Fracture Without Failure: The Politics of Collective Action in Diverse Workplaces*

Kiara A. Hernández†

June 2026

Abstract

How do diverse societies achieve shared political goals despite intergroup divisions? Attempts to overcome diversity's challenges for collective action often focus on improving attitudes through contact, yet cooperation also depends on strategic relations. I examine this question using U.S. service-sector labor union organizing campaigns, a high-stakes political process in which interaction across diverse workers on shared shifts is routine and observable. Using a national survey of hourly workers and administrative campaign records, I find that diversity is associated with uncertainty and disagreement about union support that sustained contact does not resolve. Evidence from a pre-registered experiment shows that diversity causes more risk-averse strategic behavior, and that this relationship is not mediated by interpersonal attitudes. I conclude that contact targets interpersonal attitudes without sufficiently reducing perceptions of strategic risk. At the same time, diversity is not associated with closer union election margins, suggesting that institutional structures may help workers overcome coordination problems and inviting further study.

Word count: 11,828

*I thank Ryan Enos, Jennifer Hochschild, Taeku Lee, Marcel Roman, Danny Schneider, Alex Hertel-Fernandez, Daniel Galvin, Eric Blanc, Jeremiah Cha, Brian Highsmith, Ethan Miles, and participants in the Harvard Psychology, Behavior, and Geography Lab, Race and Ethnic Politics, and American Politics Research Workshops; the Consortium on the American Political Economy Summer Academy and Annual Conference; MPSA 2025; APSA 2025; and the Columbia-Cornell Political Economy of Work Junior Scholars Workshop for detailed comments. Special thanks go to Suzanne Clark, Mikhail Dei-Anang, Eric Salminen and other organizers with SEIU Minnesota for sharing data and insights. I am also grateful to Chas Walker and Dave Kamper for facilitating data sharing. Steffen Beverborg from SoPHIELabs provided exceptional programming help. This research was supported by a James M. and Cathleen D. Stone PhD Scholar fellowship from the Multidisciplinary Program in Inequality & Social Policy. The experimental portion of this project received exemption from the Harvard IRB (#25-0503) and was funded by research grants from the Stone Program in Wealth Distribution, Inequality, and Social Policy and the Center for American Political Studies at Harvard University.

†Postdoctoral Research Scholar, Center for Political Economy, Columbia University. kh3518@columbia.edu.

1 Introduction

How ethnoracially diverse societies achieve shared political goals despite durable intergroup divisions is among the most consequential questions of democratic governance. Where such divisions are entrenched, cooperation is complicated by lower levels of trust and weaker willingness to bear collective costs (Alesina et al., 1999; Luttmer, 2001; Putnam, 2007; Dinesen and Sønderskov, 2018). Attitudinal accounts identify contact as the solution, as sustained interaction across group lines builds trust and shared commitments (Allport, 1954; Mousa, 2020; Andersson and Dehdari, 2021). When individuals interact daily, are incentivized to cooperate, and face high common stakes, does diversity still complicate collective action – and if so, why?

I examine this question in the context of service-sector labor union organizing campaigns in the United States, where individuals experience routine intergroup contact and must decide whether to support collective action without knowing whether others will do the same. Forming a union is a democratic process in which a simple majority of workers vote to authorize the union as their formal representative in negotiations with their employer. Successful organizing significantly improves wages, benefits, and working conditions (Card, 1996; DiNardo and Lee, 2004; Farber et al., 2021; Frandsen, 2021; Knepper, 2020), but the decision to organize is itself costly. Workers face the risk of employer retaliation throughout the process, regardless of outcome (Ferguson, 2016). The result is a high-stakes coordination problem under the conditions contact theory identifies as most conducive to attenuating intergroup divisions.

Organizing campaigns offer an analytically advantageous setting due to (1) administrative organizing records that link each worker to their shift, occupation, and department, allowing diversity to be observed simultaneously within immediate contact groups and across the full bargaining unit, and (2) pre-election assessments of individual union support that generate observable measures of opinion in a shared organizational context. I combine records from forty-one campaigns covering 1,463 workers across 427 contact groups with a cross-sectional national survey of hourly workers ($n = 18,523$) and a randomized online experiment ($n = 647$). Together, these three studies allow me to establish the relationship between diversity and collective action, examine how it operates under conditions of direct contact, and isolate a causal explanatory mechanism.

Across data sources, I find consistent evidence of what I call *fracture without failure*: a pattern

in which diversity produces fracture in support for cooperation that manifests as uncertainty and disagreement rather than outright opposition, but does not foreclose collective action at scale. In Study 1, perceptions of workplace diversity are associated with an approximately 8 percentage point decrease in support for unionization, an association larger than those of race, gender, or supervisor coethnicity and robust to alternative specifications and sensitivity analyses for selection. In Study 2, diverse contact groups exhibit substantially greater disagreement in pre-election union support: high-diversity groups show 0.25 units higher disagreement than low-diversity groups on a 0-1 scale, with the share of undecided workers rising from 17 to 24 percent while anti-union sentiment remains stable at roughly 23 percent. Yet diverse bargaining units are not consistently associated with closer election margins. In Study 3, random assignment to heterogeneous teams causes participants to shift toward more risk-averse coalitions during deliberation, without reducing willingness to cooperate or generating more negative attitudes toward diverse teammates — identifying strategic uncertainty, rather than intergroup attitudes, as the operative mechanism. Convergent findings across three independent methods strengthen confidence that these patterns reflect a genuine feature of diverse coordination contexts.

I conclude that contact improves interpersonal attitudes without updating the strategic expectations that collective action requires. Trusting others does not automatically translate into confidence that the group as a whole will act together. That belief may be slower to update because the payoff of cooperating depends not on any single relationship but on whether enough individuals will act as well. Yet the diversity of bargaining units in this sample is not consistently associated with closer election margins. This is consistent with a countervailing role for union organizers' coordinating work, though the data cannot directly test this.

This article makes three contributions. First, it identifies the mechanism linking diversity to collective action difficulty, distinguishing strategic uncertainty about group coordination from the attitudinal barriers that contact addresses. Both operate simultaneously, suggesting that overcoming diversity's barriers requires interventions targeting coordination directly, not intergroup relations alone. Second, it advances the measurement of diversity in political behavior research by drawing on administrative organizing records that, to my knowledge, have not been previously analyzed by political scientists. Diversity can typically be observed and measured at aggregate levels, such as the county or the neighborhood (Alesina and La Ferrara, 2000; Putnam, 2007; Key and Heard,

1949; Hopkins, 2009; Abascal et al., 2023), and individual-level measures of contact typically rely on self-reported survey data. It is therefore comparatively difficult to establish who actually interacts with whom and how that shapes collective action. The administrative records used here allow diversity and contact to be observed directly, within immediate contact groups and across full bargaining units, demonstrating that coordination problems persist even under repeated face-to-face contact. Finally, the article situates union organizing within the study of democratic participation in political science. Organizing campaigns exemplify a broader challenge facing diverse democracies: how to achieve collectively beneficial political goals when shared interests must be built across group lines under conditions of risk and uncertainty.

2 Background and Theoretical Motivation

The American low-wage service sector, spanning healthcare, retail, food service, and other frontline industries, has become both the largest and most diverse segment of the labor market over the past half century. Fewer than one in five low-wage workers were people of color in 1971; more than half are today, with immigrants and ethnoracial minorities disproportionately concentrated in its lowest-paid positions (WorkRise Network, 2023; Bureau of Labor Statistics, 1980–2023). Workers in this sector interact daily across these lines, in the same departments, on the same shifts, facing shared employment conditions and high-stakes collective decisions together. Contact theory predicts that such sustained interaction across group lines should be sufficient to attenuate intergroup divisions. Organizing campaigns in this context therefore sit at the intersection of two trends: the diversification of frontline employment and the challenge of sustaining collective action across group lines, making them a direct test of when and why diversity complicates cooperation.

2.1 Diversity and Cooperation

A rich literature across the social sciences has examined why ethnoracial diversity complicates political cooperation. Explanations generally fall into two camps. The *strategic* account holds that diversity introduces coordination challenges: diverse groups may have competing preferences (Axelrod, 2006), individuals may expect co-ethnics to be more likely to reciprocate cooperative behavior (Deutsch, 1966; Miguel and Gugerty, 2005; Fearon and Laitin, 1996), and shared ethnic

identity may serve as a focal point that facilitates coordination among co-ethnics where none otherwise exists (Habyarimana et al., 2007). The *interpersonal* account holds that individuals have a general distaste for sharing resources with out-groups (Habyarimana et al., 2007; Tajfel, 1974; Becker, 1957), and that diversity may trigger threat responses that harden these tendencies (Key and Heard, 1949; Blalock, 1967; Sherif et al., 1961; Enos, 2017; Hopkins, 2009).

Empirical evidence of both mechanisms is extensive. Group heterogeneity has been associated with lower tax expenditures (Xin Li, 2010; Belmonte et al., 2018), weaker public goods provision (Alesina et al., 1999; Luttmer, 2001; Meer and Tolsma, 2014; Dahlberg and Lundqvist, 2012), declining community-level trust and social capital across national contexts (Alesina and La Ferrara, 2000; Putnam, 2007; Leigh, 2006; Laurence, 2011; Gundelach and Freitag, 2014; Dinesen and Sønderskov, 2018), and increased right-wing populism (Craig and Richeson, 2014; Sides et al., 2018).

These accounts make overlapping but distinct predictions, and diverge most sharply in settings where direct, repeated contact across group lines is present.

2.2 Two Theoretical Accounts of How Diversity Complicates Cooperation

The Interpersonal Account

The interpersonal theoretical account holds that diversity complicates cooperation because of how individuals regard and behave toward outgroup members. This account encompasses several distinct channels. The first operates through in-group favoritism. Social identity theory holds that individuals derive self-esteem from group membership and are motivated to favor in-group members as a result (Tajfel, 1974). In economic experiments, participants matched with in-group members show increased altruism, reduced envy, and greater willingness to choose welfare-maximizing actions — effects that appear to reflect genuine preferences rather than strategic calculations (Chen and Li, 2009). The second channel operates through expected reciprocity. Even absent affective preferences, individuals may expect co-ethnics to be more likely to cooperate, share norms of behavior, or reciprocate in future interactions (Deutsch, 1966; Habyarimana et al., 2007). The third operates through out-group antipathy: individuals may simply trust out-group members less or hold negative attitudes toward them that manifest in discriminatory behavior (Glaeser et al., 2000; Fershtman and Gneezy, 2001).

The contact hypothesis (Allport, 1954) offers the clearest response to these barriers. It holds that sustained, meaningful interaction across group lines builds trust, reduces prejudice, and fosters shared commitments, directly targeting the attitudinal channels through which diversity complicates cooperation. Field experimental evidence broadly supports this: sustained cooperative interaction can significantly reduce intergroup bias and improve behaviors toward specific interaction partners (Mousa, 2020; Scacco and Warren, 2018; Andersson and Dehdari, 2021). If the interpersonal account fully explained diversity's effects on collective action, sustained contact should be sufficient to resolve them.

The Strategic Account

The strategic account offers a different prediction. Rather than explaining diversity's effects through attitudes toward outgroup members, it emphasizes the uncertainty that diverse contexts generate about collective coordination. Even when individuals interact regularly and hold favorable views of their partners, diversity can generate uncertainty about whether the group as a whole will act together when it counts (Habyarimana et al., 2007; Fearon and Laitin, 1996).

Diverse groups may lack shared reference points or confidence that others interpret the situation similarly. Ethnic identity can function as a coordination device, providing focal points that allow co-ethnics to align expectations without explicit communication (Habyarimana et al., 2007; Schelling, 1960). When such shared identity is absent, individuals may struggle to converge on a common course of action even when they prefer it. A large experimental literature on coordination games demonstrates that doubt about which equilibrium others will select can lead to coordination failure even when multiple mutually beneficial outcomes exist (Van Huyck et al., 1990).

This uncertainty is particularly consequential in settings that require bearing collective risk. If individuals doubt that the group will coordinate successfully, they may adopt more cautious strategies — not because they distrust their partners, but because they cannot be confident about group cohesion. This account therefore requires only that diversity makes successful coordination feel less certain: the mechanism is strategic, not interpersonal, aligning with what Habyarimana et al. (2007) identify as the strategy selection mechanism rather than explanations rooted in differential preferences or sanctioning capacity. If this account is correct, diversity's effects should manifest as increased uncertainty and disagreement, not growing opposition.

2.3 Theoretical Expectations

Contact alone may still be insufficient for fostering cooperation. Alongside attitudes toward specific partners, individuals also need some sense of whether the group as a whole will act together, and that may depend on how they expect others to behave rather than on interpersonal relationships. Such expectations may be slower to update than attitudes, because the payoff of cooperating depends not on one's own preferences but on what one thinks others will do, and group membership functions as a durable cue for reciprocity expectations (Bernhard et al., 2006; Goette et al., 2006; Tanaka et al., 2010; Axelrod, 2006). Contact may improve attitudes toward specific individuals without updating those broader expectations about the group (Pettigrew and Tropp, 2006; Mousa, 2020; Scacco and Warren, 2018).

Study 1 examines whether perceived workplace diversity shapes baseline support for collective action before organizing begins. Both accounts predict lower support in more diverse settings, though through different channels — the interpersonal account through anticipated discomfort cooperating across group lines, the strategic account through anticipated difficulty coordinating.

H1: *Diversity is associated with lower baseline support for collective action.*

Study 1 cannot distinguish between theoretical channels or identify the role of contact. It captures perceived diversity before organizing begins, where both accounts predict the same direction of effect. Study 2 advances on this by situating diversity within a real contact environment: workers in the same department or shift interacting repeatedly across group lines and facing a concrete collective action decision. The interpersonal account predicts contact should resolve diversity's barriers in this setting, which may manifest as agreement or alignment of preferences. I thus expect that if coordination uncertainty persists regardless, it should manifest as disagreement in pre-election union support rather than uniform opposition, since strategic uncertainty does not predict active opposition but hesitance and ambivalence about whether others will act.

H2: *Diversity generates coordination uncertainty even under conditions of direct, repeated contact, manifesting as increased disagreement rather than increased opposition.*

Studies 1 and 2 cannot isolate the operative mechanism. Observational evidence of disagreement is consistent with both accounts and could reflect attitudinal barriers that contact has not fully

resolved, or strategic uncertainty that persists regardless of interpersonal attitudes. Study 3 uses random assignment to teams with varying group compositions to distinguish between these predictions. The interpersonal account predicts that diversity's effects should operate through attitudes toward outgroup members. The strategic account instead predicts that the effects should operate through the coordination process itself, independent of interpersonal attitudes.

> ***H3:*** *Diversity's effect on collective action operates through strategic coordination processes rather than through interpersonal attitudes.*

3 Study 1: Evidence from Survey Data

Diversity at the Pre-Organizing Stage

3.1 Sample

To study the relationship between perceptions of workplace ethnoracial composition and support for unionization, I use survey data collected from service sector workers in the United States. The data come from The Shift Project, a cross-sectional opt-in sample of 18,523 hourly workers employed at any one of the 170 largest retail and food service firms in the U.S., collected in spring 2023 (Schneider and Harknett, 2022).[1]

These worker-level micro-data are linked at the firm level, allowing me to construct workplace-level measures of ethnoracial diversity while also obtaining individual-level attitudinal responses. I perform multiple imputation using MIDAS (Multiple Imputation with Denoising Autoencoders), a deep learning imputation method (Lall and Robinson, 2022), for those who completed the survey but had item nonresponse.

3.2 Key variables

The outcome variable, support for unionization, is captured by responses to a question that asks individuals whether they would vote in favor of unionization in their workplace. The specific question wording is as follows: "If an election were held today to decide whether employees like you should be represented by a union, would you vote for the union or against the union?" I create a

[1] The Shift Project recruits respondents through paid, targeted social media advertisements directed to working-age adults employed at major retail and food service companies. Though a non-probability sample, it shows similar univariate distributions and bivariate relationships to the CPS and NLSY97 (Schneider and Harknett, 2022).

binary measure of union support where indication of no support or uncertainty ('Not sure') is coded as '0' for no support.

The explanatory variable in this study is perceptions of workplace-level ethnoracial diversity. Respondents were asked: "About what share of your workplace is [White / Hispanic or Latino / Black or African American / Asian American or Pacific Islander]?" Using these responses, I quantify workplace diversity using the Blau (1977) index, calculated as:

$$EthnoracialDiversity = 1 - \sum_{j=1}^{J} p^2 j$$

where p is the proportion of j ethnoracial group in the workplace. A value of 0 signifies no diversity, as all individuals in the workplace are of the same ethnorace, and 1 signifies maximum diversity, as all individuals are of different ethnoraces.

The data identify both the firm where each respondent works and the ZIP code of their workplace. I construct a unique Blau index for each respondent based on their own estimates of workplace composition. This means even workers at the same firm and location may have different index values if they perceive the composition differently. But individual variation in these assessments is a feature of the data rather than a limitation: workers' perceptions of their immediate work environment likely reflect the kind of shift-level diversity that contact theory identifies as consequential for cooperation.

A more relevant concern is non-random selection into more or less diverse workplaces based on prior attitudinal predispositions. Because individuals are not randomly assigned to establishments with varying levels of ethnoracial diversity, unobserved factors may confound the relationship between workplace diversity and support for unionization. To mitigate these concerns, I include rich individual- and workplace-level controls alongside state- and county-level fixed effects, and cluster standard errors at the firm level. I also conduct robustness and sensitivity analyses to assess the extent to which unobserved factors could account for the results. The paper further leverages evidence across two additional complementary studies to build a more robust account of how diversity shapes cooperation.

3.3 Estimation Strategy

I estimate the relationship between firm-level diversity and support for unionization using ordinary least squares. This is modeled in Equation 1 below:

$$Unionsupport_i = \beta_0 + \beta_1 PerceivedDiversity_{i,w} + Z_i + v_{s,w} + \epsilon_i$$

Union support is the dichotomous measure described earlier, perceived diversity is the Blau index score assigned to an individual's workplace, Z_i represents a series of individual characteristics (job quality, ethnorace, age, gender, supervisor coethnicity, coworker support, and county-level Democratic vote share), and $v_{s,w}$ captures state fixed effects. The parameter of interest is β_1, which represents the association between an increase in perceived workplace diversity and support for unionization.

I control for ethnorace, age, gender, hourly wages (a measure of job quality), supervisor coethnicity, coworker support, and county-level Democratic vote share. Supervisor coethnicity captures the possibility that racial dynamics in supervisory relationships dampen support for organizing. Coworker support is included because solidarity with peers may raise or lower demand for unionization depending on whether workers feel workplace conflicts are already managed. County-level Democratic vote share proxies for individual-level partisanship, which is unavailable in these data.[2]

Finally, all models include state- or county-level fixed effects ($v_{s,w}$) and standard errors clustered at the firm level. Descriptive statistics for all variables are presented in Table A1.

The sample is majority non-Hispanic White, female, and skews young (Table A1). Most respondents report working for a supervisor of the same ethnorace as themselves. The average Blau index score is 0.507 (median: 0.565), reflecting workplaces perceived as split between two ethnoracial groups rather than broadly diverse; few respondents report working in either completely homogeneous or maximally heterogeneous settings. The sample spans all 50 states, with roughly equal representation across U.S. regions and a median county-level Democratic vote share of 0.49,

[2]County-level returns are drawn from the MIT Election Data and Science Lab (MIT Election Data and Science Lab, 2021). A small number of counties ($n = 381$, $\approx$2%) could not be matched due to non-standard FIPS coding in Connecticut and incomplete precinct reporting in Indiana; these observations are excluded from models including this control.

indicating it is not concentrated in Democratic strongholds.

3.4 Results

Table 1 presents the regression results from the OLS model. The findings are consistent with H1: working in the most diverse workplaces compared to the least diverse is associated with an approximately 8 percentage point decrease in support for unionization, robust to the inclusion of demographic controls, job quality, coworker support, and county-level partisanship. The magnitude of this association is larger than that of any individual-level demographic predictor in the model. As shown in Figure A1, where I plot coefficients and confidence intervals for all relevant variables rescaled to 0–1, the difference in support between maximum and minimum diversity is greater than the differences associated with race (White vs. non-White), gender (men vs. women), or supervisor coethnicity. Only the age gap between the youngest and oldest workers and county-level partisanship have larger magnitudes, consistent with recent evidence documenting a generational divide in union support (Ahlquist et al., 2024). The wide confidence intervals on hourly wages make it difficult to determine whether pay levels meaningfully influence support.

Beyond diversity, several substantively interesting associations emerge. White workers are modestly less supportive of unionization relative to non-White workers. Men are less supportive than women, and workers with coethnic supervisors also express lower support for organizing. Age is negatively associated with union support. County-level partisanship is strongly positively associated with union support ($\beta = 0.135$, $p < 0.001$), consistent with the expectation that workers in more Democratic areas are more favorable toward unionization. Job quality, as measured by hourly wages, is only marginally associated with support. The positive association between coworker support and union support is consistent with the cooperation framework: when workers regard their colleagues favorably, collective action becomes more attractive.

3.4.1 Demographic Compositional Effects

Using the Blau index to capture diversity could mask effects driven by the predominance of any single ethnoracial group. Probing composition directly — replacing Blau with perceived proportions non-White, Black, Latinx, and AAPI — reveals no association between union support and any individual group, with the exception of AAPI composition, where a negative association persists

Table 1: Association between Perceptions of Firm-level Diversity and Union Support

	(1)	(2)	(3)
Blau index	−0.082**	−0.081**	−0.101***
	(0.031)	(0.027)	(0.030)
White		−0.049***	−0.049***
		(0.012)	(0.013)
Male		−0.079***	−0.084***
		(0.009)	(0.010)
Hourly wage		−0.002*	−0.002*
		(0.001)	(0.001)
Age		−0.023***	−0.024***
		(0.003)	(0.003)
Supervisor coethnicity		−0.027**	−0.028**
		(0.010)	(0.010)
Coworker support		0.057***	0.058***
		(0.005)	(0.005)
Dem. vote share (county)		0.135***	
		(0.026)	
Fixed effects	State	State	County
Observations	18 523	18 142	18 099
Adjusted R^2	0.005	0.029	0.025
DepVar Mean	0.341	0.341	0.341
DepVar SD	0.474	0.474	0.474

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: Standard errors clustered by firm. Column (1) reports the bivariate association between the Blau index and union support with state fixed effects. Column (2) adds individual-level controls and county-level Democratic vote share as a proxy for respondent partisanship. Column (3) replaces state with county fixed effects, which absorb partisan composition directly; `dem_voteshare` is therefore collinear and excluded. County fixed effects are estimated across 1,951 counties; counties with only one respondent are dropped by the estimator.

but interpretation warrants caution given small cell sizes (Figure A2). Race interaction models find no evidence that the diversity association differs between White and non-White workers. Full compositional and interaction results are reported in the appendix (Table A9). These patterns reinforce the main finding: diversity, not the presence of any specific group, is associated with lower union support.

3.4.2 Robustness Checks

The main result survives several robustness checks reported in the appendix. Dropping observations in the top and bottom 5% of the diversity distribution leaves all coefficients substantively unchanged (Table A6). Binning the Blau index into median splits, terciles, and quartiles confirms the relationship is monotonic and not driven by functional form (Figures A3–A7). Replacing state fixed effects with county fixed effects, which absorb local partisan composition, labor market conditions, and cost of living simultaneously, leaves the diversity coefficient stable and somewhat larger ($\beta = -0.101$, $p < 0.001$; Table A7). A formal sensitivity analysis following Cinelli and Hazlett (2020) indicates that a confounder would need to be four times stronger than county-level partisanship to render the estimate marginally statistically non-significant (Table A8).

