Small-group Conformity and Political Attitudes*

David J. Hendry†

Abstract

Public responses to attitudinal questions tapping sensitive social issues are likely to paint an optimistic picture of the degree to which individuals adhere to desirable social norms. But little is known empirically about how social pressures operate at the level of interpersonal interactions. This study conducts a laboratory experiment to address the question of how even minimal social pressure leads to conformity with respect to attitude expressions about adherence to egalitarian norms. Baseline attitudinal measurements were taken of subjects, and then those measurements were used to exert social pressure in a contrived group setting. An asymmetric effect was found in which subjects who were willing to espouse an inegalitarian attitude in private were more likely to succumb to social pressure to change their expressed attitudes when faced with an opposed group opinion. For subjects who espouse egalitarian attitudes in private, social pressure to provide an inegalitarian response has little impact.

*Special thanks are due to Jim Kuklinski for invaluable advice at all stages of this project that greatly improved the final product, and to Chris DeRose for critical assistance and consulting in the development of the software used to implement the design. These individuals donated an immense amount of time to the success of this project. Various aspects of the study also benefited from conversations with Wendy Cho, Jason Coronel, Alan Gerber, Matthew Hayes, Michael Heaney, Jeff Mondak, Sunhee Park, Milan Svolik, Kaye Usry, Tarah Williams, and Cara Wong. Most importantly, the study would not have been possible without the generous permission granted to use the experimental laboratory facilities of the Department of Psychology at the University of Illinois. In securing this privilege, I am especially grateful to Dave Irwin and Mike Regenwetter. Any and all remaining errors are my own.

†Postdoctoral Research Associate, Center for the Study of American Politics, Institution for Social and Policy Studies, Yale University; david.hendry@yale.edu.

Scholarly thought on the interpersonal underpinnings of social influence processes is buttressed by compelling theoretical accounts and anecdotal evidence, but little in the way of systematic empirical demonstrations. At least one reason for the dearth of empirical research is that tapping the differences between true and expressed attitudes is exceedingly difficult in practice. In the pursuit of measuring internal attitudes, we are faced with external expressions of attitudes. In the pursuit of detecting social pressure effects, we are typically without direct observation of the theorized social influence process.

Take as one example the phenomenon of social desirability bias in the measurement of racial attitudes. If one examines, for instance, the attitudes of whites toward blacks over large swaths of time in the United States, a clear trend will emerge in which expressed negative attitudes exhibit a marked decrease (e.g., Schuman et al. 1998). In any thoughtful account, interpretations of this trend will be accompanied by statements of caution about changing social mores that govern the social acceptability of expressions of negative racial attitudes. The standard explanation is that the shared understanding of the level of social acceptability of negative racial attitudes has declined over time, and that at any given point, individuals whose true attitudes are discordant with that level perceive some degree of social pressure to mask their views.

That interpretation will not be criticized here. It is, in fact, highly plausible and quite compelling. But it is useful to note at least one phenomenon that it implies, and about which we know very little empirically. That is, whether or not the distribution of true racial attitudes changed significantly over time, the mass public nonetheless internalized the change in norms of socially acceptable public expression regarding race. But exactly how individuals came to perceive a changing social climate, as well as their own place in that change, remains empirically elusive.

Social desirability bias in the case of white racial attitudes is just one example of a general phenomenon of norm change and internalization. Similar transformations have been witnessed with respect to gender, various religious and ethnic groups, homosexuals, and so on. In each of these cases, we can reasonably surmise that there is a general understanding among the populations of most institutionalized democracies that expressions of egalitarianism constitute the appropriate social response in most circumstances. That is, norms governing appropriate public discourse are such that in order to avoid appearing deviant from one's peers, expressions of equality between

groups are a safer bet than any justifications or preferences for inequality. This implies that individuals perceive that a belief in egalitarianism is the norm in the broader population. And the development of that perception must have come through a process of learning and internalization over time.

When internally held attitudes and norms of appropriate discourse are in tension, public expressions of attitudes may not be what they seem (e.g., Noelle-Neumann 1984). And therein lies the rub for empirically oriented social scientists. Those interested in obtaining valid measures of attitudes on sensitive social issues through self reports cannot be sure that respondents are telling the truth. And those interested in the real-world dynamics of the social influence processes themselves will have difficulty observing them because their primary manifestation is likely to be in the form of everyday social interactions. In either case, without a careful research design, these phenomena will remain outside of the realm of empirical examination.

This paper examines the dynamics of social pressure effects in a laboratory setting. For now, we will leave aside questions of whether the opinions that individuals express on sensitive social matters represent their true, internally held, attitudes, and rather focus attention on the dynamics of social influence. Specifically, examinations of social pressure effects will be with respect to baseline attitudinal measures on sensitive social matters collected in an isolated setting. Isolation is not expected to eliminate social desirability bias stemming from researcher effects, but it is expected to provide a contrast with which to compare more direct forms of social pressure exerted in peer-group settings. The paper proceeds as follows. In the next section, I present a summary of social scientific research on social influence. In the subsequent section, I use the ideas from previous research to draw a set of empirical hypotheses that can be tested in a laboratory setting. Following this discussion, I describe the design and results of an experimental study that allows for testing of the hypotheses. Finally, I offer concluding remarks.

Social Influence

When it comes to salient social issues, individuals tend to be aware of whether there are norms of appropriate behavior that govern the range of expression deemed socially acceptable (e.g., Crowne and Marlowe 1960; Kuran 1995; Mendelberg 2001; Noelle-Neumann 1984). But how do they

become aware? Beginning with the basic notion that individual behavior is developed through individual-environment interactions (Tingsen 1937), it is widely recognized that important parts of that environment are mass media and the social context. Exposure to mainstream media provides individuals with information about the lines being drawn on salient political topics, the range of acceptable points of elite debate on those topics, as well as measures and interpretations of public opinion. Individual perceptions about the aggregate distribution of preferences have been found to influence both the willingness to express an opinion (Noelle-Neumann 1984) and, when willing, the direction of preferences expressed (Mutz 1998).

The major mechanism by which context imparts social influence is by limiting the number and structuring the types of opportunities for social interaction (Eulau 1986; Huckfeldt and Sprague 1987). Over long spans of individuals' lives, they will find themselves embedded within a set of institutions, networks, social groups, and so on, that for most practical purposes can be considered static (Granovetter 1985). If we can assume that a significant proportion of political talk that the typical person engages in takes place in the context of casual interactions with people who just happen to be proximate (MacKuen 1990), the impact of social context in the development of true attitudes, perceptions of the attitudes of others, and the social consequences of norm deviation would seem to be profound.

Social Influence through Interpersonal Interactions

The media provide information about society at large and the social context structures exposure to others, but the bulk of social influence takes place through direct interpersonal interactions. Through interpersonal interactions, common knowledge of socially acceptable behavior can exert social pressure on individuals to act according to behavioral norms (Festinger 1954; Kelley 1952; Sherif 1936), even when such behavior is incongruent with preferences (Festinger 1957; Kuran 1995; MacKuen 1990). Everyday social interactions, therefore, act as a type of social verification system that drives individuals in the same society toward a shared understanding of the range of behaviors that are socially acceptable (Hardin and Higgins 1996), if not actual perceptions of right and wrong.

In the political and social sciences, studies of interpersonal influence have tended to focus

on the direct effects of information filtering by opinion leaders (Katz 1957; Katz and Lazarsfeld 1955; Lazarsfeld, Berelson, and Gaudet 1948), and information exchanges between peers (e.g., Ahn, Huckfeldt, and Ryan 2010; Barabas 2004; Huckfeldt and Sprague 1987; Mutz 2002) in influencing attitude expression. Scholars working in this realm have paid less attention to normative explanations, however (but see Verhulst and Levitan 2009). Without any exchange of factual information, people still tend to draw on the attitudes and behaviors of others as indicators of the ranges of those attitudes and behaviors that are socially appropriate (Festinger 1954; Kelley 1952). In certain situations, individuals may have a sense that there are truly correct and incorrect attitudes with respect to a given social referent, and that the attitudes of one's peers serve as a benchmark against which to evaluate the appropriateness of certain opinions (Burt 1987; Festinger 1954; Kelley 1952; Levitan and Visser 2008; Visser and Mirabile 2004). In various other situations, proximate others may serve as standards for appropriate social behavior without regard to factual accuracy (Deutsch and Gerard 1955; Kelley 1952).

Social Conformity

The particular type of social influence most relevant for the present study is social conformity, that is, the tendency of individuals to publicly comply with a known or perceived group opinion. Following Sherif (1935) and Asch (1951), a long line of research in social psychology and related fields has demonstrated that social conformity is in fact empirically verifiable in the laboratory. Sherif (1935, 1936) employed an optical illusion known as the autokinetic effect (Adams 1912), in which a stationary pinpoint of light is projected onto the wall of a darkened room, and to almost all human observers, the light appears to move. Employing confederates, Sherif demonstrated that experimental subjects were quite susceptible to the power of suggestion with respect to judgments about how far the light had traveled. That is, when subjects were placed in group settings and asked to publicly express their estimations of how far the light had traveled after confederates had already done so, the judgments of subjects tended to fall in line with those of the confederates, even when the expressed judgments of confederates were wildly different than baselines established in a control condition. The Sherif studies demonstrated that when humans are asked to express judgments about relatively ambiguous phenomena, there is a strong tendency

to look to the majority opinion as a rule of thumb for the appropriate response.

