The Impact of Ideological Diversity on Educational Performance: A Pilot Study

Cecilia Elena Rouse^{1*}, Elizabeth Levy Paluck^{1*}, Jason C. Chin², Amy Wickett³,
and John-Henry Pezzuto¹

¹Princeton University

²UCLA

³Harvard University

October 18, 2019

* Corresponding authors

Acknowledgements: Funding provided by Princeton University Industrial Relations Section and the Princeton Department of Psychology. The study was approved by the Princeton University IRB Board #10025. Data and code for the project are available; code is posted at OSF at <https://osf.io/g7fwf/> and data are available from corresponding authors on request with IRB permissions.

The Impact of Ideological Diversity on Educational Performance: A Pilot Study

Abstract

What is known about the real-world educational effects of ideological diversity in the classroom? Surprisingly, given the amount of attention paid to the role of ideological diversity on higher education outcomes such as critical thinking and academic performance, scant causal evidence exists. We use a lab-in-the-field experiment to test whether the presence of ideologically more conservative students in academic discussion groups, as compared to groups of students who all slanted ideologically liberal, would improve academic outcomes in terms of the quality of each student's individual academic work. The complete population of an incoming cohort of policy graduate students ($N = 78$) took part in the experiment. Results demonstrate that students assigned to the ideologically heterogeneous discussion groups subsequently wrote individual assignments that received significantly more negative grades by a professional grader blind to experimental condition and to student identity. Survey results from participating students also suggest that students in the ideologically heterogeneous discussion groups were also significantly more likely to perceive interpersonal conflict and to dislike their group dynamics—a result that was not driven by students of a particular ideological slant. As a small pilot, this study provides questions to resolve with future research, including the role of pedagogy in managing ideological diversity, and provides a template for future experimental designs.

Early advocates for diversity motivated their claims using ideas of fairness and justice (e.g., Hurd & Plaut, 2018; Frymer & Skrentny, 2003; Richard & Kirby, 1999). In contrast to these moral rationales for promoting diversity, the past few decades have witnessed a rise in claims about the utilitarian benefits of diversity (Berrey, 2015; Wolfe and Fletcher 2013). Particularly in discussions of the labor force, utilitarian-oriented diversity advocates advance claims about the payoffs of diversity such as increased productivity, better decision-making, more ideas, more creativity, and other broader benefits like appeal to a consumer base or good positioning in the “knowledge economy.” These benefits are thought to come from many forms of diversity, including from demographic identities like race, gender, and religious identity, and from traits like thinking styles.

Utilitarian claims about the benefits of diversity are also used in higher education, where payoffs are often characterized as gains in learning, better critical thinking, and broader perspectives. For example, in a 1999 amicus brief to *Gratz v. Bollinger*, Patricia Gurin summarized research in higher education suggesting that racial and ethnic diversity is associated with students’ “...active thinking processes, growth in intellectual engagement and motivation, and growth in intellectual and academic skills” (Gurin, 1999). Gurin and colleagues (1999) argued that the presence of underrepresented students on campus transforms *all* students’ learning experiences by changing the institutional structure (e.g., by augmenting the set of courses that are relevant to teach), and by changing interpersonal interactions (e.g., by introducing a wider variety of conversation partners both inside and outside of the classroom).

More recently, claims about the benefits of diversity in higher education have been applied within a different dispute, regarding whether higher education in the United States is too

ideologically homogenous, namely too liberal. Those making this claim have argued for the value of “viewpoint diversity,” or the presence of a range views, including but not limited to political ideologies (Duarte et al., 2015). Proponents of viewpoint diversity believe that critical thinking, and ideas and research generally, improve when there is a diversity of political and ideological views present on campus.

What is known about the actual behavioral effects of diversity, in the workplace and in higher education? The literature is too vast to summarize here. But surprisingly, given the importance of the topic and the amount of attention it has received, there is a limited and somewhat mixed social science literature on the causal impact of diversity on workplace or educational outcomes like individual or team productivity, creativity, or learning. Research offers a mix of positive and negative effects for people who are assembled for the purpose of a one-time study, usually in a laboratory (e.g., McLeod, Lobel & Cox, 1996; Cherry, Kroll & Shogren, 2005; Inzlicht & Ben-Zeev, 2000; Antonio et al., 2004; Sekaquaptewa & Thomson 2002). Additionally, some studies reveal benefits of one sort of diversity but not another, such as one by Hansen, Owan, and Pan (2006), which finds task group diversity improves group productivity, but only when defined as greater proportions of women and not as greater proportions of people of color. Parsing when and why diversity yields benefits is difficult because studies define diversity differently, measure different kinds of outcomes, and show that results are not stable across different contexts, sometimes even within the same study. A final concern is that scholars rarely study the causal effect of diversity on real-world work groups within university or employment settings, where claims about diversity’s benefits have received the most attention.

