NERD and Norms: Framework and Experiments

Peter Danielson, Alex Mesoudi, and Roger Stanev†‡

We advocate and share the same theoretical framework for empirical research in ethics as exemplified in Christina Bicchieri's *The Grammar of Society*. Our research differs from Bicchieri's in our approach to experimentation: where she relies on lab experiments, we have constructed an experimental platform based on an internet survey instrument; where she relies on rational reconstructions, we do not. In this paper we focus on four contrasts in our methods: (1) we provide a space to explore ethical influence and norm transmission between participants, belief and choice revision, and reputation over time; (2) we provide ways for participants to expand the context of their and others' decisions; (3) we focus on more realistic ethical decisions than is allowed by games; and (4) we explain why Bicchieri's method of rational reconstructions presents challenges to her theory of social norms. Our methods are complementary to Bicchieri's, and together we can work toward developing more comprehensive empirically informed ethics.

1. Introduction. This paper introduces a new approach to experimental research on ethical norms, NERD (Norms Evolving in Response to Dilemmas), by contrasting it with the approach Bicchieri takes in *The Grammar of Society* (Bicchieri 2006; hereafter GS). NERD is a series of web surveys, open to the public, on the ethical issues raised by new technology, especially genomics. This might seem somewhat removed from the subject matter of GS, so we begin by showing that we share a common theoretical framework with GS, and then that our instrument allows us to run controlled experiments on social norms that can usefully inform both that general framework and the specific issues addressed by GS. GS provides an excellent contrast because (i) it is exemplary of a particular set of methods and theoretical assumptions, (ii) we share many of those theo-

†To contact the authors, please write to Peter Danielson, W. Maurice Young Centre for Applied Ethics, University of British Columbia, 227-6356 Agricultural Road, Vancouver, V6T 1Z2, Canada; e-mail: pad@ethics.ubc.ca.

‡Thanks to the NERD team for the tool and teamwork that provided us with the material for this paper, as well as for helpful comments on a draft.

Philosophy of Science, 75 (December 2008) pp. 830–842. 0031-8248/2008/7505-0028$10.00
Copyright 2008 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/594527 Published online by Cambridge University Press

retical assumptions, and (iii) we differ productively in our methodological approaches to experimentation.

2. Common Framework. We begin by emphasizing how much we share with—and commend in—Bicchieri's approach. GS exemplifies an emerging breed of behavioral science which, as Gintis (2007) argues, is unified by a beliefs, preferences, and constraints (BPC) agent model coupled with an account of human motivation informed by evolutionary theory. This empirical approach is relevant to an ethics based on social norms (Danielson 2007). Yet, as Mesoudi and Laland (2007) argue, it does not yet fully incorporate the role of social transmission, a key issue relating to norms and one that both NERD and GS address.

This common framework is a big tent. At one extreme, Binmore (1994, 1998, 2004) insists on rational orthodoxy, while at the other, Skyrms (1996, 2003) works mostly with theory and simulation. In contrast to these two extremes, we are attracted to two aspects of GS: (i) motivational pluralism: "My definition of . . . moral norms . . . takes into account the fact that there are different types of people. . . . What makes people different is the nature of their normative expectations" (GS, xi; see Danielson 2002; Kurzban and Houser 2005); and (ii) the balance of theory and experiment. Not only does Bicchieri refine her account of norms using the experimental literature, her criterion for a good theory is oriented toward devising new experimental tests: "I am not claiming here that mine is a realistic model of how we reason, but . . . it is a fairly good explanatory and predictive model, because my definitions are operational and their consequences are testable" (GS, 49).

NERD and GS share a focus on the importance of experiments for providing reliable tests of meaningful predictions. We agree that "the greatest help in understanding the effects of social norms on behavior will come from a combination of field and lab experiments that is yet to come" (GS, 63). According to GS, making meaningful predictions and testing them is essential for understanding what norms are being adopted. Bicchieri argues that we can only predict whether individual X will adopt norm R in situation S if we know what cues are present in S, and how those cues influence decisions. In order to identify cues and their effects, GS employs game theoretic models informed by the results of economic experiments. In the following section we argue that internet-based experiments such as NERD can complement GS.