But what about relatively unambiguous phenomena? Working in the tradition of Sherif (1935, 1936), Asch (1951, 1952) examined social conformity in situations in which the judgment task had clear right and wrong answers. Specifically, subjects in the Asch studies were asked to express which among a set of lines drawn on a card was identical in length to a line drawn on another card. Lines were purposely drawn so that subjects would easily be able to differentiate their lengths, and subjects making the judgment in an isolated control condition rarely expressed the wrong answer. Subjects were placed in a group setting and asked to express their judgments out loud following a group of confederates who expressed the wrong answer. Though a majority of subject responses in this group condition were correct, a surprisingly large proportion were the same incorrect response given by the group of confederates, and a majority of subjects gave the incorrect response at least once. Given little evidence that the subjects examined by Asch (1951, 1952) were unsure about the correct answer, the results point to a strong tendency for humans in public settings to conform to expectations about the appropriate social response.

While the results from the research programs spawned by Asch and Sherif are extremely useful as the largest and most prominent collection of empirical evidence on small-group conformity, it is not clear whether the type of conformity demonstrated has a direct connection to norms of public speech with respect to salient social referents (including attitudes about the equality of social groups), which are the focus of the present study. In particular, the typical decision-making task under examination in these studies is highly arbitrary, and cannot reasonably be expected to carry any social significance for the subject outside of the laboratory. In one way, this can be seen as a strength of the experimental designs in that the researcher is able to eliminate subjects' preconceived notions about particular referents. In another way, however, if we consider those preconceptions to be worthy of study in their own right, and if we consider the clear qualitative difference between expressed judgments on arbitrary tasks and expressions of opinion on salient—and sometimes sensitive—attitudinal measures, existing work on social conformity leaves open the question of how small-group pressures might operate on norms of public speech regarding egalitarianism. To individuals within the current social climate of the United States, at least one social rule that should be apparent is the notion of equality between individuals, without regard

to race, ethnicity, religion, gender, and so on. Regardless of one's true preferences, it is clear that publicly espousing an inegalitarian view carries with it greater risk of social costs than espousing an egalitarian view.

Development of Hypotheses

The lines of research spawned by Asch and Sherif offer a jumping-off point from which to develop a set of expectations about the impact of social pressure on individual expressions of political attitudes. As discussed previously, the takeaway point from these studies is that in ambiguous and unambiguous decision-making settings, experimental subjects exhibit a strong tendency to look to the group majority as an indicator of the appropriate social response. Individuals may not internalize the perceived group opinion as their own, but avoidance of appearing deviant provides a strong incentive to publicly comply, nonetheless.

Again, however, though researchers in social psychology and related fields have tended to couch the phenomena that they examine in terms of general descriptions of social norms, decision-making, opinion formation, and so on, it is not completely clear that the types of behavioral tasks that they examine provide a close analogue to expressed attitudes on sensitive social issues. On the one hand, attitudes, unlike lines drawn on index cards, are by their nature matters of opinion without a clear right and wrong answer. If individuals, no matter how strongly they hold their own opinions, are able to see that certain political questions are coupled with a range of potential responses across which reasonable people might disagree, then the impact of social pressure to conform to some perceived group opinion may be less than in unambiguous decision-making settings. On the other hand, the lack of a clearly correct answer, as with the autokinetic effect, may lead individuals to be even more susceptible to the power of suggestion by immediate peers. Because previous research does not offer clear guidance in this respect, I will begin the process of investigation by situating the expression of political attitudes on sensitive social issues within the general framework of decision-making tasks examined by previous research on social conformity. This decision leads to the following:

> *The General Social Conformity Hypothesis: When asked to first express an attitude privately, and then publicly in the face of a unanimously opposed majority, individuals*

will tend to alter their privately expressed attitudes in order to publicly comply with the unanimous group attitude.

In other words, we can imagine a situation in which an individual is compelled to express her attitude with respect to a social issue, and where she has some sense that this expression of opinion will be kept private. Then later, she is again compelled to express her opinion on that same issue, but this time in a group situation in which (1) there is no sense of privacy, and (2) her fellow group members have unanimously given an opinion that is opposed to what she stated in private. In this situation, the *General Social Conformity Hypothesis* predicts that the individual will change the attitude that she previously expressed in order to conform to the group opinion.

Also, note what the *General Social Conformity Hypothesis* is not saying. Specifically, there is no inherent assumption that individuals who express attitudes in the absence of specific knowledge about the opinions of their peers are free from all social pressure. It is in fact expected that individuals carry with them into the laboratory some notion of social expectations, and that being compelled to express an attitude to a researcher is often enough to invoke conformity pressures. Knowledge of the distribution of expressed attitudes among a proximate group of peers who will witness the individuals' attitude expressions, however, is expected to be stronger, and to manifest itself through dissipation or reversal of any biases that might be exhibited in isolated response settings.

Additionally, as discussed previously, the inherently subjective nature of attitudes means that most people will understand that there is no clear right or wrong answer. On most sensitive social issues, however, there is a sense of what is and what is not appropriate to say publicly. So while individuals could genuinely feel that reasonable people might disagree, they also carry the looming feeling that publicly appearing deviant from the dominant opinion in the population may carry with it social costs. Further, it has been argued convincingly elsewhere that when people hold conflicting cognitions, they experience cognitive dissonance, that is, a sense of psychological discomfort that leads them to try to reduce the conflict (Festinger 1957). Analogous to the idea of conflicting cognitions, individuals may also experience psychological discomfort from holding internal attitudes that are in conflict with the known appropriate social response. If so, we would expect individuals to experience more social pressure to change their privately expressed attitudes when, after having given a private response that is inconsistent with the appropriate social

response, they are faced with a unanimous group opinion that is consistent with the appropriate social response. Borrowing the terminology of cognitive dissonance theory, this leads to the following:

> *The Social Dissonance Hypothesis: When asked to first express an attitude privately, and then publicly in the face of a unanimously opposed majority, individuals will be more likely to alter their privately expressed attitudes in order to publicly comply with the unanimous group attitude if the attitude that they expressed in private runs counter to norms of social acceptability.*

The *Social Dissonance Hypothesis* simply predicts a more nuanced set of circumstances leading to conformity than the *General Social Conformity Hypothesis*. Revisiting the hypothetical scenario put forth above in which an individual is compelled to express an attitude privately and then publicly in the face of unanimous opposition, imagine further that the particular attitude being expressed is with respect to a sensitive social issue, and that most people have an accurate sense of the range of attitudes deemed socially appropriate with respect to that issue. The *Social Dissonance Hypothesis* predicts that if the individual expressed an attitude in private that falls outside of the range of social acceptability, she will be under more pressure to conform to the opinion of the unanimously opposed group than if she expressed an attitude that falls in line with what is deemed socially appropriate. For example, assuming that individuals in the United States recognize attitudes expressing inegalitarian viewpoints as generally socially unacceptable, those who express an inegalitarian opinion in private are more likely than those who express an egalitarian opinion in private to submit to interpersonal pressures to express the opposite in public, according to the *Social Dissonance Hypothesis*.

Even if we do find that either under general or limited circumstances, individuals have some systematic tendency to conform to the opposed opinion of a group of peers, public compliance in and of itself does not indicate whether the change in responses is due to internalization of a new attitude (e.g., through learning or updating) or preference falsification to avoid appearing deviant from the group. On the one hand, Sherif (1936) and Verhulst and Levitan (2009) demonstrated that even brief interactions can lead to continued influence weeks, or even months, later. Also, cognitive dissonance theory would predict that the very act of providing a particular response publicly would lead an individual to be more likely to internalize that attitude (Festinger 1957). Based on these considerations, I propose the following:

The Attitude Internalization Hypothesis: When asked to first express an attitude privately, then publicly in the face of a unanimously opposed majority, and then again privately, individuals who altered their responses between the first private setting and the group setting will maintain the attitude expressed in the group setting in the second private setting.

However, and as has been discussed, it is also plausible to think that expressions of attitudes on sensitive social issues do not operate in exactly the same manner. For instance, in a large body of work, Kuran (1995), has used formal logic, anecdotal accounts, and qualitative analyses of political events to argue convincingly that individual public compliance with a belief that is incongruent with internal attitudes is a theoretically plausible phenomenon. At the individual level, one can think of agents earning a type of expressive benefit by stating publicly their true preferences with respect to social referents and incurring a type of expressive cost by stating publicly a view that goes against the majority opinion (Kuran 1995). The primary mechanisms leading to these costs and benefits are negative and postive reinforcement by other group members. Under very general circumstances, it can be shown that norms of public compliance that run counter to some or all of a group's members' internal attitudes can persist if the costs of going against the group are sufficiently high. Based on this discussion, I offer the following:

The Preference Falsification Hypothesis: When asked to first express an attitude privately, then publicly in the face of a unanimously opposed majority, and then again privately, individuals who altered their responses between the first private setting and the group setting will alter them again in the second private setting in the direction of the original response.

Predictions about whether any attitude changes witnessed as a result of social pressure are real attitude changes (the *Attitude Internalization Hypothesis*) or artificial public compliance (the *Preference Falsification Hypothesis*) are expressed as competing hypotheses because the previous literature offers competing expectations.

The Real-Time Social Pressure Experiment

The research programs spawned by Asch and Sherif provide a powerful framework for testing questions of social conformity that has produced decades of research in social psychology and related fields, but has thus far been almost completely ignored by political scientists (for the only known example of a similar design with respect to political attitudes, see Verhulst and

Levitan 2009). The design objective was to situate the study within the tradition of previous research on social conformity, but to push the envelope into territory not yet explored by previous researchers—that is, public expressions of political attitudes as they relate to egalitarian norms.[1] Using a within-subjects design, a pre-treatment measurement of subjects was used to establish a baseline measurement of the types of responses to potentially sensitive social issues that subjects were willing to give in isolation. These were then used in a fictitious group setting to generate social pressure in the opposite direction of the baseline measurement. Following that, subjects were measured once more to obtain evidence on whether the response offered in the public setting was internalized.