While this large and fragmented body of work does not offer unified findings as of yet, there are a few general lessons. The first is that diversity can positively affect attitudes and behaviors under various circumstances, but when diversity does not produce benefits for all, it can be difficult to experience and to manage well. Some psychological studies show that diverse group discussions exert strain on group members who oppose majority viewpoints (Van Dyne & Saavedra, 1996), minority members can experience tokenism (Torchia, Calabro & Huse, 2011), and stereotypes or perceptions of “sides” in diverse communities can lead to interpersonal conflict and reactive devaluation of colleagues’ work (e.g., Sackett, DuBois, & Noe, 1991). Second, across contexts, mixed findings on the benefits of diversity seem to be associated with variations in the structure of the interactions, including the structure of group tasks, or in general how differences are handled as part of the group’s process. For example, “devil’s advocate” approaches to discussion in diverse groups seem to be less successful than encouraging anyone to dissent, and in general more structure is better than less structure (Nemeth, Personnaz, Personnaz, & Goncalo, 2004; Sunstein & Hastie, 2015; Wanous & Youtz, 1986).

Some political science work has borne out the claim that viewpoint diversity aids critical thinking. One study suggested that ideologically heterogeneous (vs. homogenous) political discussions stimulate cognitive elaboration and information-seeking (Levendusky, Druckman, & McClain, 2016). However, the political science literature primarily focuses on the outcome of attitude strength, finds mixed results on the effects of discussion group composition, and measures these attitudinal outcomes outside of the workplace or the classroom (e.g., Farrar et al., 2009; Robison, Leeper, & Druckman, 2018)

While accumulated findings seem suggestive of lessons for higher education, it is notable that there are no field experiments testing the causal impact of diversity in post-secondary schooling. A number of quasi-experimental or observational studies find evidence that diversity affects educational and attitudinal outcomes (Gurin, 1999; Milem, 2003, p. 126-169). However, concerns about omitted variable bias (e.g., students with particular traits choosing more diverse campuses or class topics) open these studies to critique. There have been only a few econometric tests that take advantage of random roommate assignment (Sacerdote, 2011) or class-based random assignment to a group task (Hansen, Owan, & Pan, 2006) to test claims about learning and diversity. These studies find modest positive effects of roommate demographic difference on grade point average (Sacerdote, 2011), and some positive (regarding gender diversity) and some null (regarding racial diversity) effects of peer composition in classroom task groups on class performance (Hansen, Owan, & Pan, 2015). These well-identified findings represent a very small proportion of studies in this literature, and of the types of claims made about diversity in higher education. Further, to date we are aware of no completed studies testing the effect of “viewpoint diversity,” specifically a mix of ideologically conservative and liberal views.

For all of these reasons, rigorous research is needed on the as-yet untested question of whether viewpoint diversity, and specifically ideological diversity, improves individual educational outcomes. In response, we planned a lab-in-the-field experiment. To our knowledge, this is the first experiment of its kind to test the impact of ideological diversity on academic outcomes. The experiment tested whether the presence of ideologically more conservative students in academic discussion groups, as compared to groups of students who all slanted ideologically liberal, would improve academic outcomes in terms of the quality of each student’s individual academic work. In the same study, we measured students’ experience of these

ideologically heterogeneous vs. homogeneous discussions, to understand the impact of ideological diversity on students' subjective experience, and also as a way to contextualize any effects on academic outcomes. As a small-n study of one entire cohort of a small professional school, we consider the present study to be a template for further research on this important topic. While we did not pre-register our predictions for the experiment, one of us, the dean of a policy school, published an opinion piece in advance of the study advocating for ideological diversity as a means for students to come to more considered and thoughtful decisions as policymakers (Rouse, 2017). Thus, for at least one member of the research team, expectations are on record as optimistic for the effects of viewpoint diversity on student outcomes.

Method

We examine the effect of ideological diversity in an educational setting where political ideology is highly task-relevant: at a policy school. In an afternoon-long experiment, we randomly assigned the entire entering class of graduate students at a policy school in the Northeast to ideologically homogenous or ideologically heterogeneous group discussions about a policy issue. Prior to the discussion on the same day, students completed an academic exercise in which they wrote a policy memo. Next, they were assembled by researchers into groups of three to discuss their memo's policy recommendations. After the discussion, students were separated and asked to individually rewrite their policy memo. Both the pre- and post-discussion memos were later graded for argument quality, writing quality, and overall quality by a writing professor from the policy school who was blind to condition.

Critically, students completed surveys after writing each of their memos, allowing us to elicit responses about their immediate experiences within the academic exercise. Analyzing the

results of these surveys, we examine the mechanisms by which ideological diversity affects the process and personal experiences of group discussions, and subsequent critical thinking about policy.

Materials and Methods

Participants and procedure.