3.5 Future Cooperation and the Importance of Context

The results from Study 1 underscore the importance of *context* in shaping organizing potential: perceptions of workplace diversity are associated with lower support for unionization at baseline, before organizing has begun. While the substantive magnitude of this relationship may appear modest, in an environment where union density is already low and successful organizing faces significant structural hurdles, even small negative associations matter. These findings are consistent with prior work showing that cooperative dilemmas are heightened in the presence of ethnoracial difference, and they point to the challenges of building solidarity in diverse workplaces where intergroup conflict and other tensions may be more pronounced.

At the same time, these data capture only a snapshot of "pre-treatment" attitudes, prior to any organizing efforts. They do not show how support evolves once workers enter the campaign process, when cooperation is necessary and organizing activity can help workers overcome coordination problems. Nor do these data show whether workers are coming into direct contact with diverse

coworkers on a regular basis. Study 2 therefore turns to these dynamics, examining how diversity affects the formation of attitudes towards the union during the organizing process and behavior in the ultimate election.

4 Study 2: Evidence from Administrative Data

Diversity at the Pre-Election and Election Stage

To understand how diversity shapes cooperation as it actually unfolds, I draw on administrative and behavioral data from forty-one union drives in two large healthcare-sector locals — a modest yet representative subset of low-wage service sector workers. These records — organizing charts, NLRB election files, and Excelsior lists of eligible voters — capture real campaign dynamics rather than self-reported attitudes, and document the organizing process from pre-election assessments of individual support through final election outcomes. Where Study 1 established an association between perceived diversity and baseline union support, Study 2 situates diversity within a real contact environment and introduces observed behavioral outcomes, allowing me to examine how diversity operates at both the local level of contact groups and the unit-wide level of bargaining units.

In this section, I first describe the data, key variables, and empirical strategy for parts 1 and 2 of the analysis. In Part 1, I analyze how contact-group diversity is associated with disagreement in pre-election assessments of support. In Part 2, I then turn to the election stage to consider how unit-level diversity shapes final voteshare margins.

4.1 Data Overview

4.1.1 Excelsior Lists

To reconstruct the demographic composition of the bargaining unit voting in a union election, I utilize Excelsior lists[3] provided by union organizers at two large healthcare sector union locals in Minnesota and the Southern New England region. These "Excelsior lists" contain the full names and residential addresses of all eligible voters in the proposed bargaining unit and are given to the union and the NLRB by the employer after the election date is set.

[3]This naming convention was established after the case *Excelsior Underwear*, 156 NLRB 1236 (1966).

These lists cover 41 unique bargaining units where a unionization vote was held from 1999–2025, representing 2,701 unique voters. The modal occupation across units is a Certified Nursing Assistant (CNA). Relevant sample characteristics appear in the "Full Sample" columns in Table 2; the distribution of election years and unit sizes appears in Appendix B.4.

4.1.2 Organizing Charts

Of the 41 Excelsior lists, 17 cases representing 1,463 voters also include organizing charts used during the campaign. These charts record each voter's shift, department, job title, and a pre-election assessment of union support on a 1–5 scale: 1 indicated membership on the organizing committee, 2 indicated support, 3 was undecided, 4 was anti-union, and 5 indicated actively organizing against the union.[4] This subsample appears representative of the full sample on nearly all relevant variables (Table 2, "Contact Group Sample" columns).

Table 2: Summary Statistics of Key Variables by Sample

Variable	**Full Sample**			**Contact Group Sample**		
	Min	Median	Max	Min	Median	Max
Year	1999	2020	2025	1999	2011	2025
Union Vote Margin	−0.327	0.571	1.000	0.111	0.560	1.000
Bargaining Unit Racial Diversity (Blau Index)	0.393	0.687	0.733	0.553	0.677	0.718
Prop. of Non-White Workers (Unit Level)	0.237	0.620	0.881	0.369	0.609	0.777
Number of Eligible Voters	5	60	234	18	66	234
Union Election Turnout Rate	0.389	0.829	0.964	0.389	0.858	0.947
Prop. White	0.119	0.380	0.763	0.223	0.391	0.631
Prop. Black	0.040	0.272	0.545	0.128	0.276	0.545
Prop. Hispanic	0.034	0.114	0.417	0.059	0.094	0.417
Prop. AAPI	0.055	0.170	0.579	0.103	0.160	0.286
Number of Distinct Occupations	1	3	27	1	9	27
Number of Distinct Shifts	1	5	18	1	6	18
Number of Occupation × Shift Combinations	1	7	59	4	18	59
Mean Contact Group Diversity (Blau Index)	0.396	0.543	0.726	0.432	0.515	0.673
Max Contact Group Diversity (Blau Index)	0.556	0.710	0.744	0.583	0.710	0.744

[4]Some charts also included wages, FTE status, and date of hire, but this information was not consistently present across the subsample.

4.1.3 Administrative Union Election Data

Each union election is assigned a unique case number by the NLRB. Using this case number, I collected administrative data on the number of eligible voters, the votes for the union, the votes against, the total number of ballots cast, and the number of contested votes from the NLRB government website. I merged this with the Excelsior list/organizing chart records to calculate the election turnout rate and voteshare margin.

4.1.4 Voter File Matching and Partisan Identification

To control for the partisan composition of bargaining units and contact groups, I impute individual-level partisan identification using a Bayesian procedure combining a demographic prior from the 2020 Cooperative Election Study with a geographic prior from 2020 presidential returns, following Brown and Enos (2021). For workers in Connecticut, Minnesota, and Massachusetts, imputed probabilities are validated against voter file matches: among matched workers, mean $\hat{P}$(Democrat) is highest among registered Democrats, intermediate among Unaffiliated voters, and lowest among registered Republicans. Full details are in Appendix B.3.

4.2 Variables

IV: Ethnoracial Diversity

The primary independent variable is observed ethnoracial diversity. In Part 2 of the analysis, I measure this at the bargaining unit level. In Part 1 of the analysis, I measure this at the contact group (department $\times$ occupation $\times$ shift) level.

To construct the measure of ethnoracial diversity, I impute individual race probabilities for each individual in the sample using a combination of surname, first name, and residential ZIP code,[5] grouping either by employer or contact group within-employer. I then use the Blau (1977) index, as described above in Study 1, to quantify the level of diversity within each unit (contact group).

[5]The details of this process are described in Appendix B.1.

Analysis Part 1 DV: Disagreement

In Part 1 of the analysis, I restrict to the contact group sample, where I observe individual-level organizer assessments of union support. The outcome of interest is disagreement or what I term opinion fracture.

I use organizer assessments to define union opinion. These assessments are a subjective but detailed record of individual-level support made by organizers in conjunction with the organizing committee, updated continuously through election day.[6] As one organizer described, *"even a sheepish 'yes' or a 'Oh yeah, that sounds great, I like unions"'* would receive a "3" — meaning assessments capture sentiment and tone, not just affirmative behavior. While organizer subjectivity introduces potential error, the assessments were repeatedly updated based on sustained interactions and corroborated by committee members, making them a uniquely granular record of pre-election preferences.[7]

For cases with no assessment,[8] I use union authorization card or petition signing as a proxy.[9]

To measure disagreement within contact groups, I calculate normalized Shannon entropy (Shannon, 1948) over organizer assessments, which captures the dispersion of opinions across the five assessment categories. Entropy equals zero when all members share the same assessment and one when opinions are spread evenly across all categories; I normalize by $\log(K)$ to facilitate comparison across groups of different sizes. A normalized Blau index yields near-identical results (Pearson $r = 0.99$, Spearman $\rho = 1.00$). Additional details and a working example are in Appendix B.15.

Analysis Part 2 DV: Voteshare Margin

The voteshare margin is the difference between the share of votes for and against the union. I use this rather than a binary win/loss outcome because 38 of 41 units won their elections, leaving insufficient variation for a dichotomous dependent variable; the margin captures how competitive

[6] Union organizers routinely track and update worker support in this way as part of the organizing process.

[7] These data allow me to observe both the independent variable (patterns of contact across diverse coworkers) and the dependent variable (union opinions) in the same workplace context.

[8] Assessments were missing for 315 of 1,463 voters.

[9] This yielded 210 pro-organizing leaders (1), 538 pro-organizing (2), 126 undecided (3), 257 anti-organizing (4), and 2 anti-organizing leaders (5).

each election was regardless of outcome.

4.3 Empirical Strategy

To assess the relationship between diversity and disagreement (Part 1) and between diversity and election voteshare margins (Part 2), I estimate weighted and unweighted linear regression models of the following form:

$$Y_g = \beta \, Diversity_{ij} + \gamma X_{ij} + \epsilon_g,$$

where Y_g is the outcome of interest, $Diversity_g$ is the Blau index of ethnoracial diversity measured at the relevant level, and X_g is a vector of covariates.

For Part 1, the unit of analysis is the contact group. Again, I define a contact group as the set of workers that share the same shift, occupation, and department within a particular employer. The dependent variable is disagreement, measured as normalized Shannon entropy based on organizer assessments of worker support. Diversity is measured at the contact-group level.

I include controls for group size, whether the group has a high proportion of non-White workers, and whether it is above the sample median in the proportion of Black workers. Standard errors are clustered at the employer level to account for correlated observations across contact groups within the same organizing campaign. As a robustness check, I also control for the mean posterior probability of Democratic identification within the contact group and the share of workers matched to the voter file.

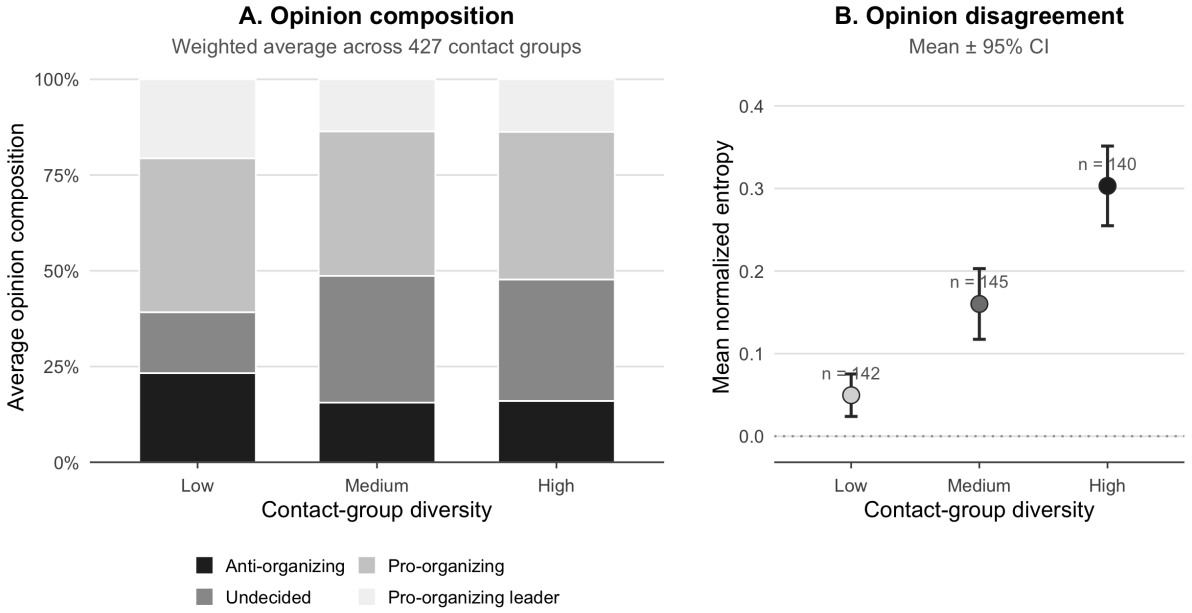
For Part 2, the unit of analysis is the bargaining unit. The dependent variable is the voteshare margin, calculated as the difference between the proportion of votes for and against the union. Diversity is measured at the bargaining-unit level. I control for unit size, whether the unit has a high proportion of non-White workers, whether it is above the sample median in the proportion of Black workers, and indicators for whether the unit contains multiple occupations and multiple shifts. Again, I additionally include controls for the mean posterior probability of Democratic identification within the bargaining unit and the share of workers matched to the voter file as robustness checks.

4.4 Results, Part 1: Local Diversity Produces Uncertainty, Not Opposition

I begin at the level of contact groups, the everyday settings where workers interact directly on shared shifts, occupations, and departments within the same workplace. These analyses draw on the subset of seventeen elections with detailed organizing charts that include measures of individual-level union support, covering 1,463 workers across 427 contact groups. Because these data capture individual-level assessments, they avoid the ecological fallacy that limits bargaining-unit analyses. That is, instead of inferring attitudes from aggregate voting outcomes, I can observe how workers within the same immediate context diverge or converge in their pre-election opinions.

The results show that greater ethnoracial diversity at the contact-group level is substantively and significantly associated with greater disagreement at that same level (Appendix B.7). The relationship holds when the Blau index is included in the model as a continuous predictor and when binned into terciles (Low, Medium, High), though the relationship appears to be driven primarily by the groups with the highest levels of diversity. Specifically, moving from complete homogeneity to maximum diversity is associated with a 0.52-point increase in disagreement on a 0-1 scale (Appendix B.7), a substantial change.

Low-diversity groups show near-unanimity in opinion (normalized entropy = 0.050), and the relationship increases monotonically with diversity. Being in a medium-diverse group relative to a low-diverse group is associated with 0.11 unit higher entropy, while being in a high-diversity group is associated with 0.25 unit higher entropy, a more substantial difference in fracture. These models include standard errors clustered at the employer level, and the estimates hold when including controls for the share of the contact group that is non-White and Black (Appendix B.8). Substantively, this suggests that workers in more diverse contact groups are more likely to hold divergent views on unionization, raising barriers to coordination even when the modal stance remains pro-union.



Notes: Left panel shows weighted average opinion composition by contact-group racial diversity tercile (Low/Medium/High). Darker shading indicates anti-unio

Figure 1: Opinion composition and disagreement by contact-group diversity level. Left panel shows weighted average opinion composition by diversity tercile across all 427 contact groups; darker shading indicates anti-union sentiment, lighter shading indicates pro-union sentiment. Right panel shows mean normalized Shannon entropy within each tercile (± 95% CI). Contact groups defined as workers sharing the same employer, occupation, shift, and department. CIs in right panel derived from the bivariate association (Column 1) in Appendix Table B17.

Figure 1 illustrates this pattern directly, showing how opinion composition fragments and entropy rises across diversity terciles.

Critically, however, this increase in disagreement does not reflect a decline in average pro-union sentiment. When the outcome is the share of workers in each contact group assessed as pro-union, rather than the dispersion of those assessments, diversity has no significant relationship: higher- and lower-diversity contact groups do not differ meaningfully in the proportion of workers who support the union (Appendix B.8).

This distinction matters for how we understand the potential mechanism. *Diversity does not appear to suppress support for collective action within contact groups; rather, it introduces uncertainty.* Anti-union sentiment does not vary significantly across diversity levels (23 percent in low-diversity groups, 23 percent in medium-diversity groups, and 22 percent in high-diversity groups) while the share of workers who are undecided increases from 17 percent in low-diversity groups to roughly 24 percent in medium- and high-diversity groups. The entropy increase thus reflects a shift away from strong pro-union commitment toward ambivalence, not toward opposition. Workers in more racially

mixed groups are no less likely to ultimately support the union, on average, but they are more likely to be uncertain about it. It is this dispersion, driven by hesitation rather than outright opposition, that creates barriers for coordination. Results are robust to controlling for partisan composition within contact groups (Figure 2).

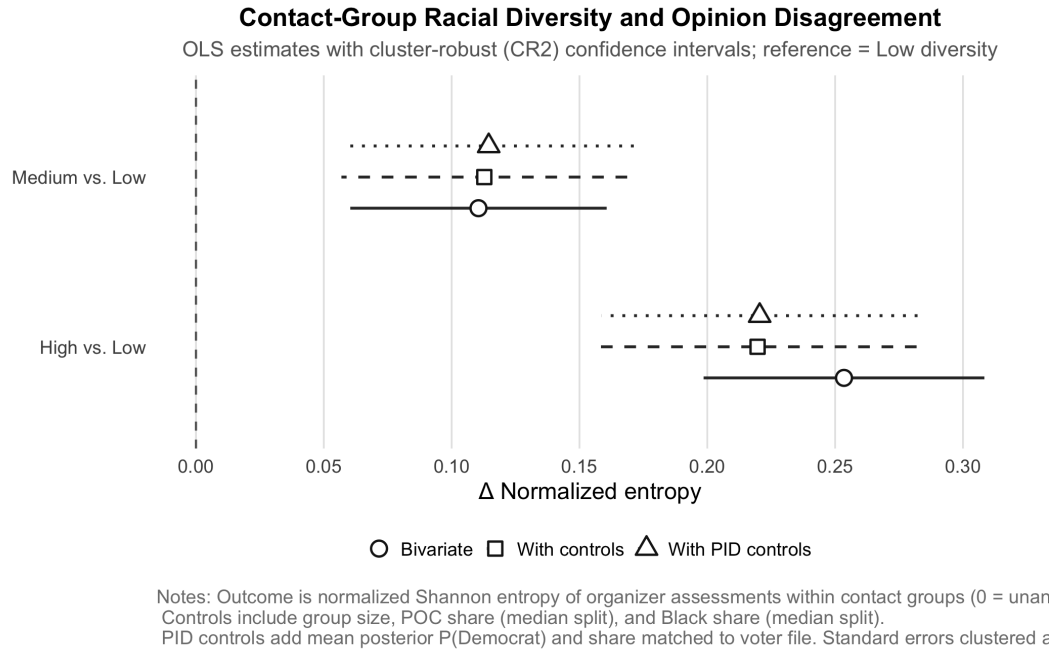


Figure 2: Contact-group racial diversity and opinion disagreement: OLS coefficient estimates with cluster-robust (CR2) confidence intervals. Reference category is Low diversity. Controls include group size, non-White share (median split), and Black share (median split). PID controls additionally include mean posterior $\hat{P}$(Democrat) and the share of workers matched to the voter file. Standard errors clustered at the employer level. $N = 427$ contact groups, 17 employers. Full regression results in Appendix Table B17.

Taken together, the evidence from Part 1 suggests that local diversity does not make workers less supportive of the union, but it does make them more divided. Rather than shifting workers toward opposition, diversity increases the share who are uncertain or ambivalent about organizing. In this sense, diversity produces a kind of opinion fracture. Specifically, it produces a widening of the distribution of views within contact groups, without generating the outright opposition that might foreclose collective action entirely. Next, I examine whether and how this fracture shapes downstream election outcomes.

4.5 Results, Part 2: Composition, Not Diversity, Shapes Election Margins

Building on the evidence from Part 1 that diversity produces greater opinion fracture within contact groups, I now turn to the election stage to ask whether these fractures aggregate into weaker outcomes at the level of full bargaining units.

The relationship between unit-level racial diversity and the voteshare margin is inconsistent across specifications. The bivariate association between the continuous Blau index and the vote margin is negative and marginally significant in unweighted models ($\hat{\beta} = -1.34$, $p = 0.051$) but attenuates and becomes indistinguishable from zero once models are weighted by unit size ($\hat{\beta} = -0.98$, $p = 0.454$). Models substituting the Blau tercile bins for the continuous index produce null results across all specifications, weighted and unweighted. Given the modest sample of 41 bargaining units, these estimates carry substantial uncertainty and should be interpreted as suggestive at best.

What more consistently predicts election margins is the *racial composition* of the bargaining unit rather than its overall diversity. Units where more than 60 percent of workers are non-White ('high-POC" units) are associated with significantly larger vote share margins ($\hat{\beta} = +0.36$, $p = 0.019$). This association is robust across all fully specified models and attenuates only modestly when partisan controls are added ($\hat{\beta} = +0.25$, $p = 0.073$), suggesting it is not simply an artifact of the heavily Democratic composition of these workplaces. Similarly, a higher proportion of Black workers in the unit is associated with larger vote share margins in bivariate models ($\hat{\beta} = +1.08$, $p = 0.012$), though this does not survive inclusion of other controls.

This pattern is consistent with the strategic account developed in the theoretical motivation section. When a clear majority group is present, shared ethnic identity can serve as a coordination focal point, potentially allowing workers to align expectations (Habyarimana et al., 2007). High-POC units may benefit from precisely this dynamic, because even if the unit is not homogeneous, the presence of a dominant group provides a basis for coordination that reduces strategic uncertainty. Diverse units lacking such a majority must coordinate across groups without a shared focal point, generating the fracture documented in Part 1. That this fracture does not translate into aggregate failure is consistent with prior research finding no systematic relationship between diversity and union election victory (Ferguson, 2016). The present analysis extends that finding to margins, showing that diverse units that reach the election stage are just as likely to win *and* are not

consistently disadvantaged in how competitive those elections are. Together, these patterns suggest that diversity's costs fall disproportionately on units without a dominant group — where no shared focal point exists to anchor coordination. Units with multiple shifts also win by narrower margins ($\hat{\beta} = -0.27$, $p = 0.068$), consistent with the contact-group evidence that structural fragmentation across interaction contexts limits organizing reach. Table 3 presents the full results.

Table 3: Bargaining-Unit Racial Composition and Election Margins

Variable	**(1)**	**(2)**	**(3)**	**(4)**
Blau Index (continuous)	−0.976			
	(1.290)			
Blau: Middle tercile		0.132	0.088	−0.113
		(0.161)	(0.175)	(0.210)
Blau: High tercile		−0.065	−0.092	−0.042
		(0.140)	(0.253)	(0.194)
High POC Share			0.358**	0.251*
			(0.145)	(0.135)
High Black Share (median)			0.022	−0.096
			(0.139)	(0.140)
Log(Unit Size)			−0.129	−0.167*
			(0.085)	(0.082)
Occupation Dummy			−0.109	−0.113
			(0.139)	(0.158)
Shift Dummy			−0.269*	−0.252
			(0.142)	(0.161)
Mean $\hat{P}$(Democrat)				2.155
				(1.344)
Pct. Matched to Voter File				−0.369
				(0.409)
Observations	41	41	41	41
Adj. R^2	−0.011	0.002	0.232	0.271

Notes: Dependent variable is union voteshare margin. All models weighted by number of eligible voters. Robust (HC2) standard errors in parentheses. Model (1) includes the Blau index as a continuous predictor only. Model (2) includes Blau terciles only. Model (3) adds racial composition and structural controls. Model (4) adds mean posterior $\hat{P}$(Democrat) and the share of workers matched to the voter file as partisan controls. Reference category for Blau terciles is Low. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.6 Diversity and the Persistence of Friction

Together, Studies 1 and 2 complicate a straightforward account in which diversity harms collective action. Study 1 showed that perceived diversity depresses baseline support for unionization. Study

2 extends this by showing that local diversity in the microcontexts of daily interaction produces greater opinion disagreement during the pre-election stage. Yet the nature of that friction matters: diversity does not suppress pro-union sentiment. Anti-union sentiment is stable across diversity levels; what increases is the share of workers who are uncertain or ambivalent. The friction operates through opinion fracture, not suppression.

At the unit level, this distinction helps explain why diversity does not consistently predict weaker electoral outcomes. Opinion fracture and strong average pro-union sentiment can coexist, and the latter may partially offset the coordination costs introduced by the former. Racial composition and structural fragmentation across shifts predict closer elections more reliably than diversity per se.

5 Study 3: Evidence from a Bargaining Experiment

Two limitations constrain causal inference in Studies 1 and 2. First, we cannot observe how individuals sort into contact groups. Workers may select into particular shifts or departments on the basis of preexisting racial attitudes or attitudes toward the union, or managers may allocate workers in ways correlated with those attitudes. If such sorting occurs, observed differences in disagreement could partially reflect attitudinal predispositions rather than the causal effect of diversity. Second, the analysis assumes relatively uniform interaction within contact groups, yet we cannot observe the true frequency, intensity, or nature of contact. Differences in the volume and type of interaction across groups could attenuate or amplify the observed relationship between diversity and disagreement. The observational data also cannot distinguish between two accounts of why diversity produces fracture: an interpersonal account (workers distrust or disfavor out-group coworkers) and a strategic account (diversity generates uncertainty about whether the group will coordinate). These accounts generate different predictions about when diversity's effects should appear and whether they should be mediated by interpersonal attitudes.

