

Chapter 6

The Relational Foundation of Norm Enforcement

Christine Horne

What Are Norms?

Norms are rules, about which there is some degree of consensus, that are socially enforced. Norms therefore overlap with, but are distinct from, internal states such as values or morals. The key element of norms that distinguishes them from internal states is their social nature—the fact that they are enforced externally by individuals.

Further, norms are not simply descriptive. That is, norms are not equivalent to the sum of behaviors in a group—the frequency or the typicality of a behavior. Patterns of behavior may provide information about what the norms are. But, in order for norms to exist, sanctioning must occur.

This conceptualization of norms means that in order to fully understand them, we have to explain why they are enforced. It is not obvious why people punish. Sanctioning is costly. It can take time and effort. It can be embarrassing and provoke retaliation. So why do it? Existing explanations focus on the characteristics of behavior (see, for example, Coleman, 1990) and the human brain—finding evidence, for example, that harmful behaviors make us angry and motivate us to punish (Fehr & Gächter, 2002). These explanations contribute to our understanding of sanctioning but still leave unanswered questions. If all that mattered was the characteristics of behavior and the anger it provoked, then we would expect the same behavior to be treated the same way in every time and place. But, it is not. Just as the state enforces laws more in some neighborhoods than in others, so do people enforce norms more in some groups and contexts than others. Harmful behaviors are sometimes ignored, and harmless behaviors are sometimes punished. In some social environments, even cooperative behavior may be sanctioned

C. Horne (✉)

Department of Sociology, Washington State University, Pullman, WA, USA
e-mail: chorne@wsu.edu

(Herrmann, Thöni, & Gächter, 2008). Something more than simply the harm caused by a behavior or emotional reactions to it must be driving enforcement.

The view of norms as social (rather than as something that exists inside individuals' heads or as simply patterns of behavior) further implies that they are a group-level rather than an individual-level phenomenon. They emerge in groups and are maintained by groups. Thus to understand norms, we need to study not just sanctioning, but sanctioning within groups.

Evidence for the Relational Foundation of Norms

I argue that understanding of social relationships is essential for explaining norm enforcement. I have developed a set of theoretical predictions and tested them in a series of laboratory experiments. Below I describe the theory and the experimental evidence that shows how social relationships affect sanctioning and how sanctioning can affect social relationships.

Social Relationships and Norm Enforcement

I focus on one key characteristic of social relationships—dependence—as well as on two other social factors—metanorms and metanorm expectations. For each of these three factors I present the theoretical argument, briefly describe the experiments testing the theory, and summarize the results.

Dependence and Sanctioning Benefits

Dependence refers to the extent to which an individual values his or her relationship with another person and the goods that he or she can get from that relationship (Emerson, 1962, 1972; Molm, 1997; Molm & Cook, 1995).¹ The more that individuals in a group depend on one another, the more interdependent they are and the more cohesive the group (Emerson, 1962, 1972).

What is the connection between dependence and sanctioning? Researchers frequently assume that people will punish behaviors that cause harm (see, for example, Coleman, 1990). On this view, enforcing norms produces direct benefits (a reduction in antisocial behavior) for all those affected by the target behavior. Those direct benefits provide an incentive to sanction. But when group members are interdependent, enforcing norms can also produce indirect benefits (Horne, 2004).

¹The definitions of dependence and cohesion used here are drawn from Emerson and Molm. Their work is part of a larger body of research on exchange developed by theorists Blau (1964), Homans (1974), and Kelley and Thibaut (1978).

This occurs when the gains that people experience as a result of deviance being discouraged increase their ability to exchange. In other words, people benefit directly when harmful behavior is punished. They also receive indirect benefits when they interact with others who have gained from the punishment of harmful behavior. This is because when individuals are dependent on those around them, their well-being is tied to the well-being of those others. They benefit when those with whom they interact have something to offer. If the other has few resources, the individual will not be able to gain much through exchange with that other—no matter how much he or she needs what the other has, the other will not have much to give. Thus individuals benefit when they are not personally victimized; they also gain when their neighborhood is safe and secure. Individuals prosper when they are not cheated; they also profit when levels of cheating are relatively low.

The fact that an individual's well-being is connected to the well-being of others means that the gains associated with sanctioning may be larger than they appear on their face (Horne, 2004). It also means that as patterns of social relationships change, the benefits associated with sanctioning shift, and sanctioning behaviors shift as well—even if the target behavior and the harm that it causes remain the same (Horne, 2008).

Further when individuals are dependent on others, they may enforce norms that benefit those others, even if they personally would prefer a different norm to be in effect. In some situations, everyone agrees on the harm caused by a behavior. In others, people have different interests in a behavior (or different understandings of its consequences). When this occurs, dependence relations can lead people to enforce norms they do not prefer.

