

Putting social rewards and identity salience to the test: Evidence from a field experiment with teachers in Philadelphia

Syon P. Bhanot^{*‡}, Gordon Kraft-Todd^{†‡}, David Rand^{†‡}, Erez Yoeli[‡]

Abstract: We partnered with the School District of Philadelphia (SDP) to run a randomized experiment testing interventions to increase teacher participation in an annual feedback survey, an uncompensated task that requires a teacher's time but helps the educational system overall. Our experiment varied the nature of the incentive scheme used, and the associated messaging. In the experiment, all 8,062 active teachers in the SDP were randomly assigned to receive one of four emails using a 2x2 experimental design; specifically, teachers received a lottery-based financial incentive to complete the survey that was either "personal" (a chance to win one of fifteen \$100 gift cards for themselves) or "social" (a chance to win one of fifteen \$100 gift cards for supplies for their students), and also received email messaging that either did or did not make salient their identity as an educator. Despite abundant statistical power, we find no discernible differences across our conditions on survey completion rates. One implication of these null results is that from a public administration perspective, social rewards may be preferable since funds used for this purpose by school districts go directly to students (through increased expenditure on student supplies), and do not seem less efficacious than personal financial incentives for teachers.

Keywords: Behavioral policy, Identity, Social Incentives, Field experiment, Education

Academic researchers in economics and psychology have increasingly explored the use of light-touch behavioral interventions to influence individual behavior. These efforts have led to the development of a number of tools to address the many barriers to behavior change, including framing manipulations in messaging (Kahneman & Tversky, 1981), harnessing social incentives (Bandiera et al., 2010, Ashraf et al., 2014, Kraft-Todd et al., 2015, Bursztyn & Jensen, 2015), and making personal identities salient (Steele & Aronson, 1990, Spencer & Castano, 2007, Carr & Steele, 2010, Benjamin et al., 2010, Cohn et al., 2015, Benjamin et al., 2016, Kessler & Milkman, 2016). However, while behavioral interventions have shown significant promise, the body of evidence supporting their applicability across a variety of real-world contexts remains relatively thin.

In the policy domain, behavioral interventions have increasingly been considered as tools to promote contributions to public goods specifically. This is often motivated by the mixed evidence on the efficacy of (small) financial incentives in field experiments (Kraft-Todd et al., 2015). Furthermore, even when they are effective, small financial incentives at the individual level can aggregate together to be quite costly when utilized at scale. Thus, if the expected impacts of financial incentives are modest, and the expected costs non-trivial, the cost effectiveness of financial incentives can be very low. This reality has encouraged academics and policymakers to consider utilizing social incentives for behavior change, with the argument that they can be cost effective tools for addressing policy challenges (Allcott, 2011, Ashraf et al., 2014, Hallsworth et al., 2017). Recent research suggests that social incentives can be particularly efficacious in the context of prosocial behavior (for a review, see Kraft-Todd et al., 2015, who discuss the inconsistent evidence supporting the use of financial incentives to promote prosocial behavior, contrasting it with the more promising impacts of social motivators in this area).

* Swarthmore College;

‡ Applied Cooperation Team, Yale University

† Psychology Department, Yale University;

Address correspondence to Syon Bhanot at
(sbhanot1@swarthmore.edu)

Copyright: © 2018. The authors license this article under the terms of the Creative Commons Attribution 4.0 International License.

One alternative approach would be to offer monetary “social rewards” rather than personal rewards as an incentive for certain behaviors (Anik et al., 2013). An example of this might be offering someone a \$50 donation in their name to charity in exchange for doing a certain behavior, instead of \$50 in cash for that behavior. This approach offers a number of benefits. First, it leverages individuals’ prosocial motivations and avoids some of the potential pitfalls of financial incentives, particularly in the domain of prosocial behaviors. For example, one might worry that personal financial incentives for collectively-beneficial behavior might “crowd out” intrinsic motivation (Gneezy & Rustichini, 2000). Second, social rewards might be preferable from the perspective of the policymaker, even if they are not more efficacious than personal incentives, because money used for social rewards can improve social welfare more directly than money spent on private rewards. This is especially true when the social reward is something like a charitable contribution or increased funding for public goods.

Another form of behavioral intervention that is increasingly common in the literature is the use of identity salience manipulations as a way to change behavior. This research suggests that making a specific “identity” salient at the moment of decision making might encourage individuals to conform to that identity by engaging in the behaviors associated with it (Benjamin et al., 2010, Cohn et al., 2015, Benjamin et al., 2016, Kessler & Milkman, 2016). By this logic, when seeking to encourage prosocial decisions, it might be effective to make an identity that is aligned with prosocial choices more salient before asking for prosocial action. For example, in Kessler and Milkman (2016), the authors test the impact of making a “giving identity” salient in a charitable donation solicitation by reminding treatment subjects of their previous donations. They find that this manipulation significantly increased giving rates, suggesting that reminders of one’s identity might have tangible impacts on prosociality. Furthermore, evidence suggests that identity salience interventions, in addition to being efficacious on their own, may be useful in enhancing other behavioral interventions, especially ones intended to enhance prosociality. Consistent with this, many Get Out the Vote Letters begin with an identity salience manipulation such as, “You are a voter!” (Bryan et al., 2011).

In this article, we describe an intervention involving an email campaign designed to test the impact of social versus personal rewards and an identity salience manipulation on prosocial behavior change amongst school teachers—namely, encouraging teachers to complete a 30-minute annual survey that benefits the school district by providing information that helps improve the educational system. All 8,062 teachers¹ in the School District of Philadelphia (SDP) were included in the study sample. Our randomized design allows us to compare the impact of a monetary reward implemented in the form of supplies for teachers’ students (social) versus money for the teacher (personal). Furthermore, we test whether or not making a teacher’s professional identity as an educator more salient influences response rates, and how this identity salience manipulation might interact with the efficacy of social versus personal rewards. Our article contributes to the growing body of literature using randomized evaluations to learn more about how actors in the educational system make decisions and how behavioral manipulations can shape educational outcomes (Fryer et al., 2012, Kraft & Rogers, 2015, Gehlbach et al., 2016, Levitt et al., 2016). Furthermore, our results speak to literature in public administration on the motivations underlying public service, and how interventions based on behavioral science might be more effective than pay-for-performance when it comes to motivating public sector employees (Ritz et al., 2016, Grant, 2008).