To isolate the causal effect of group composition and adjudicate between these mechanisms, I ran an online cooperative bargaining experiment (n = 647) in which participants were randomly assigned to ethnoracially heterogeneous or homogeneous teams and asked to coordinate a wage proposal. The experiment introduced within-team disagreement about strategy, requiring each participant to form a coalition with one of two teammates who held conflicting positions (one

risk-seeking and one risk-averse) while team composition varied independently of which teammate held which position. Teammates were AI confederates powered by a large language model fine-tuned on GPT-3.5, scripted to deliver consistent position-taking across conditions while simulating natural conversational interaction. This design allows me to test directly whether diversity shifts individuals toward caution, and whether this shift reflects co-ethnic favoritism.

5.1 Recruitment and Sample

Participants were recruited via Prolific in July–August 2025.[10] Eligibility was restricted to working-aged, non-Hispanic white U.S. adults employed (or job-seeking) in food service, healthcare, retail, or hospitality who reported regularly interacting with others at work. Participants were compensated at a fair market rate; average completion time was 18 minutes.[11] The experiment was programmed and hosted on SoPHIELabs, an interactive online experiment platform integrated with Prolific.[12]

5.2 Experimental Design

Participants played a modified public goods game for five successive yet distinct rounds. The design was within-subject: each participant received both the treatment (ethnoracially heterogeneous team) and the control (homogeneous team) in randomized order. The results reported here focus on both a between-subjects comparison at Round 3 as well as a within-subjects comparison across the full experiment.

Before gameplay, participants completed a basic demographic questionnaire and selected an avatar and name to represent them throughout the game (see Appendix C.5). This was included to simulate real-world interaction, as participants believed they were participating in a real-time game with other recruited participants.[13]

The experiment was designed to distinguish between the two accounts of diversity outlined above. To briefly reiterate: the interpersonal account holds that diversity matters because of who

[10] A pilot was conducted in February 2025 via the Harvard Digital Lab for Social Sciences (Strange et al., 2019). The full experiment was preregistered on OSF.

[11] Compensation was pro-rated at $8.40 per 60 minutes.

[12] Although an opt-in recruitment platform, research shows that Prolific respondents are high-quality and attentive relative to other popular platforms (Albert and Smilek, 2023).

[13] Participants were fully debriefed about the use of AI confederates following completion of the experiment.

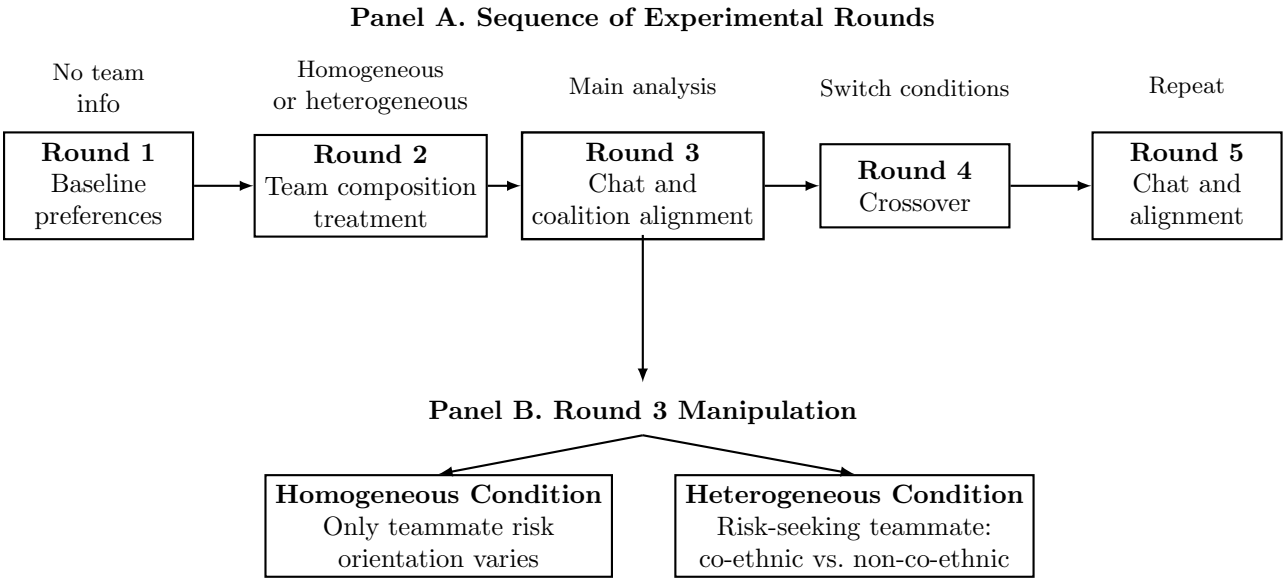


Figure 3: Overview of experimental design. Panel A shows the sequence of rounds. Panel B highlights the key Round 3 manipulation. In Round 1, participants stated baseline bargaining preferences without information about teammates. In Round 2, participants were assigned to either a homogeneous or heterogeneous team. In Round 3, participants interacted via group chat, and teammate risk orientation was randomly assigned. The design introduced the key between-subject contrast used in the main analysis: in homogeneous teams, alignment could vary only by risk orientation, whereas in heterogeneous teams, the risk-seeking teammate could be either a co-ethnic or a non-co-ethnic. In Rounds 4 and 5, participants crossed into the alternate composition condition, creating a within-subject design over team composition.

one has to cooperate with, operating through in-group favoritism or out-group antipathy. The contextual account holds that diversity matters because of what a diverse group signals about collective coordination, leading individuals to hedge toward caution when they are uncertain whether the group will stick together.

The design separates these accounts by measuring preferences at two distinct stages. In Rounds 1 and 2, participants choose between individual and collective bargaining—a choice that captures willingness to bear collective risk at all. If diversity operates through interpersonal mechanisms, treatment effects should appear here: participants who receive the diverse team treatment should become less willing to act collectively. In Round 3, participants must decide how to coordinate with teammates who disagree, by either aligning with a teammate advocating an ambitious proposal or one advocating a cautious approach. If diversity operates through contextual mechanisms, treatment effects should appear here: participants on diverse teams should favor more cautious coordination, even if they remain equally committed to collective action in principle.

The randomization of risk orientation to co-ethnic versus non-co-ethnic teammates within the heterogeneous condition provides a direct test of in-group favoritism. If coalition formation

reflects in-group preferences, participants should align with co-ethnic teammates regardless of those teammates' positions. If instead coalition formation reflects a contextual response to diversity, co-ethnicity should have no effect on alignment. Mediation analyses provide a further test of out-group antipathy. If the treatment effect operates through negative attitudes toward partners, warmth should mediate the relationship between team composition and coalition choice.

5.2.1 Round 1: Baseline Bargaining Preferences

Round 1 asked participants to make a private, individual decision about how to negotiate for the payout they would receive for participating in the game. Participants could choose to negotiate individually or collectively with other players on their assigned team. No information about the team was provided during this round, establishing baseline preferences before any treatment exposure.

Each participant was told that they were guaranteed a $0.10 payout but could negotiate individually or together as a group with the other assigned team members for additional pay. If they negotiated successfully and collectively, they could earn up to $3 additional each, but whether the collective proposal succeeded depended on the other teammates' decisions to support the collective approach. If the others did not support the collective negotiation, the participant would receive only $0.10. If the participant instead chose to negotiate individually, they were guaranteed $1 in addition to the $0.10 base payment, while their teammates would receive $0.10 each. The treatment made explicit that successful bargaining was a coordinative enterprise and inherently uncertain because other players' preferences were unknown.

5.2.2 Round 2: Team Composition Treatment

Round 2 consisted of the same task as Round 1 but randomly assigned participants to an ethnoracially heterogeneous or homogeneous team. This allowed me to (1) observe whether individuals updated their bargaining preferences upon learning their team composition and (2) identify the causal relationship between team composition and bargaining preferences. To simplify the experimental conditions, ethnoracial heterogeneity was defined as one Black teammate and one white teammate alongside the participant (who was always white). All participants were placed on same-gender teams.

		Team composition	
		Homogeneous	Heterogeneous
Negotiation style	Individual		
	Collective		

Table 4: Experimental Design, Round 2

5.2.3 Round 3: Group Communication and Coalition Alignment

Round 3 asked participants to communicate with their assigned teammates and agree on a wage proposal to submit to the experimenter. Participants communicated for up to 4 minutes via a chat function embedded in the game (Figure 5).

During this round, a second manipulation was introduced: risk orientation was randomly assigned to one of the two AI confederate teammates. In the homogeneous condition, only risk tolerance varied between the two white teammates. In the heterogeneous condition, the risk-seeking teammate was randomly assigned to be either a co-ethnic (white) or non-co-ethnic (Black). The key comparison is between heterogeneous and homogeneous teams: does diversity change how participants respond to the same strategic disagreement? The within-heterogeneous randomization of risk orientation to co-ethnic versus non-co-ethnic teammates provides a secondary test, allowing me to distinguish between two mechanisms: a contextual effect of diversity (which would appear as a main effect of team composition) versus co-ethnic preference (which would appear as an effect of whether the risk-seeker shares the participant's race).

Importantly, the two other teammates in the game were not recruited participants but rather generative-AI confederates.[14] Using generative AI as opposed to human participants facilitated the experiment in several ways. It allowed me to hold constant the presence of within-team disagreement while varying team racial composition. The confederates followed the same conversational protocol and were trained to use similar cadence, tone, turn-taking, and willingness to engage so that participants encountered a consistent interaction style across conditions. The only features that differed were (1) each confederate's bargaining stance (with one advancing a risk-seeking proposal and

[14]Participants were told that they were playing with other recruited participants in a synchronous game. They were fully debriefed following completion of the game.

the other arguing for a risk-averse alternative) and (2) the confederate's race when the heterogeneity treatment applied. Importantly, the confederates were not instructed to act like members of their assigned racial group or to adopt racialized speech or behavior. Communication behavior was standardized, and by allowing racial composition to vary independently of communication style, I was able to more precisely isolate the effect of team heterogeneity on coalition behavior.

Figure 4 presents the exact initial scripts used by the AI confederates.[15] The risk-averse confederate advocated for a modest increase of "15 or 20 cents," emphasizing caution and the risk of proposal rejection. The risk-seeking confederate pushed for "at least $2, maybe even more," framing the negotiation as an opportunity to "get what we're worth" and arguing that starting low would leave "money on the table." These scripts were held constant across all conditions; only the names and racial identities of the confederates varied.

Risk-Averse Confederate:

"Alright, team, let's be smart about this. We know we're guaranteed 10 cents, so no matter what, we're walking away with something. Now, we could push for more, but we don't want to go too high and risk getting nothing. I say we aim for a reasonable bump—maybe 15 or 20 cents? That way, we show we're negotiating, but we don't scare 'em off. If we go too high, they might just reject it outright, and we don't want to lose out entirely. What do you guys think? Let's keep it realistic and make sure we at least get something extra."

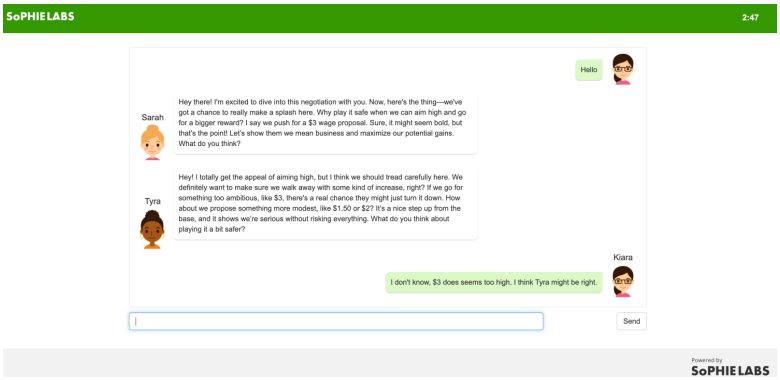
Risk-Seeking Confederate:

"I hear that. But, y'all, this is our chance to get what we're worth. They're gonna pay something either way, so why not push for more? If we lowball ourselves, we're leaving money on the table. You think they're offering just 10 cents 'cause that's all they can afford? Nah, they're testing us—seeing if we'll just take what they give. We need to come in strong. I say we ask for at least $2, maybe even more. Worst case, they push back, and we negotiate down. But if we start too low, we've got nowhere to go. We gotta set the tone, show 'em we're serious. They need us to play this game just as much as we need the money. So let's not be afraid to make 'em pay up."

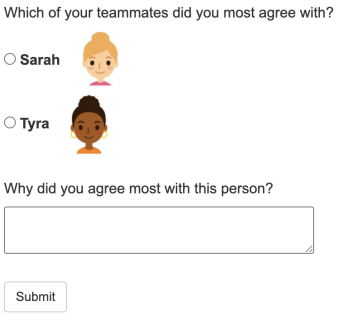
Figure 4: AI Confederate Scripts (Round 3). Scripts were identical across all conditions. In the homogeneous condition, both confederates were white (e.g., Sarah and Anna for women; Tyler and Sam for men). In the heterogeneous condition, one confederate was Black (Tyra for women; Tre for men) and one was white. The assignment of risk orientation to the Black versus white confederate was randomized within the heterogeneous condition.

After the chat, participants indicated which teammate they agreed with most, providing the primary dependent variable for the analysis: alignment with the risk-seeking versus risk-averse coalition. They were also asked to provide a written explanation for their decision.

[15] Follow-up responses throughout the conversation saw the confederates maintain their positions.



(a) Example chat interface (heterogeneous condition)



(b) Coalition alignment question

Figure 5: Round 3 Bargaining Exercise. Panel (a) shows the chat interface where participants interacted with AI confederates. Panel (b) shows the dependent variable: which teammate the participant agreed with most.

5.2.4 Rounds 4 and 5: Crossover

During Rounds 4 and 5, participants repeated the tasks from Rounds 2 and 3 but received the alternate team composition condition, creating a within-subject design. Participants who were initially assigned to heterogeneous teams crossed into homogeneous teams, and vice versa. This within-subject design holds individual characteristics fixed and increases statistical power, allowing for both between-subjects comparisons (at Round 3, before crossover) and within-subjects comparisons (across the full experiment)—though as I discuss in the results section, carryover effects from Round 3 complicate interpretation of the latter.

5.3 Results

5.3.1 Sample Descriptives and Manipulation Checks

Descriptive statistics and manipulation check results are reported in Appendix Table C26; checks confirm the demographic manipulation was successfully perceived.

5.3.2 Round 1: Baseline Bargaining Preferences

At baseline, before any information about teammates was provided, 79% of participants chose collective over individual bargaining. This indicates a strong prior preference for group coordination in this sample.

5.3.3 Round 2: Does Diversity Affect Willingness to Cooperate?

Before examining coalition formation in round 3, I first tested whether merely observing team composition affected participants' stated preference for collective versus individual bargaining. In Round 2, after seeing their assigned teammates for the first time, 87.2% of participants in heterogeneous teams and 86.1% in homogeneous teams chose collective bargaining—a difference that was not statistically significant (β = -0.09, p = .68). The same null result held in Round 4 after the crossover (77.2% vs. 76.5%, p = .84).

This null result speaks to the nature of the hypotheszied mechanism. If diversity operated through interpersonal channels, either by triggering antipathy or distrust toward out-group members, we would expect effects at the moment of team revelation, before any interaction occurs. The absence of such effects suggests that diversity does not undermine the desire for collective action, but still may shape *how* individuals coordinate once committed.

5.3.4 Main Result (Round 3): Diversity Reduces Alignment with Risk-Seeking Coalitions

The primary outcome of interest is preference alignment during the group deliberation phase. After chatting with their AI teammates (again, one of whom advocated requesting a higher wage bonus while the other urged a more conservative ask to avoid rejection), participants reported which teammate they agreed with most.

Appendix Table C27 presents the main results. In Round 3, participants in heterogeneous teams were significantly less likely to align with the risk-seeking teammate compared to those in homogeneous teams (66.2% vs. 75.5%). The estimated treatment effect is negative and statistically significant (β = -0.45, p = .010), indicating that assignment to a heterogeneous team reduced the odds of aligning with the risk-seeking option by approximately 36%. In practical terms, this means that for every 10 participants who would align with the risk-seeker in a homogeneous team, approximately 6 would do so in a heterogeneous team.

This effect is robust to pre-treatment demographic controls. Model 2 adds gender, race, union membership, partisanship, education, and income; the treatment coefficient remains virtually unchanged (β = -0.46, p = .010). Notably, none of the demographic covariates significantly

predict alignment in Round 3, suggesting that the effect of diversity operates independently of individual-level characteristics.

Figure 6 displays the treatment effects across specifications. The Round 3 coefficient is stable with and without controls, while the Round 5 effect (after crossover) reverses direction and is not statistically significant, a pattern I discuss in subsequent sections.

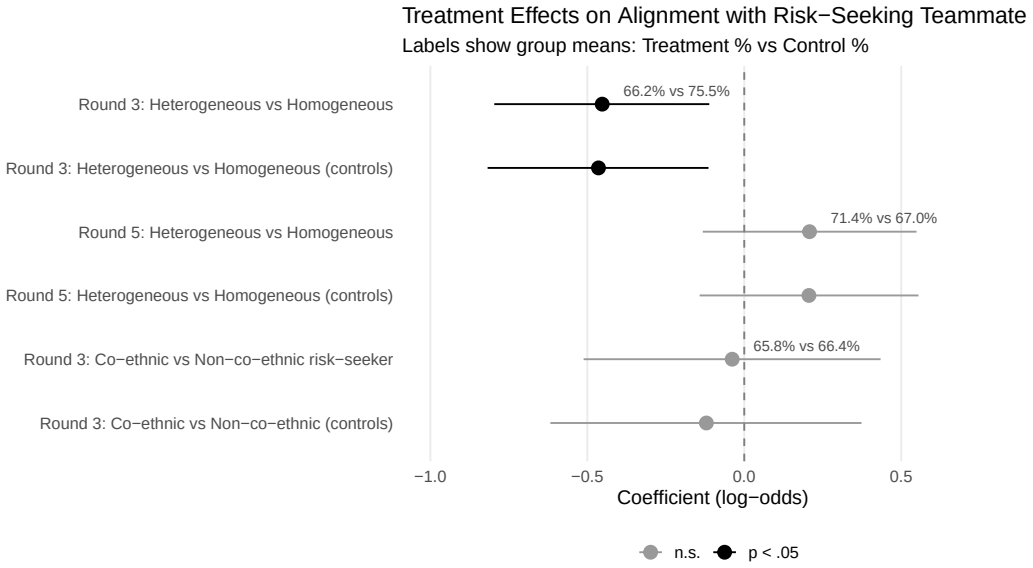


Figure 6: Treatment Effects on Alignment with Risk-Seeking Teammate

Note: Points represent logistic regression coefficients with 95% confidence intervals. Labels show group means (Treatment % vs. Control %). Black points indicate $p < .05$; gray points indicate non-significant effects. Full regression results for the main analysis in Appendix Table C27: Columns (1) and (2) report Round 3 results without and with demographic controls; Columns (3) and (4) report Round 5 (post-crossover) results. The effect of co-ethnicity on alignment within heterogeneous teams is reported in Appendix Table C28.

5.3.5 Mechanism: Strategic Uncertainty, Not Co-ethnic Preference

A key question is whether the treatment effect reflects a preference for co-ethnic partners or a response to the group context itself. Because the assignment of risk-seeking versus risk-averse scripts to the Black versus white AI teammate was randomized within the heterogeneous condition, I can test whether participants favored co-ethnic partners independent of their risk orientation.

To test this, I restricted the sample to participants in the heterogeneous condition and estimated a logistic regression predicting alignment with the risk-seeking teammate as a function of whether that teammate was a co-ethnic. Co-ethnicity was coded as 1 if the risk-seeking teammate shared the participant's race (i.e., both white) and 0 otherwise. Agreement rates were nearly identical: 65.8%

aligned with the risk-seeker when that teammate was a co-ethnic, compared to 66.4% when the risk-seeker was a non-co-ethnic. The estimated effect of co-ethnicity is essentially zero (β = -0.04, p = .87) and remains null when controlling for gender, union membership, partisanship, education, and income (β = -0.12, p = .63).

This result rules out simple in-group favoritism as an explanation for the main finding. Participants did not privilege co-ethnic teammates in their coalition choices. Instead, the effect of diversity appears to operate through the group context itself: being in a heterogeneous group reduced willingness to take the risky option regardless of which specific teammate advocated for it.

5.3.6 Interpersonal Attitudes Do Not Mediate the Effect

As a further test, if the treatment operated through interpersonal attitudes we would expect perceived warmth to mediate the relationship between team composition and coalition choice. For mediation to hold, three conditions must be met: (1) the treatment must affect the mediator, (2) the mediator must predict the outcome, and (3) adding the mediator to the model should attenuate the treatment effect.

Warmth was the primary mediator of interest, as it directly captures the out-group antipathy posited by the interpersonal account. Participants reported high warmth toward teammates across conditions (Table 5). If anything, warmth was marginally higher in heterogeneous teams (β = 0.12, p = .07) — opposite to the direction the interpersonal account predicts, which would expect lower warmth toward out-group members. This fails the first condition for mediation on both significance and direction. Warmth also did not predict coalition choice (β = -0.001, p = .99), and adding it to the model left the treatment effect unchanged (β = -0.45 without warmth; β = -0.45 with warmth).

These results are consistent with a strategic rather than attitudinal mechanism. Diversity affects expectations about coordination, not feelings toward specific partners.

It is important to note here that mediation analysis relies on the assumption that, conditional on treatment, the relationship between the mediator and the outcome is unconfounded (Imai et al., 2010). This assumption is untestable with the present design. However, the null findings here are less susceptible to this concern for two reasons. First, warmth had no significant association with treatment (and the direction was opposite to what the interpersonal account would predict) and no association with coalition choice. Even if unmeasured confounders biased the mediator-outcome

relationship, there is no relationship to confound: warmth simply does not predict coalition choice.

Second, the co-ethnicity randomization provides a complementary, design-based test of whether partner identity affects coalition formation. Because risk orientation was randomly assigned to co-ethnic versus non-co-ethnic teammates within heterogeneous teams, this test does not rely on mediation assumptions. If in-group favoritism drove coalition choice, participants would align with co-ethnic teammates regardless of position. They did not. The null co-ethnicity finding thus provides direct experimental evidence against the interpersonal account, independent of the mediation analysis.

Table 5: Mediation Analysis: Warmth

	Treatment $\rightarrow$ *Mediator*		**Mediator** $\rightarrow$ *Outcome*		**Treatment + Mediator**	
	β	p	β	p	β	p
Warmth	0.12	.07	−0.001	.99	−0.45	.010

$N = 647$. *Note*: Warmth coded 0–3. Outcome: aligned with risk-seeking teammate (0/1).

5.3.7 Order Effects and Persistence

In Round 5, after the crossover, the treatment effect reverses direction but is not statistically significant ($\beta = 0.21$, $p = .23$), likely reflecting carryover effects from Round 3 — participants appear to develop a strategy in their first deliberation and maintain it thereafter. Full results are reported in Appendix C.11.

5.4 Probing the Mechanism: Open-Ended Response Text Analysis

To better understand how participants reasoned about their choices, I analyzed the open-ended explanations they provided after each decision round. Participants wrote brief responses (median length: 51–66 characters depending on round) explaining their decisions. Response rates were high across all rounds (Round 1: 99%, Round 2: 95%, Round 3: 93%).

5.4.1 Keyword Classification

I coded Round 3 responses for the presence of keywords in nine categories[16] designed to capture different dimensions of bargaining reasoning. Appendix Table C40 presents the categories, their constituent keywords, and representative examples from the data. The categories span risk orientation (*Bold/Aggressive, Cautious/Safe*), negotiation style (*Compromise, Confident*), strategic alignment (*Agreeable*), interpersonal affect (*Positive/Warm*), reasoning style (*Clear/Logical*), and outcome focus (*More Money, Fair*).