3. NERD: From Surveys to Experimental Platform. NERD is an experimental platform built on a public opinion survey instrument. These Web-based surveys pose realistic ethical dilemmas about genomic technology, aiming at stress testing participants' moral norms. Figure 1 shows a sam-

https://doi.org/10.1086/594527 Published online by Cambridge University Press

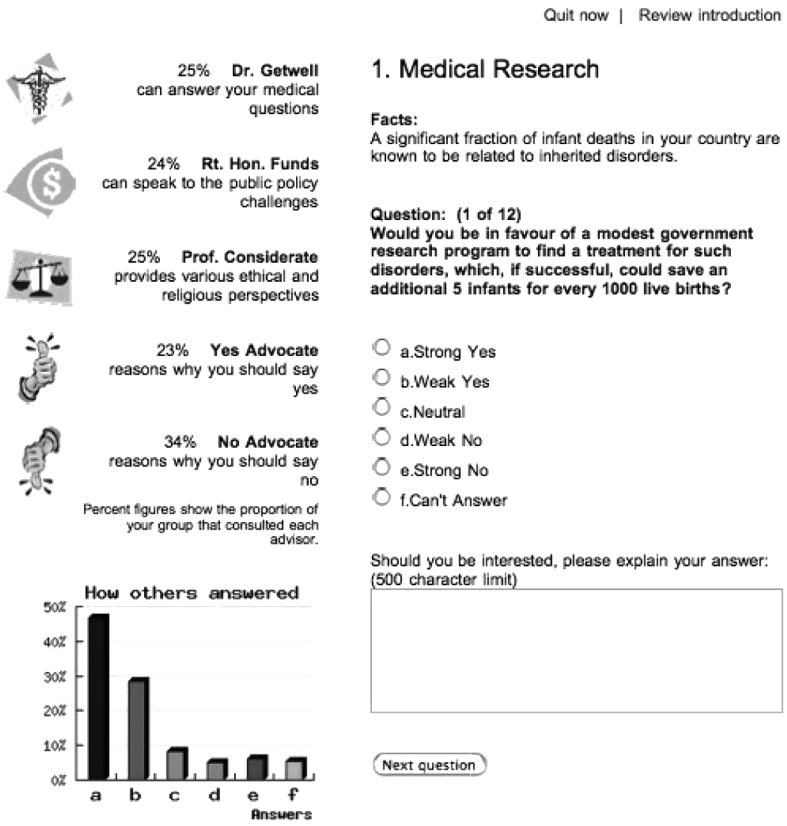


Figure 1. A survey question with feedback.

ple question from a survey concerning the application of genomics to human health, based on a real case observed in Cyprus—the introduction of genetic screening for an inherited disorder, beta-thalassemia (Bornik and Dowlatabadi 2004; Ahmad et al. 2006). Subsequent questions elicited responses to increasingly difficult ethical decisions, including genetic screening of adults and fetuses for the disorder, restrictions on marriage based on genetic profile, and selective abortion. "Advisors" provide information regarding medical, policy, and ethical aspects of the scenarios, while aggregate responses of other participants can be shown.

This survey instrument meets the needs of research clients for first-level answers about public evaluations of technologies. In contrast to many public participation instruments, it allows well-informed (unlike most surveys) and anonymous (unlike focus groups) participation. For example,

https://doi.org/10.1086/594527 Published online by Cambridge University Press

we found strong public acceptance of genomics for human health yet strong rejection of its use for food production (in another survey concerning genomics in salmon aquaculture). More relevant, however, is NERD's use as an experimental platform for exploring the issues raised by GS, particularly, the cues potentially eliciting moral norms (e.g., experts' opinions, past participants' responses, renegotiating ambiguous meaning with others via social feedback, group identity as triggering norm-abiding behavior).

Participants can be exposed to subtly different cues embedded in scenarios. We can look for different reactions to animal welfare scenarios when the species variable is manipulated (e.g., pig vs. rat). We have also begun to explore the effects of social influence. In one study, half of the participants were randomly assigned to a feedback group and were shown a graph of the real-time distribution of group responses and the percentage of the group who consulted each "advisor" (Figure 1). The other half received no group feedback. Figure 2 shows the feedback and nonfeedback medians for six questions. There was no significant difference for the first two questions, but, for the last four, feedback produced significantly stronger agreement than the nonfeedback controls, possibly indicating an "informational cascade" (GS, 196), one that strengthens an already existing effect.

