All SROs are not created equal: Heterogeneity in school policing

Hanna Han and Katherine Kwok *

November 7, 2025

Click here for the most recent version

Abstract

The expansion of police presence in U.S. schools over the past two decades raises questions about the trade-off between student safety and exposure to the criminal justice system. In this paper, we study how school context moderates the effect that School Resource Officers (SROs) have on in-school offenses and school discipline. Using a difference-in-differences design that leverages a federal SRO hiring grant, we estimate bounds on the variance of the treatment effect distribution to detect treatment effect heterogeneity in SRO hiring and discover meaningful differences in how hiring SROs impacts the number of students suspended or referred to law enforcement. To explain the detected treatment effect heterogeneity, we consider observable school characteristics such as student demographics and school staff and unobservable differences in how schools implement their SRO program, measured by training text models to extract themes from grant narratives that describe the problems facing the school. We find that schools whose SRO programs focus on building community have decreased suspensions and referrals or arrests than schools that focus on the SRO's fulfillment of traditional law enforcement duties. Our results highlight how a school's context shapes the impact that SROs have in schools.

*Department of Economics, University of Wisconsin-Madison (email: hhan82@wisc.edu). We thank Jeff Smith, Jesse Gregory, and Corina Mommaerts for their guidance and feedback. We are grateful to the School Resource Officers in school districts across Wisconsin whose insights about their work provided much needed nuance. We also thank the seminar participants at the University of Wisconsin-Madison Public Student Seminar, Empirical Micro Seminar, and Public Seminar; Institute for Research on Poverty Summer Workshop; and Association for Public Policy and Analysis Fall Conference for their helpful comments.

1 Introduction

Police officers have become a familiar sight on school campuses within the U.S. in the past two decades. As of the 2019-20 school year, there were 23,426 School Resource Officers (SROs), who are sworn law enforcement officers that work within school grounds to promote safety (Davis, 2023). While having an SRO on campus can increase the sense of safety among students, teachers and school staff (Brown, 2005; Cuellar et al., 2017; Wood and Hampton, 2021), SROs can also increase students' exposure to the criminal justice system, both directly via higher arrest rates (Owens, 2017) and indirectly by increased suspension and expulsion rates (Sorensen et al., 2023; Weisburst, 2019) which are associated with higher odds of arrest and incarceration (Monahan et al., 2014; Mowen and Brent, 2016; Bacher-Hicks et al., 2024). This is a cause for concern as juvenile criminal justice involvement results in a myriad of negative long-term outcomes such as reduced likelihood of high school completion (Gottlieb and Wilson, 2019; Hjalmarsson, 2008) and college enrollment (Kirk and Sampson, 2013), and increased probability of incarceration as an adult (Aizer and Doyle Jr., 2015).

Something that would help school districts and policymakers better navigate the known trade-offs of SRO programs than mean effects alone would be understanding how program and school context contribute to the observed SRO effect. For instance, Fisher et al. (2022) finds that SROs stationed in schools with a majority of black students perceive threats to school safety as stemming from within-school misconduct (e.g., student fights, bullying) while SROs in schools with a majority of white students perceive the threats as being external (e.g., intruder). Differences in threat perception likely translates to differences in the specific tasks of the SRO (e.g., patrolling hallways to keep an eye on students vs. patrolling outside to ensure school grounds are secure), which has implications for how much of an impact the SRO may have on students. Qualitative research indicates wide-ranging variation in program context, such as the SRO's job responsibilities (Davis, 2023; McKenna et al., 2016; Ryan et al., 2018) and level of involvement in student discipline (Curran et al., 2019), and in school context (Ghavami et al., 2021; Lynch et al., 2016). However, the localized nature of SRO programs due to limited federal guidance on their implementation (Counts et al., 2018) and the resulting lack of centralized administrative data

means we are unable to observe the variation across programs and its impact directly.

In this paper, we use observable school characteristics and novel grant narrative data to proxy for the variation in program and school context in order to explore the treatment effect heterogeneity in SRO programs. We use the Civil Rights Data Collection (CRDC) from the 2011-12 to 2017-18 school years to show how SRO hiring impacts school-level disciplinary and academic outcomes. We use a staggered difference-in-differences design that leverages the 2013-2016 federal COPS Hiring Program (CHP) grant that funds SRO hiring at schools across the U.S. By comparing schools that were awarded the grant against schools that applied but were not awarded, our empirical strategy addresses a key selection issue of schools with higher levels of student misbehavior being more likely to establish an SRO program. We follow the methodology of Callaway and Sant'Anna (2021) as the timing of the CHP grant award varies. When we estimate mean effects to serve as a baseline for the heterogeneity analysis, our findings on violent offenses and larger effects for male students and students with disabilities are consistent with past quantitative works (Weisburst, 2019; Sorensen et al., 2021, 2023).