Participants

Subjects were recruited from undergraduate political science courses at the University of Illinois at Urbana-Champaign in exchange for extra credit. Subjects were recruited under the auspices of taking part in a study on "group and individual decision making" that was carried out from January 2012 through April 2012.

The Behavioral Task

To relate the experimental design to the preceding discussion, a behavioral task was chosen to tap subjects' levels of adherence to egalitarian norms. In the United States, as in many other multicultural societies, it is without controversy to claim that in most public social settings, norms of appropriate social behavior dictate compliance with a norm of individual and group equality, without respect to race, ethnicity, religion, gender, and so forth. That is, individuals' true attitudes with respect to, say, religious equality, are expected to vary—sometimes greatly—in reality, but the range of appropriate public expression of religious attitudes lies wholly on the equality side of the scale. With respect to specific identifiable groups, however, tolerance of dissent from the norm of equality can vary at any given time, subject to contemporaneous social circumstances. Therefore, in order to examine differences between subjects who are willing to express an inegalitarian attitude and those who are not, a target group had to be chosen

[1] Verhulst and Levitan's (2009) focus is with respect to political attitudes in general.

that would elicit a relatively high level of variation in subject responses. Based on a variety of polling data, as well as a current social climate in which the threat of violence by Islamist groups looms large in the Western public conscience, expressed attitudes about Muslims and Islam by non-Muslims was chosen as the behavioral task of interest for examining social pressure.[2]

Table 1 presents the questionnaire items that were used as behavioral measures to tap individuals' willingness to express an inegalitarian attitude. For each question, a determination was made as to what constituted the egalitarian and inegalitarian responses, and these decisions are indicated in the latter columns of the table. The number of response options available to subjects on the questionnaire is also indicated. Further, each question is given an identification number (first column). In much of the discussion that follows, it will be useful and efficient to refer to questionnaire items by their identification number rather than their full question wording.[3]

[Table 1 about here.]

Procedure

Overview At a computer terminal, subjects provided responses to survey questions first in isolation, then in a simulated group composed of other participants simultaneously taking part in the study, then again in isolation. To simulate the experience of participating in a computer-networked group interaction, subject photographs (headshots) were used in conjunction with original software[4] to mimic certain aspects of social networking websites, with which a large proportion of subjects were expected to be familiar. Because the experiment relied on a deception,

[2]The incentive for subjects to participate in the study was extra course credit for one of their political science courses, and therefore, in the name of fairness, all students in a given course were offered the opportunity to participate. That is, Muslim students were not excluded from the recruitment pool. However, a survey question asked students for their religious affiliations. In all analyses that follow, self-identified Muslims are excluded.

[3]For the full question wording, see Section A of the Online Appendix.

[4]Software to implement the design was written by the author in consultation with a professional developer. Given the dynamic requirements of the experimental design, and to ensure cross-platform compatibility for future development, the use of web forms was chosen as the most straightforward implementation method. Page content, layout, actual dynamics, simulated dy-

following their participation subjects were questioned as to whether the deception was convincing. Subjects who indicated that the deception was not convincing, or that they were convinced for only a portion of the time, were excluded from all analyses.

Preliminaries Up to 10 subjects at a time gathered in a waiting room adjoined to the set of individual, private, enclosed computer terminals. After filling out preliminary paperwork, subjects were escorted one at a time into a terminal in which the computer was equipped with a camera. At this terminal, subjects used a computer interface to enter identifying information and upload a headshot of themselves to a server. Depending on the number of subjects attending a given session, the preliminary picture upload portion of the study took approximately 5 to 15 minutes to complete, during which time the subjects not uploading a photo at any given moment were allowed to converse with one another. After the final photo was taken, all subjects were simultaneously read a set of instructions outlining the structure of the experiment, specifically noting that they would first answer a set of survey questions as individuals, then be assigned to a group of other participants and perform some tasks in groups, and then perform more tasks as individuals.[5]

Individual Pre Period Following the instructions, subjects took seats at enclosed computer terminals where they encountered a set of instructions on how to begin the individual questionnaire. After first selecting and confirming the identifying information provided during the picture upload process, subjects answered a series of survey questions in typical computer-interface fashnamics, and database queries were handled using HTML, CSS, PHP, JavaScript, and MySQL, respectively. More detailed information about software development is available from the author upon request.

[5]A slight variation in instructions was employed to generate different levels of social pressure. Specifically, a subset of subjects was told that following the experimental protocol on the computer, they would perform some more tasks with their fellow group members in a face-to-face setting. Ultimately, this variation was found to have no explanatory power for the outcomes examined. All analyses that follow pool across these experimental conditions. Specific instructions are provided in Section B of the Online Appendix.

ion. The primary items of interest for the purposes of hypothesis testing are questions that tap non-Muslim subjects' attitudes with respect to Muslims and Islam. To assist in masking the intended purpose of the study, in between questions of primary interest subjects were asked questionnaire items that tapped other political attitudes, as well as demographic and personality characteristics. All questionnaire items measure concepts worthy of study in their own right in a typical mass survey, but the primary purpose of the non-Muslim, non-Islam questions in this study was to serve as distraction between repeated tapping of the primary concept of interest.

Group Period Following the initial individual-level survey items, subjects reached a waiting screen where they were told that their group was being formed for the group period. After a brief delay, each subject was taken to a screen where their own picture was shown alongside two other participants taking part in the study at the same time.[6] Subjects were told that the two participants appearing alongside them would comprise their groups and that the group period of the study would consist of answering questions similar to those asked in the *Individual Pre Period*, the difference being that all group members would be able to witness all other group members' responses. Additionally, subjects were told that the order in which group members would answer each question would be randomly chosen. As subjects navigated through the questions during the group round, the order in which group members would answer each question was made known to the subjects by displaying group members' pictures in order from left to right, with group members' responses displayed below their pictures.

In reality, the groups were fictitious constructions. Pictures of other participants taking part in the experiment and pre-programmed dynamic responses actually played the role of confederates. That is, for any given subject, group members' supposed responses were either predetermined or a function of the subject's previous responses in the *Individual Pre Period*. Specifically, intermixed within a set of distraction items that subjects answered in the first or second position in the group, and for which they faced a mixture of agreement and disagreement from their fellow group members, subjects encountered a series of questions about Muslims and Islam that they had previously been asked in the initial round of questioning. For these items (the questions of

[6]Section C of the Online Appendix presents screenshots of the experimental protocol.

primary interest), subjects were asked to answer last (or, in one case, in the middle position) in the face of a group whose responses were unanimously in diametric opposition to the response given by the subject in the initial round of questioning. For example, suppose that in the *Individual Pre Period* a subject responded to a question by choosing the egalitarian (inegalitarian) response. Then, in the *Group Period*, they would be asked to respond in the face of a unanimous majority choosing the inegalitarian (egalitarian) response. In this way, the responses that subjects supply in isolation provide a baseline measure of attitudes from which to gauge their propensities to deviate from that baseline when confronted with an opposed group. This feature of the design parallels the basic approach employed by Asch (1951, 1952, 1956) and Sherif (1935, 1936) with respect to arbitrary decision-making tasks, but applies it to the measure of political attitudes that should be expected to carry some social significance for the subjects outside of the laboratory.

[Table 2 about here.]

Following a series of questions about Islam and Muslims in which subjects were compelled to answer after other group members while facing a unanimously opposed majority, subjects then encountered a set of similar questions in which they were compelled to answer in the first group position. Similar to the design executed by Jacobs and Campbell (1961), the repeated instances of answering in the last group position allows the subject the opportunity to learn about the supposed distribution of opinion in the group with respect to the dimension of interest. Then the instances that follow in which the subject answers in the first group position allow for hypothesis testing about the effect of social pressure to conform when that pressure is indirect rather than direct. The questions of interest used to examine social pressure effects, as well as the group dynamics with respect to these questions, are presented in Table 2. Question identification numbers, as well as question wording and response options in Table 2 correspond exactly to those in Table 1.

Individual Post Period Following the *Group Period*, a subset of subjects were again asked to answer a set of questionnaire items in isolation under the condition of anonymity. Other than the specific set of distraction questions asked, this period directly mirrored the *Individual Pre Period*. The purpose of this second isolation period was to allow for hypothesis testing about the nature of the responses given across the study (importantly, whether responses given in the group setting

were carried over into the post-group setting). Survey questions were followed by a debriefing that informed subjects of the deception employed in the study, and then a set of questions that probed for suspicion about the deception.

Results

To get an initial sense of subject responses on the questions of primary interest for the testing of social pressure with respect to egalitarian norms, Figure 1 presents the distribution of expressed attitudes for the questions in Table 1 given in the *Individual Pre Period*. The most striking feature of these data is that subjects overwhelmingly tend to give egalitarian responses to the questions. For all questions, a greater proportion of respondents initially gave the egalitarian response than otherwise. And for most questions, this was overwhelmingly so. The proportion of subjects willing to express an inegalitarian preference ranged from 2% in the case of Q2 to 48% in the case of Q7. Further, the responses to Q6 and Q7 suggest that the types of questions that are most likely to elicit inegalitarian attitude expressions are those that cue subjects to consider hypothetical situations in which their own group is approaching the status of a numerical minority. The overall picture from the *Individual Pre Period* is that the subject population exhibits a relatively low initial willingness to express attitudes of an inegalitarian nature.

[Figure 1 about here.]