In other words, even when there is a rule forbidding a harmful behavior and people disapprove of it, punishment of that behavior will vary. When social relationships are strong sanctioning will increase; when they are weak sanctioning will decline (Horne, 2001, 2008). We cannot assume that people will react negatively to harmful behavior or that a normative rule will be enforced consistently across social settings. Instead, the extent to which norms are enforced varies with the structure of social relationships (Horne, 2007, 2008). Accordingly, we would expect norms to grow and fade as social relationships shift.

To test these ideas I conducted two experiments using a norms game (Horne, 2008). Four subjects played a public goods game in which they had opportunities to contribute to one or more group funds. Individuals who contributed bore the costs, but all group members benefitted. Thus each individual hoped that others would contribute but also experienced the temptation to free-ride. Following each contribution decision points in the group fund were distributed to group members. Then all group members engaged in exchange—each person made decisions about how many of their points they wanted to keep, and how many they want to give to each other group member. They could adjust the number of points they gave to others based on whether those others had contributed to the group fund. The difference in what subjects gave to those who donated to the group compared with those who had the opportunity to do so but did not constituted a sanction. If participants gave more points to those who contributed than to those who did not, then they were enforcing a norm-favoring contribution. Participants played this game for a large number of rounds.

I manipulated the benefits associated with contributing to the group (and accordingly, the benefits group members would receive if sanctions encouraged people to contribute). When benefits were low, an individual's contribution produced only a few (two) points for each group member. When benefits were high, the individual's contribution produced a larger number of points (six). Participants gained more from others' contributions in the large than small benefit condition and therefore presumably had more interest in seeing that failures to contribute were sanctioned.

I also manipulated the extent to which group members were dependent on each other. I did this by varying the value of points that individuals received from others compared to the value of points in the individuals' own personal accounts. In the low-dependence condition, points that an individual received from others were worth the same as their own points. Participants could do just as well on their own as they did interacting with others. In the high-dependence condition, points that the individual received from others were worth three times their own points. Participants earned more points if they exchanged with others than if they did not.

The results showed that the size of the benefits associated with donating to the group fund (or the harm to group members when donations were not made) did not, in and of itself, affect sanctioning (Horne, 2008). That is, sanctions were not necessarily stronger when the consequences of the target behavior were larger. Rather, the consequences of the behavior interacted with the level of interdependence in the group such that sanctioning was greatest in groups in which the consequences of donating to the group were large *and* group members were highly interdependent.

Further, variation in the structure of dependence relations within a group affected patterns of sanctioning (Horne, 2008). In some conditions, subjects had conflicting interests in the group funds. In those conditions, individuals who were dependent on other group members tended to enforce norms that benefitted those others, rather than the norms they personally preferred. (Interestingly, although subjects in interdependent groups enforced norms preferred by others, they did not necessarily follow them.)

The results are consistent with the argument that interdependence among group members magnifies the benefits of sanctioning, in turn strengthening enforcement. They demonstrate that social relationships matter for norm enforcement.

Metanorms

Norm enforcement differs from punishment. Anybody can punish anyone for anything. But norm enforcement requires some element of consensus within the group. Consensus might arise if the target behavior affects all group members in the same way. If everyone has the same interest in a behavior, we would expect to see all those in the same situation react similarly. But consensus may have a more social component. Individuals care about what others think of them. They want others to cooperate with them. To encourage them to do so, the individual needs to demonstrate that he or she is a good person with whom to interact. Therefore, people will try to behave in ways that will maintain relationships and lead others to treat them positively rather than negatively.

What can people do to demonstrate that they are good exchange partners and good group members? One thing they can do is follow norms. When individuals obey group norms, they demonstrate that they know how to behave. Their actions provide evidence of their reliability and trustworthiness.

In addition to following norms, people can enforce them. If an individual punishes behavior that other group members would like to see punished, then he or she is demonstrating that he or she understands what the group norms are. Further, he or she is providing evidence of his or her commitment to the norm. He or she is establishing that he or she is willing to bear personal costs to enforce it (Posner, 2000). And he or she is showing that he or she is not just a poser—imitating others for the sake of popularity (Centola, Willer, & Macy, 2005). People demonstrate a commitment to honesty when they themselves are honest, but they also demonstrate that commitment when they punish deceit, blow the whistle on bad behavior in the workplace, and so forth.

Individuals who want to be treated well need to consider how their actions—including their sanctioning behavior—will be seen by those around them. When thinking about enforcing a norm, they will take into account the costs (potential retaliation, emotional discomfort, and so forth) and the benefits (including a reduction in deviant behavior). But they will also consider how others are likely to view their sanctioning activity. They anticipate potential reactions. In other words, they pay attention to metanorms (Horne, 2001).