Participants were 78 graduate students (49% male; 51% female) from the entering class of a policy school in the northeastern United States. Of participants, 26 (33%) were international students, 15 (19%) of whom were studying in the United States for the first time. The sample's ethnic composition was diverse: of participants, 11 (14%) were Hispanic or Latinx, eight (10%) were black, 13 (17%) were Asian or Pacific Islander, eight (10%) were of "other" racial or ethnic background, and 38 (49%) were white. Finally, 13 (17%) participants identified as slightly or more conservative on economic issues, each of whom was assigned to one of the ideologically heterogeneous groups.

Students' political ideologies, and other demographics, were elicited by survey three weeks prior to the exercise. Specifically, they were asked the liberal-conservative self-identification question adapted directly from the American National Election Studies (ANES) Survey (2016). In addition to their general political ideology, students were also asked about ideological positions on social and economic issues. For the purposes of conditional random assignment to groups, we defined conservative students as those who self-identify as "slightly conservative," "conservative," "extremely conservative," or "moderate; middle of the road" on the ANES general ideology question and at least slightly conservative on economic issues, as

defined by the question “In terms of economic issues, how liberal or conservative are you?” which was answered on a 7-point Likert scale with 1 being “extremely liberal,” 2 being “liberal,” 3 being “slightly liberal,” 4 being “neither,” 5 being “slightly conservative,” 6 being “conservative,” and 7 being “extremely conservative.” This is formally defined by equation (1):

$$cons = i \geq 4, e \geq 5 \text{ where } 1 \leq i, e \leq 7 \quad (1)$$

Eighty-three students were randomly assigned to 28 group discussions of three people in one of two conditions: either ideologically homogenous (control) or ideologically heterogeneous (treatment). Ideologically heterogeneous discussion groups had one conservative student (the student body tends liberal). Groups were assigned on site on the day of the experiment, after the authors took attendance for the exercise (there are a total of 95 students in the cohort; 84 attended the experiment on the day; six were unable to stay for the duration of the exercise for reasons unrelated to the experiment and were excluded from analysis for a 7.1% attrition rate). In total, the 78 students who completed the entire experiment were divided into 13 ideologically heterogeneous discussions and 15 ideologically homogeneous discussions.

Measures.

Pre-discussion memo (memo 1). All participants were asked, three weeks in advance, to attend an afternoon exercise prior to the start of the academic term. Upon entry, all participants received a prompt for a policy memo, which they wrote individually in the 15 minutes allotted. The prompt asked participants to make recommendations for the Prime Minister of Canada heading into the Paris Climate Accords, on the issue of climate reparations: whether rich developed nations should take greater financial responsibility for climate change than developing

nations (see S1 Appendix 2 for memo prompts and all surveys described below). We attempted to select a policy area that was not unique to the U.S. and would not lead to viewpoints that broke along other demographic lines, e.g., race or gender, as they might in a predictably pronounced manner in discussions regarding immigration, affirmative action, or abortion.

Pre-discussion survey. After the allotted 15 minutes, experimenters collected these memos and distributed a survey (“pre-discussion survey”) asking students to rate their knowledge of climate change and energy policy (which was relevant for the policy memo they would write), confidence in their policy memo, comfort debating policy, and measures of extraversion. Most of the variables were measured using a 7-point Likert scale, where 1 reflected poor/unlikely and 7 reflected exceptional/very likely. One question inquiring how long the participant had worked on energy policy used four response categories ranging from “never” to “more than 3 years.” The pre-discussion survey also included two items about traits regarding liking for talking and for being the center of attention, both taken from the Big Five Extraversion Factor (Goldberg et al., 2006), for a total of nine items.

Group discussions. After 10 minutes allotted for completing the pre-discussion survey, the experimenters collected survey responses and then randomized students to groups of three, conditional on ideology. Participants were given their group assignments without a mention of the rationale for their composition, and directed to nearby classrooms where they discussed with their randomly-assigned group members the recommendations they made in their policy memos. Directions for discussion were kept to a minimum, neither encouraging nor discouraging debate: “You will have about 20 minutes to discuss your briefing document—your recommendation and the arguments used to support it—with your group. There are no guidelines for the discussion.

Please use the 20 minutes to discuss the recommendations and why you made them with your group members.”

Post-discussion memo (memo 2). Following discussion, participants were directed back to the auditorium, and provided the opportunity to individually re-write their memos under the instructions: “Earlier you drafted a briefing document that consisted of bullet points for the Prime Minister of Canada to review as preparation for his meeting with a global coalition of other heads of state where the Paris Climate Accord will be discussed. You now have the opportunity to redraft it ...” They were given another 15 minutes to write, and were not allowed to refer to the first memo they drafted before the discussion. Experimenters then collected the redrafted memos and distributed a post-discussion survey.

Post group discussion survey. The post group discussion survey (“post-discussion survey”) asked 22 questions about participants’ individual experiences in their group discussions, how their group discussions functioned, the content of their discussions, and how they would estimate the ideological composition of their groups and of the school’s graduate student population. After an allotted 15 minutes, experimenters collected the post-discussion survey, concluding data collection. The authors then debriefed the participants as a group, transitioning into a discussion of student concerns about diversity on campus.