There are useful implications of our specific manipulations for the management of public education systems. First, if identity salience manipulations motivate teachers to act prosocially, such techniques would be useful levers for school districts looking to motivate teacher behavior change. Second, if social rewards are (at least) as impactful as personal financial rewards in motivating teachers to take prosocial actions, it would support the substitution of social rewards in place of personal rewards for teachers where the latter currently exist, as social rewards arguably better serve the educational aims of school districts given that the money expended on these rewards often go directly to the students.

Contrary to our expectations, our main findings are null results: social and personal rewards had roughly the same impact on survey completion, and identity salience had no meaningful effect on survey completion. Taking advantage of this large

dataset, we also conduct non-experimental, exploratory analyses to determine which school characteristics predict survey completion, in order to detect potential moderator variables for our treatments. We find that teachers at schools with higher parent and student satisfaction and with higher teacher ratings are more likely to complete these surveys, while teachers at schools where students have fewer behavioral issues (e.g. lower suspension rates and higher attendance rates) are less likely to complete these surveys.

Based on these exploratory results, we conduct additional analyses of our treatment effects that provide some potentially useful (although far from definitive) insights. We find some evidence that the identity salience manipulation actually lowered survey completion rates at schools with relatively lower-rated teachers (which was determined using an SDP metric on teacher effectiveness), while having a more positive effect on survey completion rates at schools with highly-rated teachers (determined using the same SDP metric). These findings suggest that identity manipulations like ours may be better suited for use at schools with better teachers. Finally, we find minimal evidence of an interaction between our identity salience manipulation and past survey completion by teachers. That is, the point estimate is positive, but also very small, suggesting that our identity manipulation did not trigger a meaningful “consistency” motivation (Gneezy et al., 2012, Freedman & Fraser, 1966, Mullen & Monin, 2016) for teachers who completed the survey last year.

Taken together, these results suggest that the identity manipulation we designed may not be of practical use outside of higher quality schools, but that social rewards may be preferable to personal rewards as a tool to motivate teachers. That is, while both reward types motivated teachers equally well, social rewards directly increase student welfare more than personal rewards. More broadly, our findings suggest that while the use of behavioral interventions in public administration shows promise, there is room for further testing and development of effective, scalable interventions that motivate public sector employees (Ritz, et al., 2016).

Experiment Design

Implementing Partners

We worked with two institutional partners to plan, design, and implement the randomized experiment

we present here. First, we developed the intervention directly with the School District of Philadelphia (SDP), who served as the implementing partner. Second, we received institutional support from the Philadelphia Mayor’s Office, through GovPHL and the Philadelphia Behavioral Science Initiative (www.phillybsi.org), a broader effort to integrate behavioral science into public policy through collaborations between academic researchers and city policymakers.

Subjects, Context, and Design

The study’s sample was the full population of teachers employed by the School District of Philadelphia—a total of 8,062 teachers. Every spring, Philadelphia teachers receive an email from the school district with a link to an “end-of-year survey,” designed to elicit their feedback about various issues affecting schools. This survey is one of the primary channels through which the school district learns about what is happening in schools across the city. Therefore, increasing engagement with this survey is a key policy priority for the city. From an academic perspective, completing the survey can be thought of as a “cooperative” behavior on the part of the teacher; that is, it involves the teacher bearing a personal time-and-effort cost to generate a collective benefit. Our intervention involved the integration of a randomized experiment into the standard annual procedure of sending emails to teachers about the survey, namely through manipulations of the messaging content of the email and the rewards used to motivate survey completion.

We randomly assigned teachers to one of four treatment groups,² with each group receiving a different type of email (and three ensuing email reminders with consistent messaging). Our intervention used a 2x2 factorial design, in which we varied the type of reward used (“personal” vs. “social”) and whether or not a “teacher identity” was made salient in the language in the email.

The experimental conditions are shown in Table 1. Subjects in the personal rewards conditions were offered the chance to win one of fifteen \$100 Barnes and Noble gift cards for personal use, while subjects in the social rewards treatment were offered the chance to win one of fifteen \$100 Office Depot gift cards to purchase school supplies for their students. To distinguish between treatment groups, subjects who received the personal reward and no identity manipulation will hereafter be called the “Standard” group, those who received

Table 1
Experimental conditions (2x2 design)

	Personal Financial Incentive	Social Financial Incentive
Standard Email Language	“Standard” (2,025 subjects)	“Social Rewards Only” (1,999 subjects)
Identity Salience Email Language	“Identity Only” (2,010 subjects)	“Social Rewards + Identity” (2,028 subjects)

the social reward and no identity manipulation will be called the “Social Rewards Only” group, those who received the personal reward and the teacher identity manipulation will be called the “Identity Only” group, and those who received both the social reward and identity manipulation will be called the “Social Rewards + Identity” group.

Note that there was no pure “control” group that did not receive emails (or incentives of some form) as part of the intervention. The reason for this was two-fold. First, the SDP did not want to offer some teachers a financial reward to complete the survey and not offer other teachers a similar financial reward (though variation in the nature of the reward was deemed acceptable). Second, from an experimental design perspective there was a concern about spillover effects (through word of mouth) if we had a control group with no incentives. Also note that the emails sent to teachers included instructions to take the survey through an employee portal. Copies of the exact emails sent are included in the Appendix.

The initial treatment emails were sent on April 4, 2017, with the follow up reminder emails sent on April 25, May 11, and May 25. The outcome variable we measured was survey response for teachers at the individual level. There are no measurement concerns with this metric, because it came directly from reliable administrative data and represents an explicit measurement of the policy objective from the perspective of the SDP.

Power and Minimum Detectable Effects

Our 2x2 research design was structured to test 3 different comparisons (social vs. financial rewards, overall; teacher identity salience vs. no identity salience, overall; and the interaction of social rewards and teacher identity salience). Given that the sample size in our study was fixed, we conducted an ex-ante power analysis for each of these three pairwise

comparisons both without and with a Bonferroni correction for the three comparisons to determine the minimum detectable effects (MDEs) from our intervention. Based on conversations with the SDP, an effect size of roughly five percentage points was agreed upon as being practically significant for policy and therefore formed the basis of our assessment of MDEs. Note that using the Bonferroni correction of the p-value is widely regarded as providing very conservative estimates of power (Anderson, 2008).

We used data from the 2015 and 2016 teacher surveys to come up with an ex-ante expected completion rate for the survey in 2017. Specifically, 54% of SDP teachers completed the survey in 2016 and 57% completed the survey in 2015, so we used an estimate of 54% for our power calculations. The results are presented in Table 2. The MDEs were close to our five percentage point target for practical significance for the SDP.