The first part of our analysis involves showing evidence of meaningful treatment effect heterogeneity in SRO hiring. To do this, we follow Heckman et al. (1997) and use classical statistical bounds established by Fréchet (1951) and Höffding (1970), or "FH bounds," to identify the bounds on the variance of the distribution of treatment effects and to test the null hypothesis of a common treatment effect across schools. The strength of using FH bounds is that it captures meaningful treatment effect heterogeneity that does not depend on our choice or the availability of observable school characteristics to use as moderators, which would be the case if we were exploring heterogeneity by estimating an interaction effect. We extend the literature on FH bounds by applying it to a difference-in-differences design by estimating the counterfactual outcome distribution of the treated group using the Changes-in-Changes framework from Athey and Imbens (2006). Changes-in-Changes is analogous to difference-in-differences, but the method identifies the distribution of treatment effect, not just the average treatment effect, by capturing the change over time in the control group to estimate the counterfactual distribution of the treated group. We find evidence of meaningful treatment effect heterogeneity at a p-value < 0.001 for suspensions and grade retention in schools treated in 2013 and 2014, and for

expulsions and law enforcement actions in schools treated in 2013. We also capture heterogeneity in the direction of the treatment effect and find that a vast majority of treated schools have an increase in suspensions (with no schools having a zero effect) while nearly all treated schools have a zero effect in expulsions. There is less of a definitive trend in the direction of treatment effect for law enforcement actions and grade retention.

Next, we examine how SRO program and school context moderate the effect of SRO hiring on student outcomes. To capture how SRO programs may be implemented differently across schools, we train text models on CHP grant narratives to identify common themes. In the narratives, grant applicants describe the problems facing their school and outline potential responsibilities of the SRO. We also consider school context, focusing on student demographics and the presence of specific school staff such as school counselors after a review of the qualitative literature on SROs (Curran et al., 2019; Fisher et al., 2022). Estimating the treatment effect separately for schools that were associated with a given moderator and for schools that were not, we find statistically significant differences between ATTs of the moderator schools and of the comparison schools for multiple moderators, which suggests program and school context can explain aspects of observed treatment effect heterogeneity.

First, we find schools whose grant narrative contained themes that focused on building community and trust have a decrease count of suspensions, law enforcement actions, and chronic absenteeism. Schools whose grant narrative discussed themes that focused on traditional law enforcement also had decreases in suspensions and law enforcement actions, but schools focused on community had bigger reductions in these outcomes. The results suggest schools that focus on community are less likely to use more formal disciplinary measures or that SROs are able to use softer skills such as mentoring to reduce student misbehavior. An additional implication is schools that focus on community may create a better school environment for students, which has the effect of increasing attendance.

Second, we find schools whose grant narrative view threats to school safety as being internal experience a decrease in expulsions while schools that viewed threats as being external experience an increase. One possible explanation for these results is that schools whose SRO is

focused on reducing within-school student misbehavior may focus more on duties such as patrolling hallways that makes him more visible to students, possibly deterring misbehavior. The large decrease in chronic absenteeism, a value that corresponds to a decrease of 92.4% of the pre-treatment mean, suggests such schools may be increasing students' sense of safety, resulting in the decrease in chronic absenteeism. On the other hand, schools that view threats as being external may have students who are involved in criminal activity outside of school, increasing the likelihood that there are students who are committing misconduct that is severe enough to warrant expulsion.

Third, schools who have a majority of students that are black or Hispanic had an increase in students who retained a grade—corresponding to an 88.0% increase from the pre-treatment mean —than schools with more mixed student demographics. Schools with a grant narrative that focused on law enforcement also had a larger increase in grade retention than schools that did not contain such themes. These results may stem from differences in students' sense of safety at school after SROs are stationed, which may impact their learning. Theriot and Orme (2015) find that black students felt less safe at school when there was an SRO. In addition Curran et al. (2021) find that interactions with SROs increased students' sense of danger; it may be the case that in schools that focus on traditional law enforcement responsibilities, the SRO is more visible to students and also interacts with students more frequently or in qualitatively different ways.

Our paper contributes to a growing body of quantitative work on the impact of SROs on student outcomes (see Sorensen et al. (2023) for a review). For disciplinary outcomes, the quantitative literature finds that SRO presence increases arrests, referrals to law enforcement, and exclusionary discipline such as suspensions and expulsions, especially for students of color (Gottfredson et al., 2020; Owens, 2017; Sorensen et al., 2021, 2023; Weisburst, 2019). These studies evaluate treatment effect heterogeneity in SRO programs through several different channels: students characteristics (gender, race and ethnicity, household income, or disability status); school level (e.g. middle or high school); number of SROs hired; SRO tenure; and the type of duties completed by the SRO (Gottfredson et al., 2020; Swartz et al., 2016; Zhang, 2019). Studies that evaluate heterogeneity from program context, however, are qualitative and based on small

samples. Our paper contributes to the literature by considering how differently these programs can be implemented across schools and quantifying treatment effect heterogeneity in SRO hiring on student outcomes, moderated by aspects of program implementation.