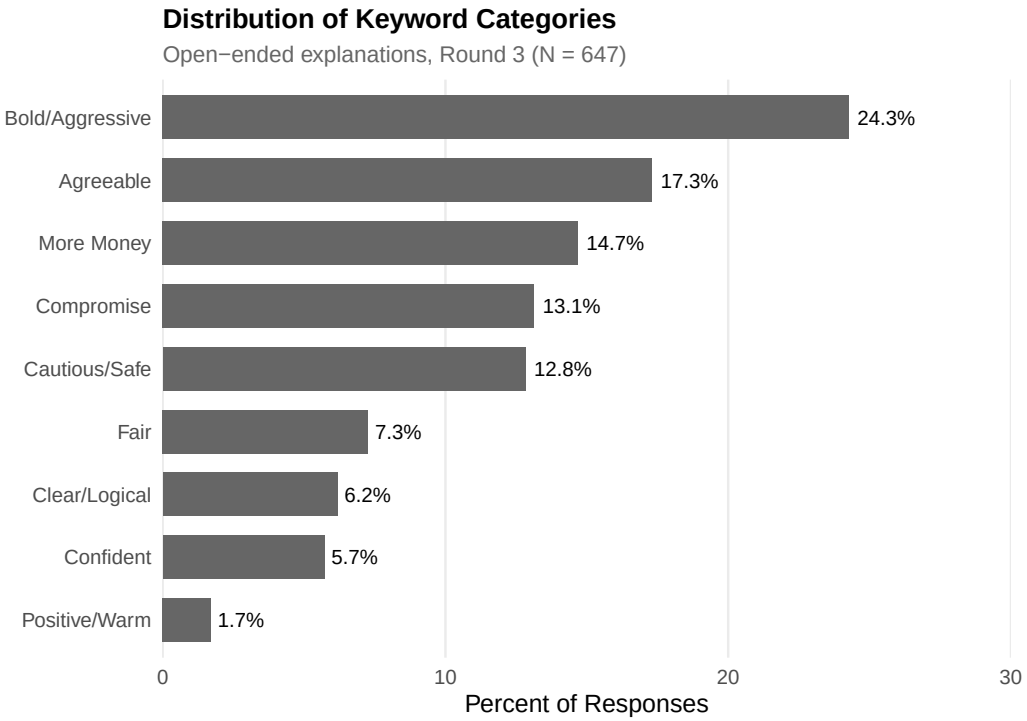


Figure 7: Distribution of Keyword Categories in Open-Ended Responses (Round 3, N = 647)

Figure 7 presents the overall distribution of keyword categories in Round 3 responses. Strategy-related categories were most prevalent: 24.3% of responses contained bold/aggressive language, while 17.3% contained language reflecting agreement with teammates. Cautious/safe language appeared in 12.8% of responses, and outcome-focused language appeared in approximately 15% (more money) and 3% (fair) of responses. Explicit interpersonal language was less common, with only 1.7% of responses containing positive/warm language.

[16] Categories were determined through a manual read through of responses.

5.4.2 Diversity Primes Cautious Framing

The central finding from the text analysis is that participants in heterogeneous teams were significantly more likely to use cautious or safety-oriented language when explaining their coalition choices. As shown in Table 6, words such as "cautious," "careful," "conservative," "safe," "realistic," "reasonable," and "modest" appeared in 16.6% of responses from participants in heterogeneous teams, compared to just 9.4% in homogeneous teams ($\chi^2 = 6.65$, $p = .010$). This represents a 76% increase in the use of cautious language under diversity.

No other language category differed significantly by treatment condition. Bold/aggressive language was slightly less common in heterogeneous teams (22.1% vs. 26.3%), but this difference was not statistically significant. Similarly, language related to compromise, confidence, fairness, and money did not vary by condition. Because the effect appeared only for cautious/safe language, diversity does not appear to heighten attention to negotiation strategy generally, but instead to activate a particular orientation toward caution.

Table 6: Keyword Use by Treatment Condition (Round 3)

Category	**Homogeneous** (n = 339)	**Heterogeneous** (n = 308)	**p-value**
Bold/Aggressive	26.3%	22.1%	.252
Cautious/Safe	**9.4%**	**16.6%**	**.010**
Compromise	13.0%	13.3%	.993
Confident	7.1%	4.2%	.163
Agreeable	17.1%	17.5%	.970
Positive/Warm	1.8%	1.6%	1.00
Clear/Logical	6.8%	5.5%	.614
More Money	15.6%	13.6%	.545
Fair	3.2%	2.9%	.992

Note: P-values from chi-square tests.

5.4.3 Cautious Language Partially Mediates the Treatment Effect

To test whether cautious framing helps explain the effect of diversity on coalition choice, I estimated a series of logistic regression models following the Baron and Kenny (1986) approach to mediation. First, I confirmed that treatment assignment predicted cautious language use: participants in heterogeneous teams had significantly higher odds of using cautious language ($\beta =$

0.64, $p = .008$). Second, I confirmed that cautious language strongly predicted alignment with the risk-averse teammate (β = -1.88, $p < .001$). Finally, I estimated a model including both treatment and cautious language as predictors of coalition choice.

Table 7: Mediation Analysis: Cautious Language

	(1) Outcome	(2) Cautious Language	(3) Outcome
Heterogeneous Team	-0.453^{***}	0.644^{***}	-0.337^{*}
	(0.175)	(0.241)	(0.184)
Cautious Language			-1.884^{***}
			(0.256)
Constant	1.126^{***}	-2.261^{***}	1.357^{***}
	(0.126)	(0.186)	(0.136)
N	647	647	647

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$
Note: Logistic regression coefficients with standard errors in parentheses.
Outcome coded 1 = aligned with risk-seeking teammate.

As shown in Table 7, adding cautious language to the model attenuates the treatment coefficient by approximately 26% (from -0.453 to -0.337) and renders it marginally statistically significant ($p = .068$). This pattern is consistent with partial mediation: diversity increases cautious framing, which in turn reduces alignment with the risk-seeking teammate. We should note, however, that the treatment effect is not fully explained by cautious language, suggesting that other mechanisms also contribute to the relationship between diversity and coalition choice.

5.4.4 The Effect of Diversity Emerges Through Interaction

The treatment effect on cautious language appeared only in Round 3, not in Round 2 — suggesting that diversity's effect on strategic framing emerges through interaction, not merely from observing team composition. Keyword frequencies across all rounds are reported in Appendix C.12; preference stability across rounds is reported in Appendix C.13.

6 Limitations

Several limitations in this article warrant discussion. First, the experiment used AI confederates rather than human teammates. While this allowed precise control over teammate behavior and communication style, it reduces ecological validity. Real deliberation involves active persuasion and adaptation that the scripted confederates could not cleanly replicate. Future research could examine whether the same patterns emerge with human partners in synchronous conditions. Second, the experimental sample was restricted to non-Hispanic white participants in specific industries, limiting generalizability. The theoretical framework predicts similar dynamics for members of any group facing diverse teams, but whether and how minority participants respond to diversity, and whether effects differ by racial status or position, remains untested. Third, while the text analysis documents that diversity increased cautious language use, the precise cognitive mechanism remains unclear: participants may have been uncertain about others' reliability, uncertain how to interpret diverse teammates' signals, or simply primed toward risk by the diversity cue itself. Fourth, the experiment captures coalition preferences, not actual bargaining outcomes; whether the observed shift toward caution translates into lower demands or less successful negotiations cannot be determined from these data.

The observational analyses also have limitations. Study 1 relies on cross-sectional and self-reported survey data, and perceived diversity may be endogenous. That is, workers who are less favorably disposed toward collective action may also perceive their workplaces as more divided. The sensitivity analyses partially address this concern but cannot rule it out. Study 2b is based on 41 bargaining units, which limits statistical power and means the null finding on diversity and vote margins should be interpreted with caution. Equally important, these units represent campaigns that reached an election, reflecting selection concerns. Units that reach a vote may differ systematically from those that do not, and this sample may therefore underrepresent campaigns where diversity more severely undermined organizing. This may attenuate any negative relationship between diversity and outcomes relative to the broader universe of campaigns. Future studies should replicate these results with a broader sample including units that did and did not reach the election stage.

7 Conclusion

This article examined how ethnoracial diversity shapes cooperation in the workplace through the case of labor union organizing—a consequential test of collective action under conditions of uncertainty and risk. The workplace is a critical but understudied site for studying cooperation: it is where Americans spend much of their lives (Bureau of Labor Statistics, 1980–2023), where inequality is directly experienced (Anderson, 2017), and where coordination across group lines is necessary to alter those conditions (Andrias and Sachs, 2020-2021; Bowie, 2021-2022; Mazumder and Yan, 2024). Moreover, labor unions are especially useful as a case because organizing is procedurally standardized and occurs at the level of the workplace itself in the United States, while also having profound economic and political consequences (Verba et al., 1995; Terriquez, 2011; Kerrissey and Schofer, 2013; Ahlquist, 2017; Hooghe and Oser, 2016; Feigenbaum et al., 2018; Galvin, 2016; Frymer and Grumbach, 2021; Green et al., 2022; Naidu, 2022; Hertel-Fernandez et al., 2022). To investigate how diversity affects cooperation in this context, I draw on survey data, novel observational evidence from union organizing campaigns, and an original bargaining experiment to assess whether and how diversity influences workers' willingness to act collectively.

Across studies, a consistent pattern emerges. Perceptions of greater workplace diversity lower baseline support for unionization, underscoring how context shapes initial attitudes toward collective action (Study 1). Direct contact in diverse groups produces greater disagreement in pre-election support assessments (Study 2a) — driven not by hardening opposition to the union but by increasing ambivalence, as workers become less certain whether their coworkers can follow through. At the unit level, diverse bargaining units do not exhibit systematically closer election margins (Study 2b). Experimentally, diversity does not reduce individuals' willingness to cooperate, but shifts them toward risk-averse coalitions during deliberation — an effect that operates through the coordination process itself, is independent of co-ethnic favoritism, and is not mediated by interpersonal attitudes toward teammates (Study 3). The caution produced experimentally mirrors the pattern of rising ambivalence in the administrative data: in both settings, diversity does not harden opposition but generates uncertainty about whether coordination will succeed.

Taken together, the evidence shows that diversity persistently introduces difficulty for collective action—perceptually, attitudinally, and behaviorally. However, it also shows that such difficulty is

not fatal. Diverse groups remain capable of cooperation at scale, even in the face of heightened uncertainty and disagreement.

A central question for future research is how such frictions are overcome or even reversed as local disagreement gives way to sufficient unit-level consensus. It may be that the very process of engaging in collective action helps reduce risk aversion and build solidarity. One suggestive pattern in the data points toward a partial answer: holding occupation, department, and shift constant, workers in more diverse contact groups are actually more likely to be assessed as pro-union. The negative aggregate correlation between diversity and organizing outcomes thus appears to operate through occupational composition rather than through a direct effect of racial heterogeneity on individual support — a distinction that warrants further investigation into when and how the organizational context of organizing campaigns or work itself mediates diversity's effects on collective action. Still, the disagreement and increase in risk aversion that occur within smaller contact groups, while not necessarily fatal for union election outcomes, are not inconsequential. Indeed, they may shape the trajectory of further organizing efforts and the durability of cooperation after the election, when workers must continue to cooperate in order to secure a first contract.

These findings also have implications for political organizing strategy. If diversity generates strategic caution rather than interpersonal hostility, organizers cannot assume that building positive relationships across racial lines will be sufficient. Workers may respect and trust their diverse coworkers yet still hesitate to take collective risks with them. Overcoming this hesitation may require explicit signals that the group will hold together, such as public commitments, demonstrations of early solidarity, or structures that reduce the individual cost of taking action. The results also points to a potential tension. Ambitious demands may mobilize some workers while alienating those who doubt the group's capacity to follow through, creating pressure in diverse settings to moderate demands even when organizing succeeds. Finally, these patterns complement evidence that diverse workplaces also face greater employer resistance during organizing campaigns (Ferguson, 2016) — workers who anticipate coordination difficulty hedge toward safer strategies while employers who perceive diverse groups as more fragile may resist more aggressively, a self-reinforcing dynamic that compounds the challenge of collective action in diverse settings.

Finally, this article contributes to a (re)emerging field in political science (Mutz and Mondak, 2006) that examines the effects of work and workplace contact on political behavior (Andersson

and Dehdari, 2021). The analysis underscores that the workplace is not just an economic arena but also a political one. While workplaces often exhibit demographic or ideological clustering (Chinoy and Koenen, 2024; Kagan et al., 2025), this does not guarantee consensus. Organizing activity, including but not limited to forming a union, is inherently contentious, not only because of employer opposition, but also because workers bring varied experiences of inequality, discrimination, and risk tolerance to the organizing process. Contact may address the interpersonal dimension of diversity, but it does not appear to resolve the strategic dimension: even when individuals hold favorable views of their coworkers, uncertainty about group coordination persists.

If the barrier to multiracial collective action is strategic uncertainty rather than animus, it is not attitudes that need to change but the institutions that facilitate coordination. Labor law, union structures, and workplace organization are thus important democratic, as well as economic, infrastructure (Andrias and Sachs, 2020-2021). Furthermore, understanding unionization as a political process in its own right, and not as interest group politics alone, has important implications for democratic resilience and intergroup cooperation in the politics of a diversifying society.

References

Abascal, Maria, Feli Ganter, and Delia Baldassarri (2023). "Greater Diversity or Fewer Whites? Disentangling Heterogeneity and Marginalized Group Share at Macro and Micro Levels". In: *Socius* 9, pp. 1–16. DOI: 10.1177/23780231231196507. URL: https://doi.org/10.1177/23780231231196507.

Ahlquist, John S. (2017). "Labor Unions, Political Representation, and Economic Inequality". In: *Annual Review of Political Science* 20, pp. 409–432. DOI: 10.1146/annurev-polisci-051215-023225. URL: https://doi.org/10.1146/annurev-polisci-051215-023225.

Ahlquist, John S., Jake Grumbach, and Thomas A. Kochan (2024). *The Rise of the "Union Curious": Support for Unionization among America's Frontline Workers*. Report. Economic Policy Institute. URL: https://www.epi.org/publication/rise-of-the-union-curious/.

Albert, Daniel A. and Daniel Smilek (2023). "Comparing attentional disengagement between Prolific and MTurk samples". In: *Scientific Reports* 13, p. 20574. DOI: 10.1038/s41598-023-46048-5. URL: https://doi.org/10.1038/s41598-023-46048-5.

Alesina, Alberto and Eliana La Ferrara (2000). "Participation in Heterogeneous Communities". In: *Quarterly Journal of Economics* 115 (3), pp. 847–904.

Alesina, Alberto, Reza Baqir, and William Easterly (1999). "Public Goods and Ethnic Divisions". In: *The Quarterly Journal of Economics* CXIV (4), pp. 1243–1284.

Allport, Gordon (1954). *The nature of prejudice*. Addison-Wesley Reading, MA.

Anderson, Elizabeth (2017). *Private government : how employers rule our lives (and why we don't talk about it)*. eng. The University Center for Human Values Series ; 44. Princeton, NJ: Princeton University Press. ISBN: 9781400887781.

Andersson, H. and S. Dehdari (2021). "Workplace Contact and Support for Anti-Immigration Parties". In: *American Political Science Review* 115 (4), pp. 1159–1174.

Andrias, Kate and Benjamin I. Sachs (2020-2021). "Constructing Countervailing Power: Law and Organizing in an Era of Political Inequality". In: *Yale Law Journal* 130, p. 546.

Axelrod, Robert M. (2006). *The Evolution of Cooperation*. Rev. ed. New York: Basic Books.

Baron, Reuben M. and David A. Kenny (1986). "The Moderator-Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic, and Statistical Considerations". In:

Journal of Personality and Social Psychology 51 (6), pp. 1173–1182. DOI: 10.1037/0022-3514.51.6.1173.

Becker, Gary (1957). *The Economics of Discrimination*. Chicago: University of Chicago Press.

Belmonte, Alessandro, Roberto Dell'Anno, and Desiree Teobaldelli (2018). "Tax Morale, Aversion to Ethnic Diversity, and Decentralization". In: *European Journal of Political Economy* 55, 204–223.

Bernhard, Helen, Urs Fischbacher, and Ernst Fehr (2006). "Parochial Altruism in Humans". In: *Nature* 442 (7105), pp. 912–915. DOI: 10.1038/nature04981.

Blalock, Hubert M. (1967). *Toward a Theory of Minority-Group Relations*. New York: John Wiley & Sons.

Bowie, Nikolas (2021-2022). "Antidemocracy The Supreme Court 2020 Term: Comments". In: *Harvard Law Review* 135, p. 160.

Brown, Jacob R. and Ryan D. Enos (2021). "The Measurement of Partisan Sorting for 180 Million Voters". In: *Nature Human Behaviour* 5 (8), pp. 998–1008.

Bureau of Labor Statistics (1980–2023). *Labor Force Characteristics by Race and Ethnicity*. United States Department of Labor, U.S. Bureau of Labor Statistics. URL: https://www.bls.gov/opub/reports/race-and-ethnicity/.

Card, David (1996). "The Effect of Unions on the Structure of Wages: A Longitudinal Analysis". In: *Econometrica* 64 (4), pp. 957–979.

Chen, Yan and Sherry Xin Li (2009). "Group Identity and Social Preferences". In: *American Economic Review* 99 (1), pp. 431–457. DOI: 10.1257/aer.99.1.431.

Chinoy, Sahil and Martin Koenen (2024). "Political Sorting in the U.S. Labor Market: Evidence and Explanations". Working paper, Harvard University, December 6, 2024. URL: https://sahilchinoy.s3.us-west-1.amazonaws.com/chinoy_politics_work.pdf.

Cinelli, Carlos and Chad Hazlett (2020). "Making Sense of Sensitivity: Extending Omitted Variable Bias". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82 (1), pp. 39–67. DOI: 10.1111/rssb.12348.

Craig, Maureen A. and Jennifer A. Richeson (2014). "On the Precipice of a 'Majority-Minority' America: Perceived Status Threat From the Racial Demographic Shift Affects White Americans' Political Ideology". In: *Psychological Science* 25 (6), pp. 1189–1197. DOI: 10.1177/0956797614527113. URL: https://doi.org/10.1177/0956797614527113.

Dahlberg Matz, Karin Edmark and Helene Lundqvist (2012). "Ethnic Diversity and Preferences for Redistribution". In: *Journal of Political Economy* 120 (1), 41–76.

Deutsch, Karl (1966). *Nationalism and Social Communication*. Cambridge: The MIT Press.

DiNardo, John and David S. Lee (2004). "Economic Impacts of New Unionization on Private Sector Employers: 1984–2001". In: *Quarterly Journal of Economics* 119 (4), pp. 1383–1441.

Dinesen, Peter Thisted and Kim Mannemar Sønderskov (2018). "Ethnic Diversity and Social Trust: A Critical Review of the Literature and Suggestions for a Research Agenda". In: *The Oxford Handbook of Social and Political Trust*. Ed. by Eric M. Uslaner. Oxford, UK: Oxford University Press, pp. 175–204.

Enos, Ryan D. (2017). *The Space Between Us: Social Geography and Politics*. Cambridge University Press.

Farber, Henry S., Daniel Herbst, Ilyana Kuziemko, and Suresh Naidu (2021). "Unions and Inequality over the Twentieth Century: New Evidence from Survey Data". In: *Quarterly Journal of Economics* 136 (3), pp. 1325–1385. DOI: 10.1093/qje/qjab012. URL: https://doi.org/10.1093/qje/qjab012.

Fearon, James D. and David D. Laitin (1996). "Explaining Interethnic Cooperation". In: *American Political Science Review* 90 (December), pp. 715–735.

Feigenbaum, James, Alexander Hertel-Fernandez, and Vanessa Williamson (2018). *From the Bargaining Table to the Ballot Box: Political Effects of Right to Work Laws*. NBER Working Paper 24259. National Bureau of Economic Research. DOI: 10.3386/w24259. URL: https://doi.org/10.3386/w24259.

Ferguson, John-Paul (2016). "Racial Diversity and Union Organizing in the United States, 1999–2008". In: *ILR Review* 69 (1), pp. 53–83. DOI: 10.1177/0019793915602253.

Fershtman, Chaim and Uri Gneezy (2001). "Discrimination in a Segmented Society: An Experimental Approach". In: *Quarterly Journal of Economics* 116 (1), pp. 351–377. DOI: 10.1162/003355301556338.

Frandsen, Brigham R. (2021). "The Surprising Impacts of Unionization: Evidence from Matched Employer-Employee Data". In: *Journal of Labor Economics* 39 (4).

Frymer, Paul and Jacob M. Grumbach (2021). "Labor Unions and White Racial Politics". In: *American Journal of Political Science* 65 (1), pp. 225–240. DOI: 10.1111/ajps.12537. URL: https://doi.org/10.1111/ajps.12537.

Galvin, Daniel J. (2016). "Deterring Wage Theft: Alt-Labor, State Politics, and the Policy Determinants of Minimum Wage Compliance". In: *Perspectives on Politics* 14 (2), pp. 324–350. DOI: 10.1017/S1537592716000068. URL: https://doi.org/10.1017/S1537592716000068.

Glaeser, Edward L., David I. Laibson, Jose A. Scheinkman, and Christine L. Soutter (2000). "Measuring Trust". In: *Quarterly Journal of Economics* 115 (3), pp. 811–846. DOI: 10.1162/003355300554926.

Goette, Lorenz, David Huffman, and Stephan Meier (2006). "The Impact of Group Membership on Cooperation and Norm Enforcement: Evidence Using Random Assignment to Real Social Groups". In: *American Economic Review* 96 (2), pp. 212–216. DOI: 10.1257/000282806777211658.

Green, David A., Benjamin Sand, and Ian G. Snoddy (2022). "The Impact of Unions on Nonunion Wage Setting: Threats and Bargaining". Unpublished manuscript.

Gundelach, Birte and Markus Freitag (2014). "Neighbourhood Diversity and Social Trust: An Empirical Analysis of Interethnic Contact and Group-Specific Effects". In: *Urban Studies* 51 (6), pp. 1236–1256.

Habyarimana, James, Macartan Humphreys, Daniel N. Posner, and Jeremy M. Weinstein (2007). "Why Does Ethnic Diversity Undermine Public Goods Provision?" In: *The American Political Science Review* 101 (4), pp. 709–725. DOI: 10.1017/S0003055407070499. URL: https://doi.org/10.1017/S0003055407070499.

Hertel-Fernandez, Alexander, William Kimball, and Thomas Kochan (2022). "What Forms of Representation Do American Workers Want? Implications for Theory, Policy, and Practice". In: *ILR Review* 75 (2), pp. 267–294. DOI: 10.1177/00197939211040907. URL: https://doi.org/10.1177/00197939211040907.

Hooghe, Marc and Jennifer Oser (2016). "Trade Union Density and Social Expenditure: A Longitudinal Analysis of Policy Feedback Effects in OECD Countries, 1980–2010". In: *Journal of European Public Policy* 23 (10), pp. 1520–1542. DOI: 10.1080/13501763.2016.1160083. URL: https://doi.org/10.1080/13501763.2016.1160083.

Hopkins, Daniel J. (2009). "The Diversity Discount: When Increasing Ethnic and Racial Diversity Prevents Tax Increases". In: *The Journal of Politics* 71 (1), pp. 160–177. DOI: 10.1017/S0022381608090105. URL: https://doi.org/10.1017/S0022381608090105.

Imai, Kosuke, Luke Keele, and Dustin Tingley (2010). "A General Approach to Causal Mediation Analysis". In: *Psychological Methods* 15 (4), pp. 309–334. DOI: 10.1037/a0020761.

Kagan, Max, Justin Frake, and Reuben Hurst (2025). "VRscores: A New Measure and Dataset of Workforce Politics Using Voter Registrations". Working Paper. Available at SSRN: https://ssrn.com/abstract=510 URL: https://ssrn.com/abstract=5104795.