Table 2
Minimum detectable effects

	MDE (Percentage Points)	
Sample Size	Without Bonferroni	With Bonferroni
8062	0.0310	0.0358
4031	0.0438	0.0505

Notes: This table shows the minimum detectable effects (MDEs) in percentage points, assuming power of 80 percent, when the full sample of 8,062 teachers is considered and when only 4,031 are considered (half the sample, for analysis of treatment effects, within a given condition, of the other condition).

Randomization

The method used in this intervention was a simple randomization carried out by the researchers using Stata and transferred into Excel. The randomization was carried out at the individual teacher level and then shared with the SDP for implementation.

Data Collection and IRB Issues

All required data is regularly collected by the SDP. We obtained IRB approval from the Swarthmore College IRB to receive the data from the school district and conduct the data analysis.

We received anonymized administrative data at the teacher-level from the SDP on survey completion, along with treatment assignment and some basic information about each teacher (namely, whether they taught in the SDP in the academic year prior, whether they completed the survey in the academic year prior, and what school they taught in). We also used public data on school performance metrics and characteristics from the SDP Open Data Initiative³ to gather details at the school level, which we merged with the teacher-level data to supplement our analysis.

Hypotheses

Our hypotheses are based on the large body of literature on the efficacy of identity salience manipulations and personal versus social incentives, drawing from both rational and behavioral models of decision making. We had three primary hypotheses. First, we hypothesized that the social rewards treatment would outperform the personal rewards (in line with Kraft-Todd et al., 2015). Note that a model that assumes pure self-interest would predict the opposite (namely that personal financial incentives would outperform social incentives). Second, we hypothesized that the identity salience manipulation would increase prosocial behavior, in the form of survey response by teachers (Kessler & Milkman, 2016). Third, we hypothesized that the identity salience manipulation would amplify the effectiveness of social rewards versus personal rewards. We also had one secondary hypothesis *ex-ante*, namely that the identity salience manipulation would have a greater impact on individuals who completed the survey in 2016, as it serves to reinforce the importance of behavioral consistency (Gneezy et al., 2012).

We also developed one *ex-post* hypothesis based on exploratory data analysis around the predictors of survey completion. Specifically, having

found that teachers from higher-quality schools (i.e., schools with more highly-rated teachers and higher parent/student satisfaction in SDP performance metrics), were more likely to complete the survey in general, we hypothesized that both the social rewards and identity salience manipulations would be more impactful for teachers from these schools. The logic for this was that our manipulations would have a greater impact on teachers with more prosocial preferences and/or a stronger commitment to their careers as educators, relative to other teachers. We are able to provide suggestive evidence for this hypothesis.

Analytical Approach

In order to assess the success of the intervention, we use simple linear probability model regressions (OLS regressions with a binary outcome). We also provide results from logit models in the Online Supplement, for robustness. Because of randomization, and the very limited data available at the teacher level based on the data sharing agreement with the SDP, the specifications are quite straightforward, and are as follows:

$$y_i = \beta_0 + \beta_1 \cdot \text{SocialRewards}_i + \beta_2 \cdot \text{2016Completion}_i + \beta_3 \cdot \text{2016Ineligible}_i + \gamma + \varepsilon \quad (1)$$

$$y_i = \beta_0 + \beta_1 \cdot \text{Identity}_i + \beta_2 \cdot \text{2016Completion}_i + \beta_3 \cdot \text{2016Ineligible}_i + \gamma + \varepsilon \quad (2)$$

$$y_i = \beta_0 + \beta_1 \cdot \text{IdentityOnly}_i + \beta_2 \cdot \text{SocialRewardOnly}_i + \beta_3 \cdot (\text{SocialRewards} + \text{Identity})_i + \beta_4 \cdot \text{2016Completion}_i + \beta_5 \cdot \text{2016Ineligible}_i + \gamma + \varepsilon \quad (3)$$

$$y_i = \beta_0 + \beta_1 \cdot \text{Identity}_i + \beta_2 \cdot \text{GoodTeachers}_i + \beta_3 \cdot (\text{Identity}_i \cdot \text{GoodTeachers}_i) + \beta_4 \cdot \text{2016Completion}_i + \beta_5 \cdot \text{2016Ineligible}_i + \varepsilon \quad (4)$$

$$y_i = \beta_0 + \beta_1 \cdot \text{SocialRewards}_i + \beta_2 \cdot \text{GoodTeachers}_i + \beta_3 \cdot (\text{SocialRewards}_i \cdot \text{GoodTeachers}_i) + \beta_4 \cdot \text{2016Completion}_i + \beta_5 \cdot \text{2016Ineligible}_i + \varepsilon \quad (5)$$

$$\begin{aligned}
y_i = & \beta_0 + \beta_1 \cdot \text{Identity}_i \\
& + \beta_2 \cdot \text{2016Completion}_i \\
& + \beta_3 \cdot (\text{Identity}_i \cdot \text{2016Completion}_i) \\
& + \beta_4 \cdot \text{2016Ineligible}_i + \gamma + \varepsilon
\end{aligned} \quad (6)$$

Note that γ in specifications 1, 2, 3, and 6 refer to school fixed effects, while the variables “2016Completion” and “2016Ineligible” together constitute a control for survey completion behavior of a given teacher in the previous year (teachers could either have done the survey the previous year, not done it the previous year when eligible, or not been eligible to do it the previous year because they did not work at the SDP at the time).

Specifications 1-3 are all simple measurements of average treatment effects, using OLS regressions. Specification 1 estimates the causal impact of social rewards as a main effect, with controls for last year’s survey completion behavior and school fixed effects. Specification 2 does the same but for the identity salience manipulation. Specification 3 includes all experimental conditions, and therefore includes the interaction condition involving both manipulations. Note that we also run the analysis from specifications 1 and 2 separately for teachers from schools of various overall performance levels (or “tiers”), as reported publicly by the SDP, to assess how measures of school quality