We also contribute additionally to a subset of this literature, which uses the COPS grant applications as the basis of its empirical strategy in estimating SRO impact (Gottfredson et al., 2020; Owens, 2017; Sorensen et al., 2023; Weisburst, 2019). Past studies have used grant data to determine the treated group and define the study sample, but have not leveraged grant narrative data in their estimation. These texts contain valuable information about why a law enforcement agency requested SRO funding, the possible duties of the hired SROs, and descriptions of how SROs may be allocated across schools. We use the narrative data to characterize previously unobservable heterogeneity in approach to school discipline across school districts and examine its effects on student disciplinary outcomes.

Lastly, we provide a methodological contribution to the literature that uses classical statistical bounds to explore treatment effect heterogeneity (Heckman et al., 1997; Djebbari and Smith, 2008; Buhl-Wiggers et al., 2024). In the case of these studies, treatment assignment is randomized, which makes the estimation of FH bounds more straightforward. We develop a method that applies FH bounds estimation to a difference-in-differences context using the Changes-in-Changes framework from Athey and Imbens (2006). This broadens the application of FH bounds and the documentation of meaningful treatment effect heterogeneity to wider contexts than randomized control trials.

The paper proceeds in the following order: Section 2 contains background information about SROs and the federal grant leveraged in this study and Section 3 presents the data. The analysis begins in Section 4 by estimating the mean effect of SRO hiring on student outcomes. Section 5 moves onto detecting treatment effect heterogeneity using classical statistical bounds. Section 6 explains how relevant SRO program and school contexts are measured and Section 7 takes these measures and shows how they moderate the impact of effect of SRO hiring on student outcomes. Section 8 concludes.

2 Institutional background

2.1 SROs and their roles

School resource officers (SROs) are sworn law enforcement officers appointed to work within school grounds to promote safety. These officers can be assigned to schools of all levels although more middle and high schools have SROs than elementary schools (Wang et al., 2020). Although exact duties may vary, the work of SROs generally falls into the categories of counseling, teaching, and law enforcement (Gottfredson et al., 2020). As of the 2019-20 school year, around 23,426 SROs were working in schools across the U.S. (Davis, 2023).

In general, local law enforcement agencies, including local police departments, school district police departments, and sheriff's offices, work with school district administrators and staff to create and design SRO programs. The involved parties usually sign a memorandum of understanding (MOU) defining the roles and responsibilities of SROs in a particular school district or school (Ghavami et al., 2021). Although SROs often work closely with school staff on a day-to-day basis, they still report to and take directions from supervisors in their respective law enforcement agencies (McKenna et al., 2016).

SROs can play many different roles on school grounds. First, they act as law enforcement on campus: they may patrol hallways, respond to incidents or calls for services, issue warnings, and make arrests. Second, SROs can also be hired to teach school safety programs (e.g. DARE) and mediate interpersonal conflicts between students. Finally, SROs may also act as informal counselors and mentors to students, coach sports teams or lead other extracurricular activities, and perform chaperone duties on field trips (Davis, 2023).

SRO programs have become increasingly commonplace over the past two decades. According to the School Survey on Crime and Safety, which is a nationally representative survey of public schools, the percentage of schools where security staff (including SROs) are present every week increased from 42 to 61% between the 2005-06 and 2017-18 school years (Wang et al., 2020). The proliferation of SRO programs is often attributed to the rise of community-oriented

policing since the 1980s and the Columbine shooting in 1999 (McKenna et al., 2016).

Funding for SRO programs can come from several different sources. At the federal level, the U.S. Department of Justice funds SRO hiring/rehiring through grants under the Community Oriented Policing Services (COPS) program. We detail these federal grants in the next section. Some states also provide funding for SRO hiring, such as North Carolina (Sorensen et al., 2021). According to a survey of SROs during the 2019-20 school year, only a small fraction (4.3%) of SRO programs are funded by federal funds (Davis, 2023). Among the remaining programs, 78.6% are funded by the school district and 10.2% are funded by state/local grants and taxes. To our knowledge, this was the first national survey of SRO programs and there is no nationally representative data on SRO funding sources before this. Our study period is 2011 to 2018, so the distribution of funding sources is likely similar to the survey statistics from 2019.

2.2 How the CHP grant works

The COPS Hiring Program (CHP) is a grant established as a part of the U.S. Department of Justice's Community Oriented Policing Services (COPS) program. It provides funds to law enforcement agencies to create new positions or retain existing positions that are at risk of being eliminated due to budget reductions, to aid in community policing initiatives. The original grant under the COPS program was the Cops In Schools (CIS) grant, designed specifically to disburse funding to law enforcement agencies to hire SROs and was active from 1999 to 2005. The CHP grant program launched in 2010, as the successor of the CIS grant program.[1] The CHP program has a broader reach than its predecessor as it accepts applications for community policing initiatives outside of school-based policing. While in a given year only a small proportion of agencies fund their SRO programs using federal funds, the COPS program has played a large role in the current prevalence of SRO programs across the U.S. as it provides short-term funding that can serve as "proof of concept" for school districts that may be interested in seeing whether an SRO program is a worthwhile inclusion as a budget item. Between 1999 and 2013, COPS provided funding for 7,200 new SRO positions nationwide (James and McCallion, 2013).

