

Experimental Evidence on the Complementarity of Union Membership

Morgan Foy*

June 2025

Abstract

This paper examines whether an individual's union membership status is causally related to the perceived membership rate in their workplace. In an information provision experiment in Wisconsin, I randomize teachers into receiving an estimate of the union membership rate in their school district. I find that treated teachers substantially update their prior beliefs toward the provided estimate. Likewise, membership exhibits complementarity — treated teachers adjust their planned union participation in accordance with their revised beliefs. My estimates suggest that a 10 percentage point change in beliefs about the district membership rate leads to a 3.7 percentage point shift in planned participation.

*University of California, Berkeley, Haas School of Business (email: morgan_foy@berkeley.edu). For helpful comments and suggestions, I thank Sydnee Caldwell, Ernesto Dal Bó, Miguel Ortiz, Ricardo Perez-Truglia, Reed Walker, and Guo Xu. I gratefully acknowledge financial assistance from the Institute for Research on Labor and Employment and Xlab at UC-Berkeley. This research was approved by the UC-Berkeley IRB office (2022-04-15229) and registered in the AEA RCT Registry (0013327).

I. Introduction

There has been a long legal debate about whether workers in union-covered workplaces should be required to pay union dues. In the US, about half the states have right-to-work laws prohibiting compulsory union dues and likewise the Supreme Court has ruled that mandatory dues in the public sector violate the First Amendment (*Janus v. AFSCME* 2018). The argument in favor of compulsory mechanisms centers on the free-rider problem: workers will shirk on paying membership dues even if they benefit from and are in favor of unionization (Olson 1965).

However, union membership has not necessarily disappeared in right-to-work environments. This raises the question of how unions sustain membership despite the free-rider incentive. One explanation is that social customs play a role: by participating in the union, workers gain a reputational benefit from their colleagues (Booth 1985, Naylor and Cripps 1993). Despite its importance, there is little empirical work that studies whether social mechanisms — and in particular whether *others'* participation — affect individual membership decisions.

In this study, I ask whether individuals' union membership status is affected by an exogenous shock to their beliefs about what share of their co-workers are union members. The social customs model suggests that individual and group participation decisions should be complementary: if the group membership rate increases, then individuals will be more likely to be union members and vice versa. Conversely, recent experimental work on collective action problems related to protest participation and election campaigning find evidence of strategic substitutes (Cantoni et al. 2019; Jarke-Neuert, Perino, and Schwickert 2023; Hager et al. 2023). In these cases, people were *less likely* to participate when more members of a group participated.

Union membership is a core example of a collective action problem. Yet, there is little empirical work examining whether group membership rates, such as within a firm or school district, affect individual membership decisions because of two key challenges. First, it requires individual-level data on union membership within specific organizations. Second, even with detailed membership data, it is difficult to isolate exogenous group shocks that do not also affect an individual by definition. For instance, quasi-experimental variation in co-workers' membership status, e.g., due to a policy change, is also likely to affect an individual's membership choice, an example of a reflection problem (Manski 1993).

To overcome these issues, I conduct a pre-specified information provision experiment with a sample of teachers in Wisconsin to understand whether perceptions of others' union participation affect one's own participation. My study combines administrative data on individual union membership with an online survey to measure prior and posterior beliefs. In the survey, I first elicit prior beliefs by asking teachers to predict the membership rate in their school district. Measuring prior beliefs is crucial because the information treatment could have offsetting effects if some

teachers update positively and others update negatively. Half of the sample is then randomized into receiving a truthful estimate of their school district's union membership rate. I construct district-specific membership rates by linking Wisconsin personnel data to campaign finance records that reveal which teachers pay union dues. After the information treatment, I estimate posterior beliefs by asking all participants to predict next year's membership rate. This allows me to see whether teachers have misperceptions about the average membership rate in their district, whether they correct those misperceptions in response to the treatment, and whether a change in beliefs ultimately affects the likelihood of being a member themselves.

Almost 2,500 teachers participated in the experimental survey during the spring of the 2023–24 school year. Importantly, there is considerable variation in prior beliefs both below and above the estimated district membership rates. Therefore, following Cantoni et al. (2019), I estimate a pooled regression, in which treated respondents who underestimated are coded as 1 and treated respondents who overestimated are coded as -1. This specification tests whether treated respondents adjusted their beliefs or union participation towards or away from the provided estimate.

The study has three main findings. First, while many teachers accurately predict the district membership share, the mean gap between prior beliefs and the truthful estimate is 16 percentage points. This implies that there is considerable variation in prior beliefs and room for correcting misperceptions. Second, the information provision induced a strong first stage — teachers in the treatment group are significantly more likely to update their posterior beliefs in the direction of the provided estimate. On average, treated teachers reduce the gap between their prior beliefs and the truthful estimate by 9 percentage points. Third, I ask teachers to estimate the likelihood that they will be a union member next year. Again, the treatment group is more likely to adjust their responses in accordance with the provided estimate. This indicates that planned participation is a complement to one's beliefs of others' participation. For instance, when teachers believe fewer co-workers are members, they are less likely to state that they will be members themselves. In future work, I will examine whether the treatment affected individuals' actual union membership status in the following school year.

Next, I estimate the results in a two-stage least squares (2SLS) specification, using treatment status as an instrument for the change in posterior beliefs. The 2SLS model is informative because it can be converted to a behavioral elasticity, capturing how an individual's membership status responds to changes in beliefs about others' participation. At one extreme, an elasticity of 1 implies that individual membership rates change one-to-one with the group membership rate. At the other extreme, an elasticity of 0 indicates that an individual's decision is unaffected by changes in the group membership rate. In my 2SLS design, I estimate a coefficient of 0.37, implying that a 10 percentage point change in beliefs leads to a corresponding 3.7 percentage point change in planned union participation. Given that baseline membership rates and beliefs about others' membership

are approximately equal, this coefficient can be roughly interpreted as an elasticity. Interestingly, I find a very similar estimate when naively regressing union status on the leave-one-out change in district membership rates using the administrative data. These findings reinforce the idea that a change in others' union participation (or the perception thereof) has a substantial effect on an individual's decision to participate in the union.

I explore further mechanisms by examining three pre-specified sources of heterogeneity. First, I test whether results vary by average membership rates at the district level. For example, effects could be more pronounced in places with a near-even membership split or in highly unionized districts. However, I find that the results are similar in below- and above-median unionized districts. Next, I examine whether results differ by whether an individual was a member at baseline. This is a useful exercise because one concern is that treated respondents may update not only on the group rate but also on the quality of the union. If respondents also updated their beliefs about union quality, effects should be stronger for non-members, who are likely less informed about the local union. However, I find similar results for members and non-members, suggesting that teachers primarily responded to changes in perceived membership rates rather than the quality of the union. Lastly, I examine whether the treatment effects vary by a baseline measure of pro-sociality. I find suggestive evidence that pro-social individuals respond more to the treatment, indicating that those inclined toward cooperative behavior are particularly sensitive to changes in perceived union participation.

Related literature: This paper contributes to work relating to the determinants of union membership and collective action problems more broadly. In terms of union membership, there is a long literature discussing the free-rider problem in settings where organizations cannot mandate union dues (Olson 1965, Lumsden and Petersen 1975, Ellwood and Fine 1987). Building on Akerlof (1980), Booth (1985) and Naylor and Cripps (1993) theorize that a “social customs” model of unionization can explain why membership rates are positive in open-shop settings. However, to my knowledge, there are no empirical papers that test whether social customs affect individual union membership. I address this gap by studying one specific social mechanism: the relationship between individual membership decisions and perceptions of the group membership rate.

I also build directly on recent experimental work exploring the relationship between individual and group participation in collective action problems. Prior studies have examined this in the context of protest participation (Cantoni et al. 2019; Neuert, Perino, and Schwickert 2023), election canvassing (Hager et al. 2023), and voter turnout (Hager et al. 2022).¹ These studies generally find evidence of strategic substitutes; for instance, individuals are less likely to turnout when they believe there is an increase in others' participation. I contribute to this work by examining strate-

¹In a non-experimental design, González (2020) examines how networks shape the decision to protest in Chile.

gic considerations in union membership, a key example of a collective action problem. Notably, in contrast to these prior studies, I find evidence of complementarity in union membership.² One possible explanation for this difference is that social mechanisms play a greater role in union membership where all workers are part of a specific organization. In contrast, free-riding may be easier when it comes to protesting or election participation.

Lastly, an advantage of my setting is that I observe survey respondents' stated intentions and also their actual future membership status. This will allow me to see whether the treatment affected membership in administrative data over a longer-run time frame.

II. Context and Administrative Data

A. Context

Recent legislation at the state and federal levels has prohibited public-sector labor unions from collecting “fair-share fees” from non-members, a payment which resembles normal membership dues. For instance, Wisconsin passed a law in 2011 called Act 10 that effectively made the state “right-to-work” for most public-sector workers. This law meant that teachers could choose whether they wanted to continue paying union dues or not. In 2018, the US Supreme Court ruled that no public-sector unions could automatically mandate membership because it violated the First Amendment (*Janus v. AFSCME* 2018).