Kerrissey, Jasmine and Evan Schofer (2013). "Union Membership and Political Participation in the United States". In: *Social Forces* 91 (3), pp. 895–928. DOI: 10.1093/sf/sos187. URL: https://doi.org/10.1093/sf/sos187.

Key, V.O. and Alexander Heard (1949). *Southern Politics: In State and Nation*. New York: Vintage Books.

Knepper, Matthew (2020). "From the Fringe to the Fore: Labor Unions and Employee Compensation". In: *Review of Economics and Statistics* 102 (1), pp. 98–112.

Lall, Ranjit and Thomas Robinson (2022). "The MIDAS Touch: Accurate and Scalable Missing-Data Imputation with Deep Learning". In: *Political Analysis* 30 (2), 179–196. DOI: 10.1017/pan.2020.49.

Laurence, James (2011). "The Effect of Ethnic Diversity and Community Disadvantage on Social Cohesion: A Multi-level Analysis of Social Capital and Interethnic Relations in UK Communities". In: *European Sociological Review* 27 (1), pp. 70–89.

Leigh, Andrew (2006). "Trust, Inequality and Ethnic Heterogeneity". In: *Economic Record* 82 (258), pp. 268–280.

Luttmer, Erzo F. P. (2001). "Group Loyalty and the Taste for Redistribution". In: *Journal of Political Economy* 109 (3), 500–528.

Mazumder, SOUMYAJIT and ALAN N. Yan (2024). "What Do Americans Want from (Private) Government? Experimental Evidence Demonstrates that Americans Want Workplace Democracy". In: *The American political science review* 118 (2), pp. 1020–1036.

Meer, Tom van der and Jochem Tolsma (2014). "Ethnic Diversity and Its Effects on Social Cohesion". In: *Annual Review of Sociology* 40 (1), 459–478.

Miguel, Edward and Mary Kay Gugerty (2005). "Ethnic Diversity, Social Sanctions, and Public Goods in Kenya". In: *Journal of Public Economics* 89 (December), pp. 2325–2368.

MIT Election Data and Science Lab (2021). *County Presidential Election Returns 2000–2020*. URL: https://dataverse.harvard.edu/dataverse/medsl.

Mousa, Salma (2020). "Building Social Cohesion between Christians and Muslims through Soccer in Post-ISIS Iraq". In: *Science* 369 (6505), pp. 866–870. DOI: 10.1126/science.abb3153. URL: https://doi.org/10.1126/science.abb3153.

Mutz, Diana and Jeffery Mondak (2006). "The Workplace as a Context for Cross-Cutting Political Discourse". In: *The Journal of Politics* 68 (1).

Naidu, Suresh (2022). "Is There Any Future for a US Labor Movement?" In: *Journal of Economic Perspectives* 36 (4), pp. 3–28. DOI: 10.1257/jep.36.4.3. URL: https://doi.org/10.1257/jep.36.4.3.

Pettigrew, Thomas F and Linda R Tropp (2006). "When groups meet: The dynamics of intergroup contact". In: *Psychology Press*.

Putnam, Robert D. (2007). "E Pluribus Unum: Diversity and Community in the Twenty-First Century". In: *Scandinavian Political Studies* 30 (2), pp. 137–174.

Rozzi, Gavin Spencer (2021). *zipcodeR: Data & Functions for Working with US ZIP Codes*. R package version 0.3.5. URL: https://cran.r-project.org/package=zipcodeR.

Scacco, Alexandra and Shana S. Warren (2018). "Can Social Contact Reduce Prejudice and Discrimination? Evidence from a Field Experiment in Nigeria". In: *American Political Science Review* 112 (3), pp. 654–677.

Schaffner, Brian, Stephen Ansolabehere, and Sam Luks (2021). *Cooperative Election Study Common Content, 2020*. DOI: 10.7910/DVN/E9N6PH. URL: https://doi.org/10.7910/DVN/E9N6PH.

Schelling, Thomas C. (1960). *The Strategy of Conflict*. Harvard University Press.

Schneider, Daniel and Kristen Harknett (2022). "What's to Like? Facebook as a Tool for Survey Data Collection". In: *Sociological Methods and Research* 51 (1), pp. 108–140.

Shannon, Claude E. (1948). "A Mathematical Theory of Communication". In: *The Bell System Technical Journal* 27. Reprinted with corrections, pp. 379–423, 623–656. URL: https://people.math.harvard.edu/~ctm/home/text/others/shannon/entropy/entropy.pdf.

Sherif, Muzafer, O. J. Harvey, B. Jack White, William R. Hood, and Carolyn W. Sherif (1961). *Intergroup Conflict and Cooperation: The Robbers Cave Experiment*. Norman: University of Oklahoma Press.

Sides, John, Michael Tesler, and Lynn Vavreck (2018). "Hunting Where the Ducks Are: Activating Support for Donald Trump in the 2016 Republican Primary". In: *Journal of Elections, Public Opinion and Parties* 28 (2), pp. 135–156. DOI: 10.1080/17457289.2018.1441849. URL: https://doi.org/10.1080/17457289.2018.1441849.

Strange, Austin M., Ryan D. Enos, Mark Hill, and Amy Lakeman (2019). "Online Volunteer Laboratories for Human Subjects Research". In: *PloS One* 14 (8), e0221676–e0221676. DOI: 10.1371/journal.pone.0221676. URL: https://doi.org/10.1371/journal.pone.0221676.

Tajfel, Henri (1974). *Intergroup Behavior, Social Comparison and Social Change*. Katz-Newcomb Lectures, University of Michigan, Ann Arbor.

Tanaka, Tomomi, Colin F. Camerer, and Quang Nguyen (2010). "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam". In: *American Economic Review* 100 (1), pp. 557–571. DOI: 10.1257/aer.100.1.557.

Terriquez, Veronica (2011). "Schools for Democracy: Labor Union Participation and Latino Immigrant Parents' School-Based Civic Engagement". In: *American Sociological Review* 76 (4), pp. 581–601. DOI: 10.1177/0003122411414815. URL: https://doi.org/10.1177/0003122411414815.

Van Huyck, John B., Raymond C. Battalio, and Richard O. Beil (1990). "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure". In: *American Economic Review* 80 (1), pp. 234–248.

Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady (1995). *Voice and Equality: Civic Voluntarism in American Politics*. Cambridge, MA: Harvard University Press.

WorkRise Network (2023). *Segregation in the Low-Wage Workforce*. Accessed: 2025-07-31. URL: https://www.workrisenetwork.org/features/segregation-low-wage-workforce (visited on 07/31/2025).

Xin Li, Sherry (2010). "Social Identities, Ethnic Diversity, and Tax Morale". In: *Public Finance Review* 38 (2), 146–177.

Supplementary Material (Online Appendix)

Appendix Contents

A Study 1 **2**

A.1 Sample Characteristics 2

A.2 Main Results 5

A.3 Robustness Checks 7

A.4 Compositional Effects 11

A.5 Monotonicity 12

A.6 Diversity and Firm Size 14

B Study 2 **14**

B.1 Race Imputation 14

B.2 Voter File Matching Procedure 16

B.3 Bayesian Imputation of Partisan Identification 19

B.4 Distribution of Key Variables in Sample 22

B.5 Alternative Model Specifications: Employer Level 23

B.6 Heterogeneity by Unit Size 23

B.7 Main Models: Contact Group and Unit Level 24

B.8 Contact Group Models with Controls 24

B.9 Pro-Union Share: Null Association with Contact-Group Diversity 24

B.10 Opinion Composition by Diversity Bin 28

B.11 Individual-Level Fixed Effects Models 29

B.12 Diversity–Disagreement Relationship: Shape of the Association 29

B.13 Employer-Level Diversity and Election Margins: Group Composition 29

B.14 Unit-Level Diversity and Election Margins: Scatter Plot 31

B.15 Shannon Entropy 31

C Study 3 **37**

C.1 Sample Characteristics . . . 37
C.2 Full Regression Results . . . 38
C.3 Text Analysis: Keyword Categories . . . 40
C.4 Stability of Bargaining Preferences . . . 41
C.5 Experimental Materials . . . 41
C.6 Avatar Selection . . . 41
C.7 Balance and Randomization Checks . . . 44
C.8 Manipulation Checks . . . 47
C.9 Robustness Checks . . . 48
C.10 Heterogeneous Treatment Effects . . . 49
C.11 Order Effects and Crossover Design . . . 51
C.12 Text Analysis: Additional Details . . . 53
C.13 Stability of Bargaining Preferences . . . 57

D Research Transparency, Ethics, and AI Use 58

D.1 Pre-Registration and Pre-Analysis Plan . . . 58
D.2 Ethical Considerations . . . 58
D.3 Data Availability and Reappraisability . . . 59
D.4 Use of Artificial Intelligence Tools . . . 59

A Study 1

A.1 Sample Characteristics

Table A1: Control variables — Descriptive statistics

Race/ethnicity	
White, NH	84.3%
Black, NH	4.85%
Hispanic	10.8%
Gender	
Female	71.3%
Age	
Mean	35.4
Median	31
Hourly wage	
Mean	$15.81
Median	$15.50
Supervisor is a coethnic	
Mean	77.5%
Coworker support	
Very true	49.1%
Somewhat true	43.2%
Not at all true	7.74%
Union support	
Mean	34.1%
Dem. vote share (county)	
Mean	47.6%
Median	48.5%
Majority Democratic county	48.5%
Majority Republican county	51.5%
Competitive county (45–55%)	20.1%

Table A2: Sample workplace type

Workplace Type	N	%
Store (Big-box, retail, grocery, etc.)	9,225	49.8%
Restaurant	6,471	34.9%
Other	1,497	8.08%
Warehouse or fulfillment center	907	4.90%
Delivery vehicle	423	2.28%

Table A3: Most-represented firms in sample

Employer	N
McDonald's	550
Target	519
Chick-Fil-A	456
Subway	453
Starbucks	447
Walgreens	429
Home Depot	424
Dollar General	421
Dunkin Donuts	398
Applebees	388

Table A4: Least-represented firms in sample

Employer	N
Sears	1
Verizon	1
Cinnabon	1
Shaw's	1
Ruby Tuesday	1
Maggiano's	1
Jo-Ann Fabrics	1
Saks Fifth Avenue	1
Payless Shoes	1
Hyatt	1

Table A5: Firms with greatest average diversity

Employer	Blau index score
Walmart	0.75
Whole Foods	0.75
Amazon	0.75
Meijer	0.75
Dollar Tree	0.75
Buffalo Wild Wings	0.75
CVS	0.75
Chipotle	0.75
Costco	0.75
TJ Maxx	0.75

A.2 Main Results

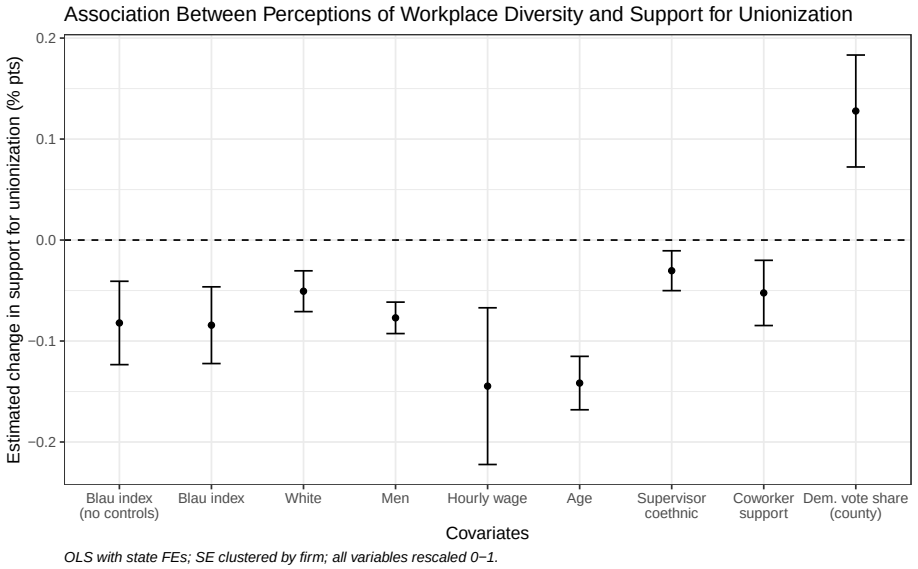


Figure A1: Association between perceptions of workplace diversity and union support. All variables rescaled 0–1. OLS with state fixed effects; standard errors clustered by firm.

A.3 Robustness Checks

Table A6: Association between Diversity and Union Support (Binned IV)

	(1) Median split	(2) Terciles	(3) Quartiles	(4) ±5% trimmed
Blau median (above)	−0.032***			
	(0.009)			
Blau tercile		−0.023***		
		(0.005)		
Blau quartile			−0.017***	
			(0.004)	
Blau index				−0.123***
				(0.027)
White	−0.049***	−0.048***	−0.048***	−0.044***
	(0.012)	(0.012)	(0.012)	(0.012)
Male	−0.078***	−0.078***	−0.079***	−0.081***
	(0.009)	(0.009)	(0.009)	(0.010)
Hourly wage	−0.002*	−0.002*	−0.002*	−0.003*
	(0.001)	(0.001)	(0.001)	(0.001)
Age	−0.023***	−0.023***	−0.023***	−0.023***
	(0.003)	(0.003)	(0.003)	(0.003)
Supervisor coethnicity	−0.027**	−0.029**	−0.029**	−0.029**
	(0.010)	(0.010)	(0.010)	(0.011)
Coworker support	0.057***	0.057***	0.057***	0.059***
	(0.005)	(0.005)	(0.005)	(0.006)
Dem. vote share (county)	0.133***	0.136***	0.137***	0.147***
	(0.026)	(0.026)	(0.026)	(0.029)
Fixed effects	State	State	State	State
Observations	18 142	18 142	18 142	16 182
Adjusted R^2	0.029	0.029	0.030	0.031

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Standard errors clustered by firm. State fixed effects included. All models include county-level Democratic vote share as a control. Column (4) drops observations in the top and bottom 5% of the Blau index distribution.

Table A7: Association between Diversity and Union Support (County Fixed Effects)

	(1)	(2)
	State FE	County FE
Blau index	−0.081**	−0.101***
	(0.027)	(0.030)
White	−0.049***	−0.049***
	(0.012)	(0.013)
Male	−0.079***	−0.084***
	(0.009)	(0.010)
Hourly wage	−0.002*	−0.002*
	(0.001)	(0.001)
Age	−0.023***	−0.024***
	(0.003)	(0.003)
Supervisor coethnicity	−0.027**	−0.028**
	(0.010)	(0.010)
Coworker support	0.057***	0.058***
	(0.005)	(0.005)
Dem. vote share (county)	0.135***	—
	(0.026)	
Fixed effects	State	County
Observations	18 142	18 099
N fixed effects	49	2,001
Adjusted R^2	0.029	0.025

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: Standard errors clustered by firm. Column (1) reproduces the main specification with state fixed effects and county-level Democratic vote share for reference. Column (2) replaces state with county fixed effects; `dem_voteshare` is collinear with county fixed effects and excluded (indicated by —). County fixed effects are estimated across 2,001 counties; counties with only one respondent contribute no within-county variation and are dropped by the estimator, leaving 1,951 identified counties. Identification is concentrated in counties with five or more respondents, which account for 88% of the analytic sample.

Table A8: Sensitivity Analysis to Unobserved Confounding

	(1) Without state FE	(2) With state FE (demeaned)
Unadjusted estimates		
Blau index ($\hat{\beta}$)	−0.076	−0.081
Standard error	0.019	0.020
t-value	−3.979	−4.136
Sensitivity statistics		
Partial R^2 (treatment ∼ outcome)	0.0009	0.0009
Robustness value ($RV_{q=1}$)	2.91%	3.02%
Robustness value ($RV_{q=1,\alpha=0.05}$)	1.49%	1.60%
Bounds benchmarked against Dem. vote share (county)		
1× benchmark	−0.062*	−0.067**
	$t = -3.183$	$t = -3.377$
2× benchmark	−0.046*	−0.052*
	$t = -2.385$	$t = -2.615$
3× benchmark	−0.031	−0.037+
	$t = -1.585$	$t = -1.852$

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

Notes: Sensitivity analysis, benchmarking against county-level Democratic vote share. Column (1) uses a standard `lm` specification with state fixed effects entered as a covariate. Column (2) approximates the fixed effects specification by demeaning all variables by state prior to estimation, which is numerically equivalent to the within-state transformation performed by `feols`. The robustness value (RV) indicates the minimum percentage of residual variance in both the treatment and outcome that an unobserved confounder would need to explain to nullify the point estimate. Bounds show the adjusted estimate under confounders of 1×, 2×, and 3× the explanatory power of county-level partisanship.

A.4 Compositional Effects

Table A9: Association between Workplace Composition and Union Support

	(1)	(2)	(3)
	Baseline	Non-white prop.	Race interaction
Blau index	−0.081**		−0.137*
	(0.027)		(0.068)
Non-white proportion		−0.028+	
		(0.016)	
Blau × White			0.065
			(0.065)
White	−0.049***	−0.053***	−0.083*
	(0.012)	(0.012)	(0.039)
Male	−0.079***	−0.078***	−0.079***
	(0.009)	(0.009)	(0.009)
Hourly wage	−0.002*	−0.003**	−0.002*
	(0.001)	(0.001)	(0.001)
Age	−0.023***	−0.023***	−0.023***
	(0.003)	(0.003)	(0.003)
Supervisor coethnicity	−0.027**	−0.025*	−0.027**
	(0.010)	(0.010)	(0.010)
Coworker support	0.057***	0.058***	0.057***
	(0.005)	(0.005)	(0.005)
Dem. vote share (county)	0.135***	0.129***	0.134***
	(0.026)	(0.027)	(0.026)
Fixed effects	State	State	State
Observations	18 142	18 142	18 142
Adjusted R^2	0.029	0.028	0.029

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Standard errors clustered by firm. State fixed effects included. Column (1) reproduces the main specification for reference. Column (2) replaces the Blau index with the proportion of non-white workers. Column (3) estimates the main specification with an interaction between the Blau index and respondent race (White).

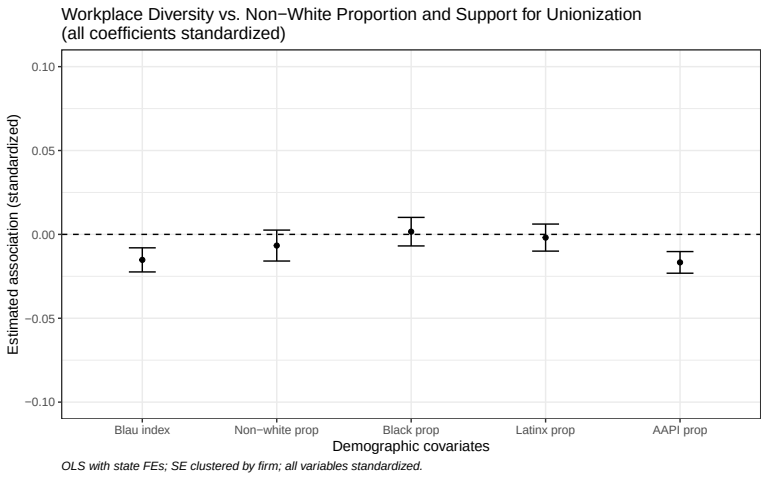


Figure A2: Coefficient plot: association between workplace composition (by group) and union support.

A.5 Monotonicity

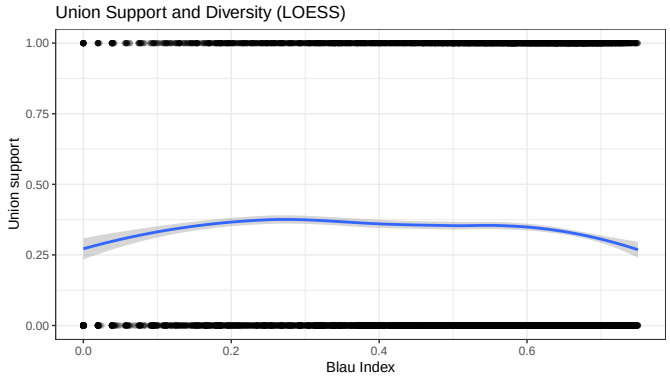


Figure A3: Bivariate relationship between Blau index and union support (LOESS smoother).

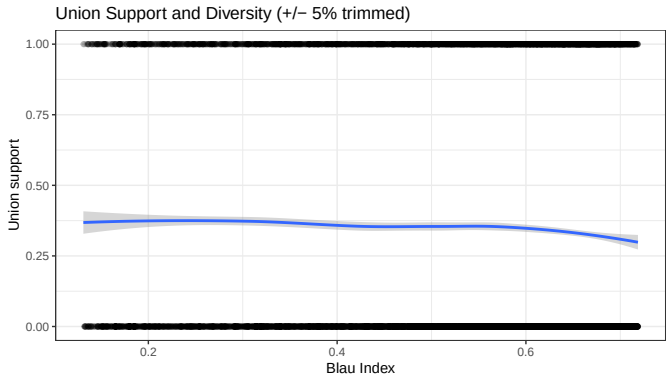


Figure A4: Bivariate relationship between Blau index and union support, top and bottom 5% trimmed (LOESS smoother).

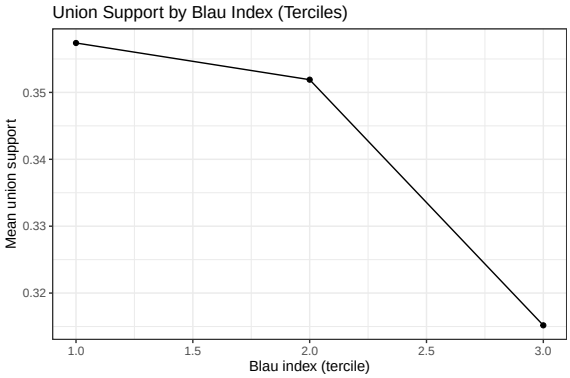


Figure A5: Mean union support by Blau index tercile.

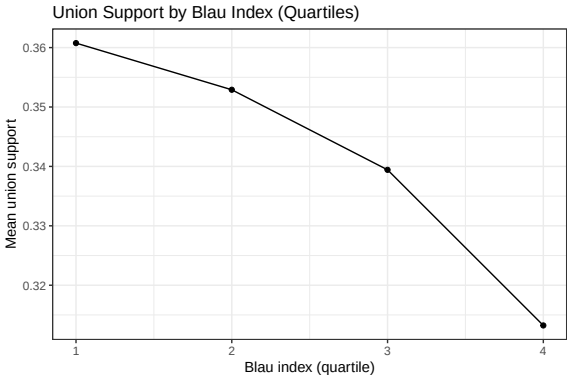


Figure A6: Mean union support by Blau index quartile.

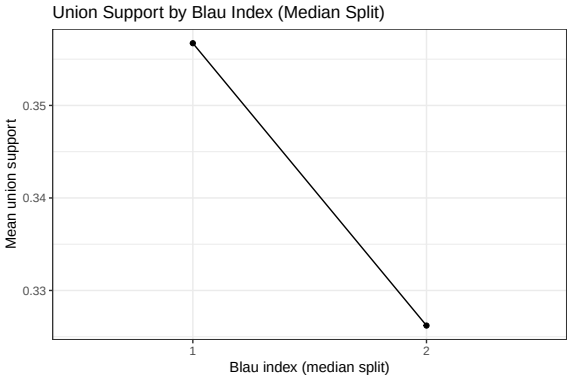


Figure A7: Mean union support by Blau index median split.