[1]The CHP grant began awarding funding in 2010, but only began awarding grants for school-based policing in 2013. There was also a pause in awards in 2018 and 2019 as the grant program was not funded in those fiscal years.

All local, state, and tribal law enforcement agencies are eligible to apply to the CHP grant. It covers up to 75% of direct personnel costs such as salary and fringe benefits, up to $125,000, over three years within a five-year period. Agencies are required to provide a minimum of a 25% match for the personnel costs unless a waiver is approved, and to show that they will be able to retain the position for 12 months following the end of the grant period. Agencies that apply for school-based policing must submit a signed Memorandum of Understanding (MOU) with their partner organization, which is often the school district the agency will be working with, within 90 days of the grant award. The CHP grant application period for a given fiscal year is usually in the spring and the grant is awarded by October 1. For example, the grant application period for the 2016 CHP grant ran from April 25 to June 23, 2016 and the grant disbursement period began on October 1, 2016. This means that the start of the grant disbursement period overlaps with the start of a school year; thus, for example, a school district that applied for and was awarded the 2016 CHP grant would, depending on how quickly they were able to fill the position, have an SRO starting from the 2016-17 school year.

The CHP grant is awarded based on three factors: the strength of the agency's community policing initiatives, which the agencies detail in their application material; the fiscal need of the agency, which is measured using changes in the law enforcement budget and the poverty and unemployment rate within the agency's jurisdiction; and the crime rate within the agency's jurisdiction. The COPS grant office assigns each applicant a score for each of the three factors mentioned above with the final score being the sum of the three scores.[2] Both the state the agency is located in and the size of the population the agency serves also play a role due to the statutory requirements of the CHP grant.[3] The variation in grant award for a given grant year is plausibly exogenous as grant allocation depends on the mix of states and population sizes represented within the pool of applicants and on the total amount of funding allocated to the CHP grant for that grant year. While the trends in student disciplinary incidents within a school may be correlated with the trends in economic conditions and crime of its surrounding

[2]The final score can also include additional points based on a miscellaneous category, which changes across grant years.

[3]The requirement states that each state or territory with eligible applicants must receive 0.5% of total grant funds, and that half of the funds must be disbursed to agencies serving populations larger than 150,000 and the other half to agencies serving less than 150,000. See Sorensen et al. (2023) for specifics about the grant award process.

community, increasing the likelihood that an agency is awarded the grant, an agency cannot predict the number or characteristics of other agencies that also choose to submit an application that grant year and the amount of funding the Department of Justice chooses to allocate to the CHP grant.

3 Data

The main data we use to construct our sample and treatment is the 2013-2016 COPS Hiring Program (CHP) grant application data, obtained from the U.S. Department of Justice through Freedom of Information Act (FOIA) requests. The data include basic information about the agency that submitted the application such as the agency name, city, state, and grant-specific information such as the number of SRO positions requested, year of application, name of partner organization (typically the local school district), and grant award status. We also have grant narrative data, which are long-form text responses that describe the problem the agency will address using the grant funds and the agency's approach to the problem.

For school-level data, we use the Civil Rights Data Collection (CRDC). The CRDC is a biennial survey that has been collected since 1968 to measure the extent to which students across the U.S. have equal access to educational opportunities. The CRDC contains information about student enrollment, presence of certain school staff, and student disciplinary incidents. We use surveys from the 2011-12, 2013-14, 2015-16, and 2017-18 school years as they include the universe of public schools in the U.S.[4] We supplement the CRDC with the Common Core of Data (CCD), which contains additional school-level characteristics such as the school level, school locale, and school region.

We use the following disciplinary outcomes reported by the CRDC: *formal law enforcement actions*, which are the number of students who receive tickets, citations, or referrals to law enforcement and school-based arrests; *suspensions*, in-school and out-of-school; and *expulsions*. We also use disciplinary offenses such as *gun offenses*, which are the number of incidents

[4] In prior survey waves, the CRDC contained only a sample of public schools in the U.S.

of physical attacks, threats of attacks, and robberies involving firearms and explosives as well as firearm or explosives possession; and *other offenses*, which are the number of incidents of sexual assault, rape, or attempted rape as well as physical attacks, threats of attack, and robberies with no weapons or weapons other than firearms or explosives. It is important to note that the disciplinary outcomes enumerate the number of students who received a specific disciplinary action—a student with multiple suspensions within a school year would be recorded as having received one suspension for that school year —while the disciplinary offenses are counts of incidents. For academic outcomes, we use *chronic absenteeism*, which is the number of students who have missed at least 15 days of school in a given school year; and *grade retention*, which is the number of students who repeat the same grade in the following school year.

We transform the outcomes so that they are counts (or incidents) per 100 students in order to factor in student enrollment, and winsorize at the 99th percentile due to extreme values in the long right tail of the distribution (see Appendix B.2 for details on the construction of the outcomes). The CRDC also contains the outcomes (excluding offenses) disaggregated by gender, race and ethnicity, disability status, and Limited English Proficiency (LEP) status, and we use these to present heterogeneous treatment effects by student subgroups.