Metanorms are a particular kind of norm that regulate sanctioning (Axelrod, 1985; Coleman, 1990). Like norms, they are socially enforced. The incentives provided by metanorms are selective—given only to the sanctioner. While the benefits of sanctioning are experienced by everyone, thus tempting people to free-ride, metanorms produce consequences only for the person who imposes the punishment. So, only the person who actually sanctions is rewarded.

Why do people enforce metanorms? Why do they reward sanctioners? Because no one, including sanctioners, wants just a fair-weather friend. Everyone maintains relationships that support them at some times but make demands on them at others. If people want to maintain relationships, then they stick with them through the profitable times as well as those times when the other has little to offer. The same is true of relationships with sanctioners. Individuals provide support to the sanctioner because they value the relationship. If they fail to be supportive, and some other group member remains loyal, then in the future the sanctioner is likely to defect to this more faithful acquaintance. The motivation to maintain ties is stronger when a relationship is valued. The more dependent people are on the sanctioner, the more they will want to support him or her. Thus dependence between group members increases the support given to sanctioners; it strengthens metanorms. In turn, metanorms affect norm enforcement (Horne, 2001, 2004).

To test these ideas, I conducted four experiments using a metanorms game (Horne, 2001, 2004, 2007; Horne & Cutlip, 2002). In this game, a computer-simulated thief stole from group members. Each time an individual was the victim of a theft, he or she could decide how to respond—whether to punish the thief or not. Group members also had opportunities to exchange with one another.

They could express their approval or disapproval of a victim's sanctioning decision by giving that person more or fewer points. The more points that participants gave to victims who punished the thief relative to those victims who did not punish, the stronger the metanorm-favoring punishment.

As in the norms game, I manipulated how dependent group members were on each other. I also manipulated the costs and benefits of sanctioning to see if metanorms could encourage people to sanction even when doing so imposed more costs than benefits on the group.

The results showed that when people were interdependent, they gave more support to sanctioners. They gave larger rewards to those who punished relative to those who did not. That is, interdependence strengthened metanorms (Horne, 2001, 2004, 2007; Horne & Cutlip, 2002).

In turn, when metanorms were strong, people were more likely to sanction. Groups with stronger metanorms had higher rates of punishment (Horne, 2001, 2004, 2007; Horne & Cutlip, 2002).

Further, metanorms encouraged people to sanction even when the costs of doing so were high and the benefits low (Horne, 2007). In fact, as the costs of sanctioning increased, the rewards given to sanctioners increased as well. Even when punishment was so costly that it produced an aggregate loss for the group, people who were highly dependent on each other rewarded such punishment. This encouragement in turn increased the rates of sanctioning in the group.

These results show that the same behavior that causes the same harm will be treated differently depending on the structure of social relationships. A behavior may be punished in one social environment and not in another. This is not because people are any less disapproving of the behavior. Rather, it is because the social relationships that support sanctioning are weak. Further, in groups in which members are dependent on each other, people may provide support to sanctioners that encourages them to punish deviance even when doing so is both individually and collectively irrational.

Metanorm Expectations

Rewards encourage sanctioning. But even anticipation of others' likely reactions may affect punishment decisions. Because people want to be rewarded, they try to determine what behaviors others would like to see punished and the punishment efforts that others will view positively. Thus, in addition to actual rewards and punishments, people's expectations about what sanctioning behaviors others are likely to approve also drive sanctioning (Willer, Kuwabara, & Macy, 2009).

The problem for the individual is that it is not always clear exactly what others want. In forming expectations therefore, people rely on a number of clues. One clue is the harm caused by a behavior. It is reasonable to think that if a behavior hurts others, those others would like to see it punished. Another clue is the frequency of behavior (Horne, 2009b). If the individual sees many others engaging in a particular

behavior, he or she might well conclude that others approve of the behavior and would disapprove of aberrations. People may also rely on the characteristics of the setting (Horne, 2010). It is widely known, for example, that informal control of criminal and deviant behavior varies across neighborhoods. Neighborhoods with certain characteristics (high poverty, low stability, and so forth) have lower levels of informal control and higher rates of crime. Why is this? One possibility is that people in those neighborhoods do not expect that others will support their sanctioning efforts.

I conducted several experiments that test the arguments that existing patterns of behavior and characteristics of the setting create expectations about others' potential reactions and that those expectations affect sanctioning decisions (Horne, 2009b, 2010).

The first two experiments test the argument that people use the typicality of behavior as a clue to help them anticipate others' potential reactions (Horne, 2009b). I created an expectations game in which each subject participated in a group with seven simulated actors. The actors took turns making a particular choice—the choice between X and W. This decision was as arbitrary as it sounds. Subjects literally had to choose between the two letters. The X–W choice had no consequences in and of itself. It had no association with status, aesthetic judgment, norms, or any other evaluation outside the lab. The point was to create an artificial behavior with no or as little as possible existing social meaning. The only factor that might make X or W more socially salient was the number of actors in the group who chose it.