Analysis Strategy

We tested the effects of ideologically diverse groups on assessments of student writing employing linear models with robust clustered standard errors at the level of the discussion

group. Thus we model the “score” or assessment of the writing of individual i in discussion group j (Y_{ij}) using models such as those in equations (2)-(4):

$$Y_{ij} = \beta_0 + \tau D_{ij} + \varepsilon_{ij} \quad (2)$$

$$Y_{ij} = \beta_0 + \tau D_{ij} + \boldsymbol{\theta} \mathbf{X}_{ij} + \varepsilon_{ij} \quad (3)$$

$$Y_{ij} = \beta_0 + \tau D_{ij} + \boldsymbol{\theta} \mathbf{X}_{ij} + \gamma P_{ij} + \varepsilon_{ij} . \quad (4)$$

In these equations, the coefficient τ represents the causal effect of the ideological treatment represented by the binary D_{ij} in which $D_{ij} = 1$ refers to participants in the ideologically diverse condition, and $D_{ij} = 0$ refers to participants in the ideologically homogenous condition; β_0 is the intercept and ε_{ij} represents the error term. Thus, in equation (2) τ represents the unconditional difference in means between the two types of groups. In equation (3) we condition on a vector of covariates ($\mathbf{X}_{ij}$) which includes a student’s sex, whether the student served in the military, is non-white, and is a first time in U.S. student. Equation (4) also includes the “score” the student on the pre-discussion essay (P_{ij}).

We also sought to understand the impact of assignment to an ideologically heterogenous group on the quality of the group discussion. Because there were so many individual items in our survey, we sought to group individual items. To do so, we conducted exploratory factor analysis using oblimin rotation which resulted in 4 factors named “Agreement in Group Discussion,” “Frustration in Group Discussion,” “Restraint in Group Discussion.” and “Enthusiasm in Group Discussion” (henceforth referred to as indices) which cumulatively explain

54.4% of variance. See Appendix 1 for additional detail on how the factor analysis was conducted.

Results

A key feature of our study was that we randomly assigned students to treatment status, thereby addressing potential issues of self-selection into ideologically homogenous or heterogenous groups. Table 1 shows descriptive statistics for the analysis sample by treatment status, whether the group was a “control” (meaning it included three liberal students) or “ideologically diverse” (meaning it included one conservative and two liberal students). The first panel of variables reflects the demographic characteristics of the sample. Notably, as shown in column (3), there was a 20 percentage point difference in the proportion male in the treatment groups compared to the control groups, a difference that was significant at the 5 percent level of significance (see column (4)). Similarly, there was a 16 percentage point difference in the proportion of the group that indicated it was their first time in the U.S.; this difference was also statistically significant at the 10 percent level. The treatment and control groups also differed in the proportion of the group indicating that they were currently also in the U.S. military forces (a statistically significant 11 percentage point difference). Combined, a test of the joint significance of all of the demographic variables had a p-value of 0.05 (column (5)). These differences reflect the fact that conservative students in the sample were more likely to be male, to be international, and to be in the military.

The second panel shows the means, by treatment status, of questions we asked after the participants completed the first draft of their policy memo but pre-treatment—i.e., before they participated in their assigned group discussions. In this panel, there are no discernable

differences between the treatment and control groups (on average) in terms of prior familiarity with climate change or energy policy, confidence in their completed policy memo or in debating policy, thoughts on whether or not their classmates would agree with them, and the measures of introversion. While the signs of many of the differences suggest that the control groups had more prior relevant policy experience, and more confidence in their memo, the joint difference of all of these measures was not statistically different from zero ($p = 0.77$, as shown in column (5)).

Finally, the bottom panel of Table 1 shows the mean assessments of the pre-discussion memos by treatment status. The memos were graded blind for the number of arguments and counterarguments and for the quality of their arguments, writing, and the overall quality of the memo on a scale from 1–5. Again, none of the pre-treatment differences were statistically significant (as shown in column (4)), nor were they jointly significant ($p = 0.25$, column (5)). On four of the five measures, the direction of difference is such that participants assigned to an ideologically homogenous (control) group were graded higher on their policy memos: the number of arguments, the quality of the argument, the quality of the writing, and the quality of the overall memo. For this reason, we show robustness tests for our results that control for the quality of the pre-discussion memo when testing the treatment effect on the post-discussion memo. We also show our results with and without control for the demographic differences (male, international status, military identity) for all results.

Altogether, Table 1 demonstrates reasonable balance between students assigned to ideologically diverse and ideologically homogenous (control) groups. While we observe significant demographic differences amongst participants who identify as liberal and conservative in our study, these differences track with the demographic differences in political

affiliation in the United States (Doherty, Kiley, and O’Hea 2018; Foy and Restifo 2018), and participants were comparable in terms of their pre-treatment attitudes, perceptions, and writing accomplishments. That their pre-discussion memos were assessed comparably is particularly important, given that our key behavioral outcome measure is precisely grading on the same rubric for the same task, with the only difference being that the behavior followed an ideologically diverse or homogenous discussion.

