when the social mechanisms that govern traditional interaction are transferred to electronically mediated interaction and vice versa. The current paper contributes to each of these goals: It documents that established aspects of social structure, such as social networks and social groups, influence behavior in a purely virtual environment, while at the same time providing new insights into how different components interact with each other.

The empirical opportunities of electronically mediated interaction can arise from the fact that it often leaves a trail of data that is more extensive than the one left by traditional interaction. Other opportunities arise, perhaps counterintuitively, from the fact that the number of salient social dimensions available to participants is often limited in mediated environments. For instance, the current empirical scenario offers a distilled relationship between networks and groups. Because the environment, by design, constrains each participant to one group, the boundaries between groups are spanned only by interpersonal relationships. In settings where

individuals are members of multiple social groups, a relationship may span the boundaries of one group but lie within another, and important boundary spanning may occur because an individual is a member of multiple groups. Furthermore, group membership can also in many settings be either uncertain or partial, leading to a different view of what exactly it means to span group boundaries. Such settings offer a richer, albeit more complex, scenario for examining how different components of social structure interact, and provide a promising venue for future research.

A more nuanced view of the complex interrelations between institutions, identity, and networks can yield important benefits for both organizational theory and practice. At a practical level, this understanding can inform interventions within organizational environments in material ways. Although feasible in some circumstances, interventions in interpersonal relationships can be complicated and difficult to manage (Hunter and Krackhardt, 2007). In some instances, interventions that target group memberships and identity may be more feasible than those that target the social ties themselves. A better understanding of the conditionality of network effects will enable individuals to consider their network ties in terms of how they interact with other elements of the social environment rather than focus solely on the network structure. The benefits of such understanding, however, are not limited to intervention. The ability to predict processes related to anti-normative behavior may be extremely important as well. Collingsworth and Menezes (2009) show how the pattern of email interactions between Enron employees changed markedly in the period leading up to its failure, and how such changes could potentially provide an early warning system for organizations. Collingsworth and Menezes focus on the social tension associated with the imminent troubles, but the benefit of such an approach can be

increased by coupling analysis of network structure with analysis of groups to detect social structures associated with a high risk of deviant behavior.

At a theoretical level, attention to different aspects of social structure will yield a deeper understanding of the underlying processes that drive classical results within the field of social networks. Rather than treating reciprocal interactions between individuals as substitutes to identity perceptions and cultural influences, we should treat them as tightly coupled and complementary components of social interaction in a general sense. Such a view not only provides opportunities for applying network approaches and insights to new fields of inquiry, but also has the potential to resolve some of the issues on which social networks studies give inconsistent results. It is true that individuals are embedded in social networks that influence and constrain them, but the network of interpersonal relationships is itself embedded in a fabric of other social structures, which networks research too often treats as inconsequential. A greater appreciation of this type of embeddedness will enrich both social network theory and structural sociology.

REFERENCES

- Bernhard, Helen, Ernst Fehr, and Urs Fischbacher. 2006. "Group Affiliation and Altruistic Norm Enforcement." *The American Economic Review* **96**(2): 217-221.
- Blackburn, Tom, and Eric Hammel. 2006. *The Jolly Rogers: The Story of Tom Blackburn and Navy Fighting Squadron VF-17*. St. Paul, MN: Zenith Press.
- Blumer, Herbert. 1969. *Symbolic Interactionism: Perspective and Method*. Berkely, CA: University of California Press.
- Boellstorff, Tom. 2008. *Coming of Age in Second Life: An Anthropologist Explores the Virtually Human*. Princeton, NJ: Princeton University Press.
- Breiger, Ronald L. 1974. "The Duality of Persons and Groups." *Social Forces* **53**(2): 181-190.
- Burt, R. S. 1999. "Entrepreneurs, Distrust, and Third Parties: A Strategic Look at the Dark Side of Dense Networks." In *Shared Cognition in Organizations: The Management of Knowledge*, edited by Leigh L. Thompson and John M. Levine. Mahwah, NJ: Lawrence Erlbaum Associates: 213-243.
- Burt, Ronald S. 2000. "The Network Structure of Social Capital." *Research in Organizational Behavior* **22**: 345-423.
- Burt, Ronald S. 2001. "Bandwidth and Echo: Trust, Information, and Gossip in Social Networks." In *Networks and Markets*, edited by James E. Rauch and Alessandra Casella. New York: The Russel Sage Foundation: 30-74.
- Burt, Ronald S. 2010. "Structural Holes in Virtual Worlds." *Working Paper*. University of Chicago Booth School of Business, Chicago, IL.
- Burt, Ronald S. 2011. "Network-Related Personality and the Agency Question: Multi-Role Evidence from a Virtual World." *Working Paper*. University of Chigago.
- Burt, Ronald S. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* **110**(2): 349-399.
- Coleman, James S. 1988. "Social Capital in the Creation of Human-Capital." *American Journal of Sociology* **94**: S95-S120.