Our sample contains schools linked to law enforcement agencies that applied for the CHP grant for school-based policing between 2013 to 2016. Schools are matched to law enforcement agencies using the grant narrative data (Appendix B.3 describes this process). The sample is restricted to the 92.2% of schools whose applications stated they did not have an SRO stationed in the school at the time of grant application so that the results can be interpreted as an extensive margin effect of establishing a new SRO program. The final sample contains 10,586 schools in 908 different school districts.

Table 1 contains descriptives from the 2011-12 CRDC[5] for schools that applied for but did not receive the CHP grant and schools that received the CHP grant. We find that schools awarded the grant have more students of color and students who receive free or reduced price lunch in ad-

[5]Chronic absenteeism and whether the school has an SRO were measured from the 2013-14 CRDC, and gun offenses and other offenses were measured from the 2015-16 CRDC as these questions were not asked in earlier waves. All other variables are measured from the 2011-12 CRDC.

dition to having more enrolled students on average than schools that did not receive the grant. They also have a higher number of students who experienced suspensions, law enforcement actions, chronic absenteeism, and grade retention and higher incidents of other offenses, but fewer students who experienced expulsions. These differences at the baseline may reflect disparities in community policing initiatives, unemployment rates, poverty rates, and crime rates, all of which factor into the grant award decision as discussed in Section 2.2.

4 Mean effect of SROs on student outcomes

Prior to exploring how the variation in SRO program implementation explains heterogeneity in the treatment effect, we estimate the mean effect of receiving a grant for SRO hiring on student disciplinary and academic outcomes. We do this because having baseline results gives us a point of reference that can put estimates of heterogeneous treatment effect into context, and also allows us to compare our results against existing literature.

4.1 Empirical strategy

To establish a causal relationship between SROs and student disciplinary and academic outcomes, we use a staggered difference-in-differences design that leverages the federal CHP grant which funds SRO hiring at schools. This is done to avoid endogeneity concerns that would arise from a naive comparison between schools that have SROs and those that do not, as having higher rates of detected student misbehavior could make schools more likely to establish SRO programs, leading to an overestimation of their effect on student discipline. Instead, we use the CHP grant to delineate our sample as schools that applied to the grant, which means all comparisons are between schools that were motivated to establish an SRO program, and to determine treatment status, which is assigned by the CHP grant office (Section 2.2 describes the process) and not individual schools. It is important to note that our sample captures a specific subset of schools that are interested in hiring SROs *and* have limited means to finance the hire through local funding sources, which may result in estimates that differ from one derived from a more

nationally representative sample. Furthermore, because treatment is defined as being awarded the CHP grant, our estimates capture less the effect of hiring SROs and more the effect of receiving the funding to hire SROs. Past work suggests that receiving the CHP grant does lead to increased SRO hiring (Owens, 2017; Sorensen et al., 2023), but we are unable to directly observe SRO placement due to data limitations.[6]

Recent literature on difference-in-differences has shown that the traditional two-way fixed effects (TWFE) estimator may be inconsistent when there is variation in treatment timing and treatment effect heterogeneity, as early treated units are included in the comparison group for late treated units (de Chaisemartin and D'Haultfœuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021). Thus, to sidestep potential pitfalls and to account for dynamic treatment effects, we use the framework and estimator from Callaway and Sant'Anna (2021). Under this approach, a key identifying assumption is the conditional parallel trends assumption, which states that student outcomes in treated schools and control schools would have evolved in parallel over time had treated schools not been awarded the CHP grant, conditional on covariates not affected by treatment.

Let D_{st} be the treatment indicator for school s located in school district $d(s)$ linked to a law enforcement agency that is awarded the CHP grant in year t. Omitting unit level subscripts for clarity, treated schools are assigned to a treatment group based on the year they were first treated with the indicator G_g, where $g \in \{2013, 2014, 2015, 2016\}$. In cases where the law enforcement agency applied for the CHP grant multiple times between 2013 to 2016, the school belongs to the treatment group in the first year it is awarded the grant. Once treated, treatment is sustained through the end of the study period as the grant funds SRO positions for three years and applicants are required to retain the position for one year following the end of the grant. Let the indicator $C = 1$ for schools that applied (even across multiple years) but did not receive the grant. These schools are in the "never treated" group, or $G = \infty$, which serves as the comparison for all treatment groups. Table 2 shows the distribution of the control and treated groups across years. Let Y_{gt} be the outcome of interest for schools in treatment group g in year t.

[6]While the CRDC includes a question about SRO placement in schools starting with the 2013-14 CRDC, issues with the 2015-16 CRDC questionnaire (**?**) result in missing values for 80.6% of treated schools and 71.6% of control schools. This limits our ability to capture a pre/post change in SRO presence within treated schools.

We begin by estimating the average treatment effect on the treated (ATT) of being awarded the CHP grant on student outcomes for each treatment group g and post-treatment year t pairing, conditional on pre-treatment covariates X. These estimates are referred to as *group-time ATTs*:

$$ATT(g,t) = \mathbb{E}\left[\left(\frac{G_g}{\mathbb{E}[G_g]} - \frac{\frac{p_g(X)C}{1-p_g(X)}}{\mathbb{E}\left[\frac{p_g(X)C}{1-p_g(X)}\right]}\right)(Y_t - Y_{g-1} - \mathbb{E}[Y_t - Y_{g-1}|X, C=1])\right] \tag{1}$$

where $p_g(X) = P(G_g = 1|X, G_g + C = 1)$, or the propensity score of being in treatment group g conditional on pre-treatment covariates.