These first NERD surveys isolated participants socially and temporally. Anonymous participants took one or perhaps both surveys, but we had no way to link their responses to the two surveys. NERD 2.0 will allow participants to enroll pseudonymously, enabling them to build reputations, post comments for others, return to answer other surveys, and/or add replies to earlier answers or comments. This change in survey design opens space for new experiments, some of which are described below.[1]

4. Advantages of NERD over GS's Economic Games. Experiments such as the aforementioned on feedback could plausibly be performed using standard economic games, such as those discussed in GS. However, NERD offers several advantages over lab experiments that make it a useful complement to Bicchieri's work.

4.1. Broader Participant Motivation. In the lab games of experimental economics, which GS draws heavily upon, participant motivation can be specified precisely: subjects are pretested for competence, strongly assured of anonymity, and paid well based on game performance. NERD has the advantage of tapping broader motivation more relevant to ethics. The

1. For the latest NERD 2.0 surveys, visit http://www.yourviews.ubc.ca.

https://doi.org/10.1086/594527 Published online by Cambridge University Press

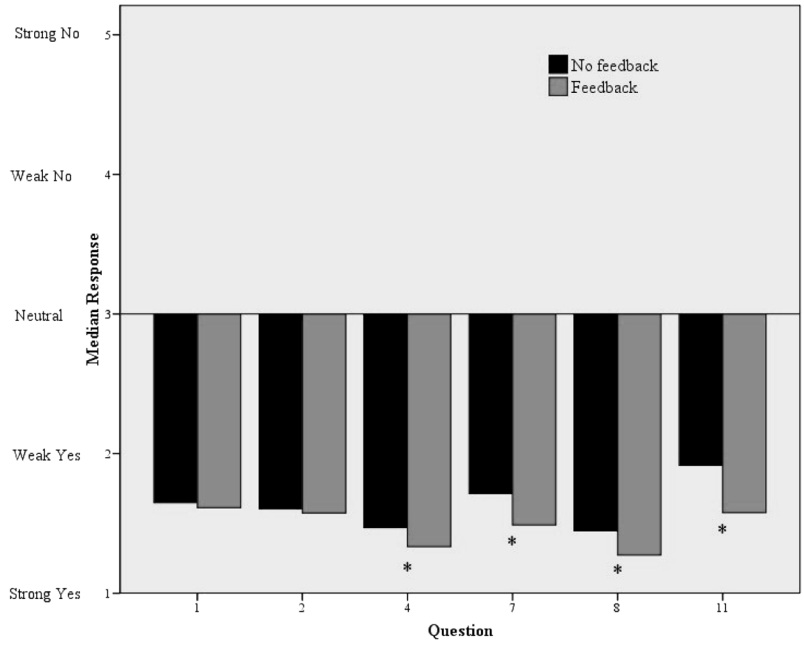


Figure 2. Feedback effect in Human Health Survey. Bars show median Likert responses, with 1 = strongly agree and 5 = strongly disagree. Mann-Whitney tests showed for questions 1 and 2 no significant difference between feedback and no-feedback conditions, and significant differences (marked with an asterisk) for question 4 ($z = -2.28$, $p < .05$), question 7 ($z = -2.78$, $p < .005$), question 8 ($z = -2.90$, $p < .005$), and question 11 ($z = -3.34$, $p < .001$).

intuition is that asking people to make realistic ethical decisions taps motivations closer to real life than does asking them to play games. Indeed, NERD surveys have so far attracted thousands of participants without any payment. Evidently, the experience of working through hard but well-structured ethical problems, and, for some, comparing their decisions with those of advisors and other participants, is itself rewarding. But we cannot know what motivates participants. A potential problem is the confounding role of fictional scenarios. For example, in one scenario we ask participants to imagine "living 30 years ago in a rapidly developing country in which health care and research are funded mainly by the government." We know very little about the extent to which participants answer for themselves rather than in the historical role set by a scenario. In contrast, for competent players, games aim to have one-dimensional payoffs, which may not apply to ethical dilemmas.

NERD provides access to a broader sample of participants than the

https://doi.org/10.1086/594527 Published online by Cambridge University Press

economic games in GS, which are typically restricted to undergraduates. To some extent, NERD is also biased toward a "convenience" sample. We can report that more people prefer health to agricultural genomics, but these were people attracted to Web surveys on these controversies. Our demographics show that our samples are not representative of, say, the Canadian public (e.g., our sample is better educated). Nevertheless, experimental manipulations such as random partitioning into control and experimental groups somewhat insulates us from these sampling problems.