The actors made their X–W choices one at a time. The subject went seventh. This meant that he saw all but one actor's choice before making his own. After the last actor made his X–W decision, everyone was able to react to each other's choices by giving them points. This time the subject went first. He had to make his sanctioning decision without knowing what anyone else would do. But, he knew that other people would be making their sanctioning decisions after him. And if others reacted negatively to him, he would have fewer points to take home at the end of the experiment.

The second experiment was the same as the first except that subjects were given information that made the X–W choice more socially meaningful. The experiment instructions said that research has revealed a surprising, yet consistent, finding—preferences for particular patterns of lines are associated with the number of friends people have. Those who prefer one category of line tend to have more friends; those who prefer the other tend to have fewer friends. The line patterns were the letters X and W. In other words, subjects had exactly the same choice to make as those participating in the first experiment. But this time they had information that their X–W choice might tell people whether they had lots of friends or only a few.

In the first experiment, the results show that behavior patterns had no effect on sanctioning. But, in the second experiment in which subjects had reason to think that their choices were socially meaningful, the typicality of behavior in the group did affect subjects' sanctioning decisions. Participants gave fewer points to those who made an atypical choice (Horne, 2009b).

I conducted a third study to examine the effects of setting on metanorm expectations and sanctioning (Horne, 2010). In this study, I showed college student participants pictures of a “good” neighborhood and a “bad” neighborhood. I told them to imagine that they were in the neighborhood and saw a crime being committed. I asked them how likely they would be to do something to try to stop the crime and, if they did so, how they thought others in the neighborhood would react.

The results showed that experimental subjects shown pictures of “bad” neighborhoods said that they were less likely to intervene to stop a crime than those shown pictures of “good” neighborhoods. These responses were completely explained by subjects’ expectations regarding how positively or negatively other residents were likely to react to their intervention efforts (Horne, 2010). That is, characteristics of the setting (the neighborhood) affected metanorm expectations. Those expectations explained subjects’ stated willingness to sanction. These results are consistent with the patterns of informal sanctioning across urban neighborhoods. Criminological research has long shown that people in “bad” neighborhoods exercise less informal control than those in “good” neighborhoods (see, for example, Bursik & Grasmick, 1993; Sampson & Groves, 1989; Sampson, Raudenbush, & Earls, 1997). The experimental results described here show that metanorms explain the link between neighborhood characteristics and sanctioning in the lab, suggesting that metanorm expectations may help to account for this link in neighborhoods.

Thus there is evidence that existing patterns of behavior and the characteristics of a setting may affect metanorm expectations and, in turn, sanctioning. There are other possible sources of clues as to the sanctioning actions that others will support. Individuals may also use others’ sanctioning behaviors as sources of information. Consider, for example, nations’ commitment to the International Criminal Court (ICC). The ICC enforces international human rights norms. When nations committed to the ICC, they made a commitment to the enforcement of those norms. Why did nations commit? One reason is that they considered the likely reactions of other countries on which they were dependent (Goodliffe, Hawkins, Horne, & Nielsen, 2012). The rhetoric of those other countries provided some information about likely reactions, but their actual commitment to the ICC provided even more. Thus a nation thinking about whether or not to commit to the ICC would consider whether other countries on which it was dependent had committed. As nations made commitments, the calculations of the uncommitted countries tied to those nations shifted. Very quickly nations’ commitment decisions in conjunction with patterns of interdependence between countries produced increasing numbers of commitments (Goodliffe et al., 2012).

The studies described above provide evidence that metanorm expectations matter. If expectations of reactions affect sanctioning decisions, then expectations may become self-fulfilling prophecies. In particular, if group members’ expectations are wrong—if they misperceive the behaviors others disapprove and would like to see sanctioned (see, for example, Perkins, Haines, & Rice, 2005)—then they will punish the wrong behaviors. But as they do so, they will create the norms they thought existed.

Norm Enforcement and Social Relationships

The series of studies described focuses on the effects of a structural feature (the characteristics of relationships within groups) on the emergence of norms (norm enforcement) within a group. But, norm enforcement also affects social relationships. That is, the causal arrow also goes the other way. As group members sanction, they hope that others will support their efforts. But they are taking a chance. They do not know for sure how others will react. As others actually reward sanctioning efforts, the relationships between the sanctioner and other group members become stronger. People place increasing value on their relationships (Horne, 2000).

Thus over time, as people enforce norms, relationships become stronger. This does not mean that the sanctioner's relationship with the deviant becomes stronger. Researchers often focus on the relationship between the sanctioner and the deviant—identifying features of that relationship that make sanctioning less likely and examining how sanctioning affects the relationship. Such work shows that strong relationships between deviants and potential sanctioners tend to dampen punishment. Here the focus is on the other group members. As group members anticipate support for punishing and as they provide support to others who sanction, they strengthen relationships with each other. Thus enforcing norms can make groups more cohesive, increasing the likelihood of future norm enforcement.