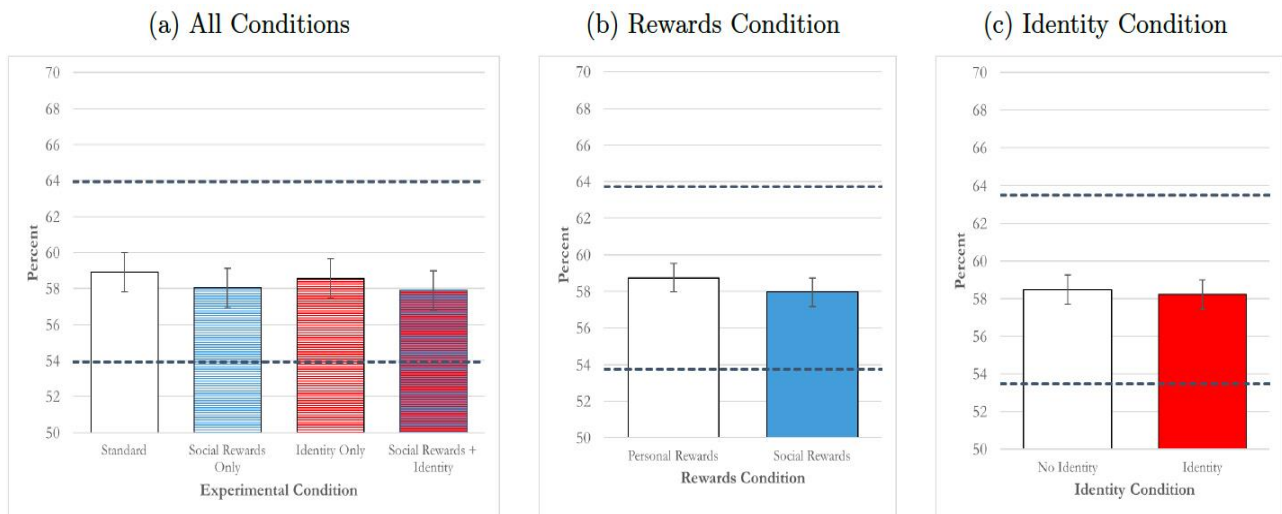
might interact with intervention efficacy (one of our secondary questions of interest).

Specifications 4-6 allow us to conduct exploratory analyses involving interactions between the experimental conditions and two important baseline variables. First, Specifications 4 and 5 interact the two main effects with a dummy variable identifying teachers who work in schools with more highly-rated teachers, a designation that we determined using a school-level metric on teacher effectiveness from the 2015-2016 SDP School Progress Report (SPR). Specifically, we used the SPR measure that reported on the percentage of teachers at a given school who received an effectiveness rating of “distinguished,” and identified schools with high teacher quality as those whose percentage of “distinguished” teachers was above the median. Second, Specification 6 interacts the identity main effect with a dummy variable for whether or not a teacher completed the annual teacher survey in 2016, the year prior to this experiment. This was done to determine if identity salience manipulations were more impactful when they aligned with the idea of “consistency” with past behavior (in this case, completing the survey last year).

Results

Figure 1 and Table 3 present the results related to

Figure 1
Survey completion by condition (std. error marked)



Notes: The dotted lines indicate ± 5 pp from the Standard treatment mean, which we determined a priori to be the threshold for a meaningful effect.

Table 3
Average treatment effects

	Identity Manipulation		Social Rewards Manipulation		All Conditions	
	(1)	(2)	(3)	(4)	(5)	(6)
	2017 Comp.	2017 Comp.	2017Comp.	2017 Comp.	2017 Comp.	2017 Comp.
Identity	-0.00252 (0.0110)	-0.00203 (0.00943)				
Social Rewards			-0.00777 (0.0110)	-0.00804 (0.00937)		
Identity Only					-0.00356 (0.0155)	-0.00108 (0.0135)
Social Rewards Only					-0.00885 (0.0155)	-0.00710 (0.0132)
Social Rewards + Identity					-0.0102 (0.0155)	-0.0100 (0.0132)
Completed in 2016		0.208*** (0.0113)		0.208*** (0.0113)		0.208*** (0.0113)
Ineligible in 2016		0.154*** (0.0173)		0.154*** (0.0173)		0.154*** (0.0173)
Constant	0.585*** (0.00777)	0.731*** (0.0571)	0.587*** (0.0078)	0.734*** (0.0570)	0.589*** (0.0109)	0.734*** (0.0574)
Observations	8062	8062	8062	8062	8062	8062
R^2	0.000	0.327	0.000	0.327	0.000	0.327
School Fixed Effects	No	Yes	No	Yes	No	Yes

Notes: This table shows the main results from this experiment, in the form of average treatment effects, using linear probability models. Specifications 1-2 show the average treatment effects of the Identity manipulation only, 3-4 show those for the Social Rewards manipulation only, and 5-6 show all treatment conditions (with the condition involving both manipulations), respectively. Regressions with and without controls are included—the controls are: 1) dummy variables for survey completion and ineligibility for the survey (meaning not employed by SDP) in the previous year (the omitted group being teachers eligible but not completing in the previous year); and 2) school fixed effects.

our main three hypotheses. As is apparent from visual inspection of Figure 1, neither the social rewards treatment nor the identity salience manipulation increased survey completion—if anything, they reduced it slightly. This is confirmed in the regression results in Table 3: in the specifications that estimate main effects with controls (columns 2 and 4), the coefficients on the identity and social rewards treatments are -0.2 and -0.8 percentage points, respectively. The results of all other specifications are quite similar. Likewise, the identity salience manipulation does not positively interact with the social rewards treatment. If anything, the interaction is slightly negative: the coefficient on

the interaction is roughly -1 percentage point (Table 3, columns 5 and 6).

To better understand how school characteristics predict survey completion, we next conduct non-experimental, exploratory analyses to detect potential moderator variables for our treatments. For this analysis, we use 21 school-level variables for which we had at least 7,000 teacher observations (representing 87% of our sample), and two teacher-level variables capturing survey completion behavior in the previous year. In Table 4, we present the results of these analyses. Column 1

Table 4
Linear probability models

		Single LPM	Multiple LPM	Stepwise LPM	Unique factor loading >.5
School-level predictors	Climate Score (SPR)	-0.063***	0.308***	0.300***	1
	% Students w Attendance>95%	-0.094***	-0.253***	-0.244***	1
	Achievement Score (SPR)	-0.051***	-0.262***	-0.238***	1,3
	% Students w/ Positive view of Teacher Quality	0.057***	0.197***	0.226***	2
	Climate Rating (Students)	0.074***	-0.188***	-0.222***	2
	Student Retention Rate	-0.083***	-0.153***	-0.154***	1
	% Students Economically Disadvantaged	-0.043***	-0.128***	-0.122***	3
	% Teacher Effectiveness of Distinguished	0.046***	0.115***	0.097***	3
	Enrollment	-0.098***	-0.101***	-0.094***	
	Teacher Attendance Rate	0.093***	0.085***	0.085***	
	% Teacher Observation of Distinguished	0.034*	0.037	0.045*	3
	Climate Rating (Parents)	0.050***	-0.057**		2
	Parent Survey Participation	0.082***	0.047*		2
	% Students in Special Education	0.094***	-0.034		1
	% Students w/ Zero In-School Suspensions	0.003	-0.030		
	Serious Incidents (per capita)	0.050***	-0.029		1
	% Students w/ Zero Out-of-School Suspensions	-0.061***	-0.018		1
	% Students Female	0.013	-0.015		
	% Students Black	0.054***	0.010		
	Suspensions (per capita)	0.064***	-0.008		1
	Progress Score (SPR)	-0.042***	-0.002		
Teacher-level predictors	Completed 2016 Survey (if eligible)	0.246***			
	Completed 2016 Survey		0.228***	0.228***	
	Ineligible for 2016 Survey	0.031	0.116***	0.114***	
	Observations	>7129	7101	7101	
	R-squared		0.121	0.117	