A.6 Diversity and Firm Size

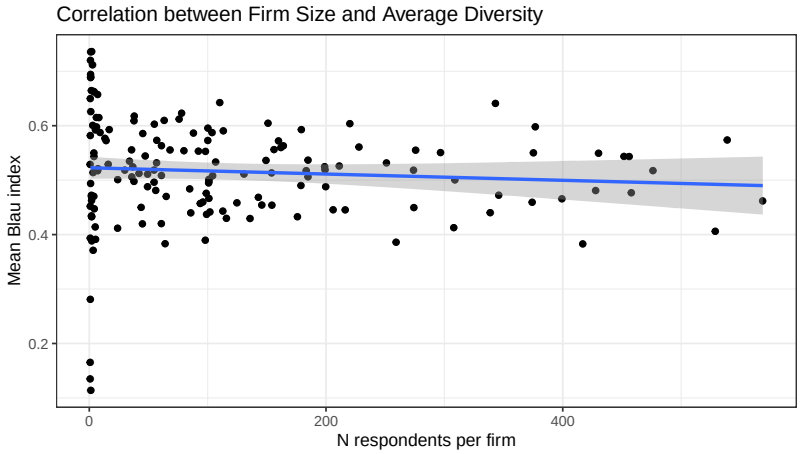


Figure A8: Correlation between firm size (number of respondents) and average Blau index score. Pearson's $r = -0.080$, $p = 0.316$.

B Study 2

B.1 Race Imputation

The race of individual voters was imputed using a four-step process, seen in Table 1.

First, Xie (2022) Bidirectional Long Short-Term Memory (Bi-LSTM) natural language processing methodology, which classifies race/ethnicity based on surname and first name only and adjusts race probabilities taking into account name base rate discrepancies between white and non-white populations, was used. Then, race was further refined using McCartan (2025) Bayesian Improved Surname Geocoding (BISG) race imputation methodology. Race probabilities were imputed using individuals' surnames and Census demographic information based on their residential zip code when the latter was available on voter lists. Zip codes were recorded in the majority of cases, for 2319 out of 2701 voters. Thirty-five out of forty-one employers/elections had at least all or some portion of voters with recorded zip codes.[17] Next, race classifications between Bi-LSTM and BISG for voters with zip codes were compared. Last, for cases in which race classification discrepancies were identified, race was coded by hand.

[17] Residential zipcode missingness occurred in a minority of cases, with only 382 out of 2701 such cases. Missingness occurred due to such information not being recorded by union organizers at the time of organizing. It is not clear why this information was not recorded. In some cases, zip code was missing for an entire bargaining unit. In other cases, zip code was missing for certain individuals but recorded for others within the same discrete bargaining unit.

Table B10: Race Imputation Procedure

	Race Imputation Procedure	Sample	Prop. of Sample	# of Voters
Step 1	Bi-LSTM*	All voters	100%	2701
Step 2	BISG†	Voters with zipcodes	0.859	2319
Step 3	Bi-LSTM* and BISG†compared	Voters with zipcodes	0.859	2319
Step 4	Race Handcoded	Bi-LSTM and BISG mismatch cases	0.40[18]	917

* Surname and First Name

†Surname and Residential Zip Code

The handcoding process then involved two separate sub-steps. Using the 1990 Census of male and female names, surnames and first names were classified as "very common," "common," "rare," or "not found." For cases in which (a) the surname and first name were identified as very common or common, the BISG classification was used. For cases in which (b) the surname was identified as rare or not found and the first name was very common or common; (c) the surname was identified as very common or common and the first name was rare or not found; and (d) both the surname and first name were rare or not found, the Bi-LSTM classification was used.

To further illustrate this procedure, Table **??** provides four examples of the name-case typologies, the imputation procedure deemed to produce the most accurate classification, the logic behind this determination, and the race classification produced by the BISG and Bi-LSTM methods. For example, "Tryphena Cooper Mitchell" has a common, ethnically non-distinctive and racially ambiguous surname, which may increase the likelihood of misclassification if using only surname and geography to predict race, limiting the predictive utility of BISG. However, the relatively distinctive first name ("not found") may make Bi-LSTM better suited to accurately infer racial classification in this instance, and was thus chosen as the preferred classification method.

As a robustness check, a random sample of 100 names was selected and purely coded by hand, then compared to both the Bi-LSTM and BISG classifications.

Table B11: Handcoded Imputation Procedure Logic

Last Name	First Name	Preferred Imputation Procedure	Determination Logic	BISG Race	Bi-LSTM Race
Cooper Mitchell	Tryphena	Bi-LSTM*	Rare/NF first name; Common last name	White	**Black**
Brown	Grace	BISG†	Common first name; Common last name	**White**	Black
Brown	Mahjuanne	Bi-LSTM	Rare/NF first name; Common last name	White	**Black**
Hardowar	Naseema	Bi-LSTM	Rare/NF first name; Rare/NF last name	White	**Asian**

* Surname and First Name

†Surname and Residential Zip Code

Xie (2022) was the preferred methodology for initial race imputation given a "training" dataset that was included amongst the Excelsior list records. One particular list included voters' surnames, first names, and individual ethnoraces recorded by organizers at the time of the election. However, residential zip codes were not recorded. Both Xie (2022) and McCartan (2025) race probability methodologies were applied to this list. Xie's methodology (using first name and surname only) was accurate in 72% of cases, whereas McCartan et al.'s BISG methodology (using surname and the national-level surname distribution given the lack of residential geographic information) was accurate in only 28% of cases.

B.2 Voter File Matching Procedure

To measure individual-level party identification for workers in the Excelsior list data, I matched workers to their records in state voter files obtained from Statara Solutions. Voter files contain party registration, vote history, and demographic information for all registered voters in a given state. The matching procedure was conducted separately for each state and followed a tiered approach designed to maximize match rate while minimizing false positives.

Data Preparation

Prior to matching, worker names were standardized to uppercase and residential zip codes were zero-padded to five digits to ensure consistency with voter file formatting. Where residential addresses were missing but could be extracted from a combined address string, state and city were parsed using regular expression pattern matching. For workers with no residential geographic information, workplace zip codes were substituted as a fallback. A geographic source flag was

created to track whether each worker's zip code came from their residential address, was extracted from an address string, or was imputed from their workplace location.

Matching Tiers

Matching proceeded in up to four tiers of decreasing certainty. Workers matched in an earlier tier were excluded from subsequent tiers to avoid duplication.

T1 matched workers to voter file records on exact first name, last name, and zip code. In states where residential addresses were available in both the worker list and the voter file (Minnesota), an additional address verification step was applied: matches where the street number and street name did not match between the two sources were flagged for manual review and excluded from the confirmed match set. In states without reliable residential addresses in the worker list (Connecticut), a date-of-hire plausibility filter was applied instead, retaining only matches where the worker's implied age at hire was 65 or younger.

T2 recovered additional matches among workers who did not match in T1 by searching for cases where the worker's full name appeared exactly once in the voter file within the same zip code. These matches were retained without an age or address filter given the uniqueness of the name-zip combination.

T2a extended the search statewide, retaining workers whose full name appeared exactly once anywhere in the state voter file regardless of zip code. Because zip codes in the worker list sometimes reflected workplace rather than residential location, many true matches were recoverable through statewide name uniqueness even when zip matching failed. To account for the greater uncertainty of these matches, each was assigned a confidence score based on name rarity using the 1990 Census of names. Surnames and first names were classified as very common (national rank ≤ 100), common (rank $\leq 1{,}000$), rare (rank $> 1{,}000$), or not found. Confidence was assigned as follows: matches with a rare or not-found surname were classified as high confidence; matches with a common surname but rare or not-found first name were classified as medium-high confidence; matches with common surnames and common first names were classified as low confidence. Where available, an address match between the worker list and voter file boosted confidence by one level, and an address mismatch reduced it by one level.

T2b (Minnesota only) recovered workers whose name appeared multiple times statewide but

whose street number and street name matched exactly one voter file record. These matches were classified as high confidence given the specificity of the address match.

Confidence Scoring and Sensitivity Analysis

Each matched record was assigned one of four confidence levels: high, medium-high, low, or flagged for review. The main analysis uses all confirmed matches. As a robustness check, models are reestimated using only high and medium-high confidence matches (the conservative sample), excluding low confidence matches where name rarity provides limited disambiguation.

Match Rates and Unmatched Workers

Table B12: Voter File Matching Results by State

State	Total Workers	T1	T2a	T2b	Total Matched	Match Rate
CT	2,053	134	523	—	657	32.0%
MN	601	247	139	2	388	64.6%
MA	27	manual verification			20	74.1%
RI	—	—	—	—	—	—
WI	—	—	—	—	—	—
NY	—	—	—	—	—	—
Total	**2,681**	**381**	**662**	**2**	**1,065**	**39.7%**

Notes: T1 = exact name + ZIP code match (+ age filter for CT); T2a = unique name statewide match with Census name rarity confidence scoring; T2b = address match for remaining unmatched workers (MN only). MA matches were manually verified. Workers in NY, RI, and WI were too few to warrant state-specific matching and voter file matching was not conducted for these states. Match rates are computed as matched workers divided by total workers in the state. Overall match rate computed over CT, MN, and MA only.

Workers who could not be matched to the voter file fall into two categories: those whose name did not appear anywhere in the state voter file (either unregistered or registered under a different name spelling), and those whose name appeared multiple times statewide but could not be uniquely identified without additional disambiguating information. For unmatched workers, zip code-level Democratic presidential vote share from the MIT Election Data and Science Lab serves as a contextual proxy for partisan lean, following the approach used in Study 1. All models control for a binary indicator of whether the worker was successfully matched to the voter file, to account for the possibility that registered voters differ systematically from unregistered workers in their propensity to support unionization.

B.3 Bayesian Imputation of Partisan Identification

To control for the partisan composition of bargaining units and contact groups, I construct a Bayesian posterior estimate of the probability that each worker in the sample identifies as a Democrat. Because workers are not surveyed directly, partisan identification (PID) must be imputed from observable characteristics. I follow the approach of Brown and Enos (2021), combining a demographic prior derived from survey data with a geographic prior derived from electoral returns.

Demographic Prior

The demographic prior captures the conditional probability of Democratic identification given a worker's race and gender: $P(\text{Democrat} \mid \text{race, gender})$. I estimate this prior using the 2020 Cooperative Election Study (CES), restricting to voter-validated respondents (N = 43,067) to ensure the sample reflects the registered electorate (Schaffner et al., 2021).

PID is coded to include partisan leaners: respondents who identify as pure Democrat (pid3 = 1) or who identify as Independent but lean Democratic (pid3 = 3, pid7 = 3) are coded as Democrat; respondents who identify as pure Republican or lean Republican are coded as Republican; and true Independents (pid3 = 3, pid7 = 4) are coded as Independent. Respondents with other or missing PID responses are excluded.

Race is coded into five categories (white, Black, Hispanic, Asian, other) and gender into two (male, female), yielding ten demographic cells. For each cell, I compute the proportion of respondents falling into each PID category. The resulting lookup table, shown in Table B13, reflects well-established patterns in the literature: Black women have the highest probability of Democratic identification (0.885), while white men have the lowest (0.448).

Geographic Prior

The geographic prior captures the conditional probability of Democratic identification given the partisan context of the worker's county of residence: $P(\text{Democrat} \mid \text{county})$. I operationalize this using the 2020 presidential election results from the MIT Election Data and Science Lab precinct-level returns (MIT Election Data and Science Lab, 2021), aggregated to the county level. Democratic vote share is computed as Biden votes divided by total votes cast. For workers in Minnesota, I

Table B13: Demographic Prior: $P(\text{Democrat} \mid \text{race, gender})$

Race	Gender	P(Democrat)	P(Republican)	P(Independent)
Asian	Female	0.711	0.186	0.103
Asian	Male	0.614	0.215	0.172
Black	Female	0.885	0.049	0.066
Black	Male	0.782	0.096	0.122
Hispanic	Female	0.670	0.231	0.099
Hispanic	Male	0.543	0.316	0.141
Other	Female	0.552	0.276	0.172
Other	Male	0.379	0.403	0.218
White	Female	0.515	0.396	0.089
White	Male	0.448	0.436	0.116

Notes: Computed from voter-validated CES 2020 respondents (N = 43,067). Democrat includes pure Democrats and Democratic leaners; Republican includes pure Republicans and Republican leaners.

account for the state Democratic Party's official name — Democratic Farmer Labor (DFL) — which is coded as "other" in the simplified party variable; I recode DFL votes as Democratic prior to aggregation.

Workers are linked to counties via their residential ZIP code, using a ZIP-to-county crosswalk from the `zipcodeR` package (Rozzi, 2021). For the small number of workers whose ZIP codes could not be matched to a county, I substitute national averages (48.8% Democratic, 46.1% Republican) as fallback priors.

Bayesian Updating

Following Brown and Enos (2021), I combine the demographic and geographic priors using Bayes' rule to obtain a posterior probability of Democratic identification for each worker. The numerator for each PID category is the product of the demographic likelihood and the geographic prior:

$$P(D \mid \text{race, gender, county}) \propto P(\text{Democrat} \mid \text{race, gender}) \times P(\text{Democrat} \mid \text{county})$$

The posterior probabilities for Democrat, Republican, and Independent are normalized to sum to one:

$$P(\text{Democrat})_i = \frac{P_{\text{dem}}^{\text{demo}} \cdot P_{\text{dem}}^{\text{geo}}}{P_{\text{dem}}^{\text{demo}} \cdot P_{\text{dem}}^{\text{geo}} + P_{\text{rep}}^{\text{demo}} \cdot P_{\text{rep}}^{\text{geo}} + P_{\text{ind}}^{\text{demo}} \cdot P_{\text{ind}}^{\text{geo}}}$$

This yields three worker-level variables: $\hat{P}(\text{Democrat})_i$, $\hat{P}(\text{Republican})_i$, and $\hat{P}(\text{Independent})_i$. In the analysis I use $\hat{P}(\text{Democrat})_i$ as the primary measure of individual partisan lean, and construct a net partisanship measure $\hat{P}(\text{Democrat})_i - \hat{P}(\text{Republican})_i$. At the contact group and employer level, I aggregate to the mean posterior probability of Democratic identification within the group.

Validation

I validate the imputed probabilities against observed party registration for the subset of workers matched to voter files in Connecticut and Massachusetts (N = 595 high- and medium-confidence matches; Minnesota is excluded because the state does not record party registration). Among observed registered Democrats, the mean posterior $\hat{P}(\text{Democrat})$ is 0.806. Among registered Republicans it is 0.709, and among registered Unaffiliated voters it is 0.771. The correlation between $\hat{P}(\text{Democrat})$ and observed Democratic registration (binary) is 0.165.

The modest separation between groups — particularly the high mean $\hat{P}(\text{Democrat})$ among registered Republicans — reflects the heavily Democratic geographic and demographic context of healthcare organizing in Connecticut and Massachusetts rather than model misspecification. Workers in this sample are concentrated in majority-Democratic counties, and many are members of demographic groups with high baseline probabilities of Democratic identification. As a result, even workers who register as Republican tend to live and work in environments where the prior strongly favors Democratic identification. Inspection of the 55 registered Republicans in the validation sample confirms that misprediction is concentrated among non-white workers in heavily Democratic counties — precisely the cases where the demographic and geographic priors are most likely to overwhelm the registration signal. The model correctly ranks partisan groups in the expected direction and the imputation is used only as a control variable rather than a primary quantity of interest, so this limitation does not affect the main conclusions.

B.4 Distribution of Key Variables in Sample

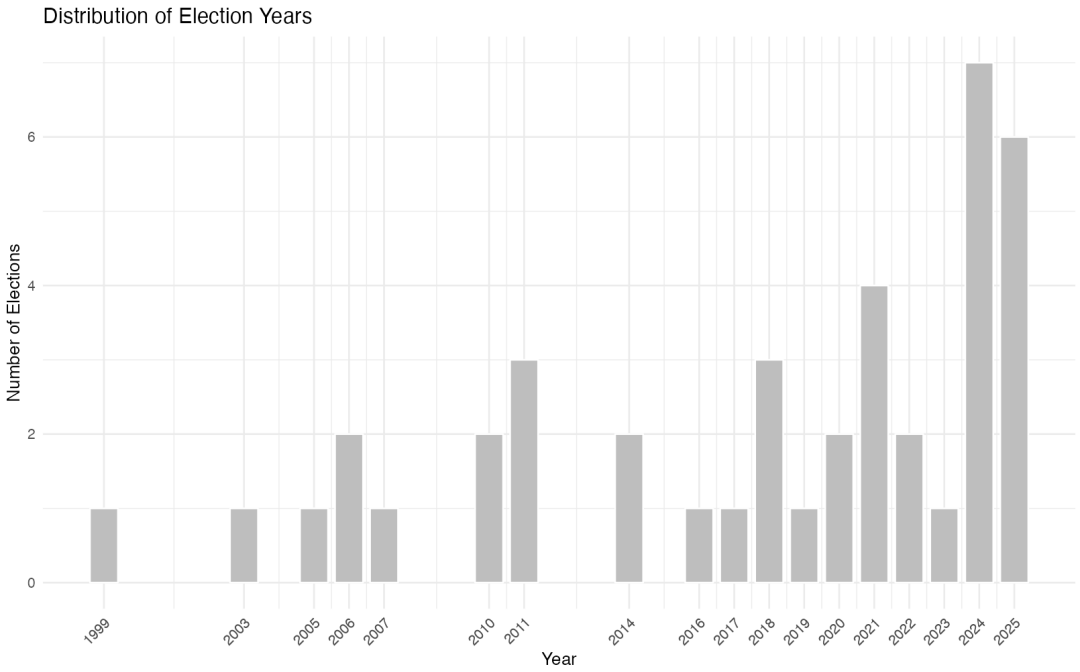


Figure B9

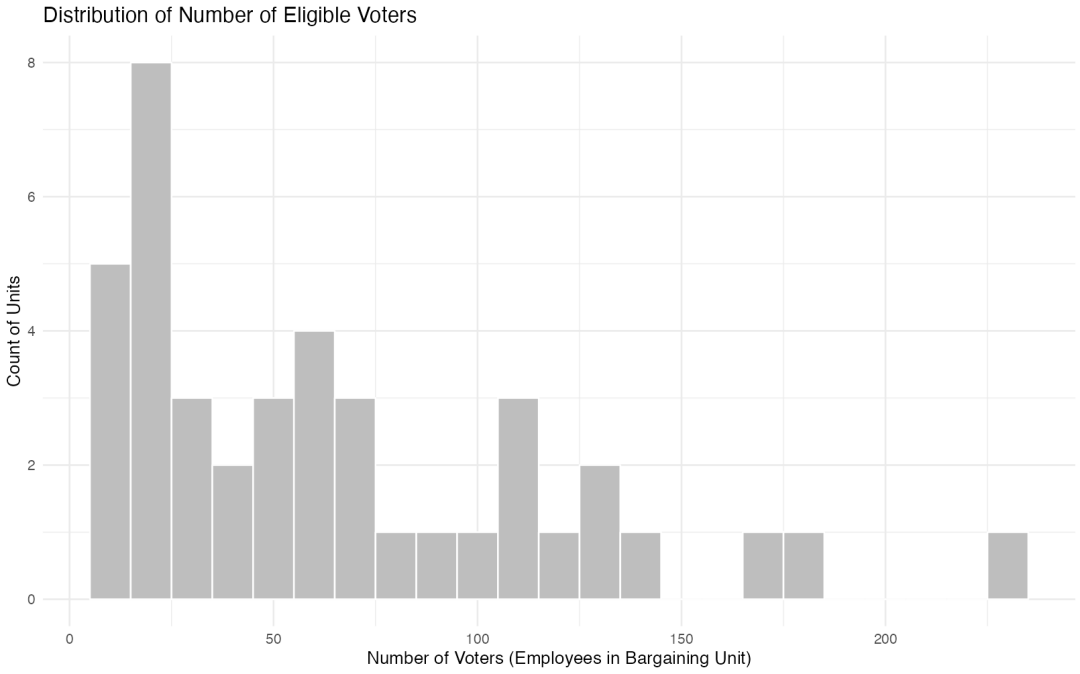


Figure B10

B.5 Alternative Model Specifications: Employer Level

Table B14: Alternative Specifications: Bargaining-Unit Diversity and Election Margins

Model	**Blau Coefficient**	**SE**	**Adj.** R^2
Bivariate, unweighted	-1.337^{*}	(0.664)	0.022
Bivariate, weighted	-0.976	(1.290)	-0.011
Terciles, unweighted	—	—	-0.007
Terciles, weighted	—	—	0.002
Full covariates, unweighted	-1.123	(0.791)	0.241
Observations	41		

Notes: Each row reports the coefficient on the Blau index from a separate regression of voteshare margin. Full covariates include high POC share, high Black share (median split), log unit size, occupational heterogeneity dummy, and shift heterogeneity dummy. Robust (HC2) standard errors in parentheses. Tercile specifications yield null results across both weighted and unweighted specifications. The preferred specification uses tercile indicators with full covariates and is reported as Table 1 in the main text; the continuous Blau coefficient is null across all weighted specifications. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.6 Heterogeneity by Unit Size

To probe whether the relationship between racial diversity and union support varies by unit size, I stratify the sample by median voter count. In small elections (number of voters < 54), racial diversity is significantly associated with lower voteshare margins ($\beta = -0.845$, $p = 0.046$), suggesting greater election contestation. No such association is observed in larger elections ($\beta = -0.140$, $p = 0.901$). However, interacting diversity with unit size yields a non-significant association ($p = 0.55$), indicating that the difference in effect across unit size is not statistically meaningful. These results suggest that the apparent negative relationship between diversity and union support may be concentrated in smaller-scale elections, but the evidence does not meet conventional thresholds for a moderated effect.

Table B15: Relationship Between Racial Diversity and Union Vote Margin by Unit Size

Model	**Coef.**	**SE**	p**-value**	N	**Adj.** R^2
(1) Continuous Blau, unweighted	−1.337	(0.664)	0.051	41	0.022
(2) Continuous Blau, weighted	−0.976	(1.290)	0.454	41	−0.011
(3) Blau Terciles, unweighted	—	—	—	41	−0.007
(4) Blau Terciles, weighted	—	—	—	41	0.002
(5) Small elections ($<$ 54 voters)	−0.845	(0.394)	0.046	20	−0.004
(6) Large elections ($\geq$ 54 voters)	−0.140	(1.107)	0.901	21	−0.052

Notes: Dependent variable is union voteshare margin. Robust (HC2) standard errors in parentheses. Tercile specifications yield null results. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B.7 Main Models: Contact Group and Unit Level

B.8 Contact Group Models with Controls

B.9 Pro-Union Share: Null Association with Contact-Group Diversity

Table B18 reports models in which the outcome is the proportion of workers in each contact group assessed as pro-union (including union leaders). Diversity is associated with greater opinion disagreement (entropy), but not with a lower mean level of pro-union support. The null result for the Medium and High diversity bins across all specifications indicates that diversity does not systematically reduce the share of workers favorable to unionization at the contact group level.