I conducted a study to test this dynamic (Horne, 2000). In particular, I examined how strengthening an alternative enforcement institution (the legal system) affected the informal controls enforced within groups. The legal system was operationalized as an agent that was supported by taxes collected from the group. The higher the taxes, the stronger the legal system. Further, the stronger the legal system, the lower the personal costs to any individual who turned to it to punish deviance. In contrast, enforcing norms personally was always directly costly to the individual. Individuals who personally punished deviance experienced costs. But when social relationships were strong, others helped to offset those costs through enforcing metanorms (rewarding those who punished). In contrast, when social relationships were weak, sanctioners received little support from others.

I found that as people used the legal system more—and bore lower personal costs for sanctioning—they also received less support from other group members (Horne, 2000). Over time, relationships weakened. People placed less value on their social relationships. In contrast, in conditions in which the legal system was weak, people engaged in more personally costly punishment and received more rewards from others. In turn, they placed more value on their relationships.

This finding is consistent with the work that shows that certain types of exchanges between actors can strengthen relationships (Lawler, 2001; Lawler, Thye, & Yoon, 2000). Norm and metanorm enforcement appear to involve interactions that similarly strengthen relationships. When people enforce norms, they hope that others will support them, but do not know for sure if they will. When others reward them for their efforts, they gain greater confidence in their relationships. Groups become more cohesive.

Summary

In sum, the structure (characteristics) of social relationships within a group affects sanctioning (Horne, 2009a). Sanctioning and support given to sanctioners in turn affect the characteristics of those relationships. That is, social structure affects individual enforcement efforts which in turn affect social structure. If this dynamic continues uninterrupted, we would expect to see groups become more and more controlling. Groups would be increasingly likely to enforce norms even if doing so provided few benefits. Norm enforcement would be strong. In contrast, if group members are not dependent on each other, if they do not value their relationships, then they are unlikely to enforce norms. It will be difficult for the group to achieve collective goals.

This dynamic of strong relationships facilitating norm enforcement which in turn strengthens relationships can be interrupted when outside institutions or events weaken people's dependence on each other. As in the study described above, increases in government involvement may weaken group members' dependence on each other for the punishment of deviant behavior. Many social institutions have the potential to weaken interdependence. When the law provides a substitute, cheaper source of control, individuals have less need of each other. The Internet weakens dependence on local social relationships for information. Employment law weakens the dependence of individuals on fellow union members. Employer-provided benefits weaken dependence on mutual benefit associations that in the past provided individuals and their families with security. Such social changes can affect the extent to which people are dependent on one another and, in turn, their sanctioning decisions.

Thus exogenous changes can have unexpected effects. If they provide a substitute for goods that people formerly worked together to provide, then they weaken people's dependence on each other. In turn, when people value their relationships less, they give less support to sanctioners. Metanorms are weaker. Norm enforcement declines. Norms lose their power.

Cumulating Theoretical Understanding of Norms

Dynamic approaches to studying norms allow for endogenous change; norms can evolve without external inputs. My work primarily focuses on how characteristics of social relationships (at the macro-level) affect the norm-related behaviors of group members (at the microlevel). I also have evidence regarding the effects of individual behaviors on characteristics of social relationships. Thus individuals are both affected by and affect the larger environment (see Andrighetto et al. and Burgemeestre et al., this volume, for alternative approaches to thinking about macro–micro-links). Even if normative rules remain constant, enforcement of those rules can change. Patterns of enforcement shift. Norms grow and fade with enforcement of the rule. Some norm change occurs endogenously; some change can be triggered by exogenous factors.

Other dynamic approaches similarly focus on the interplay between characteristics of the situation and individuals. Ostrom, Gardner, and Walker (2006), for example, identify structural features common to groups that have successfully solved collective problems (see also Janssen and Ostrom this volume, focusing on information as a key structural feature of groups). Bicchieri (2006) focuses on features of the environment that make norms salient—how the environment affects individuals' expectations regarding what others are likely to do and what those others expect, in turn affecting what the individual does. Rather than see individual internal states as immutable, Bicchieri sees them as shifting in response to the social environment. Andrighetto et al. (this volume) similarly discuss the interplay of social factors and individual internal states.

The “social norms approach” also emphasizes the effects of the larger environment on individuals. For example, research shows that patterns of behavior (such as drinking in college campuses) can affect students' perceptions of how much their fellow students drink and how favorably those students see drinking. These perceptions affect the individual student's own drinking behavior, which in turn contributes to perceptions. This pattern can be interrupted by providing students with accurate information about what their fellow students actually think about drinking.