Notes: This table shows linear probability models of school-level variables predicting teacher survey completion using single (Column 1), multiple (Column 2), and stepwise (Column 3) regression (using $p < .00434$ as removal criteria), as well as factor analysis (Column 4). Predictors and dependent measures are standardized; listed coefficients represent the change in standard deviations of survey completion for each standard deviation change in the predictor variable. A Bonferroni correction is used for 23 multiple comparisons and the variables are presented in the order of statistical significance. * $p < .00434$, ** $p < .00217$, *** $p < .000434$.

presents the standardized coefficients from individual single-variable regressions of survey completion on each of the 23 school- and teacher-level predictors. Column 2 presents the standardized coefficients from a multivariable regression of survey completion on all variables in column 1 at once.⁴ Column 3 presents the results of a stepwise regression of survey completion on the same variables, ordered by the magnitude of the effect from column 1. We also explore the underlying structure of the relationship between school characteristics and

survey completion using factor analysis with varimax (orthogonal) rotation. The analysis yielded three factors explaining 83% of the variance in all school characteristics. Column 4 presents the unique factor loadings of each of the school-level predictors when greater than 0.5. We see that the three factors map, respectively, onto metrics associated with: (1) good student behavior (student attendance, retention, etc.); (2) parent/student satisfaction (student evaluations of teachers and school climate, parent evaluations of school climate, etc.);

and (3) quality teachers (teacher effectiveness ratings at the school level, etc.). Follow-up analyses find that good student behavior negatively predicts survey completion, whereas the other two factors positively predict survey completion.⁵ See Figure 1 in the Online Supplement for details.

We use these exploratory results to guide an investigation into whether our treatments had stronger effects for certain subpopulations. Specifically, given that a teacher having completed the survey in the previous year and various school quality measures were important predictors of survey completion, we form a secondary hypothesis that our manipulations would have larger effects at schools with more highly-rated teachers and better overall performance metrics. The motivating idea

here is that our treatments might work better when teachers feel a stronger sense of prosocial commitment to their students, which may be more likely with better teachers or at better schools.

Table 5 presents one test for this hypothesis, using interaction effects between teacher quality at the school level and our manipulations to determine if there was a larger impact from the manipulations at schools with more highly-rated teachers. To measure teacher quality here, we use a dummy variable marking schools as having “high quality” teachers if the percentage of teachers at that school receiving a “distinguished” evaluation in teacher effectiveness (an SPR measure) was above the median in Philadelphia. Note that we do

Table 5
Interaction effects

	High Teach. Qual. School x Treatments		2016 Comp. x Identity
	(1) 2017 Comp.	(2) 2017 Comp.	(3) 2017 Comp.
Identity x High Teach. Qual. School	0.050** (0.022)		
Social Rewards x High Teach. Qual. School		-0.0053 (0.022)	
Identity x 2016 Completion			0.0052 (0.020)
High Teach. Qual. School	0.026* (0.016)	0.054*** (0.016)	
Identity	-0.033** (0.016)		-0.0078 (0.015)
Social Rewards		-0.0050 (0.016)	
Completed in 2016	0.24*** (0.012)	0.24*** (0.012)	0.20*** (0.015)
Ineligible in 2016	0.20*** (0.019)	0.20*** (0.019)	
Constant	0.44*** (0.013)	0.43*** (0.013)	0.76*** (0.055)
Observations	7616	7616	7130
R ²	0.056	0.056	0.340
School Fixed Effects	N/A	N/A	Yes

Notes: This table shows the results from linear probability models evaluating the interactions between the pooled treatments and two baseline characteristics: 1) specifications 1-2 interact each of the manipulations with a 2015-2016 metric for high teacher quality at the school level (a dummy variable identifying schools with an above-median percentage of teachers getting a "distinguished" rating); and 2) specification 3 interacts the identity manipulation with whether or not a teacher completed the survey in the previous year, 2016, to test for the presence of "consistency" as a motivation (note this regression omits teachers not employed by SDP in 2016). Specification 3 includes school fixed effects.

not have individual-level measures of teacher quality. Also note that Table 4 also presents the results for our secondary question of interest, regarding the importance of “consistency” as a behavioral motivation (measured using the interaction of past survey completion with the identity manipulation).

We do find a fairly large positive interaction between the identity salience manipulation and teacher quality at the school level: the coefficient on the interaction is 5.0 percentage points (Table 5, column 1). However, note that the coefficient of the identity salience manipulation (indicating the impact of the identity salience manipulation in schools with relatively lower teacher ratings) is -3.3 percentage points in this specification. Thus, the identity salience manipulation has, on net, a small positive effect in schools where teachers were more highly-rated. We find no meaningful interaction between the social rewards treatment and teacher quality at the school level, however (Table 5, column 2), or between the identity salience manipulation and prior-year survey completion (Table

5, column 3). We interpret these results as evidence that our manipulations (and identity salience in particular) were relatively less efficacious in schools where teachers received lower ratings on average, though this is more suggestive than definitive.

To further explore the link between school quality and the effectiveness of our manipulations, we present disaggregated treatment effects in Table 6. Specifically, in this analysis we disaggregate by a measure of school quality: namely, which overall “tier” of school performance the school had been placed in by the SDP. This tier classification was based on publicly-available “School Progress Report” metrics generated by the city, and ranged from the lowest tier (‘intervene’) to intermediate tiers (‘watch’ and ‘reinforce’) to the highest tier (‘model’). This analysis serves as a “second test” of whether our manipulations varied in efficacy based on school quality, broadly defined. In this case, we do not observe much difference in the estimates across schools of different quality,

Table 6
Disaggregated ATEs by school quality (SPR Tier)