Does ideological diversity influence the quality of student work?

Table 2 shows the basic results of the impact of assignment to an ideologically diverse group on the quality of the rewritten memo. The means by treatment and control status are shown in columns (1) and (2). Columns (3)-(5) show the point estimate of the difference between the two groups with robust standard errors clustered at the group level of the difference in parentheses below the point estimate and the p-value in brackets. The estimates in column (3) do not include any additional control variables (equation 2), those in column (4) include demographic controls (equation 3), and those in column (5) include demographic characteristics and measures of the quality of the pre-discussion memo (equation 4).

In general, we do not detect statistically significant differences in the number of arguments and counterarguments presented in the post-discussion memo by the treatment and control groups. By contrast, we find that those assigned to an ideologically diverse group wrote poorer quality memos, compared to those assigned to ideologically homogenous (liberal) groups, as reflected in the graded quality of their argument, of their writing, and of their overall memo. Several of the differences are significant at the 10% level, which is notable given the small sample size. Most notably, the difference in the overall quality of writing decreases by roughly a

quarter of a grade, or 0.29 on a five-point scale, when the student writes the memo after an ideologically more diverse group discussion ($p = 0.06$ when we control for demographics and the quality of the pre-discussion memo, see column (5)). Taken together, these results support the view that ideological diversity in the group has a null or a negative effect on subsequent student behavior. They do not support the hypothesis that exposure to views that differ from one's own improve critical thinking, at least as reflected in a short writing assignment administered directly after discussion.

Does ideological diversity influence the quality of student discussion?

While the evidence suggesting null or negative effects of ideological diversity on student output may be surprising from one perspective, we also know from prior literature that interactions in diverse groups can be stressful for participants. It is possible that these negative experiences are responsible for affecting behavioral outcomes like the quality of writing just after the group discussion. Specifically, an ideologically diverse discussion group could cause stress and discomfort among participants compared to being in a more homogenous group, which would in turn lead students in diverse settings to perform worse on their subsequent individual tasks. Alternatively, ideologically diverse discussions could cause *both* stress and lower performance. Our post-discussion survey results are displayed in Tables 3a–3d, detailing the treatment effect of ideologically diverse group composition on the dynamics of group discussion.

Table 3a reports the impact of assignment to an ideologically diverse group on the degree of agreement and debate in group discussion. Students assigned to an ideologically diverse group were more than half of a scale point less likely to report that the group agreed with one another; a difference that was statistically significant at the 5% level. They also reported that their group

was less likely to debate, although this finding was not conventionally statistically significant. Affirming the logic of the treatment assignment, the ideologically diverse group members reported they were more likely to hear ideas that differed from their own in their discussion (with a p-value of 0.12). These results suggest that the diverse groups did, indeed, differ in viewpoint and agreed less often, but that this disagreement was not accompanied by significantly more debate.

Table 3b shows the impact of assignment to a treatment group on questions that broadly reflect the respondent's frustration with the group discussion. These include questions regarding the extent to which participants characterized the group discussion as irritating or disrespectful, felt that their ideas were dismissed by their group or that their group did not truly listen, or that the discussion was stressful. In nearly all cases, those assigned to a treatment group reported greater frustration with the group discussion. They were more likely to report that their ideas had been dismissed by group members, that the group discussion was not respectful, that they felt others did not view their ideas as appropriate, and that others in the group did not truly listen to their ideas. Focusing on the results in column (4), four of the six differences are statistically significant at the 10% level.

Next, Table 3c reports the impact of ideological diversity on questions that reflect the extent to which participants refrained from participating in the discussion. Questions included whether the participant raised all points in their briefing document with the group, the extent to which the participant refrained from voicing ideas during discussion, and whether the discussion changed points raised in their post-discussion memo. While none of the differences were statistically significant at traditional levels, it is notable that in every case those assigned to a

treatment group reported more “distress” in the process. They were more likely to believe that others had refrained from voicing opinions, more likely to report they believed others had refrained from voicing their ideas, and were less confident in the recommendation in their redrafted briefing document. These findings also corroborate the directional finding from Table 3a that ideologically diverse groups debated less often.

Finally, the results in Table 3d reflect on the impact of ideology diversity on “enthusiasm” in group discussion. Again, we see that those in the diverse groups reflected less positively on the experience. Specifically, those assigned to a treatment group rated the quality of the discussion in their group over half of a scale point lower (with a p-value of 0.05), enjoyed the discussion less, did not believe that the discussion clarified their thinking, and were less likely to report that they would work with the group again.

Overall, Tables 3a–3d provide suggestive evidence that the process and experience of being assigned to an ideologically diverse group was stressful for students, and contain the individual items in each of the four indices. In results reported in Table 6, we estimate that these negative views were expressed by all students in the treatment groups – both those who identified more “liberal” and those who identified more “conservative.”