B. Administrative Data

I measure individual union membership by combining personnel data on all Wisconsin public-school teachers with campaign finance data that lists which teachers pay union dues. The personnel data is made publicly available by the Wisconsin Department of Public Instruction and includes a teachers' name, job assignment, pay, and basic demographic characteristics. I link this to campaign finance reports from the state chapter of the National Education Association (NEA) because the state union automatically routes a small amount of each member's dues to its political action committee (PAC).³ I link the two datasets together by a person's full name and region (see Appendix Section B for more details on the matching procedure).

This data linkage serves two purposes in the experimental design. First, I use data from the 2022–23 school year to construct an estimate of each school district's union membership rate.

²Similarly, Naidu (2022), using survey data of essential workers during COVID from Hertel-Fernandez et al. (2020), finds that workers are more likely to state that they would take part in workplace collective action given a higher hypothetical share of coworkers taking that action.

³About 97 percent of districts in the state are affiliated with the state chapter of the NEA. While other major unions also have PACs, such as the American Federation of Teachers, they do not automatically deduct a small share of membership dues for campaign finance purposes.

This is the primary information that is provided to the treatment group as described more in the next section. Second, I use the individual membership data both as a pre-experimental baseline control and as a post-experimental outcome. In other words, I will estimate how the treatment affected individual membership status controlling for past membership status.

III. Hypotheses Tested and Experimental Design

A. Descriptive Analysis

The study examines whether individual union participation is influenced by the share of one's colleagues who are members. As a proof of concept using the administrative data, I analyze whether teachers who switch districts are more or less likely to be members based on membership rates in their former and new workplaces. If union participation is solely an individual decision, then moving to a more unionized district should have no effect on their membership status. However, if union membership exhibits complementarity, we would expect someone to be more likely to be a member after moving to a more unionized district.

To test this, I examine simple event study figures of average membership rates around the time that an individual changes districts. I first group districts into quartiles of membership rates in 2017. In Figure 1, I then plot membership rates in event time for teachers who moved to a new district between 2018–20. For simplicity, I plot membership rates for those who moved from the highest and lowest membership districts into any of the four quartiles.⁴

The figures reveal some interesting patterns relating to the complementarity of union membership. For example, the green line in the left panel shows that workers who move from the highest quartile districts to the lowest quartile districts are substantially less likely to be union members after moving. The figure suggests that the average membership rate drops from about 50 percent to about 20 percent. There are similar, but smaller, drops in membership for those who moved to quartile 2 (yellow) or quartile 3 (red) districts. However, membership rates are much more stable for those who moved to other quartile 4 districts (blue).

Likewise, there is symmetry when examining those teachers who moved away from first quartile districts (right panel). For instance, the average membership rate increases from about 10 percentage points to over 30 percentage points for those teachers who moved from a first- to a fourth-quartile district (blue).

This event study analysis suggests that the membership rate in one's district plays a substantial role in an individual's decision to be a union member. Moreover, it is highly indicative of complementarity in membership, where people are more likely to be members when the average rate

⁴See Card, Heining, and Kline (2013) for a similar analysis on workers moving between high- and low-paying firms.

around them is higher and vice versa. However, this movers design is only for people who actually switched to a new workplace by definition. I therefore cannot disentangle whether this is because people update on the group membership rate upon moving, or because of other factors such as the quality of the union in a particular district. Therefore, I next describe the experimental design where I exogenously shift people's beliefs about the average membership rate in the district where they work.

B. Hypotheses Tested

The information provision experiment largely has four main objectives as outlined in the pre-analysis plan. First, I gauge to what degree teachers misperceive the share of teachers who participate in the union within their organization. Second, I ask whether treated subjects update their beliefs relative to the control group after receiving a truthful estimate of their school district's union membership rate. Third, I examine whether subjects who update their beliefs due to the intervention then change either their planned or actual union participation in the following school year. Importantly, I separately measure effects for people who overestimated and underestimated the district membership rate. This will allow me to see whether individual membership behavior is related to the group membership rate either as strategic complements (as suggested by the movers analysis) or strategic substitutes (as in the literature on protest participation). Finally, I assess whether any associated treatment effects differ by certain baseline characteristics. In particular, I examine heterogeneity by baseline union membership and pre-experimental measures of pro-sociality.

C. Experimental Design

In the spring of 2024, I sent an online Qualtrics survey via email to approximately 45,000 teachers in the state of Wisconsin. The recruitment sample included almost all public school teachers in the state with a few exceptions. First, I could not find publicly available staff emails for roughly 13 percent of the school districts. Second, I did not send the experimental survey to school districts that had no union affiliate or were not affiliated with the NEA (only 3 percent of districts). I also did not send the experimental survey to districts for which my estimated union membership rate was below 5 percent. This restriction was done to make sure that each district had at least some union presence.

The email invited teachers to participate in a 5–10 minute survey about their perceptions of their jobs and unions.⁵ For completing the survey, respondents were entered into a raffle drawing

⁵Some of the pre-experimental responses were used in Foy (2024) to understand differences in teacher perceptions across union regimes.

where 60 winners received \$50 Amazon gift cards. The first batch of email invitations was sent out on April 10 and proceeded in a staggered fashion until April 30.⁶ The survey was closed and the raffle drawing was conducted on May 14.

In the survey, I asked all respondents to estimate what share of teachers in their school district and school were members of the union last year. This response, which importantly was asked of both treatment and control groups, serves as the primary measure of prior beliefs.

After eliciting prior beliefs, I then provided the exogenous information provision. If randomized to the treatment group, participants saw the following text:

*Our best estimate is that X% of teachers in your school district were members of the teachers' union last year.**

In place of the “X”, respondents saw a number that corresponded to their district-specific membership rate. (See Appendix Section C for an example of the experimental survey.) Note that this rate was not fixed for all respondents and differed by school district. As seen in Appendix Figure A1, there is considerable variation in the district membership rates ranging from below 10 percent to above 80 percent. The asterisk pointed to a link where respondents could see more information on how the estimate was constructed.

Next, I elicited respondents' posterior beliefs by asking both treatment and control groups to predict their district and school membership rates in the next school year. By asking about next year's membership rate, I captured updated expectations without repeating the same question twice.

Finally, all respondents were asked how likely they were to be a member of the teachers' union next year with choices of very likely, somewhat likely, undecided, somewhat unlikely, or very unlikely. This question serves as the basis for the outcome variable of whether teachers planned to be in the union next year.

The survey also asked questions designed to elicit respondents' baseline pro-sociality. Following Enke, Rodriguez-Padilla, and Zimmermann (2023), I create hypothetical scenarios where respondents are asked how they would split a 10 percent budget surplus with various groups. In three separate scenarios, respondents are asked what percentage of the surplus they would keep for teacher pay versus using it for support staff pay, student program funding, and administrator pay. The administrator pay category was included to confirm that teachers were paying attention with the assumption that fewer teachers would split the surplus with their superiors. Appendix Figure A2 panel A confirms this was the case, as the average share given to support staff, students, and

⁶I take advantage of the staggered rollout in the robustness exercises in Section V to examine potential spillover effects.

administrators was 43, 40, and 17, respectively. When assessing heterogeneity, I denote teachers as “pro-social” if they split the surplus at least 50-50 with either one or two of the three groups. Appendix Figure A2 panel B shows that union members at baseline are more pro-social on average according to this measure.

The survey ended with a question to see whether participants were paying attention following Bottan and Perez-Truglia (2025). About 98 percent of teachers passed the attention check. As such, I include all respondents in the primary analysis sample, but results are robust to excluding the 2 percent who answered incorrectly.

IV. Empirical Strategy and Sample Characteristics

A. Empirical Strategy

The primary research question is whether workers update their beliefs in response to the information treatment and subsequently whether they alter their planned or actual union status. To answer this, I estimate a two-stage least squares regression equation, where the treatment serves as an instrument for the change in posterior beliefs. Breaking this down by stage, the first stage equation is

$$D_i = \alpha + \pi Z_i + X_i' \eta + u_i, \quad (1)$$

where D_i is a teacher’s posterior belief about the union membership rate in their school district and Z_i is the treatment indicator. Following Cantoni et al. (2019), I code the treatment indicator as “-1” for those respondents whose prior belief was above the truth in order to measure how treatment moved people towards the truthful estimate. The vector X_i' includes a control for prior beliefs to measure how much one updated from their initial belief as well as additional baseline controls to improve precision.

Relatedly, the reduced form equation that relates the outcome variables to treatment is

$$Y_i = \alpha + \pi Z_i + X_i' \eta + u_i, \quad (2)$$

where the terms are as before but now the outcome (Y_i) is an indicator variable for planned or actual union participation. Putting this together, the second-stage equation,

$$Y_i = \alpha + \pi D_i + X_i' \eta + u_i, \quad (3)$$

measures how much membership rates (Y_i) change due to a corresponding change in beliefs (D_i), which will be instrumented by treatment (Z_i). It should be noted that the exclusion restriction may not strictly hold because, in addition to changing beliefs, the treatment may change other

perceptions such as the belief about the quality of the union. I discuss this further in Section VI.

The primary analysis will pool respondents who underestimate and overestimate their district-specific membership rate. As such, I also include an indicator variable for having a prior above the provided estimate and the interaction between each control and the above-estimate indicator. This is necessary to include because overestimators and underestimators may react differently to treatment. To gauge this, I also present results for underestimators and overestimators separately.

B. Sample Characteristics

I first examine observable characteristics of the teachers who participated in the experiment relative to the broader teaching workforce. I then check whether the randomization successfully balanced teachers across the treatment and control groups.