4.2. Open Category Scheme. NERD has an open category scheme, allowing participants to comment on questions, the survey, and their own and others' responses. By encouraging comments, we can explore framing effects in our surveys. While lab experiments can also solicit participant comments, NERD provides more time and privacy in an environment of one's own choice. In previous surveys, textual comments were left on approximately 15% of questions; future surveys should elicit more, since future comments—under certain conditions, controlled by us—will be posted publicly.

Since NERD's open category is not schema-fixed, it allows for flexibility, permitting the identification of salient data relations not seen or expected during the planning stage of the experiment. Hence NERD can be placed on a continuum between the lab and field experiments discussed in GS: more exploratory than lab experiments, allowing us to generate and test hypotheses, yet retaining the experimental control of lab experiments.

Moreover, our design facilitates rapid experimental prototyping in response to comments. For example, comments on our first two surveys suggested that participants were frustrated by the narrative drive of the survey, where the government always chose the pro-technology option. One participant commented: "I have to point out that based on my previous answers I am strongly opposed to the scenario looking the way it does now. Yet, because things have progressed to this point I would have to say that. . . ." In response, we ran a small experiment by replacing the pro-technology frames in the human health survey with contrary government policies. A subsample took this "flipped" survey, without significantly different answers. This is, of course, good news for our overall survey, as it shows that even large differences in questions need not have a significant framing effect. More importantly, these effects can be tested quickly and rigorously, rather than simply assumed or ignored.

4.3. Social Influence and Reputation. Participants in the economic games discussed in GS are typically anonymous and participate only briefly. NERD 2.0, in contrast, allows participants to enroll pseudony-

https://doi.org/10.1086/594527 Published online by Cambridge University Press

mously, opening space for new experiments. We can experiment with participants' reputations by allowing pseudonymous participants to rate other pseudonymous participants' contributions for quality, usefulness, expertise, etc., with each participant building up a reputation score. We can then measure the effect of that reputation in ongoing social interactions, predicting that participants with high reputations will have a disproportionately large influence on others (Boyd and Richerson 1985; Henrich and Gil-White 2001).

NERD also allows controlled experiments on conformity. Classic social psychology experiments (Sherif 1936; Asch 1951) have demonstrated that participants will readily modify their judgments of ambiguous perceptual stimuli (e.g., matching lines of similar length) to conform to a group majority, even when that majority judgment clearly conflicts with the participant's individual judgment. NERD can test for similar conformity effects using more ethically significant issues than line length. NERD can also explore more precisely the effect of group structure on conformity. Different subgroups within the site can be formed with varying ratios of participants who have previously been assessed as being either "for" or "against" a specific issue, measuring exactly how large a majority must be to cause a minority to conform. Will a majority of nine influence a minority of one? Will eight influence two? Will a group composed of equal numbers of participants "for" and "against" polarize, given the lack of a majority (Sunstein 2000)? Do these ratios have the same effects in groups of 100 as they do in groups of 10? NERD allows us to systematically and severely test the effect of group structure to quantify any conformity effect, for example, by finding a threshold majority ratio that is necessary to sway a minority. Indeed, these are the kind of contextual cues that GS argues are crucial in activating social schemata, which in turn elicit social norms. One assumption in GS is that social norms are elicited when (among other conditions) "[a person] *i* believes that a sufficiently large subset of *P* [the population] conforms to *R* [the social norm] in situations of type *S*" (GS, 11). But how large is a "large subset" for different people? With NERD, we can quantitatively measure how large this subset must be for different participants.

It is also not known how conformity persists in the absence of these cues (i.e., the immediate group setting). The conforming minority might revert back to their prior individual judgment in the absence of the majority, suggesting public compliance rather than a stable change in underlying attitudes. Compared to lab-based economics and psychology experiments, NERD allows us to more easily track participants' responses before, during, and after they are placed into subgroups, in order to measure the long-term persistence of conformity. We expect long-term, stable changes in underlying beliefs to be more likely in the case of ex-

https://doi.org/10.1086/594527 Published online by Cambridge University Press

periments conducted with NERD given greater participant motivation (see Section 4.1).