Other dynamic approaches to norms focus primarily on the intersections of individual behaviors. These approaches embed behavioral assumptions in agents who then act. Individual behaviors affect others' decisions. Thus individuals are affected by the behavioral rules they are programmed to follow and the behaviors of those around them. Individual actions intersect to produce macro-level patterns of behavior (for an example of this approach see Elsenbroich, this volume). Interactions may lead to a variety of macro-level outcomes—equilibrium, continual change, cycling between different macro-level patterns, and so forth.

Thus some dynamic approaches focus on macro-level features of the environment, their effect on individuals, and the effect of individual behaviors on those macro-level features. Others focus on how the interplay of behaviors by actors following simple behavior rules produces macro-level patterns. The difference between these two approaches is that the first embeds influences on individuals in the social environment, while the second produces individual behavior by embedding assumptions in actors. This distinction is not as great as it may appear. At a conceptual level, structural constraints can be converted into internal states of agents or vice versa. For example, an assumption that individuals have a taste for conformity might produce the same kinds of behaviors as a social environment that restricts choices. An assumption that individuals have a taste for uniqueness might produce behaviors similar to those we would see in a social environment that encourages innovation. Though subtle, this distinction may nonetheless affect how researchers think about norms. At a practical level, because agent-based models highlight characteristics of actors and the distribution of actors of different types, the most obvious type of intervention is to change the characteristics of actors or their distribution. But because outcomes are the consequence of many interactions, it is difficult to predict what the outcome of a particular change would be. Further, it is easier to

change our assumptions about actors than it is to change actors themselves. In contrast, a structural approach highlights structural interventions as a way to change norms. Whatever the characteristics of individual actors, certain kinds of structures foster norm enforcement more than others.

While we have learned much about norms, there are still many unanswered questions. What might help us as we continue to study the emergence, change, and decline of norms?

Values and Expectations

Norms may be effective because they are internalized into the individual's value system. They may also be effective because they shape individuals' perceptions and expectations that in turn affect their behavior. Researchers differ in the extent to which they emphasize these two mechanisms (for a related discussion see Xenitidou & Elsenbroich, 2010). To some extent, agent-based approaches locate norms in the individual, while structural/situational approaches locate them in the larger environment.

But research is often not clear about the extent to which norms produce individual values or strategies for action (Yamagishi, Hashimoto, & Schug, 2008). For example, research suggests that cultural variation is reflected in individual internal states (see, for example, Haidt & Graham, 2007; Haidt, Koller, & Dias, 1993; Markus & Kitayama, 1991). Much work seems to assume that norms have been internalized so that individuals in different cultures adhere to different values and therefore behave differently. It is possible that individuals carry cultural tastes and preferences into the lab with them. But it is also possible that they bring expectations about others into the lab. Thus their behavior may reflect individual values or it may reflect strategies based on understanding of a society's norms (Bicchieri, 2006; Yamagishi et al., 2008).

The fact that research does not always explicitly distinguish between the two possibilities is a problem because internalized values and expectations about the social world may be the result of different causal factors and mechanisms and may have their effects through different mechanisms. Values are thought to be relatively stable and carried in the individual from one context to another. Expectations are more likely to be formed in situations and to change as the individual moves across social contexts. Researchers need to be clear about these two possibilities in order to collect data evaluating their contribution. Research questions suggested by a focus on internal states will likely be different from those suggested by a focus on the external environment. In the context of neighborhood crime, for example, criminologists have tried to explain why crime rates are higher in neighborhoods with some characteristics than others. One explanation is that norms differ. But, when researchers talk to people they do not find evidence that people in poor neighborhoods have different values than those in rich neighborhoods (Kornhauser, 1978: 214–221). As described above, however, it may well be the case that people in “bad”

neighborhoods have different expectations of their neighbors than people in “good” neighborhoods. People in both kinds of neighborhoods may disapprove of crime (have the same values), but they may have different expectations about the extent to which others disapprove of crime and will try to do something about it—and therefore different norm enforcement patterns. By distinguishing between these two possibilities, we might be better equipped to understand the relation between norms and crime across communities.

Behavior is likely due to some combination of individuals’ preferences (what the individual wants, cares about, and so forth) and individuals’ perceptions of others’ preferences (what others want and will approve of). Similarly, norm enforcement is likely affected by individuals’ views of behaviors and their perceptions of what others would like to see punished. Thus both values and expectations are likely to matter. Norms may be internalized so that they become values as well as norms. But blurring the distinction between the two makes cumulating theoretical knowledge about norms difficult.

In addition to identifying individuals’ expectations of others, we also need to do more work to understand how those expectations develop and change. While researchers have begun to look at this issue (see, for example, Bicchieri, 2006), there is still much that we do not understand.