Participation in the Study: The experimental survey was sent to 43,637 teachers across the state of Wisconsin. Of those invitations, 2,789 individuals filled out at least one question of the survey and 2,525 completed enough of the survey such that they were randomized to one of the two treatment arms. This implies a response rate of about 5.8 percent, a rate similar to other studies that conduct cold recruitment strategies.⁷ I further drop 33 respondents who were either not licensed teachers or who could not be matched to the administrative personnel data (see Appendix Section B for more details). After these restrictions, there are 2,492 individuals in the experimental study, 1,232 randomized to the treatment group and 1,260 randomized to the control group. In the main analysis, I also exclude 70 respondents who did not answer either the prior belief, posterior belief, or planned participation question.

Table 1 compares the survey participants to other educators in the administrative data within the same district. Column 1 shows the baseline mean characteristics for non-participants and column 2 shows the difference for those in the experimental study. In general, there are significant differences: those who selected into participation have more experience, were more likely to be working full time or in multiple roles, and were less likely to be math teachers. The survey sample was also more likely to be male and less likely to be non-white. Of particular interest, the survey participants were 6 percentage points more likely to be union members via the matching procedure in the fall of the 2023–24 school year (i.e., before randomization), a difference of roughly 15 percent. This is not too surprising given that the recruitment email stated explicitly that the study was about teachers’ perceptions of their jobs and unions.

⁷It was also infeasible to provide payments to everyone who participated given the large potential sample size. Most similarly, Giacobasso et al. (forthcoming) sent invitations to residents in Dallas County via mail and received a response rate of 3.6 percent. Cantoni et al. (2019) received a response rate of 19.1 percent using a sample of undergraduates from Hong Kong University of Science and Technology, though all participants were paid.

Balance Across Treatment Groups: More importantly for the validity of the design, I examine whether the treatment and control groups were balanced across observable characteristics in Table 2. In addition to the baseline characteristics from Table 1, I also test whether there are differences in the baseline rate of pro-sociality or about respondents’ prior beliefs of their district and school membership rates. Column 1 shows the average rate of each characteristic in the control group and column 2 shows the difference in the treatment group. Here, the observable characteristics are well balanced across treatment arms, indicating that the randomization was successful. Across the fourteen variables, there are only two that are statistically significant at the 10 percent level: treated respondents were slightly less likely to be male (p-value = 0.06) and slightly less likely to be non-white (p-value = 0.06). Moreover, when regressing the treatment indicator on the full set of covariates, I fail to reject an F-test that baseline characteristics are jointly the same (p-value = 0.53).

I also examine balance separately by those who underestimate or overestimate their respective district membership rates, given that the analysis will pool these two groups. Appendix Tables A1 and A2 repeat the balance analysis but splitting by whether the respondent underestimated or overestimated, respectively. Once again, covariates are well-balanced across treatment groups. Though not statistically significant, note however that the point estimate of baseline union membership for underestimators was -3.2 percentage points (p-value = 0.29), indicating that treated respondents were somewhat less likely to be union members. This will be an important control in the analysis given that pre-randomization membership will very strongly predict future planned or actual membership status.

V. Experimental Results

A. Main Results

In this section, I examine whether the treatment affected participants’ beliefs about the union membership rate in their school district before turning to whether the updating of beliefs affected individuals’ planned or actual membership decisions.

First Stage on Beliefs: Figure 2 plots the distribution of beliefs around the truthful estimate of an individual’s district membership rate. The figure shows that the modal person had accurate perceptions of membership in their district.⁸ However, there is considerable variation to the left and right of 0 as well, implying that many people both underestimated and overestimated their district

⁸The centering around 0 is not due to people guessing 50 percent when they are uncertain (Fischhoff and De Bruin 1999). Appendix Figure A3 shows that this pattern holds even when restricting to districts with a membership rate below 40 percent or above 60 percent.

rate, respectively. There was slightly greater overestimation: the average estimate was about 7 percentage points above the district rate.

I next turn to the first question regarding the experimental intervention: whether respondents in the treatment group updated their beliefs after being provided with a truthful estimate of their district membership rate. Table 3 presents the results of equation 1, estimating whether the intervention affected posterior beliefs. In panel A, I show results for the pooled sample where treatment is recoded as -1 for those who overestimate their district rate. Column 1 includes no controls except for one's prior belief.⁹ Column 2 adds a control for whether the teacher was a union member at the start of the school year. I include additional baseline controls in column 3 to improve precision.¹⁰ Across specifications, there is a persistent, statistically significant coefficient of about 8.5 percentage points (p-value < 0.001), which indicates that treated respondents were more likely to update their beliefs towards the provided estimate. This is a substantial reduction in misperceptions as the baseline absolute gap between prior beliefs and the district rate was about 17 percentage points.

Panels B and C of Table 3 present results separated by underestimators and overestimators, respectively. There was about a 5 percentage point increase in posterior beliefs for underestimators and an 11 percentage point reduction in posterior beliefs for overestimators. This again represents substantial changes in beliefs as the baseline level of misperceptions for underestimators was 12 percentage points and the baseline level of misperceptions for overestimators was 21 percentage points.

I visualize the updating of beliefs in Figure 3. Panel A plots the distribution of posteriors (blue) overlayed on the distribution of priors (gray) for the control group. The figure makes clear that the control group's posterior beliefs are very similar to their prior beliefs. Panel B shows the corresponding figure for the treatment group. Here, the posterior beliefs are much more concentrated to the left of the figure, indicating that people moved their beliefs towards the provided estimate.

To more formally estimate the degree of updating in the treatment group, I calculate learning rates following Cavallo, Cruces, and Perez-Truglia (2017). This procedure estimates the weight that treated participants put on the signal versus the weight they place on their prior belief. The estimating equation is

$$\text{Updating}_i = \alpha + \beta_1 Z_i \times \text{Perception gap}_i + \beta_2 Z_i + \beta_3 \text{Perception gap}_i + \epsilon_i, \quad (4)$$

where Updating_i is the difference between an individual's posterior and their prior belief and

⁹Additionally, the pooled sample includes a control for whether the respondent's prior belief was above the district membership share as well as an interaction between the indicator and prior beliefs.

¹⁰The vector of controls includes baseline union membership status, years of experience in the district, full-time employment percentage, number of positions, whether the teacher had an advanced degree, whether they taught reading or math, whether the teacher was male or non-white, their pro-sociality score, and fixed effects for each of the 13 union regions. See Appendix Section B for more details.

Perception gap_{*i*} is the difference between the district rate and an individual's prior belief. The coefficient of interest, β_1 , represents the degree of updating that occurs for members of the treatment group net of any updating that occurs in the control group. A coefficient of 1 would indicate that teachers in the treatment group perfectly update their beliefs towards the provided estimate. As seen in Appendix Table A3, The β_1 coefficient is 0.55, which means that teachers in the treatment group placed a weight of 55 percent on the provided signal and a weight of 45 percent on their own prior belief. A learning rate of 0.55 is similar to what has been found in other information provision experiments (see Haaland, Roth, and Wohlfart (2023) for a review).

Reduced Form Effect on Planned Participation: The first stage analysis showed that treated respondents updated their beliefs in the direction of the provided district estimate. I next examine whether the treatment affected respondents' planned participation in the teachers' union next year. It's possible, for instance, that respondents who underestimated the district rate would be more likely to say that they plan to be members in the future, a sign of strategic complementarity. On the other hand, respondents may be less likely to say they will participate in the union, perhaps because they think they can free ride on others' participation.

I examine this by estimating Equation 2 and present the results in Table 4, where once again column 1 shows results without any controls, column 2 adds a control for an individual's baseline membership status, and column 3 includes additional controls to improve precision. The outcome is an indicator variable equal to 100 if respondents said they are *very likely* or *somewhat likely* to be members of the teachers' union next year and 0 if they responded *uncertain*, *somewhat unlikely*, or *very unlikely*. Focusing on the pooled sample in Panel A, I find a positive but insignificant coefficient of 2.5 percentage points in column 1. This points to complementarity in membership but the estimate is imprecise. Recall, however, that there was slight imbalance in treatment assignment between members and non-members for underestimators, which is important as prior membership is very predictive of future membership behavior. When controlling for just prior membership in column 2, I find a similar but more precise coefficient of 2.7 percentage points (p-value = 0.04.) This indicates that planned participation is complementary to the group membership rate. In column 3, there is a very similar estimate of 3.2 percentage points when adding additional controls (p-value = 0.02).

As in Table 3, panels B and C of Table 4 split the sample by underestimators and overestimators, respectively. My preferred specification with controls in column 3 shows that there is about a 2–4 percentage point effect in the direction of the provided estimate for both underestimators and overestimators. This indicates that the pooled effect is not driven by one side of the estimate; in both cases, planned participation moves in the same direction as the updated first-stage beliefs.

The prior analysis focused on planned participation next year, which was defined as those who

answered “very likely” or “somewhat likely” to the question about the likelihood of being a future union member. However, the question was on a 5-point likert scale, meaning that an alternative framing is whether they moved up or down the scale. For instance, 12 percent of the sample said that they were “undecided.” In Appendix Table A4, I instead use a scale from 1 (very unlikely) to 5 (very likely) as the outcome variable. This analysis shows relatively more precise estimates in favor of complementarity.