Given this baseline distribution of responses, I now turn to to the question of changes in expressed attitudes under group pressure. Figure 2 presents the raw data for all of the questions asked in the *Group Period* for which subjects were compelled to respond after at least one of their fellow group members had expressed an opposed view. Here, the data points are jittered for ease of visualization and are separated by whether subjects initially provided an inegalitarian or egalitarian response. For Q5, Q6, and Q7, a large percentage of subjects who initially provided the inegalitarian response publicly complied with the egalitarian response when compelled to follow their fellow group members. For Q2, only 3 subjects initially provided an inegalitarian response, but 2 out of these three altered their responses in the face of opposed group pressure. On the flipside, across questions a very small percentage of subjects who initially provided the

egalitarian response in isolation ended up changing to express the inegalitarian response in the face of group pressure.

[Figure 2 about here.]

While Figures 1 and 2 are suggestive, Table 3 presents logistic regression models of the likelihood of changing responses between the *Individual Pre* and *Group* periods. Responses to all of the questions used to induce social pressure in the *Group Period* are analyzed individually. In addition, the final two columns of Table 3 pool the responses such that the unit of analysis is the unique subject-question combination. To account for unobserved subject characteristics, these models also include subject-level random effects. For each individual question, and for the pooled analysis, the first column presents a baseline model and the second column controls for a variety of demographic characteristics using the self-report measures used as distraction items for the social pressure experiment. Specifically, dummy variables are used to capture race, gender, and religion, while a four-level variable is used to capture the number of years in school.[7] In addition, these models control for the size of the group that gathered in the common waiting area on the date of the particular study in which the subject participated. Because the hypothesized mechanisms of subjects altering or masking their attitudes have to do with considerations about appearing deviant from the broader population or the immediate social group, the number of subjects participating in the study at one time could plausibly cue a sense of sociality among the subjects and therefore induce variation in behavior based on the size of the group. Ultimately, the characteristics of individuals do not add explanatory value beyond what is already known. Specific coefficient estimates for the controls are omitted because, except for two individual instances, across models they all fail to achieve conventional levels of statistical significance.[8]

[Table 3 about here.]

The impact of the initial response, however, is substantial, and remains so even after controlling for a host of other factors. The model for Q2 indicates that statistical significance disappears after adding controls, but the estimated effect remains marginally significant ($p = .08$). To ease interpretation of the estimated effects, Figure 3 presents predicted probabilities for each of the

[7]Summary statistics for control variables are presented in Section D of the Online Appendix.

[8]Coefficient estimates for control variables are presented in Section D of the Online Appendix.

individual question models without controls from Table 3. The plot illustrates the estimated effect on the probability of changing responses under group pressure of moving from an inegalitarian response to an egalitarian response. For all questions, this effect is substantial. Specifically, when a subject provides an egalitarian response in private, the probability of changing responses is less than .1, regardless of the specific question. On the other hand, when a subject initially gives an inegalitarian response, the probability of changing this response in the face of group pressure increases dramatically, rising to as much as .67 in the case of Q2. Similar, though slightly less dramatic, relationships are found with respect to the other questions as well. The results lead to the conclusion that when subjects initially express an inegalitarian attitude in private, the chance of changing responses in the face of group pressure increases to between 45 and 65%.

[Figure 3 about here.]

The results presented thus far show strong initial support for the *Social Dissonance Hypothesis*, and little reason to believe the *General Social Conformity Hypothesis*. All signs in the analysis thus far point to the idea that if the subject initially gives a response that runs counter to known norms in the broader population, she will be more susceptible to group pressures to provide the opposite response. Subjects who express a view in agreement with the broader population norm show little susceptibility to social pressure to publicly comply with a group norm that runs counter to the social norm in the broader population. Whether examining subject behavior with respect to specific questions or all questions pooled together, the subject's choice initially to provide the inegalitarian response, that is, the response that is clearly not the socially desirable one, turns out to be highly consequential for predicting whether she will alter that expressed attitude when faced with unanimous opposition. And this effect is large and consistent.

However, it is not clear from the analyses presented thus far whether the social pressure that subjects encounter has been internalized to any significant degree. For all of the attitudinal questions of interest analyzed in Table 3, the subject has been compelled to respond in the last position after witnessing fellow group members responding in the direction polar opposite to the subject's private responses. But, after having been socialized to the group opinion in this manner, what happens when the subject herself is called upon to be the opinion leader? That is, after having gone through several rounds of questioning in the last position, what happens when the

subject is compelled to provide a response to attitudinal questions along the same dimension, but this time in the first position in the group? Does the level of egalitarianism expressed in private have a similar impact on the propensity to change as for the direct social pressure setting? Here the evidence is more mixed. Table 4 presents the results of logistic regressions of the likelihood of changing responses between the *Individual Pre* and *Group* periods for all of the questions for which subjects were asked to answer prior to the rest of their group members.

[Table 4 about here.]

As with the analyses for the direct social pressure questions, models are presented both for each question separately and for all of the questions pooled together, with and without demographic controls.[9] Additionally, in these models a covariate is added to capture the proportion of changed responses for the direct social pressure questions in order to provide a control for individual susceptibility to social pressure prior to being placed in the first position.

The models for Q3 and Q7 indicate that neither the level of egalitarianism exhibited in the private setting, nor the number of changed responses in the face of direct social pressure, has a significant impact on the likelihood that subjects will carry any perceived social pressure into their roles as the first group respondent. In fact, the coefficient estimates for the level of egalitarianism of the initial response are in the direction opposite expectations. For the model that pools questions, the level of egalitarianism that the subject expresses in private is not a significant predictor of the propensity to change, but the proportion of changed responses under direct social pressure conforms to expectations. For Q4, however, the story is the same as that for the direct social pressure setting. In fact, even after controlling for the individual propensity to alter responses in the *Group Period*, the magnitude of the effect of level of egalitarianism expressed in private is substantially larger than for any of the direct social pressure questions. For this final question of interest, for which the subject is asked to lead the group opinion after having been socialized to the fact that the distribution of group opinion is opposed to her privately expressed view, the result is exactly as predicted by the *Social Dissonance Hypothesis*.

Subject responses when they are compelled to follow an opposed group of peers provide strong evidence in favor of the *Social Dissonance Hypothesis*, and responses when they are compelled to

[9]Coefficient estimates for control variables are presented in Section D of the Online Appendix.

lead a group of peers provide mixed evidence that this same effect carries over to a situation with a lower degree of social pressure. But the decision to change responses between the *Individual Pre Period* and the *Group Period* does not reveal the mechanism underlying the change. The design feature intended to allow for adjudication between the *Attitude Internalization Hypothesis* and the *Preference Falsification Hypothesis* is the set of questions on the same attitudinal items in the *Individual Post Period*. However, because, as we have seen, altering responses under group pressure is so overwhelmingly associated with an initial inegalitarian response, and because so few subjects actually gave an inegalitarian response in the initial private setting, we are left with a very small number of subjects with which to examine the relationship between susceptibility to change under social pressure and subsequent private behavior. Therefore, the discussion that follows will necessarily be preliminary.

[Figure 4 about here.]

Figure 4 presents, for each question asked both in the *Group Period* and the *Individual Post Period*, and for each subject that changed responses between the *Individual Pre Period* and the *Group Period*, how many returned to their original responses and how many stayed with the changed response.[10] What is immediately apparent is that when pooling across questions, the majority of subjects who changed responses between the *Individual Pre* and *Group* periods held to their changed response rather than reverting back to the responses given in private. On first glance, given that the vast majority of subjects who changed responses under group pressure were moving from a inegalitarian state toward an expressed preference more in line with egalitarianism, this result would on the face seem to have important positive implications for the ability of interpersonal interactions with unfamiliar others to produce positive social change.

However, simply examining the behavior of the pooled majority masks important variation in responses based on the level of social pressure exerted. Specifically, it is important to note that

[10]In this discussion, reverting back to the original response is actually a shorthand for reverting back to either the egalitarian or inegalitarian side of the question's response scale. For questions with more than two response options, subjects will be considered to have reverted back to their original responses if they move from an inegalitarian response to an egalitarian response and back again, or vice versa, even if the original and final responses are not exactly the same.

all of the questions for which a majority of respondents reverted back to their original responses from the *Individual Pre Period* were used to exert direct social pressure in the *Group Period*. And even though the number of subjects who changed their initial responses is quite small for formal statistical analysis, this differential effect between the direct social pressure and indirect social pressure questions is strong enough that it is borne out in a multivariate regression setting. Table 5 presents logistic regressions of reverting back to original responses, as previously, both with and without demographic controls. Whether examining all subjects who change responses between the *Individual Pre* and *Group* periods, or limiting the analysis to only those subjects who initially gave an inegalitarian response in private, the coefficient estimates indicate that the direct social pressure questions are positively related to the likelihood of reverting back to the original response. In summary, the results provide an initial suggestion that the stronger the social pressure mechanism underlying any observed changes in responses, the more likely subjects are to falsify their expressed preferences in the opposed group setting, only to revert back when given a chance to do so in isolation. This is a plausible explanation for the observed pattern, but greater statistical power is required to determine whether this is an effect that holds when individual questions are examined in isolation. A more definitive statement about the conditions under which the *Attitude Internalization Hypothesis* and the *Preference Falsification Hypothesis* are more or less likely to hold will be left for future research.

[Table 5 about here.]