Extending this experimental design can explore how participants' reputation and expertise interact with conformity. Can an informed or highly reputable minority sway the attitudes of an uninformed or low-status majority? The influence of expertise might be counteracted by a lack of similarity between experts and nonexperts, thus inhibiting imitation (Axelrod 1997). NERD allows us to systematically and severely test the effects of all of these contextual cues—reputation, conformity, expertise—and interactions between them.

5. An Objection to GS's Explanatory Model of Social Norms. According to GS, we must understand the mechanisms by which norms control our actions or there is little hope in predicting and influencing social behavior. To understand these mechanisms, Bicchieri uses rational reconstructions. These rational reconstructions are not descriptions of the real beliefs and preferences people have or how people deliberate. Rather, they provide good explanations of social norms and are reliable inasmuch as they generate meaningful predictions with testable consequences (GS, 48).

At the core of Bicchieri's rational reconstructions lies the belief-desire model of choice. This model is said to enable us to explain, predict, and test subjects' actions based on observations of those subjects' beliefs and desires. Therefore, were the beliefs to be different, we would expect behavior to change in predictable ways.

Such a model carries with it some epistemic virtues, including allowing experimenters to identify and elicit salient cues in social contexts. It is the mapping from context to interpretation—categorizing the situation as being of a certain type—and thus to subjects' beliefs and expectations that elicits a preference for a social norm (GS, 57). In other words, subjects infer from situational cues what the appropriate behavior is, what they expect others to do, and what they believe others expect them to do. Ultimately, it is the situational cue upon which the subject focuses that governs the mapping from context to interpretation and, finally, the activation of social norms.

When it comes to the existence of a social norm, a conditional norm follower is presented as if facing a Bayesian game. When the conditions of normative and empirical expectations are both met, the conditional norm follower is said to assess a higher probability to being matched with another norm follower in that game. Note that Bicchieri thinks that whether or not a Bayesian decision model provides an acceptable explanation for what we observe depends on our willingness to take "as if" models seriously, which in turn relates to the possibility of drawing interesting predictions from them (GS, 28).

https://doi.org/10.1086/594527 Published online by Cambridge University Press

However, if what we expect from experiments is the ability to distinguish artifacts from real effects and grounds for arriving at reliable data, the question should not be limited to whether interesting predictions can be drawn from models, since, most likely, they can. Our goal for an appropriate epistemology of experiments for the social sciences should include whether predictive failures (e.g., significant departures from what was expected) can assist us in drawing better and further informative predictions. That is, when predictions fail—for example, due to model misspecifications or violations of our initial assumptions—does Bayesian rational reconstruction provide us with the means for identifying what went wrong? If it cannot, the situation presents serious methodological challenges to an experimental approach driven by rational reconstructions focused on understanding social norms.

Take the following hypothetical scenario. Suppose we have a survey instrument that has repeatedly proven successful in eliciting social norms involving animal genomics and welfare scenarios, specifically in identifying instances of cooperative and competitive interactions that Canadians often employ when judging genomics expenditure. Furthermore, suppose results have been shown to corroborate Bicchieri's theory of social norms insofar as identifying salient situational cues, that is, how different ways of categorizing situations—where a good is allocated—determine the adoption of different fairness norms.

Imagine that, in one study, our instrument allows us to identify situational cues concerning the use of plant genomics—involving inexpensive and nonintrusive research benefiting human health—with situational cues very similar to those of past genomics studies. Yet, suppose, much to our surprise, that changes in the way a situation is framed—public versus private funding—has induced no change in subjects' responses and adherence to pertinent social norms, such as fairness, and we observe a stable and overwhelming opposition to funding such research.

Here is a possible Bayesian rational reconstruction of the experimental results in terms of GS' conditional norm followers:

Theory T. B's theory of social norm.

Hypothesis H. Conditional norm followers' empirical and normative expectations are met by a majority of participants.

Observation O. norm adoption, the predicted support for funding plant genomics.

Observation not-O. norm violation, an overwhelming opposition to funding plant genomics (an anomalous result).

https://doi.org/10.1086/594527 Published online by Cambridge University Press

Auxiliary hypothesis A. Schema for fairness is appropriately primed in situation *S*.

That is, the conjunction **H** and **A** would entail **O**, but what we observe is **not-O**.

Notice the difficulties in reinterpreting what might have occurred. According to GS, social norms are understood to apply to classes or families of situations, not to every possible situation or context (GS, 12). First challenge: Can we classify our new case of plant genomics as part of a previous set of genomics studies? Is the set of candidate categories underlying the social situation the same as previous genomics scenarios? Are they similar enough to warrant the same inductive potential? It is not clear how the experimenter would answer this without a framework allowing for further probing of participants.