Substantive and Abstract Norms

Researchers study norms at different levels of abstraction. Conceptualizing norms in terms of games such as the ultimatum game or the social dilemma game has produced research that has contributed greatly to our understanding of norm enforcement. There are many benefits to focusing on basic theory and abstract norms, including that doing so contributes to the cumulation of knowledge. But such approaches may limit the range of norms that we consider. It may be useful, therefore, to also study substantive norms.

For example, consider American norms governing race relations. The explicit norm is that we should be color-blind, that race does not matter. Despite this explicit norm, many Americans think that other Americans disapprove of intimate relations across racial lines (Dave Thomas Foundation for Adoption, 2002). These perceptions are inconsistent with both what people say they support and what they think the explicit norm is. So in the United States explicit norms have shifted from making distinctions between racial groups to saying that race should not matter. At the same time, there are people who are aware of the explicit norm of color-blindness, whose own views may or may not be consistent with that norm, and who believe that others do not support the explicit norm. Although the explicit norm is that race should not matter, segregation persists. Such complex norms raise questions that might not be raised if our research is limited to standard games. In this case it suggests that norms may be more or less explicit and that explicit and implicit norms may differ and may vary in how they change over time.

Lab and Field

I tested my theoretical ideas primarily using lab experiments. Experiments are very useful for testing theories because they provide strong evidence of causation. Many norm scholars rely on experiments and in particular on standard games. Others use data from the field. For example, they use computer simulations to see if they can recreate real-world conditions in a simulation outcome. To the extent that simulation outcomes are consistent with real-world patterns, there is support for the theory embedded in the simulation. Similarly, to the extent that data from the lab and the field are consistent, we can have greater confidence in our theories (see, for example, Ostrom et al., 2006).

Lab experiments can be used to explicitly tests ideas suggested by results obtained in the field. For example, researchers have found that people in different societies play standard economic games differently. Further, they have found correlations between characteristics of the society and how people in those societies play standard games. In many ways these findings reflect those of survey researchers showing that values/attitudes vary across cultures. They show that market integration increases individuals' cooperative behavior in the ultimatum game (Henrich et al., 2001). They also show that variation in antisocial punishment across culture is correlated with the rule of law and religion (Herrmann et al., 2008). But the evidence is largely correlational. Further, the reasons for these correlations are not well understood. Standardized games are useful because they provide behavioral measures and because they facilitate making comparisons across studies and cultures. They suggest important insights into cultural and structural factors that may affect norms. Researchers could build on these findings to develop theory about just how structure/culture affect norms. To test such causal theories, one could manipulate cultures and structures in the lab and observe the norms that emerge. That is, it may be useful to depart from standard games and design experiments that will test the effects of societal level factors on norms.

Lab experiments also have their limits. For example, incentive structures in the lab are usually clear. In the field they are not. Yet norms emerge amidst uncertainty, ambiguity, and conflict. To fully take advantage of field settings, however, we need to develop ways to measure norms. While measuring norms in the lab is relatively straightforward, in the field it is more challenging. It is difficult to get accurate indicators of norms that people do not want to talk about. Researchers need to be able to measure explicit norms as well as the norms that people think exist but will not admit to believing themselves. We must also be careful to develop measures that distinguish between values and norms.

Conclusion

Researchers have learned much about the emergence of norms over the last 20 years. Yet, unanswered questions remain about norm enforcement and even more about norm content and norm change. To move our understanding forward we should take

advantage of the strengths of different approaches. In order to do so, researchers need to be clear about their assumptions so that those different approaches can build on each other's knowledge. In particular, we need to be clear about the distinction between evaluations that are social and those that are internal to individuals. We need good measures of theoretical concepts that can be used across settings. Finally, it will be useful to bring multiple methods to bear on the same theoretical problems. Computer simulations will allow us to examine complex interaction processes and their outcomes. Lab experiments will allow us to empirically test causal relations and mechanisms. And applications in the field will allow us to explore the applicability of our theories across settings. We will learn more through using multiple methods than any single approach alone. Taking advantage of the strengths of multiple approaches will contribute to the development of cumulative theoretical knowledge.