Individual group-time ATT estimates are then aggregated using weights into dynamic ATT estimates, or the average treatment effect by the number of years of exposure to treatment e:

$$\theta(e) = \sum_{g \in \mathcal{G}} 1\{g + e \leq \mathcal{T}\} P(G = g|G + e \leq \mathcal{T}) ATT(g, g + e)$$

Due to the structure of the data in which the treatment groups progress yearly but the school-level outcomes progress biennially, we pool together the one- and two-year effect of receiving grant funding to hire SROs such that the dynamic ATT estimates progress biennially, or $e \in \{-4, -2, 0, 2, 4\}$. For example, the dynamic ATT estimate for two years after treatment $\theta(2)$ aggregates individual ATT estimates that capture one-year effects from the 2014 and 2016 treatment groups using the 2015-16 and 2017-18 CRDC, respectively, or $ATT(2014, 2015)$ and $ATT(2016, 2017)$; and two-year effects from the 2013 and 2015 treatment groups using the 2015-16 and 2017-18 CRDC, respectively, or $ATT(2013, 2015)$ and $ATT(2015, 2017)$. Pooling allows us to compare all estimates against the same pre-treatment period.[7]

Figure 1 shows the conditional dynamic ATT estimates for all disciplinary outcomes and grade retention.[8] Appendix Figure A1 shows the unconditional dynamic ATT estimates and Ap-

[7] It is the standard in event study plots to normalize all estimates using the last period before treatment. Our data structure prevents us from using the same comparison period to normalize estimates across all treatment groups. Specifically, while we can use one year before treatment as the comparison for the 2014 and 2016 treatment groups (using the 2013-14 and 2015-16 CRDC, respectively), the closest pre-treatment comparison for the 2013 and 2015 treatment groups is two years before treatment (using the 2011-12 and 2013-14 CRDC, respectively) as the CRDC is a biennial survey.

[8] Chronic absenteeism, gun offenses, and other offenses are excluded because these outcomes are only present in two CRDC survey waves, resulting in a pre/post comparison.

pendix Figure A2 shows conditional dynamic ATT estimates that are not pooled. Figure 1 suggests that the conditional parallel trends assumption is satisfied for disciplinary outcomes, but not for grade retention. This means we are unable to interpret the estimates for grade retention as causal as there are statistically significant differences between the treatment and control groups for this outcome; however, we are able to interpret the estimates for the other outcomes as causal.

Dynamic ATT estimates are aggregated to produce the overall ATT estimate, which captures the average treatment effect across varying lengths of exposure to treatment:

$$\theta_{overall} = \frac{1}{\mathcal{T}-1} \sum_{e=0}^{\mathcal{T}-2} \theta(e)$$

Our estimation includes the following controls using pre-treatment characteristics of schools: proportion of enrolled students who are male, black, Hispanic, disabled, have Limited English Proficiency (LEP), receiving free or reduced price lunch; school level (elementary, middle, or high); school locale (urban, suburban, town, or rural); school region (Northeast, South, Midwest, West); and the population size that the law enforcement agency serves (more or less than 150,000). We include these covariates to strengthen the plausibility of the parallel trends assumption as schools that received the CHP grant differ from schools that did not in terms of student composition and school characteristics, which may impact how student outcomes evolve over time. Standard errors are estimated using the Wild bootstrap (Wu, 1986; Cameron et al., 2008) and clustered at the school district level as treatment is determined at the agency level, and agencies are matched to school districts.

4.2 Baseline results

Tables 4 and 5 contain the post-treatment ATT estimates of impacts on student outcomes in response to receiving CHP grant funding for SRO hiring. Looking at the first row of the tables, we find that receiving the CHP grant results in an average decrease of 0.186 per 100 students who experience law enforcement action, or a 32.2% decrease from the pre-treatment mean (Table

3 contains pre-treatment means). The number of incidents of other offenses also decreased by 1.506 per 100 students, or a 76.2% reduction from the pre-treatment mean. These results suggest that students in schools that received funding to hire SROs may be deterred away from committing serious offenses, which would also decrease the need for formal law enforcement action. Furthermore, having an SRO on campus who regularly interacts with students and becomes familiar with their background may result in the SRO foregoing formal law enforcement action in favor of less severe or more informal responses to misbehavior. We do not find meaningful effects of the CHP grant award on the count of students that experienced exclusionary discipline, chronic absenteeism, grade retention, or gun offenses.