2SLS Estimates: Tables 3 and 4 showed evidence of strategic complementarity. Treated respondents who underestimated the district rate were more likely to say they would be a member next year, while those who overestimated the district rate were less likely to say they would be a member next year. In Table 5, I put these pieces together in a two-stage least squares regression, where the change in beliefs is instrumented by treatment assignment. Focusing on the results with controls in column 3, the coefficient of 0.37 indicates that a 10 percentage point increase in beliefs about the district membership rate increases the likelihood that one plans to be a member by 3.7 percentage points. Importantly, this relies on the assumption that the only thing that treatment changes is updating one’s beliefs about the district membership rate. This may not be strictly true because people may also update on other items, such as the quality of the local union. I discuss this further in Section VI.

Magnitude: The 2SLS estimates can be viewed in the context of a behavioral elasticity. For instance, an elasticity of 1 would imply that an individual’s membership decision corresponds one-to-one with their beliefs about others’ participation. Conversely, an elasticity of 0 would say that individuals are not affected at all by others’ participation. The 2SLS estimates imply that a 10 percentage point increase in beliefs about others’ participation induces a 3.7 percentage point increase in one’s own planned participation. This point estimate implies a behavioral elasticity of roughly 0.37.¹¹ While significantly less than one, this estimate represents a substantial correspondence between others’ participation and one’s own membership decision.

As a frame of reference, I formalize the non-experimental event study analysis using the administrative data (Section III) in a regression specification. For each teacher in the administrative data, I create leave-one-out district-level membership rates in each year. I then estimate the change in a person’s district membership rate from the prior year to the current year. This allows me to show the corresponding analysis in the administrative data by regressing one’s membership status on the change in district-level membership rates. This non-experimental analysis is not necessarily

¹¹More specifically, the baseline rate of actual membership (0.44) is roughly equivalent to the baseline belief of others’ membership (0.45). Therefore, the 2SLS point estimate of 0.37 can roughly be thought as an elasticity at this level of membership.

causal, because other factors could affect both the change in district rates and one's own membership decision such as a change to the quality of the union or administration. Nevertheless, I show the results from this analysis in Appendix Table A5 because, interestingly, the results are very similar to the experimental elasticity. In column 1, I find that a 1 percentage point change in district-level membership rates corresponds to a 0.37 percentage point change in one's own membership status. Results are similar when adding controls in column 2.

To summarize, I find that the experimental treatment induced a change in beliefs towards the provided district membership rate. Accordingly, treated respondents' planned participation changed in the same direction as their change in beliefs. When combined in a two-stage analysis, I find that a 10 percentage point change in beliefs induced a 3.7 percentage point change in planned participation on average.

B. Robustness

Alternative Specifications: I examine robustness to the main results in Appendix Table A6. Panels A–C present results for the first stage, reduced form, and two-stage least squares specifications, respectively. For reference, column 1 shows the baseline result with controls from column 3 of Tables 3–5. In column 2, I incorporate the fact that the survey asked teachers to estimate the union membership rates in their school in addition to their district. Not surprisingly, district and school beliefs are highly correlated. The column 2 estimates replace the prior and posterior district predictions with the corresponding school variables. Results are similar, suggesting that the treatment induced people to update not only on their district but on the union membership rates in their actual school building.

While treatment occurs at the individual level, there is still likely serial correlation in union membership within districts. Column 3 addresses this by including district fixed effects and clustering standard errors at the district level. There are slightly fewer observations, because some districts do not have both treatment and control observations, but the results are relatively more pronounced and precise.

One issue with the pooled sample is that it is not clear how to treat people in the treatment group who estimate the district rate exactly. Column 4 presents an alternative model, in which these treated respondents are re-coded as 0 and included in the control group. The logic is that these people should not change their behavior given that they are already correctly estimating the district rate. This change has little impact on the results.

Finally, columns 5–6 make two restrictions. In column 5, I drop the respondents who failed the attention check at the end of the survey. Likewise, in column 6 I drop outliers in terms of their pre-randomization prior beliefs. Specifically, I exclude 1 percent of the sample with the most extreme priors measured as the absolute distance between their prior belief and the district share.

Again, results are similar in these two subsamples.

Spillover Effects: Another concern is that treated respondents informed others, particularly control group members, about the average membership rate. For instance, since the survey went out over email, it is possible that teachers in a school took it at the same time or discussed it during their day. If anything, this would likely bias results towards zero given that this might result in the control group having more informed priors.

Nevertheless, there are a few pieces of information that suggest limited spillovers. For one, the 2,422 respondents in the main sample were from 976 unique schools, suggesting limited selection into the survey within a school. Additionally, control group participants had posteriors that were very similar to their priors (Figure 3 panel A).

It remains possible that both groups had more informed priors over time due to spillover effects. To test this, I incorporate the fact that the survey was administered over a five-week time frame. I therefore examine whether people who took the survey later had more accurate prior beliefs relative to those who took the survey earlier. In Appendix Figure A4, I regress the gap between prior beliefs and the district share on survey week fixed effects. The results show that the gap in prior beliefs was similar across weeks and I cannot reject that the gap in beliefs in any of the 2–5 weeks were different than the average gap in the first week. Overall, I find no evidence suggesting that participants shared the treatment information with each other.

VI. Mechanisms

In this section, I explore potential mechanisms by examining heterogeneity on three pre-specified dimensions: the baseline membership rate in the district, whether the individual was a union member at baseline, and whether the respondent is pro-social. Overall, I find suggestive evidence that the results are most consistent with a social effect, in which pro-social people are most likely to be influenced by the treatment. Additionally, I find no difference in results for members versus non-members, suggesting little updating on the quality of the union.

In Figure 4, I first split the sample by whether the respondent was in a district with a below- or above-median district membership share. It could be the case, for instance, that the social mechanisms causing people to join unions are more prevalent in unions with higher membership rates. However, I find similar results by baseline district membership rates.

Next, I look at whether there were notable differences by an individual's pre-randomization membership status. Here, I see similar results suggesting that members and non-members were equally likely to update their beliefs and planned participation in response to the treatment. It is particularly notable that respondents who were union members before treatment also update

because these people presumably already know of the quality of the union. Therefore, this provides suggestive evidence that people are responding to their updated beliefs about the share of members and not on other factors.

Lastly, I examine whether results differ by the pro-sociality measure that was elicited in the survey. Recall that I asked each person how they would split a 10 percent budget surplus between teacher pay and either support staff pay or student program funding. I denote one's pro-sociality score equal to 0 if they spend the majority of the surplus on teacher pay in both scenarios. Likewise, a person's score equals 1 if they split at least half of the surplus with either source and equals 2 if they split at least half in both scenarios. On the one hand, prosocial individuals may be less likely to change their behavior in response to the treatment because their actions are motivated by altruism, rather than by others' participation. Conversely, prosocial individuals may be more responsive to changes in group membership rates precisely because they are more attuned to social norms and the behavior of others.

I examine this in the bottom set of results in Figure 4, where I find evidence most consistent with the second channel in which pro-sociality is correlated with stronger updating. For instance, those with a pro-social score equal to 0 are no more likely to change their behavior in the treatment group relative to the control group. Conversely, for people with a pro-sociality score of 2, the 2SLS coefficient equals 0.70, with a 95 percent confidence interval that includes 1. This means that, for pro-social people, their planned participation moves almost one-to-one with their updated beliefs about the group membership rate. While not randomly assigned, the patterns reinforce the idea that people with pro-social motivations were the ones who responded to the treatment the most. Moreover, this aligns with the idea that sustaining union membership rates relies on pro-social individuals.

Finally, it remains possible that the self-reported findings are influenced by experimenter demand bias, where participants alter their stated behavior because of researcher expectations. This is especially conceivable for underestimators, who may feel pressure to say they will be members upon learning that the district estimate is higher than they thought. The effect of demand bias on overestimators is somewhat less clear — they may not perceive a lower estimate as a researcher endorsement.

Nevertheless, an advantage of my setting is that I can also look at long-run effects on actual membership status. This analysis will continue as data is made available.

VII. Conclusion

This study investigates whether teachers alter their beliefs and intentions regarding union membership after an exogenous change to their perceptions of others' participation in the union. I examine

this through the use of an information provision experiment, where treated respondents were given a truthful estimate of their school district's union membership rate. I find that the information provision closes about half of the 17 percentage point gap between teachers' prior beliefs and the district estimate. Accordingly, teachers in the treatment group were more likely to adjust their plans for union participation in the direction of their changed beliefs. This indicates that individuals' reported membership decisions are complementary to their beliefs of others' participation.