Discussion

The results of the social pressure experiment point to strong support for the *Social Dissonance Hypothesis*. As as is consistent with a large body of previous research, in the *Individual Pre Period*, when a pool of subjects was presented with a set of attitudinal questions about sensitive social issues tapping concepts such as religious equality, the vast majority expressed attitudes in favor of the egalitarian point of view. However, in the *Group Period*, interesting systematic variation became evident based on whether subjects had initially provided an inegalitarian response. Specifically, the minority of subjects who initially expressed support for an inegalitarian opinion were shown to be significantly more likely to alter their responses when later asked to answer

the same question in the face of a unanimously opposed group opinion. In other words, on any given question tapping non-Muslim subjects' attitudes with respect to Islam and Muslims, a large majority of them expressed support for an egalitarian view. These subjects were later put into a situation in which they faced a group opinion on the same question that unanimously supported an inegalitarian view, and were asked to publicly express their attitudes on this question again. In this case, the overwhelmingly dominant tendency was for those subjects to resist social pressure to publicly comply and stick with the responses that they originally gave privately. However, there was also a minority of subjects who were willing to express an inegalitarian view when they were initially asked a given question about Islam and Muslims. When this subset of subjects was put into a situation in which they were compelled to state their opinions on that same question publicly in front of a group of peers showcasing a unanimous egalitarian response, the vast majority of them altered their initial responses in order to publicly comply with the group. This association is quite strong, and carries across specific questions and slight variations in the experimental protocol. After answering several questions in the last group position, subjects were then assigned to the first group position and asked to answer several more questions along the same dimension. For one of these questions, the relationship between the level of egalitarianism expressed privately and subject behavior in the *Group Period* remained consistent with the relationship established in the direct social pressure setting.

This is exactly what was posited by the *Social Dissonance Hypothesis*. Borrowing the language of cognitive dissonance theory (Festinger 1957), social dissonance refers not to internal inconsistencies within the individual, but rather to inconsistencies between an opinion held by a subject and a known appropriate social response. Specifically, it is expected that on matters of egalitarianism with respect to various groups in society, individuals who are willing to express an inegalitarian view in private will almost surely be aware that the attitude that they have expressed runs counter to the expected appropriate social response in the broader population. Because individuals have a preference for not appearing deviant from the broader population, the dissonance between the inegalitarian view expressed and the assumed social view is expected to be a salient thought in the minds of subjects when they encounter that same question again a short time later. If the immediate social group then confronts the subject with a unanimous view

that falls in line with that of the broader population, the subject's preference for not appearing deviant will have a tendency to overwhelm any preferences for stating what she truly believes, or even for consistency with a view that was expressed only a short time prior. This preference is then likely to lead to public compliance with the expressed group opinion. This expectation was strongly supported by the experimental results.

With respect to the specific form that the public compliance takes (attitude internalization vs. preference falsification), the evidence is weaker because of the relatively small number of subjects that altered their private responses in the face of group pressure. Pooling the responses of all subjects who initially changed, the evidence indicates that the majority of them carried the changed response into the final private setting. However, there is also systematic variation in this tendency based on whether the attitudinal measure in question was used to exert direct social pressure in the group setting. Specifically, subjects who changed responses when they were compelled to follow an opposed group were significantly more likely to revert back to their original private responses when given the opportunity to do so. Though this relationship was not anticipated, it seems plausible to believe that greater social pressure leads to a greater propensity to falsify beliefs in the face of group pressure. But future work will be required to further unpack this explanation.

Conclusion

In the contemporary industrialized world, norms of appropriate social behavior dictate that on questions of race, ethnicity, gender, religion, and so on, support for egalitarian principles is the only acceptable mode of mainstream social and political behavior. And not surprisingly, citizens—especially those with greater awareness about and knowledge of the political system—express high levels of support for egalitarian principles across issue domains. But it remains a constant struggle for empirical researchers to determine the truth and falsehoods underlying those responses. The main source of the trouble is the understanding that humans have a strong desire to not appear to be deviant from those around them, especially when those around them are able to witness their behavior. It is that desire that leads to the thought that in an open setting, when asked to publicly express a view on a sensitive social issue like racial equality, some substantial portion of

citizens are likely to falsify their true preferences in favor of the actual or perceived distribution of preferences in the group. To the extent that many individuals feel this way, widespread preference falsification might be expected to ensue.

This paper explored various aspects of the relationship between interpersonal interactions and the behaviors that individuals are willing to exhibit in public settings. Leaving aside the question of whether the opinions that individuals express on sensitive social matters represent their true beliefs, the experiment examined some dynamics of social influence by comparing baseline measures of political attitudes to measures of those same attitudes when social pressure was employed to push subjects away from the baseline. The asymmetric effect found with respect to the propensity to alter responses is consistent with the explanation that after knowingly providing a response that is considered socially undesirable, subjects generally experience some level of social discomfort. This discomfort becomes enhanced and more salient when faced with a unanimous majority expressing the socially desirable, egalitarian response. Apparently, the psychic cost of simultaneously deviating from the group as well as the broader societal norm is too great for these subjects, and overwhelms any desire to express their true attitudes or be consistent. Further, though the majority of subjects who altered their initial private responses under group pressure stayed with their group responses when they were asked the same question again in private, the tendency to revert back to the original private response after having changed was increasing in the level of social pressure exerted. Future research will bring more evidence to bear on this relationship.

The experimental results presented here have clear implications for the transition in norms of public speech that has been witnessed in multicultural societies in the twentieth century. Take, for example, norms of public speech with respect to race in the United States. In a previous era, norms of appropriate social behavior dictated that blacks could be, and perhaps should be, treated unequally by whites in public settings. By the end of the twentieth century, the opposite was true: any public rhetoric expressing a preference for racial inequality would be met with social censure and generally severe reputational costs. Scholars who examine American public opinion on race are well aware that self reports of racial attitudes may actually overstate the extent to which the public adheres to the norm of racial equality. And the results presented

here may provide a mechanism. Knowledge of the appropriate social response leads adherents to the previous norm to alter their responses in public settings. This suggests that the transition in norms of public rhetoric that has been witnessed may not have been accompanied by a large-scale change in attitudes. Future avenues of research will delve deeper into this question.

References

Adams, Henry F. 1912. "Autokinetic Sensations." *Psychological Monographs* 14(59): 1–45.

Ahn, T.K., Robert Huckfeldt, and John Barry Ryan. 2010. "Communication, Influence, and Informational Asymmetries among Voters." *Political Psychology* 31(5): 763–787.

Asch, S. E. 1951. "Effects of Group Pressure Upon the Modification and Distortion of Judgments." In *Groups, Leadership and Men*, ed. Harold Guetzkow. New York: Russell and Russell pp. 177–190.

Asch, Solomon E. 1952. *Social Psychology.* New York: Prentice Hall.

Asch, Solomon E. 1956. "Studies of Independence and Conformity: I. A Minority of One Against a Unanimous Majority." *Psychological Monographs* 70(9).

Barabas, Jason. 2004. "How Deliberation Affects Policy Opinions." *American Political Science Review* 98(4): 687–702.

Burt, Ronald S. 1987. "Social Contagion and Innovation: Cohesion versus Structural Equivalence." *American Journal of Sociology* 92(6): 1287–1335.

Crowne, Douglas P., and David Marlowe. 1960. "A New Scale of Social Desirability Independent of Psychopathology." *Journal of Consulting Psychology* 24(4): 349–354.

Deutsch, Morton, and Harold B. Gerard. 1955. "A Study of Normative and Informational Social Influences Upon Individual Judgment." *Journal of Abnormal and Social Psychology* 51(3): 629–636.

Eulau, Heinz. 1986. *Politics, Self, and Society.* Cambridge, MA: Harvard University Press.

Festinger, Leon. 1954. "A Theory of Social Comparison Processes." *Human Relations* 7(2): 117–140.

Festinger, Leon. 1957. *A Theory of Cognitive Dissonance.* Evanston, IL: Row, Peterson.

Granovetter, Mark. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." *American Journal of Sociology* 91(3): 481–510.

Hardin, Curtis D., and E. Tory Higgins. 1996. "Shared Reality: How Social Verification Makes the Subjective Objective." In *Handbook of Motivation and Cognition: The Interpersonal Context*, ed. Richard M. Sorrentino, and E. Tory Higgins. New York: Guilford Press pp. 28–84.

Huckfeldt, Robert, and John Sprague. 1987. "Networks in Context: The Social Flow of Political Information." *American Political Science Review* 81(4): 1197–1216.

Jacobs, Robert C., and Donald T. Campbell. 1961. "The Perpetuation of an Arbitrary Tradition Through Several Generations of a Laboratory Microculture." *Journal of Abnormal and Social Psychology* 62(3): 649–658.

Katz, Elihu. 1957. "The Two-Step Flow of Communication: An Up-To Date Report on an Hypothesis." *Public Opinion Quarterly* 21(1): 61–78.

Katz, Elihu, and Paul F. Lazarsfeld. 1955. *Personal Influence: The Part Played by People in the Flow of Mass Communications*. New York: Free Press.

Kelley, Harold H. 1952. "Two Functions of Reference Groups." In *Readings in Social Psychology*, ed. Guy E. Swanson, Theodore M. Newcombe, and Eugene L. Hartley. Revised ed. New York: Henry Holt and Company pp. 410–414.

Kuran, Timur. 1995. *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. Cambridge, MA: Harvard University Press.

Lazarsfeld, Paul F., Bernard Berelson, and Hazel Gaudet. 1948. *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign*. New York: Columbia University Press.

Levitan, Lindsey Clark, and Penny S. Visser. 2008. "The Impact of the Social Context on Resistance to Persuasion: Effortful versus Effortless Responses to Counter-attitudinal Information." *Journal of Experimental Social Psychology* 44(3): 640–649.

MacKuen, Michael. 1990. "Speaking of Politics: Individual Conversational Choice, Public Opinion, and the Prospects for Deliberative Democracy." In *Information and Democratic Processes*, ed. John A. Ferejohn, and James H. Kuklinski. Urbana, IL: University of Illinois Press pp. 59–99.

Mendelberg, Tali. 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton: Princeton University Press.