A key question, then, is the extent to which these negative feelings about the experience affected the estimated impact of assignment to an ideologically diverse group on the quality of student writing. We explore this question in Tables 4a–4c by controlling for student experiences in the impact regression. Given our small sample size, we were mindful of having too few degrees of freedom as we attempted to control for important factors.

Results in Table 4a explore how much the group dynamic impacted the treatment impact on the number of arguments and counterarguments presented by students in their post-discussion memos. Columns (1) and (6) replicate the impacts presented in Table 2, column (4). As such the regressions control for demographics as well as the quality of the pre-discussion memo. The subsequent columns also control for the indices described above. Controlling for measures reflecting the degree of agreement (columns 2 and 7) substantially reduces the size of the estimated impact of being assigned to the treatment group on the number of arguments and counterarguments. Similarly, the level of frustration (as reflected in index 2) appears to be negatively correlated with the number of arguments in the post-discussion memo and explains much of the treatment impact when the outcome is the number of arguments. In contrast, the level of frustration is basically unrelated to the number of counterarguments and does not appear to explain the impact of ideological diversity on the post-discussion memo. An interesting puzzle is that the level of enthusiasm is negatively related to the number of arguments in the post-discussion memo and is associated with a greater treatment impact.

The results in Tables 4b and 4c are largely consistent. The extent to which the groups agreed increased the quality of the post-discussion memos and mitigated the treatment impact somewhat. Similarly, as frustration and measures of “restraint” increased, the argument quality and overall quality of the post-discussion memo decreased; both measures also explain the estimated treatment impact as well. Finally, once again, we also see that as the level of measured enthusiasm increased, the quality measures decreased and the treatment impact was magnified. We are not certain how to interpret this finding regarding the measures of enthusiasm regarding the group discussion and its impact on the post-discussion memo, although we believe it is worth further investigation in a larger sample.

Discussion

The results in this pilot study must be interpreted as suggestive given the small sample size. Nevertheless, they can provide a basis upon which to build a larger study to better understand the implications of diversity in many different settings. Specifically, our results suggest that ideological diversity can negatively impact academic outcomes, as reflected in a short policy memo written soon after a group discussion. Counter to much of the theoretically hypothesized impact of ideological diversity on critical thinking, the results likely highlight some of the challenges of bringing diverse groups of individuals together. In our study, participants were more likely to perceive interpersonal conflict as their groups did not agree with one another as much, and they felt that their group members dismissed their ideas. As such, these findings continue to raise questions about the interpersonal mechanisms through which diversity aids or hinders group processes.

We note some unique features of the experiment that contribute to diversity research. First, students discussed their memo recommendations with their discussion group, but re-drafted their memos individually—an element of the design intended to match a policy education experience in which students arrive with their own policy recommendations, discuss or debate them in the classroom, and refine these recommendations during and after discussion. Second, we had a writing professor within the policy school assess participants' memos, providing an ecologically valid measure of critical thinking outcomes within this educational setting. Third, this is the first experiment, to our knowledge, to manipulate the political ideological composition of classroom discussions and examine effects on educational outcomes.

We propose several hypotheses to test in the future. Given sample size constraints, we were not able to manipulate instructions for group discussions. Previous research suggests that discussion structure, scaffolded by a leader's instructions, plays an important role in idea exchange and generation, and that this structure may mitigate some of the conflict that comes with diverse groups (Nemeth, Connell, Rogers, & Brown, 2001; Nemeth et al., 2004). The argument for ideological diversity supposes that conservative individuals will voice ideas less palatable to their liberal counterparts (and vice versa). However, without specific instructions to take turns in discussion, participants in our group discussions may not have heard the ideas of their conservative or liberal discussion partners. A 2 (ideologically homogeneous vs. heterogeneous) x 2 (structured vs. unstructured discussion) experimental design could explore whether discussion structure might attenuate some of the negative, stressful experiences of engaging with peers on a topic where they disagree.

For discussion, we selected a topic—climate policy—that did not directly break along race or gender lines in the way that other topics might, such as immigration or abortion. These dimensions of identity can inform participants' knowledge of related policy topics, shaping the balance of expertise within a discussion. These dimensions might also make debate more personal and stressful. Future research should test the effects of ideological diversity on student achievement on academic topics that vary in how directly related they are to other dimensions of diversity.

We also welcome future research to disentangle the effects of a discussion's process from the effects of individual participants' experiences of a discussion. Indeed, the function of a discussion and those who contribute to it, comprise a complex, dynamic unit of study. In light of

our findings, we look for further empirical inquiry to enrich our understanding of diversity's positive and negative effects and to sharpen our recommendations for education.