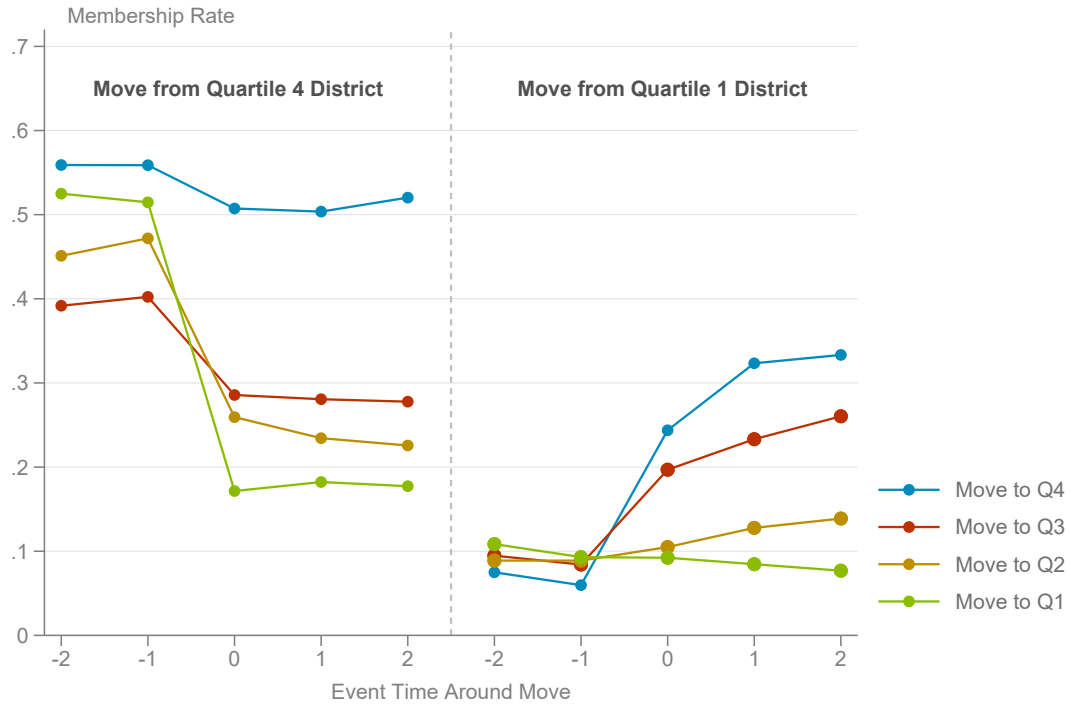
Putting this in a two-stage least squares specification, I estimate that a 10 percentage point change in one's beliefs corresponds to a 3.7 percentage point change in one's own planned union membership. This roughly translates to a behavioral elasticity of 0.37 and indicates that others' participation is a significant contributor to one's own participation in settings where union membership is not mandatory. Nevertheless, it is significantly different from 1, implying that there are other factors that contribute to individual membership decisions. In the future, I will extend the paper by examining whether the information treatment affected actual membership decisions in the following school year.

References

- Akerlof, George A. 1980. "A Theory of Social Custom, of Which Unemployment May be One Consequence." *Quarterly Journal of Economics* 94: 749–75.
- Booth, Alison L. 1985. "The Free Rider Problem and a Social Custom Model of Trade Union Membership." *Quarterly Journal of Economics* 100(1): 253–61.
- Bottan, Nicolas, and Ricardo Perez-Truglia. 2025. "Betting on the House: Subjective Expectations and Market Choices." *American Economic Journal: Applied Economics* 17(1): 459–500.
- Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128(3): 967–1015.
- Cantoni, Davide., David Y. Yang, Noam Yuchtman, and Y. Jane Zhang. 2019. "Protests as Strategic Games: Experimental Evidence from Hong Kong's Anti-Authoritarian Movement." *Quarterly Journal of Economics* 134(2): 1021–77.
- Cavallo, Alberto, Guillermo Cruces, and Ricardo Perez-Truglia. 2017. "Inflation Expectations, Learning, and Supermarket Prices: Evidence from Survey Experiments." *American Economic Journal: Macroeconomics* 9(3): 1–35.
- Ellwood, David T., and Glenn Fine. 1987. "The Impact of Right-to-Work Laws on Union Organizing." *Journal of Political Economy* 95(2): 250–73.
- Enke, Benjamin, Ricardo Rodríguez-Padilla, and Florian Zimmerman. 2023. "Moral Universalism and the Structure of Ideology." *Review of Economic Studies* 90: 1934–62.
- Fischhoff, Baruch, and Wandí Bruine de Bruin. 1999. "Fifty-Fifty=50%?" *Journal of Behavioral Decision Making* 12: 149–63.
- Foy, Morgan. 2024. "Selection and Performance in Teachers' Unions." Working Paper.
- Giaccobasso, Matias, Brad Nathan, Ricardo Perez-Truglia, and Alejandro Zentner. forthcoming. "Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending." *American Economic Journal: Applied Economics*.
- González, Felipe. 2020. "Collective Action in Networks: Evidence from the Chilean Student Movement." *Journal of Public Economics* 188: 104220.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart. 2023. "Designing Information Provision Experiments." *Journal of Economic Literature* 61(1): 3–40.
- Hager, Anselm, Lukas Hensel, Johannes Hermle, and Christopher Roth. 2022. "Group Size and Protest Mobilization across Movements and Countermovements." *American Political Science Review* 116: 1051–66.
- Hager, Anselm, Lukas Hensel, Johannes Hermle, and Christopher Roth. 2023. "Political Activists as Free Riders: Evidence from a Natural Field Experiment." *Economic Journal* 133: 2068–84.


- Hertel-Fernandez, Alexander, Suresh Naidu, Adam Reich, and Patrick Youngblood. 2020. “Understanding the COVID-19 Workplace: Evidence from a Survey of Essential Workers.” New York: Roosevelt Institute.
- Janus v AFSCME. 2018. No. 16–1466. https://www.supremecourt.gov/opinions/17pdf/16-1466_2b3j.pdf.
- Jarke-Neuert, Johannes, Grischa Perino, and Henrike Schwickert. 2023. “Free Riding in Climate Protests.” *Nature Climate Change* 13: 1197–1202.
- Lumsden, Keith G., and Craig H. Petersen. 1975. “The Effect of Right-to-Work Laws on Unionization in the United States.” *Journal of Political Economy* 83(6): 1237–48.
- Manski, Charles F. 1993. “Identification of Endogenous Social Effects: The Reflection Problem.” *The Review of Economic Studies* 60: 531–42.
- Naidu, Suresh. 2022. “Is There Any Future for a US Labor Movement?” *Journal of Economic Perspectives* 36(4): 3–28.
- Naylor, Robin, and Martin Cripps. 1993. “An Economic Theory of the Open Shop Trade Union.” *European Economic Review* 37: 1599–620.
- Olson, Mancur. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge: Harvard University Press.

Figure 1: Event Study of Membership Rates around Moves to New Districts



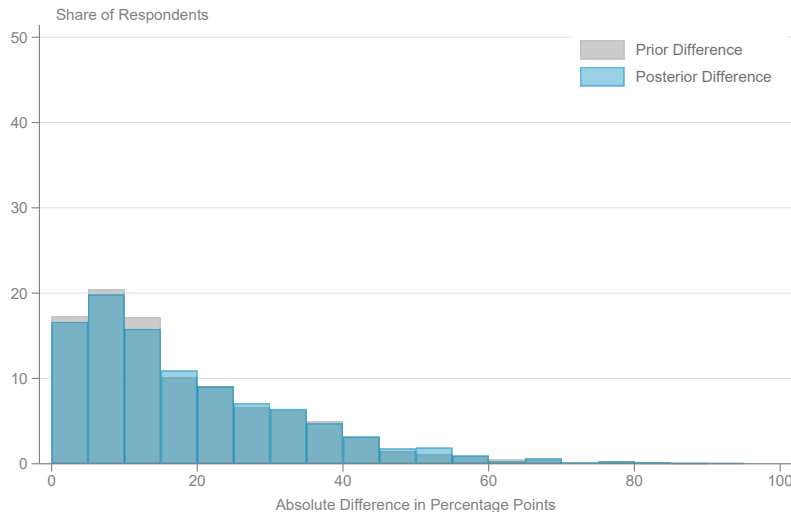
Note: The figure shows individual membership rates around the timing of a move to a new school district. The sample is teachers who moved once between 2018–2020 and have two years of data pre- and post-move. I split the sample into four quartiles of 2017 membership rates (using all teachers not just movers). I then separately examine membership rates for people who moved from a district in each of the four quartiles to a new district in any of the four quartiles. For simplicity, the figure shows only moves from quartile four (left panel) and quartile one (right panel). Individuals who moved to a new district in quartile four, three, two, or one are shown in blue, red, yellow, and green, respectively.

Figure 2: Distribution of Priors around District Estimate of Membership Rate




Note: The x-axis is the an individual's prior belief of their school district's membership rate minus the truthful estimate of the district membership rate. The y-axis is the share of respondents who fall into each 5 percentage point bin. Respondents who underestimated their district share lie to the left of the dashed line; respondents who overestimated their district share lie to the right of the dashed line.

Figure 3: Belief Updating: Gap between Prior and Truth and Posterior and Truth




(a) Control Group



(b) Treatment Group

Note: The figure shows the distribution of the absolute difference between the priors (gray) and the truthful estimate and the posteriors (blue) and the truthful estimate. Panel A shows results for the control group; Panel B shows results for the treatment group. The x-axis is the absolute difference in percentage points. The y-axis is the share of respondents in 5 percentage point bins.

Figure 4: Heterogeneity in 2SLS Estimates



Note: The figure plots two-stage least squares estimates from equation 3 by various subgroups. The outcome variable is planned participation in the union next year. The first set of results splits the sample by whether membership is above or below the median district membership rate (median = 37 percent). The second set of results splits the sample by whether the respondent was a union member in the fall of 2023, before the randomization experiment. The last set of results splits the sample by whether the respondent received a low, medium, or high pro-sociality score. See Appendix Section B for details. The gray bars represent 95 percent confidence intervals.