As discussed earlier, NERD's exploratory approach to experimentation mitigates the problem of categorization by allowing participants to form their own response categories. Not only can NERD treat questions with different sensitivities, it also permits open question instruments, giving participants the means to express themselves more freely than if they were prompted with categories for responses.

Second challenge: Since the operational definition of a social norm is defined relative to a particular population—that is, a behavioral rule can be a social norm for population P and not for population P'—is our convenience sample representative of population P or P'? Are we dealing with a representative sample of Canadian voters, animal rights sympathizers, Web enthusiasts, or other?

Although the task of identifying target population(s) is an important and often difficult part of any investigation, NERD has the advantage of being exposed to a much larger and broader participant sample than standard economic experiments, thus enabling a wide range of meaningful data analyses and methods such as random or stratified sampling and cluster analysis. For example, within a framework of probability sampling, we can quantify how likely it is that our sample data are representative of a wider population. Since a single probability sample is never guaranteed to be representative with regard to the characteristics of interest, what NERD can do is quantify how often samples can meet particular criteria of representativeness. Populations can be split and randomized, and/or divided into strata—for example, by gender or education—thus increasing precision.

Third challenge: Here is a plausible reconstruction of what the experimenter observes. Since there is an overwhelming amount of evidence corroborating **T** and consequently presumably **H**—due to previous runs

https://doi.org/10.1086/594527 Published online by Cambridge University Press

on genomics studies—whereas there is little evidence for supporting the truth of **A**, that is, hardly more evidence for than against it, and ceteris paribus, that unless **A** is false, there is no other plausible explanation for anomaly **not-O** occurring; then the situation indicates that **A** is false.

Let's look at a Bayesian numerical reconstruction. Contrary to the claims of Bayesian epistemologists, Bayesian priors are always a point of contention. For Bicchieri, the source of priors is left open, but the important thing is that once a script is activated the agent supposedly plays a Bayesian game in which priors are given. One possibility of taking the "as if" model seriously is to consider the experimenter as using Bayesian reconstruction throughout. That is, suppose $P(H) = 0.9$, relatively high, since we have seen a strong corroboration from past observations; and $P(A) = 0.5$, since we feel fifty-fifty about it, and there is no more evidence for than against the fairness norm in situation S. Bayes's theorem then implies that we assess and compare the posterior probabilities $P(H|\text{not-}O)$ and $P(A|\text{not-}O)$.

The sole application of Bayes's theorem as often employed in Bayesian statistics is not the point of contention here. What is at issue is the pragmatic relevance of using such probability calculus in reconstructing experimental results so as to elucidate compliance with social norms.

Taking the "as if" model seriously: If the assumptions are such that participants, conditional norm followers, have degrees of belief that can be expressed reliably as probabilities, and the experimenter must suggest a way to fit these beliefs together coherently so as to specify in which sense one may say that norms are rational, or compliance with a norm is irrational, then it is not clear to what extent the experimenter has given a plausible explanation of the observed anomalous result.

Further suppose that, at a later time, due to an experimenter's error in misspecifying the target population, the anomalous result was skewed. That is, suppose that following further statistical analysis—including this time nonrespondents—we find that the conditional norm followers' empirical expectations were never really met by a majority of participants in the first place. We notice that the majority of participants believed that a sufficiently large subset of the relevant population did not conform to behavioral rule R in situation of type S (i.e., **H** is false). If by Bayesian rational reconstructions of social norms we mean that once an experimental result is identified then we can show how conditional norm followers' probabilities can justify that result, then it is unclear what has been learned.

The fact that experimental results can be reconstructed so as to accord with rational reconstructions of social norms, by reassigning different prior probabilities, is no reason to think that if an experimenter had started out with the view of conditional norm followers as playing Bayesian games

https://doi.org/10.1086/594527 Published online by Cambridge University Press

with those other priors in the first place, that the experimenter would have reached the right experimental results.[2]

Notice that our objection is not contingent on whether or not people reason like Bayesian epistemologists. The objection is that a Bayesian game—functioning as a vehicle for rational reconstructions of people's compliance to social norms—faces serious obstacles when treated as a good explanatory and predictive model for learning about people's norm adherence under failed predictions. That is, in our example, a Bayesian reconstruction does not tell us which experimental assumption is responsible for the failed prediction.