Mutz, Diana. 2002. "Cross-Cutting Networks: Testing Democratic Theory in Practice." *American Political Science Review* 96(1): 111–126.

Mutz, Diana C. 1998. *Impersonal Influence: How Perceptions of Mass Collectives Affect Political Attitudes*. Cambridge: Cambridge University Press.

Noelle-Neumann, Elisabeth. 1984. *The Spiral of Silence: Public Opinion—Our Social Skin.* Chicago: University of Chicago Press.

Schuman, Howard, Charlotte Steeh, Lawrence Bobo, and Maria Krysan. 1998. *Racial Attitudes in America: Trends and Interpretations.* Revised ed. Cambridge, MA: Harvard University Press.

Sherif, Muzafer. 1935. "A Study of Some Social Factors in Perception." *Archives of Psychology* 27(187).

Sherif, Muzafer. 1936. *The Psychology of Social Norms.* New York: Harper and Brothers.

Tingsen, Herbert. 1937. *Political Behavior: Studies in Election Statistics.* London: King.

Verhulst, Brad, and Lindsey Clark Levitan. 2009. "Conformity in Groups: The Effects of Groups on Expressed Attitudes." Paper presented at the Annual Meeting of the International Society of Political Psychology, July 14, Dublin, Ireland.

Visser, Penny S., and Robert R. Mirabile. 2004. "Attitudes in the Social Context: The Impact of Social Network Composition on Individual-Level Attitude Strength." *Journal of Personality and Social Psychology* 87(6): 779–795.

Figure 1: Individual Responses to Questions about Muslims and Islam in *Individual Pre Period*

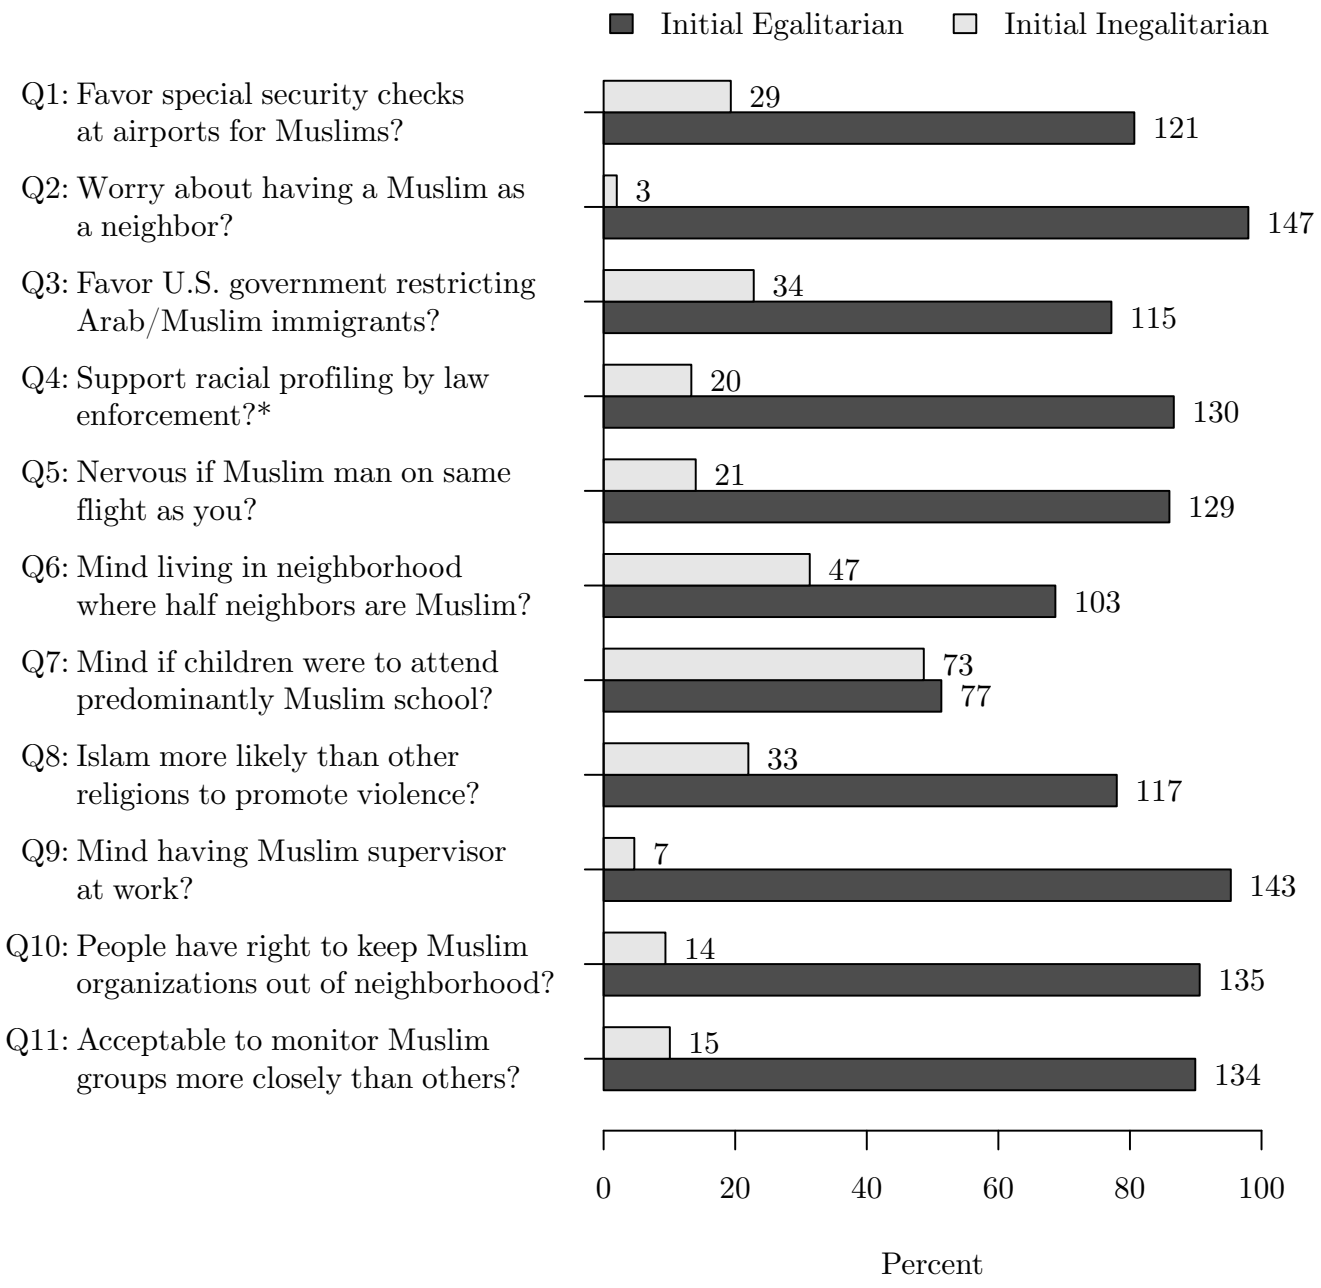

Note: Questions are listed from top to bottom in the order that they appear in the *Individual Pre Period*. ID numbers are used to refer to particular questions in the text.
*Question does not specifically mention Islam or Muslims.

Figure 2: Individual Responses Across *Individual Pre Period*, *Group Period*, and *Individual Post Period* in the Social Pressure Experiment

Q5: Nervous if Muslim man on same flight as you?

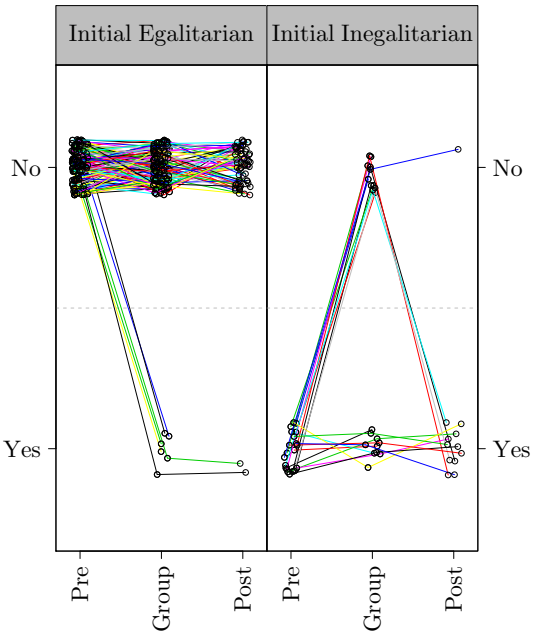

Q6: Mind living in neighborhood where half neighbors are Muslim?

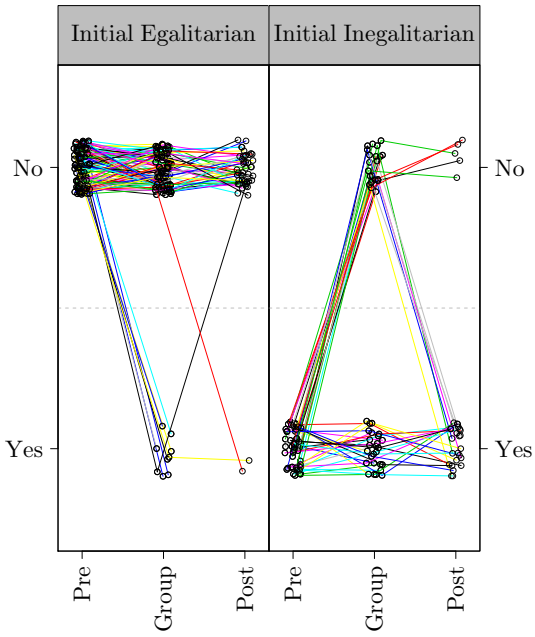

Q1: Favor special security checks at airports for Muslims?