Table 1: Selection into Survey Experiment

	(1) Non-Participants	(2) Participant Diff.	(3) P-value	(4) Observations
Experience	14.61	0.546*** (0.193)	0.01	56,837
FTE	99.56	0.831*** (0.203)	<0.01	56,837
Salary	64,729.19	-110.656 (286.408)	0.70	56,837
Number Positions	1.51	0.096*** (0.020)	<0.01	56,837
Advanced Degree	0.51	0.014 (0.010)	0.17	56,837
Reading Teacher	0.34	-0.012 (0.009)	0.21	56,837
Math Teacher	0.34	-0.033*** (0.009)	<0.01	56,837
Male	0.23	0.039*** (0.009)	<0.01	56,837
Non-White	0.07	-0.016*** (0.005)	<0.01	56,837
Member	0.37	0.057*** (0.010)	<0.01	56,837

Note: The sample is all teachers working in school districts that received the survey. Each row shows the results from regressing the stated baseline variable on an indicator variable for whether the individual took the survey. Column 1 presents the mean value for non-participants. Column 2 shows the coefficient and associated standard errors for the indicator denoting participation in the survey. Column 3 presents the associated p-value. The number of observations are in column 4. All regressions use district fixed effects and standard errors are clustered at the district level. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table 2: Balance

	(1) Control	(2) Treated Diff.	(3) P-value	(4) Observations
Experience	15.07	0.074 (0.387)	0.85	2,492
FTE	100.50	-0.029 (0.395)	0.94	2,492
Salary	64,646.86	-174.826 (558.606)	0.75	2,492
Number Positions	1.60	0.018 (0.040)	0.66	2,492
Advanced Degree	0.52	0.009 (0.020)	0.67	2,492
Reading Teacher	0.33	-0.021 (0.019)	0.27	2,492
Math Teacher	0.29	0.017 (0.018)	0.35	2,492
Male	0.29	-0.034* (0.018)	0.06	2,492
Non-White	0.06	-0.017* (0.009)	0.06	2,492
Member	0.44	-0.019 (0.020)	0.33	2,492
Pro-Social	0.94	0.011 (0.032)	0.73	2,492
Prior Belief, District	45.55	-0.779 (1.089)	0.47	2,442
Prior Belief, School	44.16	-0.787 (1.194)	0.51	2,441
Attrited	0.03	-0.009 (0.007)	0.17	2,492

Note: Each row shows the results from regressing the stated baseline variable on treatment status. Column 1 presents the mean value in the control group. Column 2 shows the coefficient and associated standard errors on the treatment status indicator. Column 3 presents the associated p-value. The number of observations are in column 4. There are fewer observations for the two prior belief variables because some respondents did not fill out this question. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table 3: First Stage Effect on Beliefs

	(1)	(2)	(3)
	Outcome = Posterior Beliefs		
<i>Panel A: Pooled</i>			
Treatment (Recode)	8.589*** (0.466)	8.568*** (0.461)	8.716*** (0.454)
Observations	2,422	2,422	2,422
<i>Panel B: Underestimators</i>			
Treatment	5.094*** (0.702)	5.108*** (0.705)	5.344*** (0.686)
Observations	1,080	1,080	1,080
<i>Panel C: Overestimators</i>			
Treatment	-11.402*** (0.613)	-11.351*** (0.599)	-11.419*** (0.595)
Observations	1,342	1,342	1,342
Baseline Membership	No	Yes	Yes
Additional Controls	No	No	Yes

Note: The outcome variable is respondents' estimates of the district membership rate next year. Panel A is the pooled sample of all participants, where treatment is recoded as -1 for those who overestimated their district share as described in Section IV. Panel B is the subsample of participants who underestimated their district share; Panel C is the subsample of participants who overestimated their district share. All specifications control for prior beliefs, column 2 controls for respondents' prior membership status, and column 3 includes additional controls as described in the text. The specifications in Panel A also include an indicator variable for one's prior beliefs being above the district membership rate and an interaction between the above district membership rate indicator and each control variable. Robust standard errors in parentheses. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table 4: Reduced Form Effect on Planned Participation

	(1)	(2)	(3)
	Outcome = Planned Participation		
<i>Panel A: Pooled</i>			
Treatment (Recode)	2.466 (1.883)	2.738** (1.350)	3.229** (1.368)
Observations	2,422	2,422	2,422
<i>Panel B: Underestimators</i>			
Treated	0.116 (2.689)	1.729 (1.906)	2.350 (1.916)
Observations	1,080	1,080	1,080
<i>Panel C: Overestimators</i>			
Treated	-4.357* (2.620)	-3.550* (1.893)	-3.933** (1.928)
Observations	1,342	1,342	1,342
Baseline Membership	No	Yes	Yes
Additional Controls	No	No	Yes

Note: The outcome variable is an indicator variable equal to 100 if the respondent said they were “very likely” or “somewhat likely” to be a member of the teachers’ union next year. Panel A is the pooled sample of all participants, where treatment is recoded as -1 for those who overestimated their district share as described in Section IV. Panel B is the subsample of participants who underestimated their district share; Panel C is the subsample of participants who overestimated their district share. Column 1 includes controls for prior beliefs, column 2 adds a control for respondents’ prior membership status, and column 3 includes additional controls as described in the text. The specifications in Panel A also include an indicator variable for one’s prior beliefs being above the district membership rate and an interaction between the above district membership rate indicator and each control variable. Robust standard errors in parentheses. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table 5: 2SLS Effect of Change in Posteriors on Planned Participation

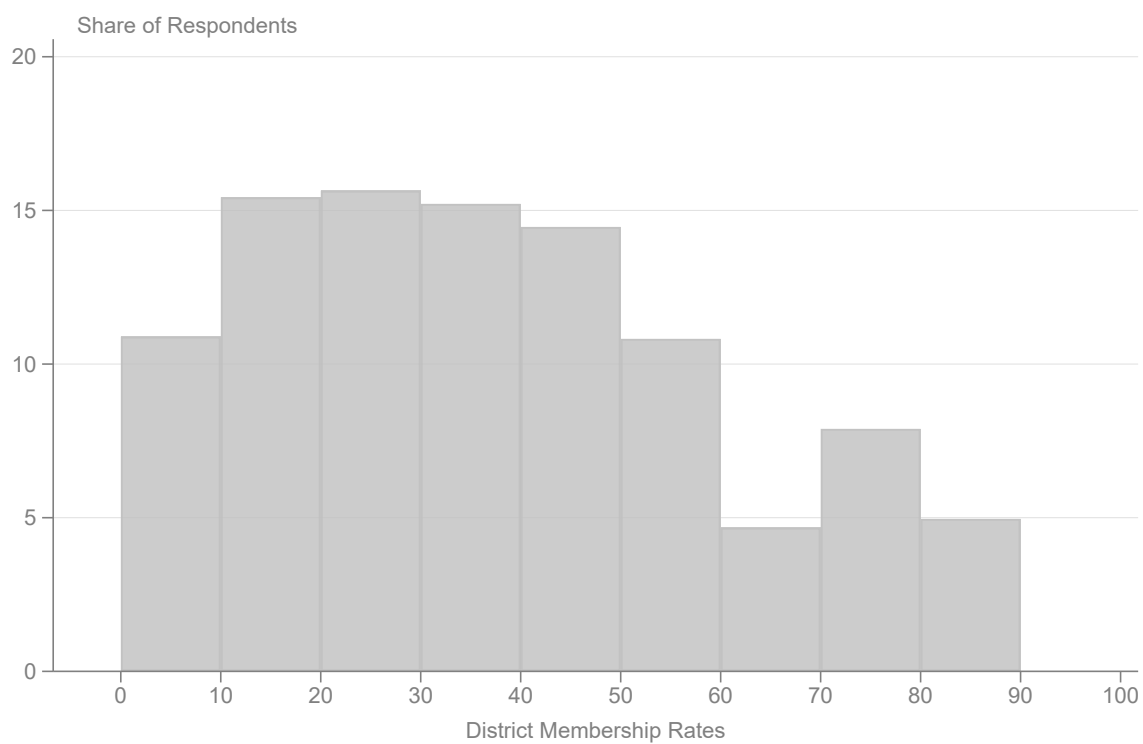
	(1)	(2)	(3)
	Outcome = Planned Participation		
<i>Panel A: Pooled</i>			
Posterior Beliefs	0.287 (0.216)	0.320** (0.156)	0.370** (0.154)
Observations	2,422	2,422	2,422
<i>Panel B: Underestimators</i>			
Posterior Beliefs	0.023 (0.527)	0.338 (0.368)	0.440 (0.350)
Observations	1,080	1,080	1,080
<i>Panel C: Overestimators</i>			
Posterior Beliefs	0.382* (0.225)	0.313* (0.165)	0.344** (0.165)
Observations	1,342	1,342	1,342
Baseline Membership	No	Yes	Yes
Additional Controls	No	No	Yes

Note: The outcome variable is an indicator variable equal to 100 if the respondent said they were “very likely” or “somewhat likely” to be a member of the teachers’ union next year. Panel A is the pooled sample of all participants, where treatment is recoded as -1 for those who overestimated their district share as described in Section IV. Panel B is the subsample of participants who underestimated their district share; Panel C is the subsample of participants who overestimated their district share. Column 1 includes controls for prior beliefs, column 2 adds a control for respondents’ prior membership status, and column 3 includes additional controls as described in the text. The specifications in Panel A also include an indicator variable for one’s prior beliefs being above the district membership rate and an interaction between the above district membership rate indicator and each control variable. Robust standard errors in parentheses. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Appendix


A. Additional Figures/Tables

Figure A1: Distribution of District Membership Rates




Note: The x-axis is the district membership rate that was used in the experimental information provision. The y-axis shows the percentage of survey participants in each 10 percentage point bin.