	School Progress Report Category							
	(1) Intervene	(2) Intervene	(3) Watch	(4) Watch	(5) Reinforce	(6) Reinforce	(7) Model	(8) Model
Identity	-0.00806 (0.0154)		-0.00270 (0.0153)		-0.00264 (0.0244)		0.0225 (0.0422)	
Social Rewards		-0.00871 (0.0153)		-0.00413 (0.0153)		-0.00951 (0.0242)		-0.0134 (0.0413)
Completed in 2016	0.219*** (0.0181)	0.219*** (0.0181)	0.190 (0.0187)	0.190*** (0.0187)	0.248*** (0.0301)	0.248*** (0.0302)	0.0896** (0.0449)	0.0909* (0.0449)
Ineligible in 2016	0.189*** (0.0250)	0.189*** (0.0250)	0.153 (0.0302)	0.153*** (0.0302)	0.121*** (0.0518)	0.121** (0.0516)	-0.0727 (0.107)	-0.0683 (0.1080)
Constant	0.727*** (0.0583)	0.727*** (0.0583)	0.946 (0.0255)	0.947*** (0.0256)	0.785*** (0.0370)	0.789*** (0.0356)	0.768*** (0.0726)	0.787*** (0.0752)
Observations	3000	3000	2977	2977	1220	1220	361	361
R ²	0.280	0.280	0.338	0.338	0.331	0.332	0.273	0.272
School Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows disaggregated average treatment effects by school quality, using the 2015-2016 School Progress Report score tier category (defined by the SDP). Linear probability model results are shown, with qualitatively-similar margin estimates from logit regressions presented in the Online Appendix. It does this for each of the main effects (identity and social rewards), for each of the four possible score categories, from the lowest-performing schools (“Intervene”) to the highest-performing schools (“Model”). All regressions include school fixed effects and dummy variables for survey completion and ineligibility for the survey (meaning not employed by the SDP) in the previous year (the omitted group being teachers eligible but not completing in the previous year).

except that the point estimate for the identity intervention at the best schools ('model' schools) is roughly 2-3 percentage points greater than for other schools. The difference is not statistically significant, however, due primarily to the relatively small number of teachers who come from "model" schools in the sample. However, this observation does add some additional support to the idea that school quality may positively impact the effectiveness of identity manipulations.

Discussion

In this article, we report the results of a randomized experiment testing the impact of social rewards and an identity salience manipulation on encouraging contributions to a public good: teachers completing an annual survey. We have three main results. First, we find that social incentives work just as well as personal incentives at motivating teachers to be prosocial. Second, we find no evidence that the identity salience manipulation increased prosocial behavior by teachers, nor did it meaningfully improve the efficacy of social rewards as an incentive (relative to personal rewards). Third, exploratory analyses indicated that survey completion happens at a higher rate at schools where teachers are rated more highly, and we found suggestive evidence that our identity manipulation in particular may have been somewhat more effective in these schools.

Our main conclusions, therefore, are as follows. First, although social incentives were not more effective than personal incentives (inconsistent with our hypothesis), the lack of difference between conditions has both theoretical and practical importance. If teachers were purely self-interested (as sometimes assumed by policymakers), then social incentives should have performed worse than personal incentives. Thus, the lack of difference between conditions implies that the teachers had at least somewhat other-regarding preferences. From a practical perspective, the lack of difference in effectiveness between social and personal incentives suggests that social incentives may be preferable from the point of view of the school district, as money spent by the district on social incentives (buying school supplies for students) directly benefits students and does so without undermining teacher motivation. This result speaks directly to a growing body of literature in public administration suggesting that interventions triggering prosocial motivations might be as or

more efficacious as pay-for-performance schemes with public sector employees (Ritz, et al., 2016, Grant, 2008).

Second, the lack of overall effect of our identity salience manipulation suggests that such manipulations, as implemented here, may not be a particularly promising approach for motivating aggregate teacher survey completion. Whether such incentives would work better for outcomes that are more obviously related to teaching (and thus to teachers' identities) remains to be seen. Furthermore, our exploratory results suggest that care should be taken regarding which sub-populations are targeted with such identity manipulations: the identity manipulation may have modestly reduced survey completion at schools where teachers receive lower ratings, while modestly increasing survey completion at schools with more highly rated teachers.

In addition to these main conclusions regarding our experimental treatments, exploratory analyses revealed interesting patterns regarding the school-level and individual-level predictors of teacher survey completion. First, teachers at schools with more satisfied parents and students were more likely to complete these surveys. Second, teachers at schools where students had fewer behavioral issues (e.g. lower suspension rates and higher attendance rates) were less likely to complete these surveys. Third, teachers at schools where teachers were rated highly were more likely to complete these surveys. One counterintuitive takeaway from these findings is that one method to increase rates of similar teacher feedback is to target schools at which students are well-behaved. Assessing the replicability of these relationships, and understanding their mechanisms, may be a fruitful direction for future research.

One might argue that our lack of significant results is driven by the fact that teachers simply did not read the emails they received carefully. While we cannot be certain, we have reason to believe that this was not the case. In particular, the SDP provided some anecdotal evidence that teachers were aware of the social/personal rewards, and were both providing feedback about and asking for more information about the timing of these rewards as the survey period came to a close. This does not necessarily mean that they read and internalized the identity salience manipulation, but it does suggest that teachers did not ignore the content of the email. Furthermore, to the extent that

the identity manipulation seemed to have been negatively impactful at schools with lower-rated teachers and positively impactful at schools with highly-rated teachers, this suggests that teachers did read the emails with sufficient care to notice the manipulation.

Another similar response to our results might be that the manipulations were simply too subtle to change behavior. While we are sympathetic to this view, we do not find it especially compelling. This intervention involved four separate emails that reinforced the treatment messaging, which is a reasonably strong manipulation when compared to other manipulations of messaging in the broader literature that uses behavioral interventions of this sort (Kessler & Milkman, 2016, Bursztny & Jensen, 2015, Bryan et al., 2011, Shang et al., 2008).

It is also plausible that teachers are already quite strongly saturated in their “teacher identity” when receiving the emails, meaning that an identity salience manipulation could not influence behavior very much. In other words, there may not have been much “room” for the manipulation to strengthen the influence of identity considerations on decision making. Here again, the fact that we do see variance in the efficacy of the identity manipulation as a function of school quality provides some evidence that this explanation is not fully satisfying. In particular, our finding that the identity salience manipulation did positively influence survey completion at schools with highly rated teachers (where teachers may identify more strongly with a teacher identity) somewhat weakens the case for already-high salience of identity across the board being the reason for the small aggregate effect.

Furthermore, one might argue that our null results on social versus personal rewards are influenced by the fact that teachers often use personal money to pay for school supplies for their students anyway, making our “social rewards” more similar to our “personal rewards” than they could have been. To the extent that this is an issue, future work might test a more unambiguous form of social reward, like a donation in the teacher’s name to a school-specific scholarship or charitable fund, instead of a gift card for school supplies.