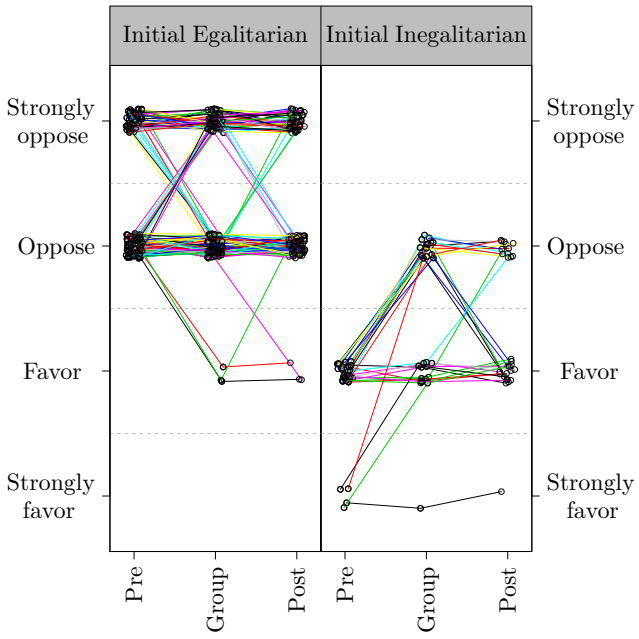

Q2: Worry about having a Muslim as a neighbor?

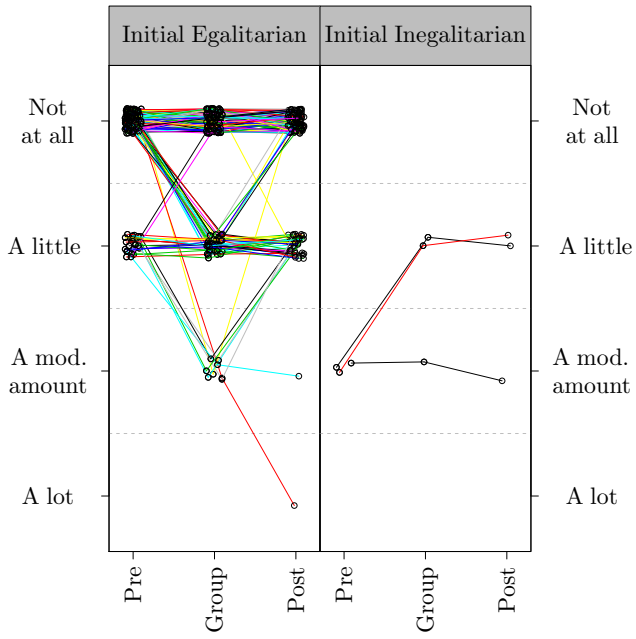

Note: Responses are separated according to whether subjects provided the egalitarian response or the inegalitarian response in the *Individual Pre Period*, indicated in the panel labels. For egalitarian and inegalitarian response choices, consult text and Table 2.

Figure 3: Predicted Probabilities of Changing Responses Between *Individual Pre* and *Group* Periods for Direct Social Pressure Questions

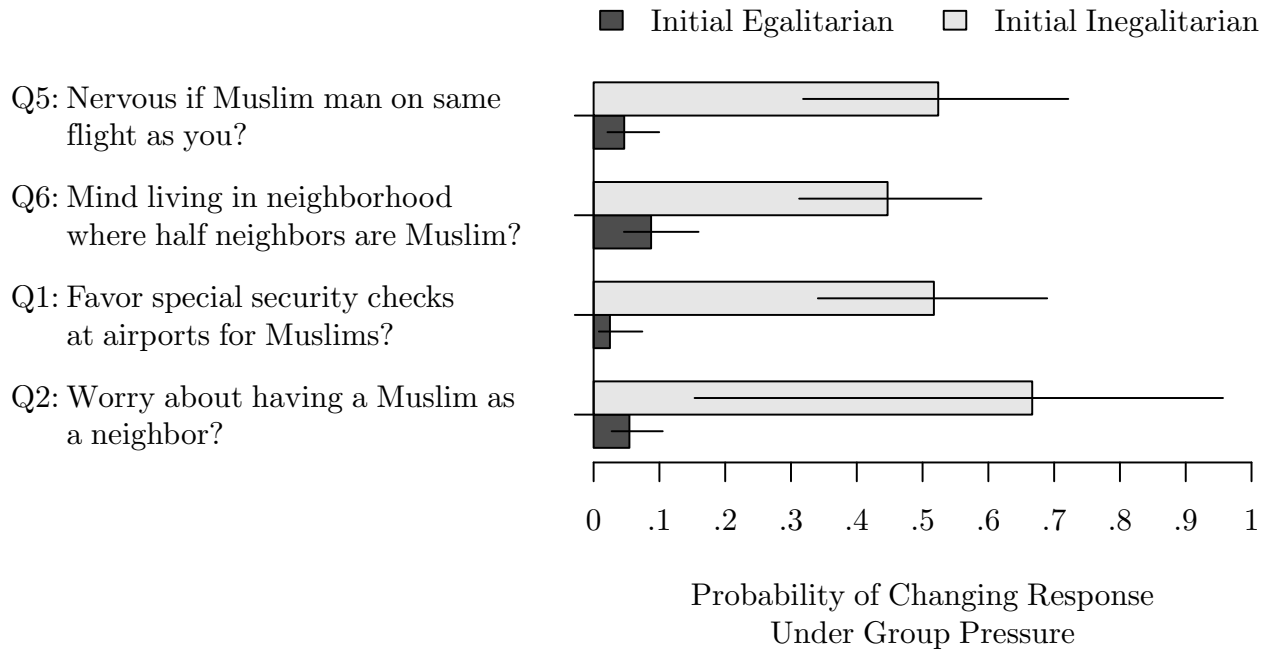

Note: Questions are listed from top to bottom in the order that they appear in the *Group Period*. ID numbers are used to refer to particular questions in the text. Line segments represent 95% confidence intervals.

Figure 4: Attitude Expression in *Individual Post Period* for Subjects who Changed Responses Between *Individual Pre* and *Group* Periods