Figure A2: Descriptive Statistics on Pro-Social Measure



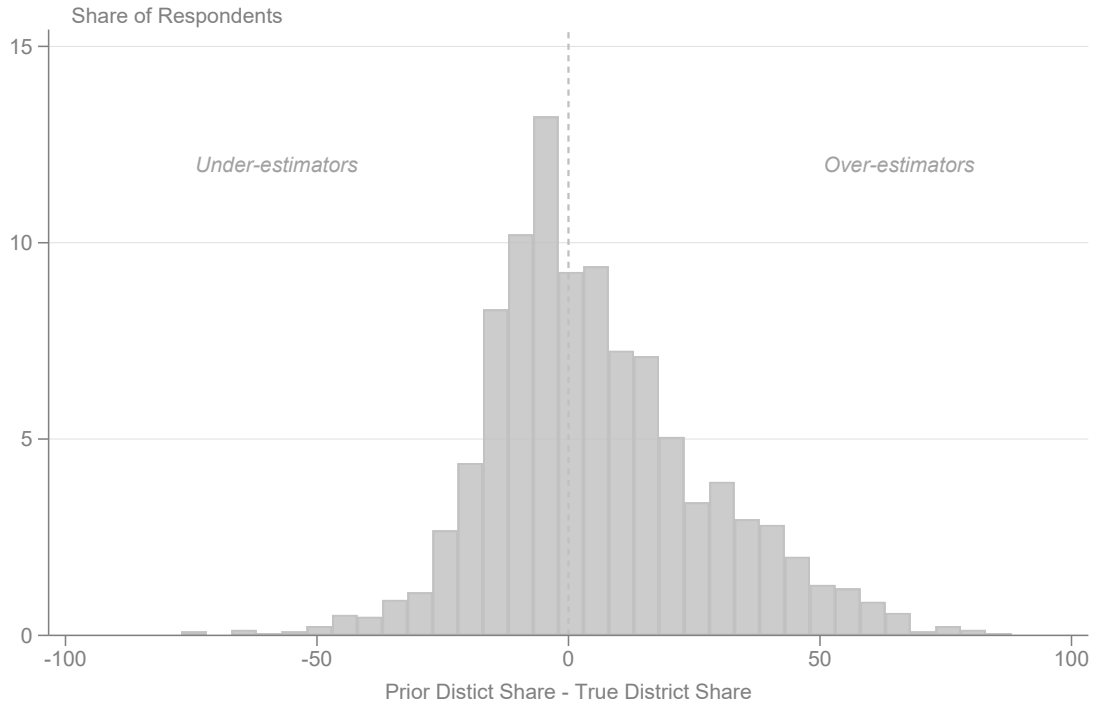
(a) Average Share Given to Each Group



(b) Pro-Social Score by Union Status


Note: Panel A shows the average percentage of the budget surplus shared between support staff pay, student program funding, and administrator pay, respectively. See Appendix Section C. Panel B shows the share of respondents within each pro-social score separated by union status.

Figure A3: Distribution of Priors Excluding Districts Near 50 Percent



Note: This is a reproduction of Figure 2, but restricted to districts where the estimated membership rate was less than 40 percent or greater than 60 percent. The x-axis is the an individual's prior belief of their school district's membership rate minus the truthful estimate of the district membership rate. The y-axis is the share of respondents who fall into each 5 percentage point bin. Respondents who underestimated their district membership rate lie to the left of the dashed line; respondents who overestimated their district rate lie to the right of the dashed line.

Figure A4: Gap Between Prior Beliefs and District Rate by Survey Week



Note: This figure presents the gap between prior beliefs and the district membership rate regressed on indicators for which week the survey was taken. The coefficients from this regression are then added to the Week 1 mean of 16.6 percentage points. Gray bars represent 95 percent confidence intervals.

Table A1: Balance, Underestimators

	(1) Control	(2) Treated Diff.	(3) P-value	(4) Observations
Experience	14.42	0.586 (0.587)	0.32	1,091
FTE	100.29	-0.272 (0.635)	0.67	1,091
Salary	63,936.89	127.805 (846.341)	0.88	1,091
Number Positions	1.59	0.055 (0.065)	0.40	1,091
Advanced Degree	0.48	0.017 (0.030)	0.56	1,091
Reading Teacher	0.39	-0.050* (0.029)	0.09	1,091
Math Teacher	0.33	0.016 (0.029)	0.56	1,091
Male	0.26	-0.014 (0.027)	0.59	1,091
Non-White	0.07	-0.010 (0.015)	0.51	1,091
Member	0.44	-0.032 (0.030)	0.29	1,091
Pro-Social	0.94	0.011 (0.048)	0.81	1,091
Prior Belief, District	28.94	-0.892 (1.470)	0.54	1,091
Prior Belief, School	29.43	-1.096 (1.686)	0.52	1,091
Attrited	0.02	-0.011 (0.007)	0.13	1,091

Note: The sample is respondents whose prior belief was below or equal to their respective district share. Each row shows the results from regressing the stated baseline variable on treatment status. Column 1 presents the mean value in the control group. Column 2 shows the coefficient and associated standard errors on the treatment status indicator. Column 3 presents the associated p-value. The number of observations are in column 4. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table A2: Balance, Overestimators

	(1) Control	(2) Treated Diff.	(3) P-value	(4) Observations
Experience	15.85	-0.409 (0.522)	0.43	1,351
FTE	100.63	0.202 (0.494)	0.68	1,351
Salary	65,506.13	-527.069 (755.039)	0.48	1,351
Number Positions	1.61	-0.018 (0.052)	0.72	1,351
Advanced Degree	0.55	0.007 (0.027)	0.80	1,351
Reading Teacher	0.29	0.000 (0.025)	0.99	1,351
Math Teacher	0.27	0.007 (0.024)	0.77	1,351
Male	0.31	-0.047* (0.025)	0.05	1,351
Non-White	0.06	-0.021* (0.012)	0.08	1,351
Member	0.46	-0.014 (0.027)	0.60	1,351
Pro-Social	0.94	0.012 (0.043)	0.79	1,351
Prior Belief, District	58.72	-0.195 (1.124)	0.86	1,351
Prior Belief, School	55.86	-0.114 (1.359)	0.93	1,350
Attrited	0.01	-0.003	0.51	1,351

Note: The sample is respondents whose prior belief was above their respective district share. Each row shows the results from regressing the stated baseline variable on treatment status. Column 1 presents the mean value in the control group. Column 2 shows the coefficient and associated standard errors on the treatment status indicator. Column 3 presents the associated p-value. The number of observations are in column 4. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table A3: Learning Rates

	Updating (Posterior - Prior)
Treated (Recode)	-0.508 (0.683)
Signal Gap	-0.006 (0.040)
Treated (Recode) \times Signal Gap	0.547*** (0.043)
Observations	2,422

Note: The table estimates the degree of learning for members of the treatment group. The outcome variable is the difference between a respondent's posterior and prior beliefs. Treated (Recode) is an indicator for whether the respondent was in the treatment group, where overestimators are coded as -1 (see equation 1). The "signal gap" is the difference between the district membership rate and one's prior beliefs. The third row shows the interaction term between these two variables. Robust standard errors are in parentheses. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table A4: Reduced Form Effect on Likelihood of Participation

	(1)	(2)	(3)
	Outcome = Likelihood Scale		
<i>Panel A: Pooled</i>			
Treated (Recode)	0.123* (0.066)	0.132*** (0.046)	0.147*** (0.046)
Observations	2,422	2,422	2,422
<i>Panel B: Underestimators</i>			
Treated	0.051 (0.095)	0.109 (0.067)	0.130* (0.067)
Observations	1,080	1,080	1,080
<i>Panel C: Overestimators</i>			
Treated	-0.181** (0.090)	-0.153** (0.064)	-0.162** (0.063)
Observations	1,342	1,342	1,342
Baseline Membership	No	Yes	Yes
Additional Controls	No	No	Yes

Note: The outcome variable is a 5-point scale of the likelihood that a respondent is a member of the teachers' union next year where 5 = very likely and 1 = very unlikely. Panel A is the pooled sample of all participants, where treatment is recoded as -1 for those who overestimated their district share as described in Section IV. Panel B is the subsample of participants who underestimated their district share; Panel C is the subsample of participants who overestimated their district share. Column 1 includes controls for prior beliefs, column 2 adds a control for respondents' prior membership status, and column 3 includes additional controls as described in the text. The specifications in Panel A also include an indicator variable for one's prior beliefs being above the district membership rate and an interaction between the above district membership rate indicator and each control variable. Robust standard errors in parentheses. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table A5: Effect of Change in Others' Membership on Own Membership, Admin Data