There are important limitations of our findings. First, we were unable to include a control group that received no incentive at all, both because of SDP priorities and because of concerns about

spillover effects if we had a control group not receiving incentives. As a result, we cannot say how much the incentives increased survey completion, but can only conclude that social rewards were roughly as effective as personal ones. Though this particular context may not be ideal for an experiment on incentives with a pure control group, future work in this area would do well to have a pure control group to measure the causal impact of incentives in general, perhaps through variations in when teachers are informed about the lottery incentives (before vs. after they complete the survey, for example).

Second, we were not able to observe which teachers actually read the email soliciting survey completion (and thus who were actually influenced by the treatment). It may be that the treatments would appear substantially more effective if we were able to focus on those who were actually treated.

Third, in terms of how our results inform the broader literature on cooperation and public goods, our outcome (survey completion) may not be ideal. Although it is true that survey completion is a public good, many of the teachers may not have actually perceived the survey as creating benefits for others. This could be due to skepticism about the effectiveness of the school bureaucracy or a perception of red tape (Dehart-Davis & Pandey, 2005; Pandey & Scott, 2002). Thus, their prosocial motivations may not have been engaged in the task. If this was indeed the case, our manipulations may have been more effective for outcomes that were more obviously prosocial.

Finally, we cannot rule out the possibility that some treatments might have improved the “quality” of the sample (i.e., the representativeness of teachers responding) or of the survey responses (i.e., the detail and thoughtfulness of the feedback) without affecting the completion rate. However, while we have no way of testing the latter possibility with our data, we believe the former is not especially likely given that we find few meaningful differences in school-level characteristics across treatments.

In sum, we found little impact of social versus personal incentives and of identity salience on teacher survey completion. Our results suggest that school districts, and public sector organizations more broadly, should further investigate the use of social incentives, as such incentives are often

preferable so long as they do not undermine motivation. The limited aggregate impact of the identity manipulation, on the other hand, contributes to a growing literature that goes beyond just identifying promising nudges to testing when those nudges actually work, and for whom they work more or less effectively.

Acknowledgment

The authors wish to thank ideas42 for their generous financial support, the research team at the Philadelphia School District, notably Adrienne Reitano and Kelly Linker, for their help in designing and implementing this experiment, and Anjali Chainani, Yuan Huang, and Nandi O'Connor from the mayor's policy office in the City of Philadelphia for their logistical support. We are also very appreciative of helpful editor and reviewer comments.

Notes

1. Note that the initial randomization consisted of 8,423 teachers, but we removed 361 teachers from the analysis who were not (for various reasons) active SDP teachers as of May 31st, 2017. This was done based on advice from the SDP and does not affect the results, since these individuals were by and large not employed by the SDP when the surveys were distributed.
2. Note that the sample for this intervention did not include charter school teachers, though they also were invited to complete the annual survey.

3. <https://www.philasd.org/performance/programsservices/open-data/>
4. Note that the regression in column 2 includes teacher-level control variables for: 1) completing the survey in 2016; and 2) not being eligible to complete the survey in 2016. These two dummies allow us to control for the three-value categorical variable capturing teacher behavior in 2016 survey completion (ineligible to complete, completed, or did not complete when eligible). This differs from the treatment of teacher-level variables in column 1, which includes single variable regressions for: 1) the ineligible dummy variable; and 2) a binary variable for whether or not a teacher completed the survey conditional on being eligible (a regression that excludes ineligible individuals). We did the analysis in column 1 in this way to make the coefficient for 2016 survey completion easier to meaningfully interpret.
5. Because the inferences drawn from any individual regression method presented in Columns 1-3 of Table 4 may suffer from bias or constitute spurious correlation, we include all three to triangulate our understanding of which school-level variables are most strongly associated with teacher survey completion. We suggest that more confidence can be put in this relationship when a variable is a significant predictor across multiple regression specifications. To that end, Table 4 presents the school-level variables in descending order of the apparent strength of relationship with teacher survey completion.

References

- Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics*, 95, 1082–1095.
- Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry preschool, and early training projects. *Journal of the American Statistical Association*, 103, 1481–1495.
- Anik, L., Aknin, L. B., Norton, M. I., Dunn, E. W., & Quoidbach, J. (2013). Prosocial bonuses increase employee satisfaction and team performance. *PLOS ONE*, 8, 1–8.
- Ashraf, N., Bandiera, O., & Jack, B. K. (2014). No margin, no mission? A field experiment on incentives for public service delivery. *Journal of Public Economics*, 120, 1–17.
- Bandiera, O., Barankay, I., & Rasul, I. (2010). Social incentives in the workplace. *The Review of Economic Studies*, 77, 417–458.
- Benjamin, D. J., Choi, J. J., & Strickland, A. J. (2010). Social identity and preferences. *American Economic Review*, 100, 1913–1928.
- Benjamin, D. J., Choi, J. J., & Fisher, G. (2016). Religious identity and economic behavior. *Review of Economics and Statistics*, 98, 617–637.
- Bryan, C. J., Walton, G. M., Rogers, T., & Dweck, C. S. (2011). Motivating voter turnout by invoking the self. *Proceedings of the National Academy of Sciences of the United States of America*, 108, 12653–12656.
- Bursztyn, L., & Jensen, R. (2015). How does peer pressure affect educational investments? *Quarterly Journal of Economics*, 130, 1329–1367.