References

- (1) Hurd K, Plaut VC. Diversity entitlement: Does diversity-benefits ideology undermine inclusion?. *Nw UL Rev.* 2017;112:1605.
- (2) Frymer P, Skrentny JD. The rise of instrumental affirmative action: Law and the new significance of race in America. *Conn L Rev.* 2003;36:677.
- (3) Richard OC, Kirby SL. Organizational justice and the justification of work force diversity programs. *J Bus Psychol.* 1999;14(1):109-18.
- (4) Berrey E. The enigma of diversity: The language of race and the limits of racial justice. Chicago, IL: University of Chicago Press; 2015.
- (5) Wolfe B, Fletcher J. Estimating Benefits from University-Level Diversity. IRP Discussion Paper No. 1410-13. Institute for Research on Poverty. 2013 Feb.
- (6) Gurin, P. Selections from *The Compelling Need for Diversity in Higher Education*, expert reports in defense of the University of Michigan. *Equity & Excellence in Education.* 1999;32(2):36-62.
- (7) Gurin P, Peng T, Lopez G, Nagda BA. Context, identity, and intergroup relations. In: Prentice DA, Miller DT, eds. *Cultural divides: Understanding and overcoming group conflict.* New York: Russell Sage Foundation; 1999. p. 133-170.
- (8) Warikoo NK. The diversity bargain: And other dilemmas of race, admissions, and meritocracy at elite universities. Chicago, IL: University of Chicago Press; 2016.

- (9) Duarte JL, Crawford JT, Stern C, Haidt J, Jussim L, Tetlock PE. Political diversity will improve social psychological science 1. *Behav Brain Sci.* 2015;38:1-13.
- (10) Levendusky MS, Druckman JN, McLain A. How group discussions create strong attitudes and strong partisans. *Research and Politics.* 2016;3(2):1-6.
- (11) Farrar C, Green DP, Green JE, Nickerson DW, Shewfelt S. Does discussion group composition affect policy preferences ? Results from three randomized experiments. *Political Psychol.* 2009;30(4):615-647.
- (12) Robison J, Leeper TJ, Druckman JN. Do disagreeable political discussion networks undermine attitude strength? *Political Psychol.* 2018;39:479-494.
- (13) McLeod PL, Lobel SA, Cox Jr TH. Ethnic diversity and creativity in small groups. *Small Group Res.* 1996;27(2):248-64.
- (14) Cherry TL, Kroll S, Shogren JF. The impact of endowment heterogeneity and origin on public good contributions: evidence from the lab. *J Econ Behav Organ.* 2005;57(3):357-65.
<https://doi.org/10/dq6n6w>
- (15) Inzlicht M, Ben-Zeev T. A threatening intellectual environment: Why females are susceptible to experiencing problem-solving deficits in the presence of males. *Psychol Sci.* 2000;11(5):365-71.
- (16) Antonio AL, Chang MJ, Hakuta K, Kenny DA, Levin S, Milem JF. Effects of racial diversity on complex thinking in college students. *Psychol Sci.* 2004;15(8):507-10.
- (17) Sekaquaptewa D, Thompson M. The differential effects of solo status on members of high-and low-status groups. *Pers Soc Psychol Bull.* 2002;28(5):694-707.

- (18) Hansen Z, Owan H, Pan J. The impact of group diversity on performance and knowledge spillover—an experiment in a college classroom. *National Bureau of Economic Research*; 2006 May 22.
- (19) Van Dyne L, Saavedra R. A naturalistic minority influence experiment: Effects on divergent thinking, conflict and originality in work-groups. *Br J Soc Psychol*. 1996;35(1):151-67.
- (20) Torchia M, Calabrò A, Huse M. Women directors on corporate boards: From tokenism to critical mass. *J Bus Ethics*. 2011;102(2):299-317.
- (21) Sackett PR, DuBois CL, Noe AW. Tokenism in performance evaluation: The effects of work group representation on male-female and White-Black differences in performance ratings. *J Appl Psychol*. 1991;76(2):263-267.
- (22) Nemeth CJ, Personnaz B, Personnaz M, Goncalo JA. The liberating role of conflict in group creativity: A study in two countries. *Eur J Soc Psychol*. 2004;34(4):365-374.
- (23) Sunstein CR, Hastie R. Garbage in, garbage out? Some micro sources of macro errors. *J Institutional Econ*. 2015;11(3):561-83.
- (24) Wanous JP, Youtz MA. Solution diversity and the quality of groups decisions. *Acad Manag J*. 1986;29(1):149-159.
- (25) Milem JF. The educational benefits of diversity: Evidence from multiple sectors. *Compelling interest: Examining the evidence on racial dynamics in higher education*. 2003:126-69.

(26) Sacerdote B. Peer effects in education: How might they work, how big are they and how much do we know thus far?. In Handbook of the Economics of Education 2011 Jan 1 (Vol. 3, pp. 249-277). Elsevier.

(27) Hansen Z, Owan H, Pan J. The impact of group diversity on class performance: evidence from college classrooms. *Educ Econ*. 2015;23(2):238-58.