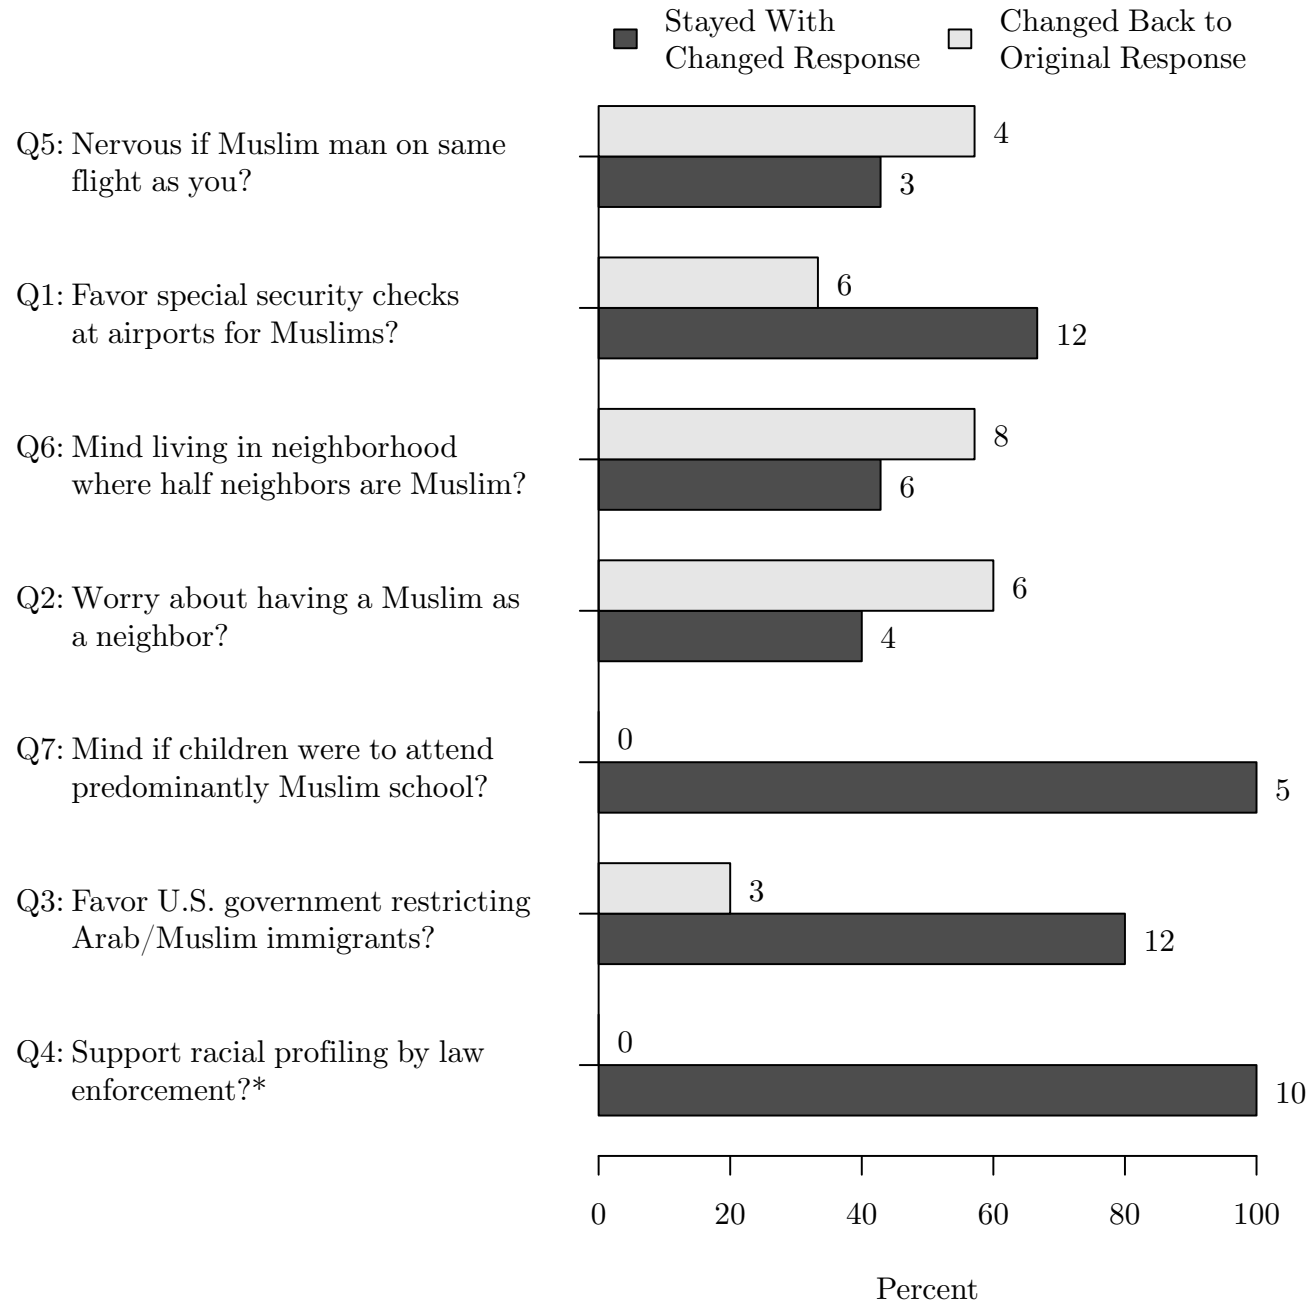

Note: Questions are listed from top to bottom in the order that they appear in the *Individual Post Period*. ID numbers are used to refer to particular questions in the text.
*Question does not specifically mention Islam or Muslims.

Table 1: Summary of Questions about Muslims and Islam Used to Examine Social Pressure

ID	Question Summary	Response Scale	Egalitarian Response	Inegalitarian Response
Q1	Favor special security checks at airports for Muslims?	4-point	Str. oppose/ Oppose	Str. favor/ Favor
Q2	Worry about having a Muslim as a neighbor?	4-point	A little/ Not at all	Mod. amount/ A lot
Q3	Favor U.S. government restricting Arab/Muslim immigrants?	4-point	Str. oppose/ Oppose	Str. favor/ Favor
Q4	Support racial profiling by law enforcement?	4-point	Str. oppose/ Oppose	Str. support/ Support
Q5	Nervous if Muslim man on same flight as you?	2-point	No	Yes
Q6	Mind living in neighborhood where half neighbors are Muslim?	2-point	No	Yes
Q7	Mind if children were to attend predominantly Muslim school?	2-point	No	Yes
Q8	Islam more likely than other religions to promote violence?	2-point	No	Yes
Q9	Mind having Muslim supervisor at work?	2-point	No	Yes
Q10	People have right to keep Muslim organizations out of neighborhood?	2-point	No	Yes
Q11	Acceptable to monitor Muslim groups more closely than others?	2-point	Violates rights	Acceptable

Note: ID Numbers are used to refer to particular questions in the text. Combined response options in the egalitarian and inegalitarian response categories are separated by slashes. For full question wording, see Section A of the Online Appendix.

Table 2: Summary of *Group Period* Dynamics for Questions about Muslims and Islam

ID	Subject Position in *Group Period*	Subject Response in *Individual Pre Period*	Peer 1 Response in *Group Period*	Peer 2 Response in *Group Period*
Q5	Last	Yes	No	No
		No	Yes	Yes
Q6	Last	Yes	No	No
		No	Yes	Yes
Q1	Middle	Str. favor/Favor	Oppose	Str. oppose
		Str. oppose/Oppose	Favor	Str. favor
Q2	Last	Mod. amount/A lot	A little	Not at all
		A little/Not at all	Mod. amount	A lot
Q3	First	Str. favor/Favor	Oppose	Oppose
		Str. oppose/Oppose	Favor	Favor
Q7	First	Yes	No	No
		No	Yes	Yes
Q4	First	Str. support/Support	Oppose	Oppose
		Str. oppose/Oppose	Support	Support

Note: ID Numbers are used to refer to particular questions in the text. For question and response summaries, see Table 1. For full question wording, see Section A of the Online Appendix.

Table 3: Logistic Regression Models of Changing Responses Between *Individual Pre* and *Group* Periods for Direct Social Pressure Questions

	Q5: 1st Group Question†		Q6: 2nd Group Question†		Q1: 3rd Group Question‡		Q2: 4th Group Question†		Pooled Direct Social Pressure Questions	
Initial Inegalitarian	3.12*	3.47*	2.13*	2.11*	3.74*	4.35*	3.55*	2.55	3.39*	3.30*
	(0.60)	(0.71)	(0.46)	(0.50)	(0.69)	(0.88)	(1.28)	(1.47)	(0.44)	(0.45)
Number of Participants		0.50*		0.15		0.08		0.19		0.23
		(0.24)		(0.17)		(0.24)		(0.25)		(0.17)
Constant	-3.02*	-6.26*	-2.35*	-3.30*	-3.67*	-6.72*	-2.86*	-4.85*	-3.80*	-5.83*
	(0.42)	(1.96)	(0.35)	(1.42)	(0.58)	(2.40)	(0.36)	(2.01)	(0.46)	(1.43)
Controls	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Subjects	150	150	150	150	150	150	150	150	150	150
Observations									600	600
Var(α_i)									2.75	2.47
Log-likelihood	-38.80	-36.15	-62.84	-59.11	-34.13	-30.76	-32.98	-29.95	-162.36	-158.66

Note: Cell entries are logistic regression coefficients with standard errors in parentheses. The binary dependent variable is 1 if the subject changed responses to a given question between the *Individual Pre Period* and the *Group Perod*, and 0 otherwise. Unit of analysis for columns 1–8 is the subject. Unit of analysis for columns 9 and 10 is the subject-question. Models in columns 9 and 10 include subject-level random intercepts, α_i, where i indexes subjects.

†Subject appears in last group position; ‡Subject appears in middle group position.

*$p < .05$, two-tailed.

Table 4: Logistic Regression Models of Changing Responses Between *Individual Pre* and *Group* Periods for Indirect Social Pressure Questions

	Q3: 5th Group Question†		Q7: 6th Group Question†		Q4: 7th Group Question†		Pooled Indirect Social Pressure Questions	
Initial Inegalitarian	-1.29	-1.73	-1.87	-1.77	4.46*	5.25*	0.10	0.00
	(0.89)	(0.94)	(0.99)	(1.00)	(1.14)	(1.56)	(0.58)	(0.57)
Proportion of Changed Responses	2.13	2.16	2.37	2.46	3.27*	4.07*	2.67*	2.59*
	(1.15)	(1.20)	(1.56)	(1.62)	(1.59)	(1.71)	(1.13)	(1.08)
Number of Participants		-0.18		-0.33		-0.45		-0.20
		(0.21)		(0.29)		(0.41)		(0.20)
Constant	-2.29*	-2.01	-2.44*	-3.50	-5.56*	-3.74	-4.17*	-3.78*
	(0.33)	(1.55)	(0.42)	(2.22)	(1.18)	(3.07)	(0.52)	(1.55)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
Subjects	150	150	150	150	150	150	150	150
Observations							449	449
Var(α_i)							2.97	2.16
Log-likelihood	-46.59	-44.08	-31.62	-28.27	-17.34	-16.16	-110.80	-108.67

Note: Cell entries are logistic regression coefficients with standard errors in parentheses. The binary dependent variable is 1 if the subject changed responses to a given question between the *Individual Pre Period* and the *Group Perod*, and 0 otherwise. Unit of analysis for columns 1–6 is the subject. Unit of analysis for columns 7 and 8 is the subject-question. Models in columns 7 and 8 include subject-level random intercepts, α_i, where i indexes subjects.

†Subject appears in first group position.

*$p < .05$, two-tailed.

Table 5: Pooled Logistic Regression Models of Changing Back to Original Response in *Individual Post Period*

	All		Initial Inegalitarian	
	Pre→*Group* Changers		*Pre*→*Group* Changers	
Direct Social Pressure Question	2.16*	2.43*	2.31*	2.73*
	(0.67)	(0.77)	(1.17)	(1.31)
Number of Participants		0.61*		0.94*
		(0.25)		(0.39)
Constant	-2.20*	-5.88*	-2.47*	-9.58*
	(0.61)	(2.19)	(1.12)	(3.46)
Controls	No	Yes	No	Yes
Subjects	48	48	31	31
Observations	79	79	46	46
Var(α_i)	0.00	0.00	0.35	0.00
Log-likelihood	-43.71	-38.89	-26.87	-21.93

Note: Cell entries are logistic regression coefficients with standard errors in parentheses. The binary dependent variable is 1 if the subject changed responses to a given question between the *Group Period* and the *Individual Post Period*, and 0 otherwise. Analysis is limited to subjects who changed responses between the *Individual Pre Period* and the *Group Period* and to questions that were asked in all three rounds of the experiment. Unit of analysis is the subject-question. Models include subject-level random intercepts, α_i, where i indexes subjects.
*$p < .05$, two-tailed.