- Coleman, James S. 1990. *Foundations of social theory*. Cambridge, MA: Harvard University Press.
- Collingsworth, Ben, and Ronaldo Menezes. 2009. "Identification of Social Tension in Organizational Networks." In *Complex Networks*, edited by Santo Fortunato, Giuseppe Mangioni, Ronaldo Menezes and Vincenzo Nicosia. Berlin: Springer **207**: 209-223.
- Csárdi, Gábor, and Tamás Nepusz. 2006. "The igraph software package for complex network research." *Software Package*. InterJournal Complex Systems, 1695.
- Emirbayer, Mustafa, and Jeff Goodwin. 1994. "Network Analysis, Culture, and the Problem of Agency." *American Journal of Sociology* **99**(6): 1411-1454.
- Entwisle, Barbara, Katherine Faust, Ronald R Rindfuss, and Toshiko Kaneda. 2007. "Networks and Contexts: Variation in the Structure of Social Ties." *American Journal of Sociology* **112**(5): 1495-1533.
- EVE Forum. 2009. "Alliance and Corporation Recruitment Center." Retrieved August 2009, from <http://www.eveonline.com/ingameboard.asp?a=channel&channelID=109585>.
- Fernandez, Roberto M., Emilio J. Castilla, and Paul Moore. 2000. "Social Capital at Work: Networks and Employment at a Phone Center." *The American Journal of Sociology* **105**(5): 1288-1356.
- Goette, Lorenz, David Huffman, and Stephan Meier. 2011. "The Impact of Social Ties on Group Interactions: Evidence of Minimal Groups and Randomly Assigned Real Groups." *American Economic Journal: Microeconomics* **Forthcoming**.
- Goffman, Erving. 1971. *Relations in Public: Microstudies of the Public Order*. New York: Basic Books.
- Granovetter, Mark. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." *American Journal of Sociology* **91**(3): 481-510.
- Granovetter, Mark. 2005. "The Impact of Social Structure on Economic Outcomes." *The Journal of Economic Perspectives* **19**(1): 33-50.
- Hechter, Michael, and Satoshi Kanazawa. 1993. "Group Solidarity and Social Order in Japan." *Journal of Theoretical Politics* **5**(4): 455-493.
- Homans, George Caspar. 1950. *The Human Group*. New York: Harcourt, Brace & World.

- Hunter, Keith O., and David Krackhardt. 2007. "Structural Facilitators of Organizational Network Change." *Paper presented at the 2007 Academy of Management Annual Meeting*. Carnegie Mellon University.
- Knorr Cetina, Karin, and Urs Bruegger. 2002. "Global Microstructures: The Virtual Societies of Financial Markets." *The American Journal of Sociology* **107**(4): 905-950.
- Lyons, Christopher J. 2007. "Community (Dis)Organization and Racially Motivated Crime." *American Journal of Sociology* **113**(3): 815-863.
- Markham, Annette N. 1998. *Life Online: Researching Real Experience in Virtual Space*. Walnut Creek, CA: Altamira Press.
- Mehra, Ajay, Martin Kilduff, and Daniel J. Brass. 1998. "At the Margins: A Distinctiveness Approach to the Social Identity and Social Networks of Underrepresented Groups." *The Academy of Management Journal* **41**(4): 441-452.
- Moody, James, and Douglas R. White. 2003. "Structural Cohesion and Embeddedness: A Hierarchical Concept of Social Groups." *American Sociological Review* **68**(1): 103-127.
- Nee, Victor, and Paul Ingram. 1998. "Embeddedness and Beyond: Institutions, Exchange, and Social Structure." In *The New Institutionalism in Sociology*, edited by Mary C. Brinton and Victor Nee. New York: Russell Sage Foundation: 19-45.
- Pettigrew, Thomas F. 1998. "Intergroup Contact Theory." *Annual Review of Psychology* **49**: 65-85.
- Pettigrew, Thomas F., and Linda R. Tropp. 2006. "A Meta-Analytic Test of Intergroup Contact Theory." *Journal of Personality and Social Psychology* **90**(5): 751-783.
- Piskorski, Mikołaj Jan, and Andreea Gorbatai. 2011. "Testing Coleman's Social-Norm Enforcement Mechanism: Evidence from Wikipedia." *Working Paper*. Harvard University.
- Polanyi, Karl. 1944. *The Great Transformation: The Political and Economic Origins of Our Times*. Boston, MA: Beacon Press.
- Polanyi, Karl. 1957. "The Economy as Instituted Process." In *Trade and Market in Early Empires*, edited by Karl Polanyi, Conrad Arensberg and Harry Pearson. Glencoe, IL: Free Press.
- Putnam, Robert D. 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon & Schuster.