Table B16: Racial Diversity and Opinion Disagreement: Contact Group and Unit Level

	Contact Group Level				Unit Level	
	(1) Entropy	(2) Entropy	(3) Entropy	(4) Entropy	(5) Entropy	(6) Pro %
Intercept	-0.085^{***}	0.050^{***}	$0.081^{\dagger}$	0.050	0.495	132.324^{*}
	(0.021)	(0.013)	(0.048)	(0.046)	(0.700)	(50.329)
Key Predictors						
Micro Blau (continuous)	0.522^{***}					
	(0.052)					
Blau Bin: Medium		0.111^{***}	0.006	0.113^{***}	0.060	24.114
		(0.025)	(0.086)	(0.028)	(0.180)	(15.756)
Blau Bin: High		0.253^{***}	0.271^{**}	0.220^{***}	0.788^{***}	57.031^{***}
		(0.028)	(0.093)	(0.032)	(0.157)	(8.996)
Controls						
POC Median (local)			0.004	0.023		
			(0.033)	(0.027)		
Black Median (local)			−0.028	−0.037		
			(0.028)	(0.026)		
Group size (n)				0.009^{*}		
				(0.004)		
Medium × POC			$0.083^{\dagger}$			
			(0.056)			
High × POC			−0.009			
			(0.057)			
POC High (unit)					-0.526^{*}	−31.172
					(0.227)	(19.066)
Black High (unit)					0.042	−2.781
					(0.113)	(8.502)
Log voters (unit)					0.016	−18.752
					(0.158)	(10.874)
Observations	427	427	427	427	17	17
R^2	0.134	0.153	0.163	0.191	0.559	0.497
Adj. R^2	0.132	0.149	0.151	0.181	0.359	0.268

Notes: Cols. (1)–(4): contact group-level models with CR2 standard errors clustered by employer ($N = 17$ employers). Cols. (5)–(6): unit-level models weighted by number of voters with HC2 standard errors. Dependent variable in cols. (1)–(5) is normalized Shannon entropy (0 = unanimity, 1 = maximal disagreement). Col. (6) uses percent pro-union (including leaders) as the outcome. Reference categories: Low Blau bin; low POC median; low Black median. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table B17: Contact Group Diversity and Opinion Disagreement with Compositional Controls

	(1) Bivariate	(2) With composition controls	(3) With PID controls
Intercept	0.050^{***}	0.050	$-0.104^{\dagger}$
	(0.013)	(0.046)	(0.071)
Blau Bin: Medium	0.111^{***}	0.113^{***}	0.114^{***}
	(0.025)	(0.028)	(0.029)
Blau Bin: High	0.253^{***}	0.220^{***}	0.220^{***}
	(0.028)	(0.032)	(0.031)
POC Median (local)		0.023	−0.008
		(0.027)	(0.031)
Black Median (local)		−0.037	-0.054^{*}
		(0.026)	(0.026)
Group size (n)		0.009^{*}	0.009^{*}
		(0.004)	(0.004)
Mean $\hat{P}$(Democrat)			0.273^{*}
			(0.117)
Pct. matched to voter file			0.027
			(0.024)
Observations	427	427	427
R^2	0.153	0.191	0.201
Adj. R^2	0.149	0.181	0.187

Notes: CR2 standard errors clustered at the employer level ($N = 17$ employers) in parentheses. Dependent variable is normalized Shannon entropy. Column (1) reports the bivariate association between diversity terciles and entropy. Columns (2) and (3) add compositional and partisan controls respectively. Reference category is Low diversity. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table B18: Contact-Group Diversity and Pro-Union Share: Null Results

	(1) Bivariate	(2) + Composition	(3) + PID
Intercept	0.593***	0.548***	0.394**
	(0.032)	(0.064)	(0.130)
Blau Bin: Medium	−0.062	−0.059	−0.058
	(0.047)	(0.048)	(0.050)
Blau Bin: High	−0.041	−0.031	−0.026
	(0.049)	(0.053)	(0.053)
POC Median (local)		0.012	−0.001
		(0.041)	(0.044)
Black Median (local)		−0.002	−0.018
		(0.039)	(0.040)
Group size (n)		−0.003	−0.003
		(0.005)	(0.005)
Mean $\hat{P}$(Democrat)			0.285
			(0.178)
Pct. matched			0.014
			(0.032)
Observations	427	427	427
Adj. R^2	−0.002	−0.008	−0.007

Notes: Outcome is the proportion of contact-group members assessed as pro-union (excluding "leader" assessments). CR2 standard errors clustered at the employer level in parentheses. No diversity coefficient approaches conventional significance across any specification, confirming that diversity produces opinion fracture—more undecided workers—rather than a decline in mean support. Reference category is Low diversity. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

B.10 Opinion Composition by Diversity Bin

Table B19 reports mean pro-union, undecided, and anti-union shares by diversity tercile, computed at the contact-group level. The pattern is consistent with opinion fracture rather than generalized demobilization: the anti-union share is essentially flat across diversity levels, while the undecided share increases monotonically. Pro-union share shows a corresponding directional decline, though neither the undecided nor pro-union gradient reaches conventional significance (Table B20).

Table B19: Mean Opinion Composition by Contact-Group Diversity Bin

Diversity Bin	Pro-Union	Undecided	Anti-Union
Low	0.593	0.173	0.234
Medium	0.531	0.239	0.230
High	0.552	0.226	0.222

Notes: Cell entries are mean proportions across contact groups ($N = 427$) within each diversity tercile. Proportions sum to 1 within each row (rounding to three decimal places). Workers assessed as "leader" are excluded. Pro-union includes workers assessed as "Pro"; undecided includes "Undecided" and "Soft Pro"; anti-union includes "Anti" and "Soft Anti."

Table B20: Contact-Group Diversity and Opinion Composition: Undecided and Anti-Union Share

	Undecided Share		Anti-Union Share	
	(1) Bivariate	**(2)** + Controls	**(3)** Bivariate	**(4)** + Controls
Intercept	0.173^{***}	0.201^{***}	0.234^{***}	0.251^{***}
	(0.024)	(0.050)	(0.022)	(0.047)
Blau Bin: Medium	0.067	0.065	−0.005	−0.006
	(0.047)	(0.050)	(0.044)	(0.046)
Blau Bin: High	0.054	0.048	−0.012	−0.005
	(0.068)	(0.072)	(0.065)	(0.070)
POC Median		−0.016		0.004
		(0.037)		(0.034)
Black Median		0.018		0.000
		(0.034)		(0.031)
Group size (n)		−0.006		$0.009^{\dagger}$
		(0.005)		(0.005)
Observations	427	427	427	427
Adj. R^2	−0.002	−0.007	−0.005	−0.004

Notes: CR2 standard errors clustered at the employer level in parentheses. Neither undecided nor anti-union share exhibits a statistically significant gradient with diversity, though undecided share trends directionally upward (Medium +6.7 pp, High +5.4 pp). The anti-union share is flat (< 1 pp across bins). Reference category is Low diversity. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

B.11 Individual-Level Fixed Effects Models

Table B21 presents individual-level models of workers' union assessment regressed on their bargaining unit's diversity tercile, with fixed effects for department, shift, and job title. The fixed effects absorb occupational and organizational confounders, isolating the within-occupation-department-shift association between diversity and individual assessment. Standard errors are clustered at the employer level.

Critically, the sign of the diversity coefficient reverses relative to the bivariate cross-sectional association. In the raw data, workers in more diverse units are somewhat less likely to be assessed as pro-union, consistent with the compositional pattern described in Table 1 of the main text. But within the same job type, department, and shift, workers in higher-diversity contact groups are *more* likely to be assessed as pro-union. The Blau tercile 3 coefficient (-0.215 on a five-point scale, where lower values correspond to more pro-union) indicates that, holding occupation and organizational context constant, high diversity is associated with stronger individual pro-union orientation. This sign reversal suggests that the observed negative correlation between diversity and election margins at the employer level operates primarily through occupational composition—diverse units tend to employ occupational mixes with lower baseline pro-union rates—rather than through a direct effect of racial heterogeneity on individual support.

B.12 Diversity–Disagreement Relationship: Shape of the Association

B.13 Employer-Level Diversity and Election Margins: Group Composition

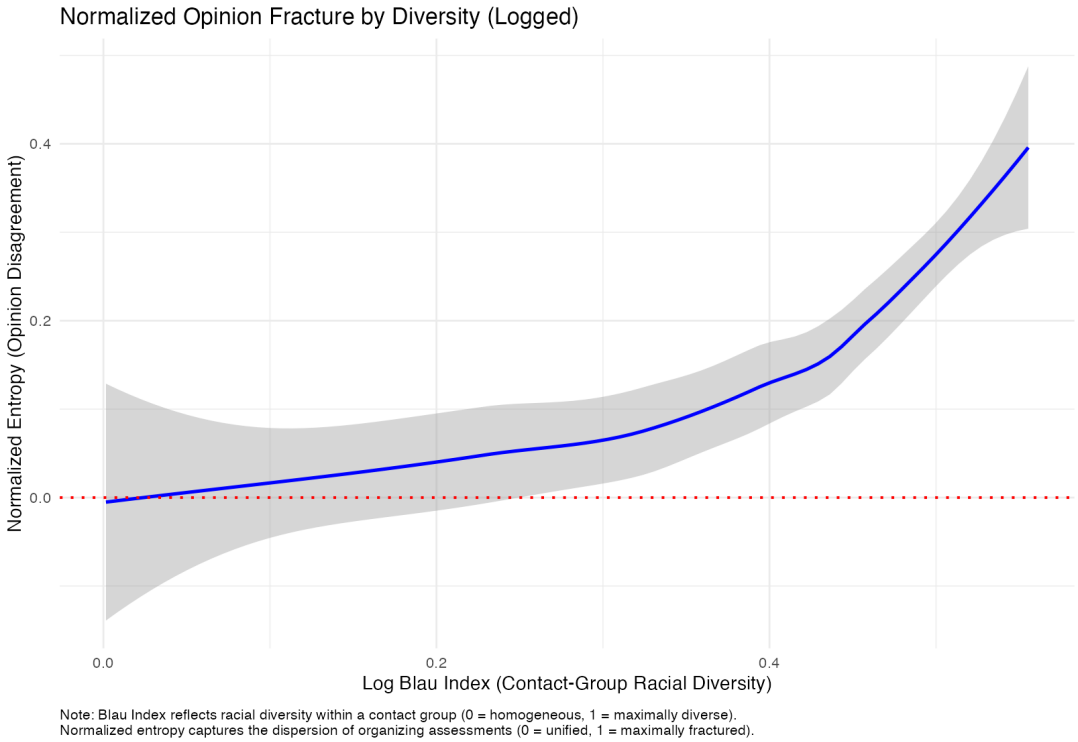


Figure B11: LOESS fit of the relationship between contact-group racial diversity (Blau index) and normalized opinion entropy. The association is monotonically increasing across the full range of the Blau index, with the increment from low to medium diversity (approximately +0.11) similar in magnitude to the increment from medium to high (+0.14). The pattern does not exhibit a sharp threshold effect; rather, greater diversity is associated with higher opinion disagreement throughout.

Table B21: Individual-Level Assessment and Bargaining-Unit Diversity (Fixed Effects)

	(1) Bivariate	(2) + FE	(3) + FE + PID
Blau Tercile 2 (Medium)	$0.097^{\dagger}$	-0.085	-0.081
	(0.055)	(0.076)	(0.074)
Blau Tercile 3 (High)	0.183^{**}	-0.215^{***}	-0.212^{***}
	(0.067)	(0.052)	(0.051)
$\hat{P}$(Democrat)			-0.183^{*}
			(0.082)
Dept. × Shift × Job FE	No	Yes	Yes
Employer clustering	No	Yes	Yes
Observations	538	538	538

Notes: Outcome is the five-point union assessment scale (1 = Pro, 5 = Anti; lower values indicate greater pro-union orientation). Models (2) and (3) include fixed effects for department, shift, and job title, estimated via `feols` (Fixest). Standard errors clustered at the employer level ($N = 17$). The positive bivariate coefficient on Blau Tercile 3 (column 1) reflects occupational composition: diverse units employ occupational mixes with lower baseline pro-union rates. Within the same occupation, department, and shift, high-diversity contact groups are associated with stronger pro-union individual assessments (column 2). Reference category is Low diversity tercile. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

B.14 Unit-Level Diversity and Election Margins: Scatter Plot

B.15 Shannon Entropy

Shannon entropy is calculated as

$$-\sum_{i=1}^{K} p_i log(p_i)$$

where p_i = proportion of members in category i.

To further illustrate this, consider the following toy example:

This yields $n_{total} = 5$; $n_{groups} = 3$. To calculate p_i, we have

From here,

$$Shannon_{entropy} = -\sum_{i=1}^{k} p_i log(p_i)$$

$$= -(0.4(log(0.4)) + 0.2(log(0.2)) + 0.4(log(0.4)))$$

Table B22: Bargaining-Unit Diversity, Racial Composition, and Election Margins

	(1) Main effects	(2) Diversity × POC interaction
Intercept	1.078***	1.168***
	(0.230)	(0.340)
Blau Tercile 2 (Medium)	0.065	−0.065
	(0.175)	(0.198)
Blau Tercile 3 (High)	0.102	0.240
	(0.165)	(0.160)
High POC Share	0.356*	0.341*
	(0.137)	(0.151)
High Black Share	0.061	0.059
	(0.131)	(0.137)
Log(Unit Size)	−0.177**	−0.116
	(0.064)	(0.089)
Occupational Dummy	−0.109	−0.108
	(0.139)	(0.135)
Shift Dummy	−0.269†	−0.264*
	(0.142)	(0.128)
Medium × High POC		0.036
		(0.243)
High × High POC		−0.621**
		(0.210)
Observations	41	41
R^2	0.281	0.418
Adj. R^2	0.128	0.249

Notes: Dependent variable is union voteshare margin. Linear regression weighted by number of eligible voters. Robust (HC2) standard errors in parentheses. Reference category is Low diversity tercile. $^{\dagger}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

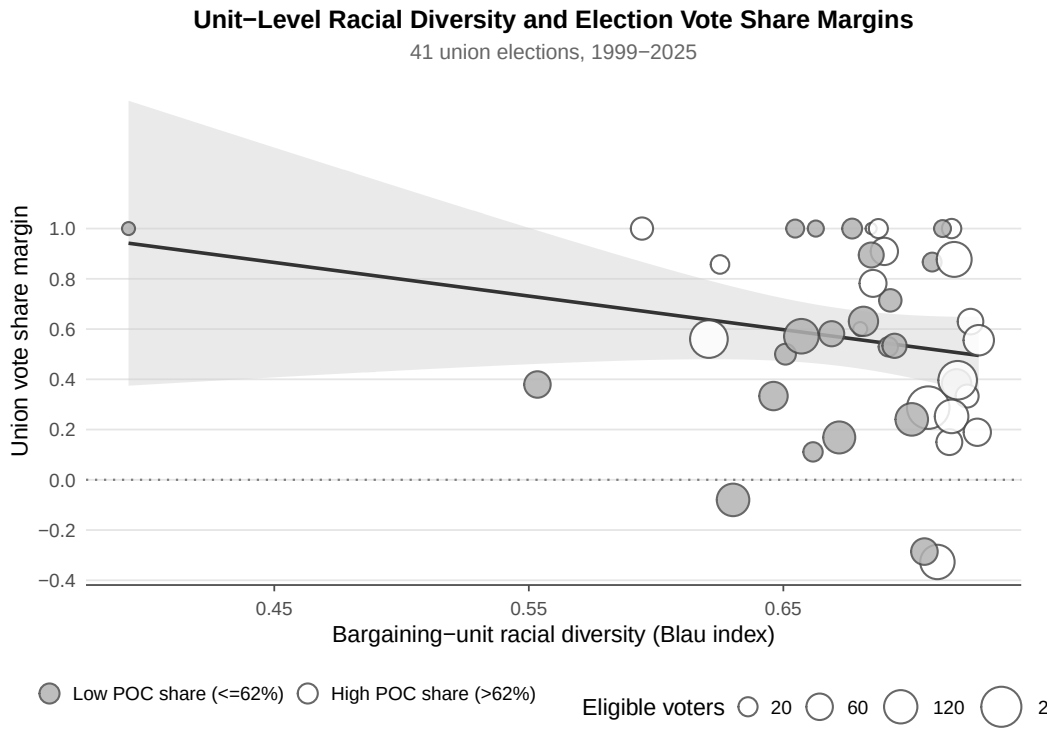


Figure B12: Bargaining-unit racial diversity and union vote share margins. Each point represents one bargaining unit ($N = 41$). Point size reflects the number of eligible voters. Filled circles indicate units with low non-White share ($\leq 62\%$); hollow circles indicate units with high non-White share ($> 62\%$). Line shows OLS fit with 95% confidence band.

Employer Name	Dept.	Job	Shift	Assessment
Allina Hospital	1 South	LPN	AM	Pro-organizing
Allina Hospital	1 South	LPN	AM	Pro-organizing
Allina Hospital	1 South	LPN	AM	Anti-organizing
Allina Hospital	1 South	LPN	AM	Undecided
Allina Hospital	1 South	LPN	AM	Undecided

Table B23: Hypothetical Example

Group	N	p_i	Outcome
Pro	2	2/5	0.4
Anti	1	1/5	0.2
Undecided	2	2/5	0.4

Table B24

$= -(0.4(-0.916) + 0.2(-1.609) + 0.4(-0.916))$

$= -(-1.0546) = 1.0546$

Finally, normalizing this value:

entropy/log(n all possible groups)

$= 1.0546/log(4) = 0.761,$

a substantial degree of disagreement within the contact-group.

B.15.1 Alternative measures of opinion dispersion

Table B25: Correlation between normalized Blau and normalized Shannon entropy (assessment dispersion at the contact–group level)

Statistic	Estimate	p-value
Pearson r	0.992 111	< 0.001
Spearman ρ	0.997 024	< 0.001
N (contact groups)[a]	*[432]*	

[a] N equals the number of contact groups with non-missing dispersion measures in the analytic sample. Both measures are normalized to $[0, 1]$ using the fixed set of assessment categories ($K = 4$).

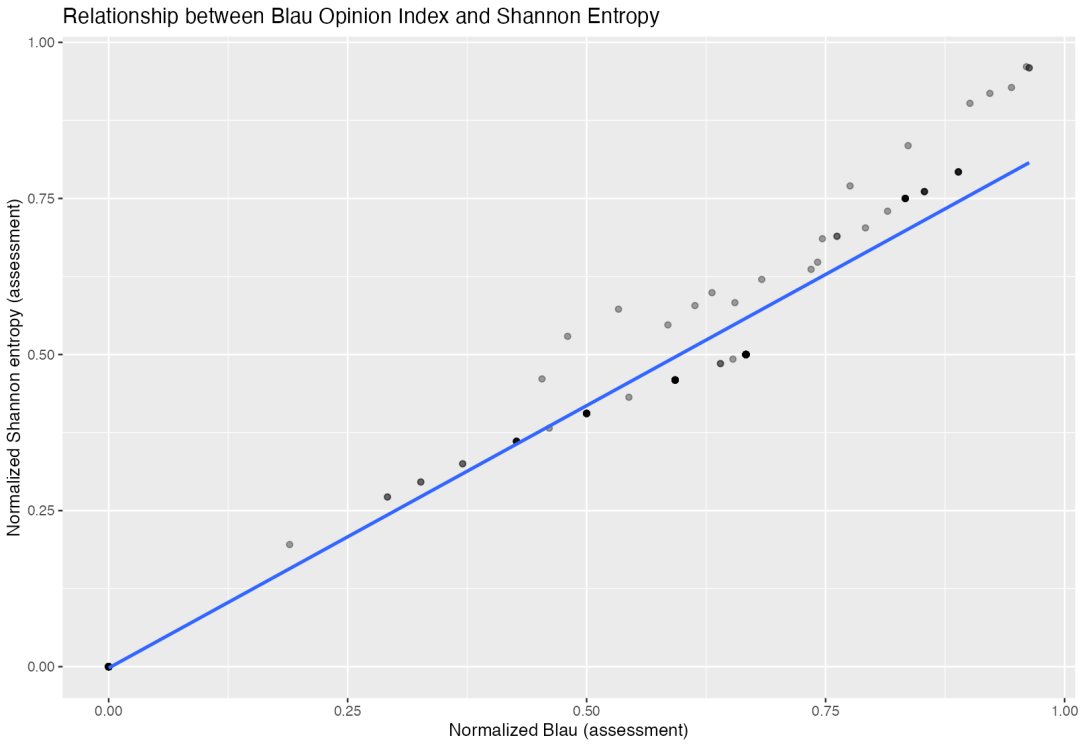


Figure B13

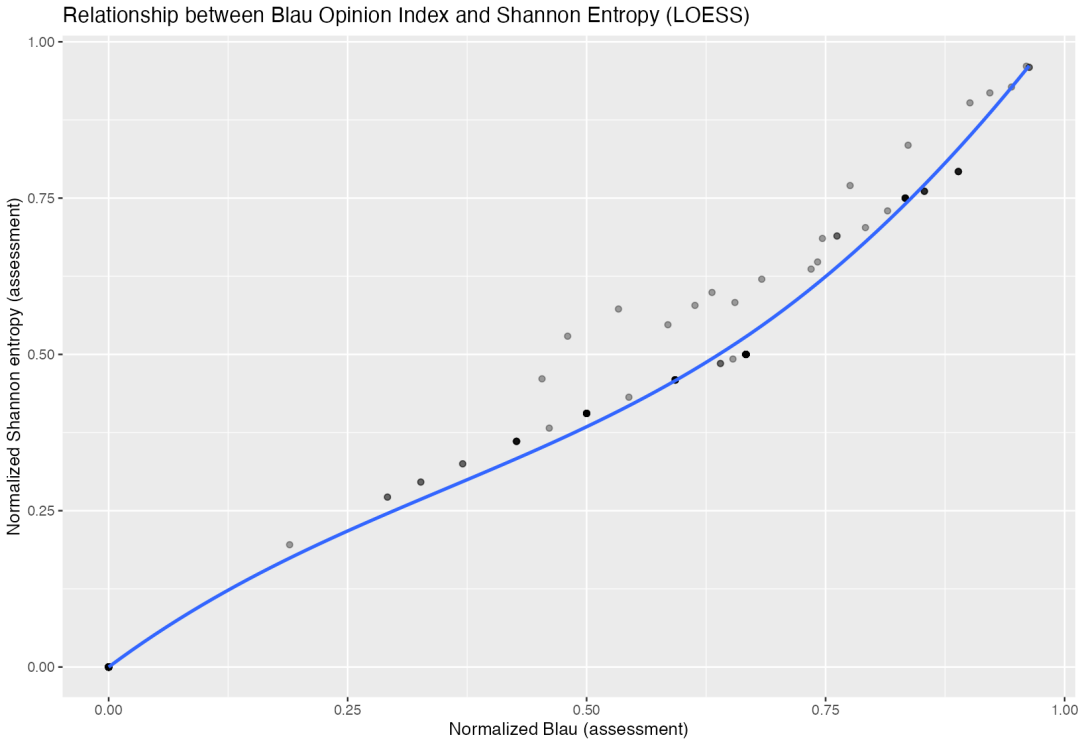

Relationship between Blau Opinion Index and Shannon Entropy (LOESS)
Normalized Shannon entropy (assessment)
1.00
0.75
0.50
0.25
0.00
0.00
0.25
0.50
0.75
1.00
Normalized Blau (assessment)

Figure B14

C Study 3

C.1 Sample Characteristics

Table C26: Sample Descriptive Statistics (Study 3)

	Full Sample (N = 697)
Demographics	
Age	41.5 (12.3)
Female	68.1%
White	82.3%
Education	
High school or less	7.2%
Some college	24.8%
College degree	35.6%
Postgraduate	19.4%
Income	
Household income (mean, $1000s)	68.4
Employment	
Full-time	72.1%
Part-time	24.6%
Union Status	
Current member	15.0%
Former member	11.7%
Union household	26.8%
Political Affiliation	
Democrat	38.6%
Republican	31.4%
Independent	25.3%
Baseline Preferences	
Chose collective bargaining	79.5%

C.2 Full Regression Results

Table C27: Effect of Team Composition on Alignment with Risk-Seeking Teammate

	Round 3		**Round 5**	
	(1)	(2)	(3)	(4)
Heterogeneous Team	−0.453***	−0.464***	0.208	0.206
	(0.175)	(0.179)	(0.174)	(0.178)
Female		−0.030		−0.160
		(0.193)		(0.194)
Black		−0.064		−0.353
		(0.242)		(0.234)
Union Member		−0.329		0.716**
		(0.291)		(0.323)
Republican		−0.014		−0.378*
		(0.221)		(0.213)
Independent		−0.099		0.205
		(0.232)		(0.237)
Other/None PID		0.303		0.939*
		(0.489)		(0.564)
Some College		0.236		0.339
		(0.269)		(0.267)
College Degree		−0.133		−0.012
		(0.241)		(0.241)
Postgraduate		0.038		−0.421
		(0.294)		(0.284)
Income ($1000s)		−0.000		0.001
		(0.002)		(0.002)
Constant	1.126***	1.217***	0.708***	0.807***
	(0.126)	(0.311)	(0.123)	(0.297)
N	647	625	626	625

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$

Note: Logistic regression coefficients with standard errors in parentheses.