Since, according to GS, social norms are embedded into scripts, the particular way a situation is framed should have a large effect on participants' expectations about others' behavior. A major difficulty is the fact that the triad "norm-script-priors" is so tightly bound that it is very difficult to distinguish potential sources of deviation. The situation is further aggravated by Bayesian reconstructions, since these are ineffective in telling us what we want to know in experiments that may involve potential sources of deviation.

7. Conclusion. We have argued that, first, we share a common framework with GS. Second, we have sketched a new way to do experiments using an internet survey instrument and outlined the key advantages of our method over the economic experiments discussed in GS. Third, we have argued that GS's method of rational reconstruction presents challenges to GS's own theory of social norms. Our platform allows us to manipulate social interaction in new ways, test hypotheses about expectations, and explore broader motivations. Our methods are complementary to the lab experiments that Bicchieri draws upon in GS and can enrich future findings that stem from the work presented in GS.

REFERENCES

Ahmad, Rana, Jennifer Bailey, Zosia Bornik, Peter Danielson, Hadi Dowlatabadi, Edwin Levy, and Holly Longstaf (2006), "A Web-Based Instrument to Model Social Norms: NERD Design and Results", *Integrated Assessment* 6 (2): 9–36.

Asch, Solomon E. (1951), "Effects of Group Pressure on the Modification and Distortion of Judgments", in H. Guetzkow (ed.), *Groups, Leadership and Men.* Pittsburgh: Carnegie, 177–190.

Axelrod, Robert (1997), "The Dissemination of Culture", *Journal of Conflict Resolution* 41 (2): 203–226.

Bicchieri, Cristina (2006), *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York: Cambridge University Press.

2. For an excellent exposition of the challenges and critique of prior probability distributions as degrees of belief, and objections to Bayesian reconstructions in general, see Chapters 3 and 10 of Mayo 1996.

https://doi.org/10.1086/594527 Published online by Cambridge University Press

Binmore, Ken (1994), *Game Theory and the Social Contract: Playing Fair*. Vol. 1. Cambridge, MA: MIT Press.
——— (1998), *Game Theory and the Social Contract: Just Playing*. Cambridge, MA: MIT Press.
——— (2004), *Natural Justice*. Cambridge, MA: MIT Press.
Bornik, Zosia, and Hadi Dowlatabadi (2004), "The Interplay of Technological Change and Social Norms: The Case of B-Thalassaemia in Cyprus." Electronic Working Papers Series, W. Maurice Young Centre for Applied Ethics, University of British Columbia, http://www.ethics.ubc.ca.
Boyd, Robert, and Peter J. Richerson (1985), *Culture and the Evolutionary Process*. Chicago: University of Chicago Press.
Danielson, Peter (2002), "Competition among Cooperators: Altruism and Reciprocity", *Proceedings of the National Academy of Sciences* 99: 7237–7242.
——— (2007), "The Place of Ethics in a Unified Behavioral Science", *Behavioral and Brain Sciences* 30: 23–24.
Gintis, Herbert (2007), "A Framework for the Unification of the Behavioral Sciences", *Behavioral and Brain Sciences* 30 (1): 1–16.
Henrich, Joseph, and Francisco Gil-White (2001), "The Evolution of Prestige", *Evolution and Human Behavior* 22: 165–196.
Kurzban, Robert, and Daniel Houser (2005), "Experiments Investigating Cooperative Types in Humans", *PNAS* 102 (5): 1803–1807.
Mayo, Deborah G. (1996), *Error and the Growth of Experimental Knowledge*. Chicago: University of Chicago Press.
Mesoudi, Alex, and Kevin N. Laland (2007), "Extending the Behavioral Sciences Framework: Clarification of Methods, Predictions and Concepts", *Behavioral and Brain Sciences* 30 (1): 36–37.
Sherif, Muzafer (1936), *The Psychology of Social Norms*. Oxford: Harper.
Skyrms, Brian (1996), *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
——— (2003), *The Stag Hunt and the Evolution of Social Structure*. Cambridge: Cambridge University Press.
Sunstein, Cass (2000), "Deliberative Trouble? Why Groups Go to Extremes", *Yale Law Journal* 110: 71–119.

https://doi.org/10.1086/594527 Published online by Cambridge University Press