- Reagans, Ray, Ezra Zuckerman, and Bill McEvily. 2004. "How to Make the Team: Social Networks vs. Demography as Criteria for Designing Effective Teams." *Administrative Science Quarterly* **49**(1): 101-133.
- Simmel, Georg. 1964. *Conflict & The Web of Group Affiliations*. New York: The Free Press.
- Tajfel, Henri, and John C. Turner. 1979. "An Integrative Theory of Intergroup Conflict." In *The Social Psychology of Intergroup Relations*, edited by William G. Austin and Stephen Worchel. Monterey, CA: Brooks/Cole: 33-47.
- Tilly, Charles. 2005. *Identities, Boundaries, and Social Ties*. Boulder, CO: Paradigm Publishers.
- Tocqueville, Alexis de. [1835] 2000. *Democracy in America*. Chicago: The University of Chicago Press.
- Uzzi, Brian. 1997. "Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness." *Administrative Science Quarterly* **42**(1): 35-67.
- Westphal, James D., and Laurie P. Milton. 2000. "How Experience and Network Ties Affect the Influence of Demographic Minorities on Corporate Boards." *Administrative Science Quarterly* **45**(2): 366-398.
- Wiesenfeld, Batia M., Sumita Raghuram, and Raghu Garud. 1999. "Communication Patterns as Determinants of Organizational Identification in a Virtual Organization." *Organization Science* **10**(6): 777-790.
- Wimmer, Andreas, and Kevin Lewis. 2010. "Beyond and Below Racial Homophily: ERG Models of a Friendship Network Documented on Facebook." *American Journal of Sociology* **116**(2).

FIGURE 1 – WITHIN-GROUP AND ACROSS-GROUP TIES




Figure 1 demonstrates the intersection between group structure and network structure. Diagram A shows the members of three social groups (circles/squares/hexagons) and their interpersonal relationships. Diagram B highlights the group boundaries, as well as the distinction between across-group ties (jagged lines) and within-group ties (straight lines).

FIGURE 2 – STRUCTURE OF RELATIONSHIPS AMONG THE LARGEST FORMAL GROUPS IN EVE ONLINE




Figure 2 shows the network structure of the 100 largest player groups (“corporations”) in EVE Online. Each node in the graph represents a single group. The across-group ties in this graph are derived directly from the ties between group members (formed through joint travel). The volume of each node is proportional to the number of members in each group. The shade of each node represents the density of ties between members within the group (darker nodes represent denser within-group networks).

FIGURE 3 – HISTOGRAM OF POD KILL PREVALENCE BY GROUPS




Figure 3 demonstrates how the likelihood that a regular kill will be followed by a pod kill varies between groups. The histogram shows the fraction of regular kills that are followed by pod kills in all groups whose members engaged in more than ten regular kills in the observation period.

FIGURE 4 – ESTIMATED COEFFICIENTS BY PLAYER EXPERIENCE




Figure 4 shows how the effect of ties varies with player experience. The figure separates the players into quintiles according to experience. The graph shows that for new players (Q1), across-group ties have no discernible effect, but having more ties within the group significantly increases the tendency to engage in anti-normative behavior. For highly experienced players (Q5), within-group ties show no effect, but for these players having more ties that span group boundaries significantly decreases the tendency to engage in anti-normative behavior.

TABLE 1 – CORRELATIONS AND SUMMARY STATISTICS

	Mean	S.D.	Min	Max
(1) Regular kill followed by pod kill (dummy)	0.14	0.35	0.00	1.00
(2) Number of actor's within-group ties (logged)	1.46	1.03	0.00	4.39
(3) Number of actor's across-group ties (logged)	2.32	1.42	0.00	6.65
(4) Actor constraint score	0.23	0.29	0.00	1.13
(5) Actor experience	34.2	19.5	0.06	103.6
(6) Size of actor's group (thousands)	0.29	2.13	0.00	48.23
(7) Actor is an isolate (dummy)	0.19	0.40	0.00	1.00

	(1)	(2)	(3)	(4)	(5)	(6)
(1) Regular kill followed by pod kill (dummy)						
(2) Number of actor's within-group ties (logged)	.011					
(3) Number of actor's across-group ties (logged)	-.013	.420				
(4) Actor constraint score	.009	-.506	-.750			
(5) Actor experience	.038	.121	.040	-.045		
(6) Size of actor's group (thousands)	-.017	-.099	-.037	.070	-.083	
(7) Actor is an isolate (dummy)	.022006	.076

Table 1 includes summary statistics and correlations of key variables.