(28) Rouse CE. A Voice in the Conversation: Why Diversity Matters at Princeton | Princeton Alumni Weekly [Internet]. Princeton University. The Trustees of Princeton University; 2017 [cited 2019Aug9]. Available from: <https://paw.princeton.edu/article/voice-conversation-why-diversity-matters-princeton>

(29) Doherty C, Kiley J, O’Hea O. Wide gender gap, growing educational divide in voter’s party identification. Pew Research Center. 2018.

(31) Foy SL, Restifo SJ. Reliably Republican? Shifts in US Veterans’ Political Party Affiliation from 1974 to 2016. *Sociological Spectrum*. 2018 May 4;38(3):174-93.

(32) Goldberg LR, Johnson JA, Eber HW, Hogan R, Ashton MC, Cloninger CR, Gough HG. The international personality item pool and the future of public-domain personality measures. *J Res Pers*.. 2006;40(1):84-96.

(33) Nemeth CJ, Connell JB, Rogers JD, & Brown KS. Improving decision making by means of dissent. *J Appl Psychol*. 2001;31(1):48-58. doi:10.1111/j.1559-1816.2001.tb02481.x

(34) Cattell RB. The scree test for the number of factors. *Multivariate Behav Res*. 1966;1(2):245-76. <https://doi.org/10/fjbdsb>

- (35) DiStefano C, Zhu M, Mindrila D. Understanding and using factor scores: Considerations for the applied researcher. *Practical Assessment, Research & Evaluation*. 2009;14(20):1-11.
- (36) Thurstone LL. *The vectors of mind*. Chicago, IL: University of Chicago Press; 1935.

Appendix 1. Factor analysis.

Exploratory factor analysis is a common data reduction analysis that accounts for as much of the total variance as possible among the variables by using latent factors (DiStefano, Zhu, & Mîndrilă, 2009). We first conducted a scree test (available upon request), finding that either three or four factors was appropriate for our data (Cattell, 1966). Solutions with three and four indices were then examined using both orthogonal (varimax) and oblique (oblimin) rotations. We determined that four indices using oblimin rotation was the best fit for the data, because of the high face validity of the indices and because orthogonal factors are rare in social science research (DiStefano, Zhu, & Mîndrilă, 2009). The complete loadings of the factors with a cutoff of 0.3 are shown in Table 5. We did not exclude any survey items because each item had a loading of greater than 0.3, although some items showed evidence of being complex factors, likely due to our limited sample size.

The four indices are as follows. Index 1 is comprised of three items reported on a 7-point Likert scale with factor loadings from -0.42 to 1 and explained approximately 10% of variance. The items in Index 1 had positive loadings for group agreement and debate amount, and a negative loading for hearing different ideas. Because this index was largely driven by agreement we labeled this index “Agreement in Group Discussion”. The questions that comprise Agreement in Group Discussion are shown in Table 3a. Index 2 comprised six items reported on a 7-point Likert scale, and had factor loadings from -0.74 to 0.80 and explained approximately 16% of variance. Index 2 has high positive loadings for adverse discussion experiences like discussion stress, discussion irritation, and dismissing ideas, and high negative loadings for discussion respect, and feeling like your group truly listened. We therefore decided to label this index

“Frustration in Group Discussion”. The individual questions in the frustration in group discussion is available in Table 3b.

Index 3 comprises six items reported on a 7-point Likert scale, with factor loadings from -0.78 to 0.70, and explains approximately 13% of variance. Index 3 has high positive loadings for refraining voice during the discussion, and high negative loadings for questions such as raising all points, and confidence presenting ideas, hence we named Index 3 “Restraint in Group Discussion”. The individual items in the Restraint in Group Discussion index are shown in Table 3c. Lastly, Index 4 comprises five questions and explains approximately 15% of variance. The questions with high positive loadings in Index 4 were all related with positive aspects of the discussion such as a quality, enthusiasm, and enjoying the discussion. We call Index 4 “Enthusiasm in Group Discussion”. Information about the individual items in the Enthusiasm in Group Discussion index is located in Table 3d. Overall, the four different indices explain approximately 54% of variance. Note that the indices correspond to the items in Tables 3a-3d in order of variance explained, and that some items have been reverse loaded.¹ Index scores used in further analysis were predicted using regression (Thurstone, 1935).

S1 Appendix 2. Memo and survey materials.

Attached to submission as separate document.

¹ Specifically, “During the discussion did you hear ideas that differed from your own?” reverse loaded onto Index 1; “To what extent would you characterize the group discussion as respectful?” and “Did you feel that others in the group truly listened to your ideas?” reverse loaded onto Index 2; and “Did you raise all the points made in your briefing document with the group?”, “How confidently did you present your ideas to the group?”, and “How confident are you in the recommendation in your redrafted briefing document?” reverse loaded onto Index 3. No items reverse loaded onto Index 4. Full loadings of all indexes are shown in Table 5.