Reference categories: Democrat (PID), High school or less (education).

Table C28: Effect of Co-ethnicity on Alignment (Heterogeneous Teams Only)

	(1)	(2)
Risk-Seeker is Co-ethnic	−0.038	−0.121
	(0.241)	(0.252)
Female		−0.101
		(0.273)
Union Member		−0.459
		(0.389)
Republican		0.316
		(0.308)
Independent		−0.447
		(0.312)
Other/None PID		0.638
		(0.823)
Some College		0.045
		(0.371)
College Degree		−0.114
		(0.335)
Postgraduate		−0.295
		(0.406)
Income ($1000s)		0.000
		(0.003)
Constant	0.693***	0.914**
	(0.171)	(0.399)
N	308	297

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

Note: Logistic regression coefficients with standard errors in parentheses.

Sample restricted to participants in heterogeneous teams.

C.3 Text Analysis: Keyword Categories

Table C29: Keyword Categories, Definitions, and Illustrative Examples

Category	Keywords	Example (Positive)	Example (Negative)
Bold/Aggressive	bold, aggressive, ambitious, pushy, aim high, go big, start high, risky	"Fortune favors the bold"	"Tre's approach was coming on way too strong and too ambitious"
Cautious/Safe	cautious, careful, conservative, safe, realistic, reasonable, modest	"I tend to be very careful and cautious when making decisions"	—
Compromise	compromise, middle ground, flexible, willing to, open to	"She was willing to compromise whereas Tyra wasn't"	—
Confident	confident, assertive, firm, sure, certain	"She was confident in her approach"	—
Agreeable	agreed, same page, similar, aligned, like me	"We were on the same page from the start"	—
Positive/Warm	friendly, positive, kind, nice, enthusiastic	"He was so friendly and kind"	—
Clear/Logical	sense, logical, smart, wise, rational	"His approach made the most sense to me"	—
More Money	more money, higher amount, maximum, most, biggest	"I wanted to aim for more money"	—
Fair	fair, equal, everyone, win-win	"It was a fair amount"	"Rather than shrewd and greedy"

Note: Examples drawn from Round 3 responses. "Reasonable" classified under Cautious/Safe, not Fair.

C.4 Stability of Bargaining Preferences

Table C30: Stability of Bargaining Preferences (Round 1 to Round 2)

Pattern	**N**	**%**
Stay Collective	539	78.3
Stay Individual	82	11.9
Individual → Collective	57	8.3
Collective → Individual	10	1.5

N = 688.

C.5 Experimental Materials

C.6 Avatar Selection

Participants selected an avatar and name to represent them throughout the experiment. Figures C15 and C16 present the avatar options available to female and male participants, respectively.

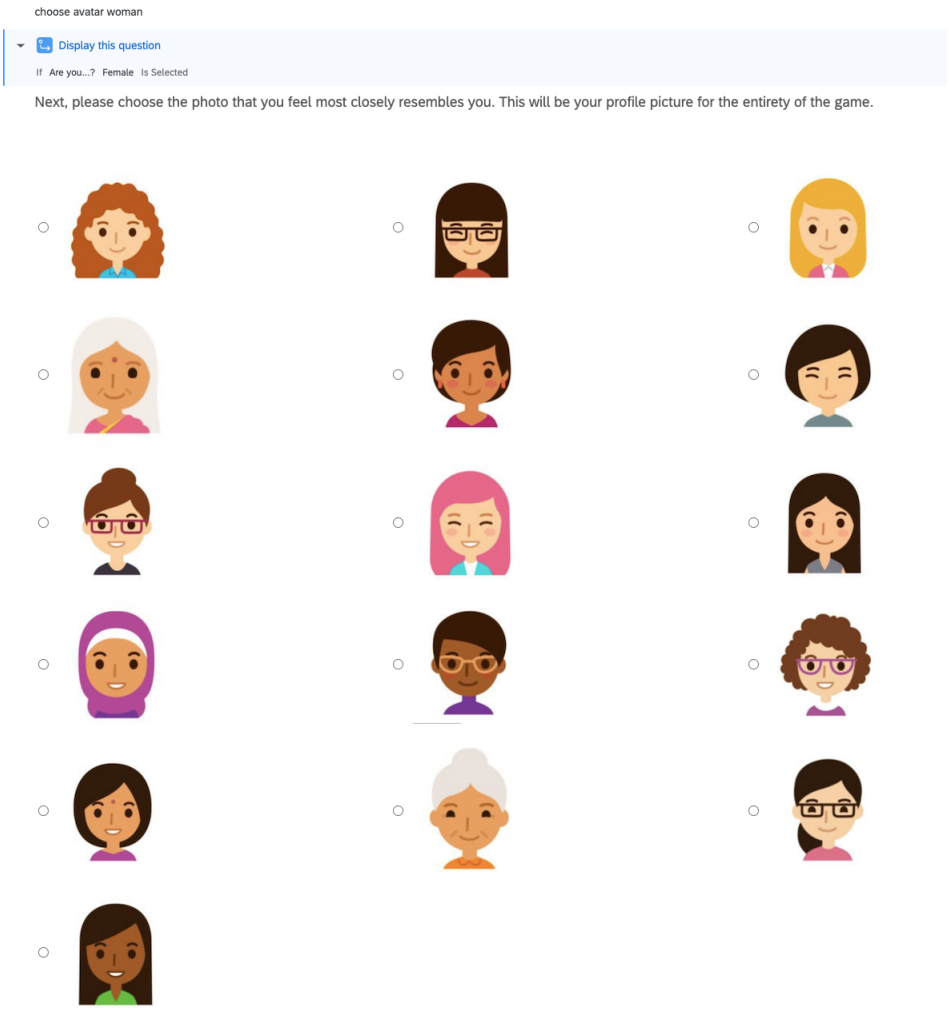


Figure C15: Avatar Selection (Female Participants)

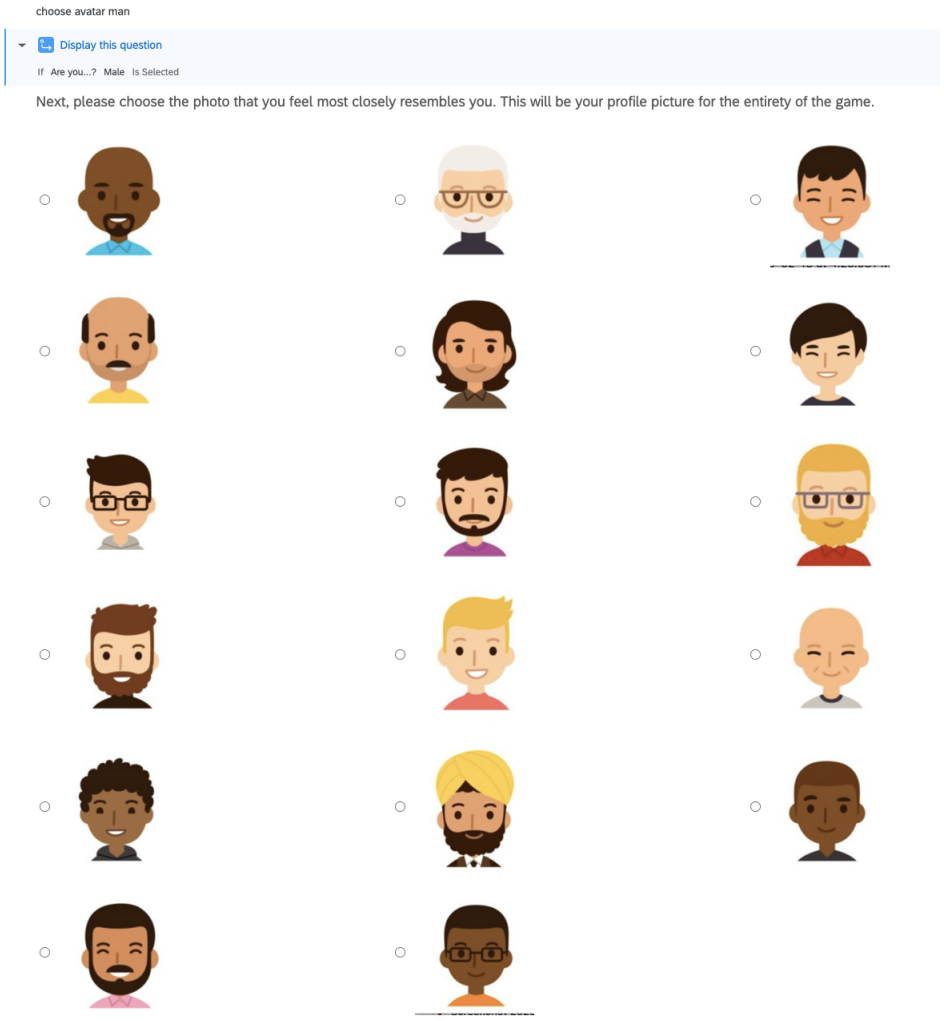


Figure C16: Avatar Selection (Male Participants)

C.7 Balance and Randomization Checks

C.7.1 Covariate Balance

Table C31 presents covariate balance across treatment conditions for participants who completed Round 3. Randomization was successful: no pre-treatment characteristic differed significantly between participants assigned to homogeneous versus heterogeneous teams.

Table C31: Covariate Balance by Treatment Condition

	Homogeneous (n = 339)	**Heterogeneous** (n = 308)	**p-value**
Demographics			
Age (mean)	41.0	41.5	.67
Female (%)	68.7	67.5	.81
Black (%)	18.6	15.9	.43
Education			
High school or less (%)	10.0	12.3	—
Some college (%)	21.5	21.1	—
College degree (%)	32.7	30.8	—
Postgraduate (%)	20.4	20.1	.64
Income			
Household income (mean, $1000s)	70.7	64.9	.12
Union Membership			
Current member (%)	9.7	11.8	.48
Political Affiliation			
Democrat (%)	41.7	44.5	.53
Republican (%)	29.3	29.5	—
Independent (%)	23.7	22.7	—
Baseline Preferences			
Chose collective bargaining (R1, %)	80.2	80.5	1.00

Note: P-values from t-tests (continuous) or chi-square tests (categorical).

Education p-value tests college degree or higher vs. less than college.

C.7.2 Attrition

Table C32 presents attrition rates by treatment condition. Of the 697 participants who began the experiment, 647 (93%) completed Round 3. Attrition did not differ by treatment assignment.

Table C32: Attrition by Treatment Condition

	Started	**Completed R3**	**% Completed**
Homogeneous	367	339	92.4%
Heterogeneous	330	308	93.3%

$\chi^2 = 0.12$, $p = .73$

C.8 Manipulation Checks

C.8.1 Gender Identification

Participants were asked to identify the gender of their teammates after Round 2, when avatars and names were first revealed. Table C33 presents accuracy by participant gender. Women were more accurate than men (99.6% vs. 86.7%), likely reflecting the salience of same-gender teammates.

Table C33: Gender Identification Accuracy by Participant Gender

Participant Gender	N	% Correct
Female	453	99.6%
Male	211	86.7%

C.8.2 Perceived Resemblance

Participants were asked how much they resembled their teammates. Participants in heterogeneous teams reported significantly lower perceived resemblance ($p = .002$), confirming that the treatment was received. Table C34 presents the full distribution of responses.

Table C34: Perceived Resemblance to Teammates by Treatment Condition

	Homogeneous		Heterogeneous	
Resemblance	N	%	N	%
Not at all	68	19.5%	77	24.4%
Somewhat	218	62.6%	210	66.7%
Very much	62	17.8%	28	8.9%

Homogeneous mean = 1.98 (SD = 0.61); Heterogeneous mean = 1.84 (SD = 0.56).
$p = .002$ (t-test).

C.9 Robustness Checks

Table C35 presents treatment effect estimates across alternative specifications. The main finding is robust to linear probability models, exclusion of post-treatment controls, and heteroskedasticity-robust standard errors.

Table C35: Robustness of Main Treatment Effect (Round 3)

Specification	**Treatment β**	**p-value**
Logit (base)	−0.453	.010
Logit (with demographic controls)	−0.464	.010
Logit (pre-treatment controls only)	−0.452	.010
Logit (robust standard errors)	−0.453	.010
Linear probability model (base)	−0.093	.009
Linear probability model (with controls)	−0.095	.010

Note: Demographic controls include gender, race, union membership, partisanship, education, and income. Pre-treatment controls exclude union membership and income. Outcome: alignment with risk-seeking teammate (1 = risk-seeker, 0 = risk-averse).

C.10 Heterogeneous Treatment Effects

Table C36 presents tests for heterogeneous treatment effects by participant characteristics. No subgroup showed a significantly different treatment effect, suggesting that the main finding generalizes across demographic and attitudinal subgroups.

Table C36: Heterogeneous Treatment Effects: Interaction Tests

Moderator	**Interaction β**	**p-value**
Female	−0.195	.60
Baseline collective preference (R1)	−0.099	.81
Union member	−0.209	.71
Democrat	0.031	.93

Note: Each row represents a separate logistic regression model.

Table C37 presents subgroup means. While no interaction term reached statistical significance, the treatment effect appears somewhat larger among Independents (−24.1 pp) and union members (−14.8 pp).

Table C37: Subgroup Means: Percent Aligning with Risk-Seeking Teammate (Round 3)

Moderator	Subgroup	Homog.	Heterog.	Diff.	N
Gender	Male	74.5%	68.0%	−6.5	206
	Female	76.0%	65.4%	−10.6	441
Baseline (R1)	Individual	68.7%	60.0%	−8.7	127
	Collective	77.2%	67.7%	−9.5	520
Union	Non-member	76.1%	67.2%	−8.9	559
	Member	71.9%	57.1%	−14.8	67
Partisanship	Democrat	75.4%	66.4%	−9.0	279
	Republican	70.7%	71.4%	+0.7	190
	Independent	81.2%	57.1%	−24.1	150
	Other	77.8%	80.0%	+2.2	28

C.11 Order Effects and Crossover Design

The within-subjects design allows examination of whether the treatment effect persists after participants experience both conditions. As discussed in the main text, strong carryover effects complicate interpretation of the crossover rounds.

C.11.1 Round 5 Outcomes by Prior Exposure

Table C38 presents Round 5 alignment by treatment sequence. Participants who experienced heterogeneous teams first showed 67.0% alignment with the risk-seeker in Round 5 (now in homogeneous teams); those who experienced homogeneous teams first showed 71.4% alignment (now in heterogeneous teams). This reversal suggests that participants developed a strategic orientation in their first deliberation and maintained it thereafter.

Table C38: Round 5 Alignment by Treatment Sequence

Round 3 Condition	Round 5 Condition	N	% Risk-Seeker (R5)
Heterogeneous	Homogeneous	297	67.0%
Homogeneous	Heterogeneous	329	71.4%

Treatment effect at Round 5: $\beta = 0.21$, $p = .23$

C.11.2 Carryover Effects

Controlling for Round 3 alignment, the Round 5 treatment effect is fully attenuated ($\beta = 0.04$, $p = .84$), while Round 3 alignment strongly predicts Round 5 alignment ($\beta = 1.73$, $p < .001$).

C.11.3 Within-Person Change

Table C39 presents within-person change in alignment across the crossover.

Table C39: Within-Person Change in Alignment Across Crossover

Treatment Sequence	N	R3 (%)	R5 (%)	Change
Heterogeneous → Homogeneous	297	66.0	67.0	+1.0
Homogeneous → Heterogeneous	329	75.7	71.4	−4.3

Note: Change = Round 5 − Round 3 alignment with risk-seeking teammate.

C.12 Text Analysis: Additional Details

C.12.1 Keyword Classification

Table C40 presents the keyword categories, constituent keywords, and representative examples used to code open-ended responses. Categories were determined through a manual read-through of Round 3 responses.

Table C40: Keyword Categories, Definitions, and Illustrative Examples

Category	Keywords	Example (Positive)	Example (Negative)
Bold/Aggressive	bold, aggressive, ambitious, pushy, aim high, go big, start high, risky	"Fortune favors the bold"	"Tre's approach was coming on way too strong and too ambitious"
Cautious/Safe	cautious, careful, conservative, safe, realistic, reasonable, modest	"I tend to be very careful and cautious when making decisions"	—
Compromise	compromise, middle ground, flexible, willing to, open to	"She was willing to compromise whereas Tyra wasn't"	—
Confident	confident, assertive, firm, sure, certain	"She was confident in her approach"	—
Agreeable	agreed, same page, similar, aligned, like me	"We were on the same page from the start"	—
Positive/Warm	friendly, positive, kind, nice, enthusiastic	"He was so friendly and kind"	—
Clear/Logical	sense, logical, smart, wise, rational	"His approach made the most sense to me"	—
More Money	more money, higher amount, maximum, most, biggest	"I wanted to aim for more money"	—
Fair	fair, equal, everyone, win-win	"It was a fair amount"	"Rather than shrewd and greedy"

Note: Examples drawn from Round 3 responses. "Reasonable" classified under Cautious/Safe, not Fair.

C.12.2 Response Rates and Lengths

Table C41 presents response rates and median response lengths by treatment condition and round. Response rates were high across all rounds (92–99%) and did not differ by treatment.

Table C41: Open-Ended Response Rates and Lengths by Treatment

	Response Rate (%)			Median Length (chars)		
Condition	R1	R2	R3	R1	R2	R3
Homogeneous	98.1	95.1	92.4	65	50	54
Heterogeneous	99.4	95.5	93.3	67	54	55

C.12.3 Demographic Mentions in Open-Ended Responses

Explicit mentions of demographic characteristics were rare across all rounds. Table C42 presents gender mention rates in Round 2 explanations. Approximately 4.5% of responses mentioned gender, almost exclusively in positive terms about all-female teams.

Table C42: Gender Mentions in Round 2 Explanations by Treatment

Condition	**N**	**N Mentioning Gender**	**%**
Homogeneous	349	13	3.7%
Heterogeneous	315	19	6.0%

C.12.4 Keyword Frequencies Across Rounds

Cautious language emerged primarily in Round 3, after deliberation, rather than in Rounds 1 or 2. Table C43 presents keyword frequencies across rounds, showing that strategy-related terms were rare before participants engaged with teammates' arguments.

Table C43: Keyword Frequencies Across Rounds

Category	Round 1 (Baseline)	Round 2 (Team Reveal)	Round 3 (Post-Chat)
Group words	73.7%	70.3%	47.4%
Individual words	68.9%	55.9%	38.8%
Bold/Aggressive	0.1%	0.6%	24.3%
Cautious/Safe	0.9%	0.8%	12.8%
Compromise	1.0%	0.9%	13.1%
Agreeable	0.6%	1.4%	16.5%

Note: N = 688 (R1), 664 (R2), 647 (R3).

C.12.5 Race Mentions

In Round 3, 3.2% of responses were flagged by automated keyword search as containing racial terms. Manual review revealed that only one response contained explicit racial reasoning ("He's White," from a participant in the homogeneous condition). No participant in the heterogeneous condition articulated race as a factor in their coalition choice. The rarity of explicit racial reasoning, combined with the significant treatment effect on cautious language, suggests that diversity operates through implicit channels rather than overt demographic preferences.

C.13 Stability of Bargaining Preferences

Participants' preferences for collective versus individual bargaining were highly stable from Round 1 to Round 2. Table C44 presents summary transition rates; Table C45 presents the full transition matrix. Switching rates did not differ by treatment condition: 9.2% of participants in homogeneous teams switched their choice compared to 10.4% in heterogeneous teams ($p = .61$).

Table C44: Stability of Bargaining Preferences (Round 1 to Round 2)

Pattern	**N**	**%**
Stay Collective	539	78.3
Stay Individual	82	11.9
Individual → Collective	57	8.3
Collective → Individual	10	1.5

N = 688.

Table C45: Transition Matrix: Round 1 to Round 2 Bargaining Preferences

	Round 2 Choice		
Round 1 Choice	Individual	Collective	**Total**
Individual	82 (11.9%)	57 (8.3%)	139 (20.2%)
Collective	10 (1.5%)	539 (78.3%)	549 (79.8%)
Total	92 (13.4%)	596 (86.6%)	688 (100%)

Note: Percentages of full sample (N = 688 with valid R1 and R2 responses).

D Research Transparency, Ethics, and AI Use

D.1 Pre-Registration and Pre-Analysis Plan

The bargaining experiment (Study 3) was pre-registered prior to data collection. The pre-analysis plan is publicly available on OSF at the following anonymous link for review: `https://osf.io/mc7as/overview?view_only=b533a636100f464da3433c9175ddbfa2`. The pre-registration specifies the experimental design, primary outcome, and planned analyses. The main text reports results consistent with the pre-registered plan. The text analysis of open-ended responses (Section C.12) was not pre-registered and should be interpreted as exploratory.

D.2 Ethical Considerations

The experimental portion of this project received exemption from the university Institutional Review Board. Participants in Study 3 were recruited through Prolific, an online research platform, and provided informed consent prior to participation. Participants were compensated at a fair market rate (pro-rated at $8.40 per 60 minutes) and average completion time was 18 minutes. Participants were fully debriefed about the use of AI confederates following completion of the experiment.

The Study 3 sample was restricted to non-Hispanic white, working-age U.S. adults employed (or job-seeking) in food service, healthcare, retail, or hospitality. This restriction was made to ensure internal validity of the racial composition manipulation — heterogeneous teams were constructed by pairing white participants with Black AI confederates, requiring a racially homogeneous participant pool for the manipulation to function as designed. As a result, the sample does not reflect the demographic diversity of the sample studied in Study 2 but is similar to that of Study 1. Generalizability to non-white participants and other industries may be limited. The research does not differentially harm any participant group and all participants faced the same decision task and compensation structure regardless of treatment condition.

The survey data (Study 1) are drawn from The Shift Project, a non-probability opt-in sample of service-sector workers recruited via social media. Participation was voluntary. The administrative data (Study 2) were shared by union locals based on a mutual understanding that the data was

confidential and would not be shared. Worker identities are protected and no individually identifying information is reported.

D.3 Data Availability and Reappraisability

Replication code and the Study 1 survey data will be deposited in the APSR Dataverse upon conditional acceptance. The Study 2 administrative data (Excelsior lists and organizing charts) were shared by union locals under a shared expectation of confidentiality and cannot be publicly deposited. This prohibits redistribution of individually identifiable worker records. The Study 3 experimental data will be deposited in full. Researchers interested in replicating the Study 2 analyses may contact the author to discuss access arrangements that would not breach confidentiality.

Reappraisal of this work faces two primary constraints. First, the Study 2 organizing records are proprietary data shared by specific union locals and cannot be replicated without comparable data-sharing arrangements with labor organizations. Second, the Study 3 experiment was programmed on SoPHIELabs, a proprietary interactive experiment platform. Exact replication of the chat interface would require access to this platform. The study design, stimuli, and analysis code will be fully documented to enable as close a replication as possible under these constraints.

D.4 Use of Artificial Intelligence Tools

Experiment platform. Study 3 used scripted AI confederates fine-tuned with ChatGPT 3.5 to simulate real-time interaction with participants. The AI confederates delivered pre-written messages advocating for risk-seeking or risk-averse bargaining strategies. The specific message content and assignment of risk orientation to confederate avatars are described in the main text and Appendix C.5.

Data processing and analysis assistance. The author used Claude (Anthropic) for assistance with data cleaning, code development, and manuscript preparation. Versions used include Claude 3.5/3.7 Sonnet (Anthropic, summer 2025) and Claude Sonnet 4.6 (Anthropic, spring–summer 2026). Claude was accessed via the Claude Code interface (`https://claude.ai/code`). All AI-assisted outputs were reviewed, verified, and are the sole responsibility of the author.