TABLE 2 – PREDICTORS OF ANTI-NORMATIVE BEHAVIOR

Variable	Model 1	Model 2	Model 3	Model 4	Model 5
Total number of actor's ties	0.141 (0.081)				
Actor constraint score	0.166 (0.091)				
Ties within group boundaries		0.197* (0.084)	0.188* (0.084)	0.061 (0.034)	0.037 (0.043)
Ties across group boundaries		-0.207** (0.080)	-0.226** (0.078)	-0.114*** (0.032)	-0.136** (0.044)
Actor constraint within-group only		0.200* (0.096)	0.200* (0.094)	0.117** (0.036)	-0.016 (0.053)
Actor constraint across-group only		-0.159 (0.084)	-0.180* (0.083)	-0.129*** (0.033)	-0.140** (0.044)
Actor experience			0.005*** (0.001)		
Size of actor's group (thousands)			-0.017 (0.010)		
Deviant group (dummy)				0.796*** (0.049)	
Ties within boundaries of deviant group (interaction)				0.049* (0.023)	
Ties across boundaries of deviant group (interaction)				-0.039** (0.015)	
Constant	-1.860*** (0.063)	-1.783*** (0.087)	-1.932*** (0.085)	-2.253*** (0.043)	-2.222*** (0.045)
Log-likelihood	-70877	-51556	-51486	-50474	-49222
Observations	175944	127073	127073	127073	127073

Robust standard errors in parentheses

* p<.05, ** p<.01, *** p<.001

Table 2 reports the predictors of anti-normative behaviors. The table shows the results of a logit regression wherein actors' individual, network, and group characteristics are used to predict the probability that they will destroy the escape capsule of an overpowered opponent. Robust standard errors of estimates, clustered by individuals' groups, are displayed in parentheses. Model 5 also includes fixed effects for individuals' groups. Because their distribution is skewed, all network variables are logged.

TABLE 3 – INDIVIDUAL-LEVEL ANALYSIS

Variable	Model 6	Model 7
Ties within group boundaries	0.197* (0.084)	0.188* (0.084)
Ties across group boundaries	-0.207** (0.080)	-0.226** (0.078)
Actor constraint within-group only	0.200* (0.096)	0.200* (0.094)
Actor constraint across-group only	-0.159 (0.084)	-0.180* (0.083)
Actor experience		0.005*** (0.001)
Size of actor's group (thousands)		-0.017 (0.010)
Constant	-1.783*** (0.087)	-1.932*** (0.085)
Log-likelihood	-4606	-4598
Observations	14757	14757

Robust standard errors in parentheses

* p<.05, ** p<.01, *** p<.001

Table 4 explores the tendency to engage in pod kills on the level of the individual. The dependent variable is the fraction of times an individual engages in pod kills. The regression model is a fractional dependent variable GLM regression. Robust standard errors of estimates, clustered on individuals' groups, are displayed in parentheses. Because their distribution is skewed, all network variables are logged.

TABLE 4 – LONGITUDINAL ANALYSIS

Variable	Model 8	Model 9	Model 10
Change in within-group ties	0.162** (0.055)	0.086* (0.040)	0.075* (0.030)
Change in across-group ties	-0.037 (0.045)	-0.060* (0.026)	0.023 (0.032)
Change in group size	0.015 (0.108)	0.067 (0.058)	-0.052 (0.056)
Constant	-0.110 (0.067)	0.101** (0.036)	-0.211*** (0.053)
R-squared	0.15	0.18	0.10
Observations	112	112	112

Robust standard errors in parentheses

* p<.05, ** p<.01, *** p<.001

Table 3 examines how a forced change in an individual's social structure affects his or her behavior. The dependent variable is the change in a dummy variable indicating whether an individual engaged in pod kills (in each period). Model 8 examines all changes in pod killing behavior. Model 9 focuses exclusively on changes larger than zero (transition into pod killing). Model 10 focuses exclusively on change smaller than zero (transition out of pod killing). All individual variables are differenced to capture within-subject variation.