- Carr, P. B., & Steele, C. M. (2010). Stereotype threat affects financial decision making. *Psychological Science*, 21, 1411–1416.
- DeHart-Davis, L. & Pandey, S. K. (2005) Red tape and public employees: Does perceived rule dysfunction alienate managers? *Journal of Public Administration Research and Theory*, 15.1, 133-148.
- Cohn, A., Marechal, M. A., & Noll, T. (2015) Bad boys: How criminal identity salience affects rule violation. *The Review of Economic Studies*, 84, 1289–1308.
- Freedman, J. L., & Fraser, S. C. (1966). Compliance without pressure: The foot-in-the-door technique. *Journal of Personality and Social Psychology*, 4.2, 195–202.
- Fryer, R. G., Levitt, S. D., List, J., & Sadoff, S. (2012) Enhancing the efficacy of teacher incentives through loss aversion: A field experiment. *National Bureau of Economics Research Working Paper Series*.
- Gehlbach, H., Brinkworth, M. E., King, A. M., Hsu, L. M., McIntyre, J., & Rogers, T. (2016). Creating birds of similar feathers: Leveraging similarity to improve teacher-student relationships and academic achievement. *Journal of Educational Psychology*, 108, 342–352.
- Gneezy, A., Imas, A., Brown, A., Nelson, L. D., & Norton, M. I. (2012). Paying to be nice: Consistency and costly prosocial behavior. *Management Science*, 58, 179–187.
- Gneezy, U., & Rustichini, A. (2000). A fine is a price. *The Journal of Legal Studies*, 29, 1–17.
- Grant, A. M. (2008). Employees without a cause: The motivational effects of prosocial impact in public service. *International Public Management Journal*, 11.1, 48–66.
- Hallsworth, M., List, J. A., Metcalfe, R. D., & Vlaev, I. (2017) The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics*, 148, 14–31.
- Kahneman, D., & Tversky, A. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458.
- Kessler, J. B., & Milkman, K. L. (2016). Identity in charitable giving. *Management Science*, Forthcoming.
- Kraft, M. A., & Rogers, T. (2015). The underutilized potential of teacher-to-parent communication: Evidence from a field experiment. *Economics of Education Review*, 47, 49–63.
- Kraft-Todd, G., Yoeli, E., Bhanot, S., and Rand, D. (2015). Promoting cooperation in the field. *Current Opinion in Behavioral Sciences*, 3, 96–101, 2015.
- Levitt, S. D., List, J. A., Neckermann, S., & Sadoff, S. (2016). The behavioralist goes to school: Leveraging behavioral economics to improve educational performance. *American Economic Journal: Economic Policy*, 8, 183–219.
- Mullen, E., & Monin, B. (2016). Consistency versus licensing effects of past moral behavior. *Annual Review of Psychology*, 67, 363–385.
- Pandey, S. K., & Scott, P. G. (2002). Red tape: A review and assessment of concepts and measures. *Journal of Public Administration Research and Theory*, 12.4, 553–580.
- Ritz, A., Brewer, G. A., & Neumann, O. (2016). Public service motivation: A systematic literature review and outlook. *Public Administration Review*, 76.3, 414–426.
- Shang, J., Reed, A., & Croson, R. (2008) Identity congruence effects on donations. *Journal of Marketing Research*, 45, 351–361.
- Spencer, B., & Castano, E. (2007). Social class is dead. long live social class! Stereotype threat among low socioeconomic status individuals. *Social Justice Research*, 20, 418–432.
- Steele, C. M., & Aronson, J. (1990). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology*, 69, 797–811.

Appendix

Figure A1: Treatment email screenshots (for first emails, sent on April 4, 2017; screenshots of three follow-up reminder emails visible in online appendix, Figures 2-5)

(a) “Standard” condition

Subject: Help Us Improve Philly Schools: Take the District-Wide Survey!

Dear Teacher,

It's time for our annual District-wide Teacher Survey. This survey is being completed by teachers across Philadelphia and will provide stakeholders with valuable information about the state of education in our city. Your input will be used to identify successes and areas where resources and support are needed.

The survey will take approximately 30 minutes to complete. All responses will be kept confidential and data will only be reported in aggregate.

We know your time is valuable, so we would like to thank you for your participation:

Complete the survey and you could win \$100!
We will randomly select 15 teachers who complete the survey and award them a
\$100 gift card to Barnes and Noble, just for you!

To complete the survey, please login to the Employee Portal (<https://sso.philasd.org>) and click the checkmark icon labeled “District-wide Teacher Survey.” All responses will be confidential, but please do not forward this email to anyone else.

As always, we greatly appreciate your participation.

(b) “Social Rewards Only” condition

Subject: Help Us Improve Philly Schools: Take the District-Wide Survey!

Dear Teacher,

It's time for our annual District-wide Teacher Survey. This survey is being completed by teachers across Philadelphia and will provide stakeholders with valuable information about the state of education in our city. Your input will be used to identify successes and areas where resources and support are needed.

The survey will take approximately 30 minutes to complete. All responses will be kept confidential and data will only be reported in aggregate.

We know your time is valuable, so we would like to thank you for your participation:

Complete the survey and you could win \$100 for your students!
We will randomly select 15 teachers who complete the survey and award them a
\$100 gift card to Office Depot, for classroom supplies!

To complete the survey, please login to the Employee Portal (<https://sso.philasd.org>) and click the checkmark icon labeled “District-wide Teacher Survey.” All responses will be confidential, but please do not forward this email to anyone else.

As always, we greatly appreciate your participation.

(c) “Identity Only” condition

Subject: Help Us Improve Philly Schools: Take the District-Wide Survey!

Dear Teacher,

It's time for our annual District-wide Teacher Survey. This survey is being completed by teachers across Philadelphia and will provide stakeholders with valuable information about the state of education in our city. Your input will be used to identify successes and areas where resources and support are needed.

The survey will take approximately 30 minutes to complete. All responses will be kept confidential and data will only be reported in aggregate.

As a proud Philadelphia teacher, you have committed your career to the advancement of education and to improving the lives of young people in the city.

It's important for educators like you to complete the survey - it helps us improve our schools and is another way to show your ongoing commitment to your students!

We know your time is valuable, so we would like to thank you for your participation:

Complete the survey and you could win \$100!

We will randomly select 15 teachers who complete the survey and award them a

\$100 gift card to Barnes and Noble, just for you!

To complete the survey, please login to the Employee Portal (<https://sso.philasd.org>) and click the checkmark icon labeled “District-wide Teacher Survey.” All responses will be confidential, but please do not forward this email to anyone else.

As always, we greatly appreciate your participation.

(d) “Social Reward + Identity” condition

Subject: Help Us Improve Philly Schools: Take the District-Wide Survey!

Dear Teacher,

It's time for our annual District-wide Teacher Survey. This survey is being completed by teachers across Philadelphia and will provide stakeholders with valuable information about the state of education in our city. Your input will be used to identify successes and areas where resources and support are needed.

The survey will take approximately 30 minutes to complete. All responses will be kept confidential and data will only be reported in aggregate.

As a proud Philadelphia teacher, you have committed your career to the advancement of education and to improving the lives of young people in the city.

It's important for educators like you to complete the survey - it helps us improve our schools and is another way to show your ongoing commitment to your students!

We know your time is valuable, so we would like to thank you for your participation:

Complete the survey and you could win \$100 for your students!

We will randomly select 15 teachers who complete the survey and award them a

\$100 gift card to Office Depot, for classroom supplies!

To complete the survey, please login to the Employee Portal (<https://sso.philasd.org>) and click the checkmark icon labeled “District-wide Teacher Survey.” All responses will be confidential, but please do not forward this email to anyone else.

As always, we greatly appreciate your participation.