References

- Axelrod, R. (1985). An evolutionary approach to norms. *American Political Science Review*, 80(4), 1095–1111
- Bicchieri, C. (2006). *The grammar of society: The nature and dynamics of social norms*. Cambridge: Cambridge University Press.
- Blau, P. M. (1964). *Exchange and power in social life*. New York: Wiley.
- Bursik, R. J., Jr., & Grasmick, H. G. (1993). *Neighborhoods and crime*. New York: Lexington Books.
- Centola, D., Willer, R., & Macy, M. (2005). The emperor's dilemma: A computational model of self-enforcing norms. *American Journal of Sociology*, 110, 1009–1040.
- Coleman, J. (1990). *Foundations of social theory*. Cambridge, MA: Harvard University Press.
- Dave Thomas Foundation for Adoption. (2002). National adoption attitudes survey. http://www.adoptioninstitute.org/survey/survey_intro.html. Retrieved on November 16 2010.
- Emerson, R. M. (1962). Power-dependence relations. *American Sociological Review*, 27(1), 31–41.
- Emerson, R. M. (1972). Exchange theory, part II: Exchange relations and networks. In J. Berger, M. Zelditch Jr., & B. Anderson (Eds.), *Sociological theories in progress* (Vol. 2, pp. 58–87). Boston, MA: Houghton Mifflin.
- Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415, 137–140.
- Goodliffe, J., Hawkins, D., Horne, C., & Nielsen, D., (2012). Dependence networks and the International Criminal Court. *International Studies Quarterly* (Forthcoming).
- Haidt, J., & Graham, J. (2007). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1), 98–116.
- Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, culture and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, 65(4), 613–628.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., et al. (2001). Cooperation, reciprocity, and punishment in fifteen small-scale societies. *AEA Papers and Proceedings*, 91(2), 73–78.
- Herrmann, B., Thöni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science*, 319, 1362–1367.
- Homans, G. C. (1974). *Social behavior: Its elementary forms*. New York, NY: Harcourt Brace and World.
- Horne, C. (2000). Community and the state: The relationship between normative and legal controls. *European Sociological Review*, 16(3), 225–243.

- Horne, C. (2001). The enforcement of norms: Group cohesion and meta-norms. *Social Psychology Quarterly*, 64(3), 253–266.
- Horne, C. (2004). Collective benefits, exchange interests, and norm enforcement. *Social Forces*, 82(3), 1037–1062.
- Horne, C. (2007). Explaining norm enforcement. *Rationality and Society*, 19(2), 139–170.
- Horne, C. (2008). Norm enforcement in heterogeneous groups: Sanctioning by majorities and isolated minorities. *Rationality and Society*, 29(2), 147–172.
- Horne, C. (2009a). *The rewards of punishment: A relational theory of norm enforcement*. Stanford, CA: Stanford University Press.
- Horne, C. (2009b). Metanorm expectations: Determining what to sanction. *Advances in Group Processes*, 26, 199–223.
- Horne, C. (2010). Unpublished data.
- Horne, C., & Cutlip, A. (2002). Sanctioning costs and norm enforcement. *Rationality and Society*, 14(3), 285–307.
- Kelley, H. H., & Thibaut, J. W. (1978). *Interpersonal relations: A theory of interdependence*. New York: Wiley.
- Kornhauser, R. R. (1978). *Social sources of delinquency*. Chicago: University of Chicago Press.
- Lawler, E. J. (2001). An affect theory of social exchange. *American Journal of Sociology*, 107, 321–352.
- Lawler, E. J., Thye, S., & Yoon, J. (2000). Emotion and group cohesion in productive exchange. *American Journal of Sociology*, 106, 616–657.
- Markus, H. R., & Kitayama, S. (1991). Culture and the self: Implications for cognition, emotion, and motivation. *Psychological Review*, 98, 224–253.
- Molm, L. D. (1997). *Coercive power in social exchange*. Cambridge: Cambridge University Press.
- Molm, L. D., & Cook, K. S. (1995). Social exchange and exchange networks. In K. S. Cook, G. A. Fine, & J. S. House (Eds.), *Sociological perspectives on social psychology* (pp. 209–235). Boston, MA: Allyn and Bacon.
- Ostrom, E., Gardner, R., & Walker, J. (2006). *Rules, games, and common-pool resources*. Ann Arbor, MI: University of Michigan Press.
- Perkins, H. W., Haines, M. P., & Rice, R. (2005). Misperceiving the college drinking norm and related problems: A nationwide study of exposure to prevention information, perceived norms, and student alcohol misuse. *Journal of Studies on Alcohol*, 66(4), 470–478.
- Posner, E. (2000). *Law and social norms*. Cambridge, MA: Harvard University Press.
- Sampson, R. J., & Groves, W. B. (1989). Community structure and crime: Testing social disorganization theory. *American Journal of Sociology*, 94(4), 774–802.
- Sampson, R. J., Raudenbush, S. W., & Earls, F. (1997). Neighborhoods and violent crime: A multilevel study of collective efficacy. *Science*, 277(5328), 918–924.
- Willer, R., Kuwabara, K., & Macy, M. (2009). The false enforcement of unpopular norms. *American Journal of Sociology*, 115, 451–490.
- Xenitidou, M., & Elsenbroich, C. (2010). Construct validity and theoretical embeddedness of agent-based models of normative behavior. *International Journal of Interdisciplinary Social Sciences*, 5(4), 67–80.
- Yamagishi, T., Hashimoto, H., & Schug, J. (2008). Preferences versus strategies as explanations for culture-specific behavior. *Psychological Science*, 19(6), 579–584.