The remainder of Tables 4 and 5 present how the impacts differ by student characteristics such as gender, race and ethnicity, disability status, and LEP status. We are unable to provide these estimates for gun offenses and other offenses as the CRDC does not report these outcomes disaggregated by student characteristics. We find a CHP grant award results in a decrease of 0.135 per 100 students who experience law enforcement action for male students while the estimate is substantially smaller and not statistically significant for female students. The effect may be concentrated in male students as they received nearly 2.5 times as many law enforcement actions in the pre-treatment period as female students. The grant award also results in a decrease of 0.069 per 100 students who experience law enforcement action among students with disability status while the estimate is not statistically significant for students without disability status. This difference may be due to the fact SROs are able to gain a more holistic understanding of a student's background unlike patrol officers who are irregularly called onto campus as needed. We do not find meaningful effects across student subgroups on the count of students that experienced exclusionary discipline, chronic absenteeism, grade retention, or gun offenses. However, we note that, while statistically insignificant, the estimates for male students across these outcomes are larger than female students as with for black students compared to white students.

Our results differ from other papers that similarly leverage federal grants for SRO hiring, which find that grant receipt has positive effects on exclusionary discipline and law enforcement actions. This likely stems from differences in the specific grant mechanism (Owens (2017) and

Weisburst (2019) use the Cops In School (CIS) grant rather than the CHP grant), study period, and the comparison group. Sorensen et al. (2023) is the closest comparison to our paper as it also uses the CHP grant as its mechanism and the CRDC for its school-level outcomes, and its use of a regression discontinuity design (RDD) means its comparison group is also schools that applied but were not awarded the grant. The difference in findings may be the result of the fact that, by virtue of RDD, they are using a hyper-local comparison group while we compare schools with greater variation in grant scores; and their study period is schools that applied for the CHP grant between 2015 to 2017 while we focus those between 2013 to 2016. Appendix Table A2 contains estimates aggregated by treatment group and, while none are statistically significant, estimates for law enforcement actions for the 2013 and 2014 treatment groups are negative while those for 2015 and 2016 treatment groups have a positive relationship, revealing differences that can exist due to study period. These results further highlight the heterogeneity present in this setting and emphasize the salience of examining treatment effect heterogeneity.

5 Evidence of treatment effect heterogeneity

The most straightforward way to show that there is treatment effect heterogeneity would be to estimate the distribution of the treatment effect, or $\Delta_s = Y_{1s} - Y_{0s}$, which comes from taking the difference in the outcome between the treated and untreated state for each school in the sample. This is not possible to do, however, as we cannot observe both states. Even under an experimental design in which schools in the control group act as the proxy for the untreated state of treated schools, it is not possible to identify the treatment effect distribution because all the permutations of mapping individual schools in the control group to those in the treated group means there is a vast range of possible treatment effect distributions. Instead, we follow Heckman et al. (1997), Djebbari and Smith (2008), and more recently Buhl-Wiggers et al. (2024), and use classical statistical bounds established by Fréchet (1951) and Höffding (1970), referred to as "FH bounds," which uses the marginal distributionsof the treated and untreated outcomes to bound their joint distribution. While the FH bounding distributions do not bound the treatment effect distribution itself, they bound some of its *characteristics*, including the impact standard

deviation which captures the variation of the treatment effect distribution.

Using the potential outcomes framework, let $F_1(Y_1)$ be the outcome distribution of schools that received the CHP grant and let $F_0(Y_0)$ be the outcome distribution of schools that did not receive the grant. Let D indicate whether the school applied for the CHP grant. Then, $F(Y_1, Y_0)$ is the joint distribution and the FH bounds can be expressed as

$$max[F_1(Y_1|D=1) + F_0(Y_0|D=1) - 1, 0] \leq F(Y_1, Y_0|D=1) \leq min[F_1(Y_1|D=1), F_0(Y_0|D=1)]$$

The upper bound distribution corresponds to a case of perfect positive dependence, or "rank preservation," in which the rank correlation between the treated and control distributions is 1. This implies a school would maintain its relative position within the distribution after receiving treatment. The lower bound distribution corresponds to a case of perfect negative dependence, or "rank inversion," in which the rank correlation between the treated and control distributions is -1. This implies a school would attain a rank inverted from its original relative position within the distribution after receiving treatment (e.g., school at the 5th percentile attaining the 95th percentile); this is an extreme case that allows us to capture the full range of possible joint distributions.

We estimate the quantile treatment effect (QTE) at each ventile, or every fifth percentile from the 5th to 95th percentiles. It is important to note that these estimates in themselves cannot be interpreted as the treatment effect *at* the given ventile (e.g., grant award affects schools at the 5th percentile by Y additional counts of the outcome) unless the rank preservation assumption is satisfied. As we do not test that the rank preservation assumption holds, the QTEs merely informs about the shape of the outcome distribution. We also estimate the Pearson correlation ρ_{Y_0,Y_1} between the treated and control distributions and the proportion of schools that had a positive treatment effect. Finally, we estimate the impact standard deviation, which captures the standard deviation of the treatment effect distribution:

$$\begin{aligned} ISD &= \sqrt{Var(Y_1 - Y_0)} \\ &= Var(Y_1) + Var(Y_0) - 2\rho_{Y_0,Y_1}\sqrt{Var(Y_1)Var(Y_0)} \end{aligned}$$