	Outcome = Membership Status	
Δ District Membership Rate	0.374*** (0.021)	0.370*** (0.020)
Lagged Membership	0.884*** (0.005)	0.878*** (0.005)
Observations	310,909	310,909
Controls	No	Yes
Number Districts	405	405

Note: The table estimates the relationship between the leave-one-out change in district membership rates and a teacher's membership status. The sample is teachers in the administrative data from 2017–2022. The outcome variable is a teacher's membership status in a given year t . The variable of interest, Δ District Membership Rate, is the change in the leave-one-out district membership rates in the school district where a teacher worked. Lagged membership is an indicator variable for a teacher's membership status in the prior school year. Both specifications include district and year fixed effects. Column 2 includes additional controls for local experience, full-time employment amount, the number of positions a teacher worked, higher education, indicators for teaching reading or math, and indicators for being male or non-white. Standard errors clustered at the district level are in parentheses. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

Table A6: Robustness Exercises

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: First Stage (Posterior Beliefs)</i>						
Treated (Recode)	8.716*** (0.454)	7.608*** (0.469)	8.947*** (0.485)	8.818*** (0.452)	8.647*** (0.455)	8.370*** (0.441)
<i>Panel B: Reduced Form (Planned Participation)</i>						
Treated (Recode)	3.229** (1.368)	3.122** (1.361)	4.191*** (1.463)	3.250** (1.372)	3.177** (1.368)	3.053** (1.361)
<i>Panel C: 2SLS (Planned Participation)</i>						
Posterior Beliefs	0.370** (0.154)	0.410** (0.175)	0.468*** (0.157)	0.369** (0.152)	0.367** (0.155)	0.365** (0.159)
Observations	2,422	2,420	2,386	2,422	2,361	2,397
School Share	No	Yes	No	No	No	No
District FEs	No	No	Yes	No	No	No
Recode Exacts	No	No	No	Yes	No	No
Drop Attention Check	No	No	No	No	Yes	No
Drop Outliers	No	No	No	No	No	Yes

Note: This table reproduces the main results with alternative specifications or samples. Panel A presents the first stage results on posterior beliefs, panel B presents the reduced form effect on planned participation, and panel C presents the 2SLS results of a change in posteriors on planned participation. Column 1 reproduces the main specification as in column 3 of Tables 3–5. Column 2 replaces the district-level priors and posteriors with the corresponding prior and posterior beliefs about union membership rates within a teacher’s school. Column 3 includes district fixed effects and clusters standard errors at the district level. Column 4 recodes the people who guessed the district share perfectly from 1 to 0. Column 5 drops those respondents who failed the attention check. Column 6 drops those observations for which the gap between the district share and their prior belief was greater than the 99th percentile. All specifications use robust standard errors and column 3 clusters standard errors at the district level. *, **, *** denote statistical significance at the 10, 5, and 1 percent levels, respectively.

B. Data Appendix

A. Measuring Union Membership

To measure union membership, I link datasets from the Wisconsin Department of Public Instruction and campaign contributions data from the Wisconsin Campaign Finance Information System. For the campaign contributions data, I use contributions to the state teachers' union (WEAC PAC) and any of its regional affiliates. The unique information common to both datasets is a person's full name and the regional WEAC chapter they work for in a given year. For the DPI data, this information is linked via district identifiers from here: <https://weac.org/region-finder-test/>.

In the DPI data, 99 percent of people are uniquely identified by name, year, and WEAC region. In the experimental sample, there are 25 individuals who are not uniquely identified by name, year, and WEAC region. For these duplicates, there are three possible scenarios: (1) The number of duplicates matches the number of unique people in the WCFIS data. (2) No one appears in the WCFIS data. (3) There is one or more people in the WCFIS data, but it is not immediately clear which person matches which teacher. In the first case, all of the duplicates are considered members, while in the second, no one is a member. In the third case, I denote membership status by which teacher's address is closest to their respective school. There are six observations that meet this latter criteria.

The campaign contributions data is updated at least once every quarter of the year. I denote someone as a member in the 2023–24 school year if they are in the contribution data in the fall of 2023 or January 2024. The survey went out in the spring of 2024, meaning that the later reports were potentially affected by the survey and treatment.

B. Experimental Analysis Sample

There were 2,789 individuals who filled out at least one question of the survey and 2,525 who completed enough of the survey such that they included in the experimental randomization. In the analysis sample, I drop 20 respondents who were not licensed teachers as support staff often have no union representation or different unions than licensed teaching staff. I also drop 13 people who could not be matched to the administrative personnel data. In the main analysis, I also exclude 70 people who did not answer either the prior belief, posterior belief, or planned participation question.

C. Data Variables

Advanced degree: An indicator variable equal to 1 if the teacher had a master's degree or higher.

Attrited: An indicator variable equal to 1 if the respondent was randomized into either treatment or control but did not answer one or more of the main questions relating to prior beliefs, posterior beliefs, or planned participation.

Experience: The number of years that a teacher has worked in their local school district.

FTE: A number corresponding to a teacher's full-time employment status, where 100 equals full-time.

Male: An indicator variable equal to 1 if the teacher was coded as male.

Math teacher: An indicator variable equal to 1 if the teacher worked as a math teacher. Teachers who teach multiple subjects, such as elementary teachers, are included.

Baseline member: An indicator variable equal to 1 if the teacher was a union member at baseline according to the campaign finance reports. I consider teachers as members at baseline if they appear in the fall or January reports from the 2023–24 school year. (The January reports cite that the contributions are from August to December 2023).

Non-white: An indicator variable equal to 1 if the teacher’s race in the administrative data was not “white.”

Number positions: The number of positions that a teacher worked. For instance, a teacher may work part-time in two different schools.

Planned participation: An indicator variable equal to 1 if the respondent answered “very likely” or “somewhat likely” to the question regarding the likelihood of being in the union next year.

Posterior Belief, District: A number ranging from 0 to 100 of a teacher’s estimate of the union membership rate in their school district next school year (2024–25). See Appendix Section C for a screenshot of the posterior belief question.

Posterior Belief, School: A number ranging from 0 to 100 of a teacher’s estimate of the union membership rate in their school next school year (2024–25). See Appendix Section C for a screenshot of the posterior belief question.

Prior Belief, District: A number ranging from 0 to 100 of a teacher’s estimate of the union membership rate in their school district last school year (2022–23). See Appendix Section C for a screenshot of the prior belief question.

Prior Belief, School: A number ranging from 0 to 100 of a teacher’s estimate of the union membership rate in their school last school year (2022–23). See Appendix Section C for a screenshot of the prior belief question.

Pro-social: A variable from 0–2 indicating a survey participant’s pro-sociality. The score is determined by whether the respondent shared at least 50 percent of the hypothetical budget surplus with any group. The score equals 1 if they shared 50 percent with one group and 2 if they shared 50 percent with at least two groups. See Appendix Section C for a screenshot of the pro-sociality question.

Reading teacher: An indicator variable equal to 1 if the teacher worked as a reading or English and language arts teacher. Teachers who teach multiple subjects, such as elementary teachers, are included.

Salary: A teacher's gross salary in dollars.

C. Survey Details

Below, I display screenshots of the Qualtrics survey questions used in the paper. Before randomization, I asked questions designed to proxy for teachers' pro-sociality. The following screenshot shows this, in which survey respondents moved the sliders to allocate a budget surplus in three hypothetical scenarios (Enke, Rodriguez-Padilla, and Zimmermann 2023).

Suppose the school district budget increased by 10 percent. If you could only spend the money on two options, how would you split the budget surplus between teacher pay on the left and the alternative choice on the right in the following three scenarios?

(The closer you move the slider to one side, the more you allocate to that option.)



The next screenshot shows the question used to elicit prior beliefs about the district and school membership rates:



What percentage of teachers in your **school district** do you think were members of the teachers' union last year?

Please enter a number between 0 and 100.

 %

What percentage of teachers in your **school** do you think were members of the teachers' union last year?

Please enter a number between 0 and 100.

 %

After eliciting prior beliefs, the treatment group was presented with the information provision seen below (control group teachers did not see this screen).



Our best estimate is that **% of teachers in your school district** were members of the teachers' union last year.*

[*Methodological notes](#)

Before the “%” sign was a number that corresponded to the district-specific membership rate. If one clicked on the “Methodological notes” hyperlink, they saw the text: “Chapter 11 of Wisconsin state statute requires that political action committees itemize the source of all funds received as part of campaign finance reporting. The Wisconsin Campaign Finance Information System makes these reports publicly available semi-annually. The Wisconsin Educational Association Council sends a small share of annual membership dues to the WEAC political action committee, unless a member requests in writing a refund of the PAC contribution. Therefore, the estimate of a district’s membership rate is constructed by linking the WEAC PAC listings with publicly available data from the Wisconsin Department of Public Instruction “All Staff Report.”

Next, all respondents (both teachers and control) saw the following questions to elicit posterior beliefs and planned union participation in the next school year:



What percentage of teachers in your **school district** do you think will be members next year?

Please enter a number between 0 and 100.

 %

What percentage of teachers in your **school** do you think will be members next year?

Please enter a number between 0 and 100.

 %

How likely is it that you will be a member of the teachers' union next year?

☐ Very likely

☐ Somewhat likely

☐ Undecided

☐ Somewhat unlikely

☐ Very unlikely