The impact standard deviation is used to detect treatment effect heterogeneity in our setting. In a setting with common treatment effects, the variance of the treatment effect distribution is zero. Thus, an impact standard deviation that is bounded away from zero makes the case that there is not a common treatment effect, but that there is treatment effect heterogeneity. We take the additional step of testing the null hypothesis of a common treatment effect using the method outlined in Appendix E of Heckman et al. (1997). The method accounts for the fact that the null is at the boundary of the parameter space of a standard deviation of $[0, \infty)$. The method involves drawing observations from the control group to construct a faux treatment group in order to estimate the impact standard deviation. As draws are made from the control group, we would expect the estimated impact standard deviation to be zero. The distribution of the impact standard deviation from 1,000 draws is then used to construct cutoff values for testing the null hypothesis.

5.1 Empirical strategy

One major difficulty in estimating FH bounds within our study context is the fact that the marginal distribution of the control group is supposed to stand in for the untreated state of treated schools. This is the case in past uses of FH bounds (Heckman et al., 1997; Djebbari and Smith, 2008; Buhl-Wiggers et al., 2024) as in these studies, treatment assignment is randomized. Within our study setting, however, treatment is not randomized; instead, we use a quasi-experimental design of difference-in-differences with the key identifying assumption of conditional parallel trends, which allows us to use the outcomes from the control group, after differencing out the change over time, as the counterfactual of the treated group. Thus, we apply this idea of using the time effect within the control group to estimate the counterfactual distribution of the treated group, or the outcome treated schools would have attained had they not received treatment, and use this as the comparison distribution when estimating FH bounds.

To estimate this counterfactual distribution, we use the Changes-in-Changes framework from Athey and Imbens (2006), which Melly and Santangelo (2015) expand on to include covariates. Changes-in-Changes is analogous to difference-in-differences, but the main difference is

that the distribution of the treatment effect is identified instead of just the average treatment effect under difference-in-differences. To illustrate how the Changes-in-Changes method works, suppose the outcome Y_s is generated by a production function

$$Y_s = h(U_s, X_s, G_s)$$

where U_s contains unobservable school characteristics that contribute to the outcome production, X_s contains covariates, and G_s is a treatment indicator. Outcomes in the periods before treatment are generated by $h^{pre}(\cdot)$ for both groups while in the periods after treatment, the outcomes of the treated group are generated by $h^{post,treat}(\cdot)$. The realization of $U_s = u$ is used to match outcome distributions of treated and control schools as U_s determines the relative rank of a school within its group's distribution. Then, the quantiles of the treatment effect distribution are

$$\tau(u) = h^{post,treat}(u, 1) - h^{pre}(u, 1)$$

The Changes-in-Changes method requires the following identifying assumptions. The first assumption is that the outcome distribution of both treated and control groups in the pre-treatment periods are generated by the same function $h^{pre}(U_s, T_s)$. This assumption allows the mapping of treated and control group outcomes in the pre-treatment period. We also assume that the functions $h(\cdot)$ are strictly monotonic. This ensures that we are able to match treated and control schools by the realized value of U_s based on the relative position of the school within the outcome distribution as U_s is unobservable. Lastly, U_s must be time invariant within the groups, meaning that its relative position within the distribution stays the same between pre- and post-treatment periods. This assumption is necessary to identify the time effect within the control group, which is used to estimate the counterfactual distribution of the treated group.

From these assumptions, for each quantile $q \in (0, 1)$, the counterfactual outcome distri-

bution of treated schools can be identified by

$$F^{CF}_{Y,11} = F^{-1}_{Y,01}(F_{Y,00}(F^{-1}_{Y,10}(q)))$$

where $F_{Y,gt}(\cdot)$ is the cumulative density function for group g in time t and $F^{-1}_{Y,gt}$ is the inverse cumulative density function. For a given q, $F^{-1}_{Y,10}(q)$ outputs the associated value, which when put into $F_{Y,00}(\cdot)$ provides the quantile in the control group at the pre-treatment period that is associated with that value. This mapping is only possible because the production function for the outcome distributions in the pre-treatment period is the same for treated and control schools. Next, we find the change in the outcome over time for that quantile in the control group using $F^{-1}_{Y,01}(\cdot)$. This new value in the post-treatment period is the counterfactual outcome of the treated school at the given q. This procedure is applied to quantiles $q \in (0,1)$ to estimate the counterfactual distribution.

We adapt our study framework in order to estimate the counterfactual distribution. Specifically, we adapt a 2×2 framework for each treatment group: the period before treatment is used as the pre-treatment period and for the post-treatment period, we take the average across all outcomes in the post-treatment periods. Figure 2 maps the observed distribution of the treated group in the post-treatment period and the counterfactual distribution of the treated group estimated using the Changes-in-Changes method for the 2014 treatment group. We note that the counterfactual distributions closely follow the shape of the observed distributions.

We estimate the FH bounds by collapsing the outcome distributions of both the observed and counterfactual distributions into percentiles then take the difference between the outcome of the two groups at each percentile. Under rank inversion, we invert the order of the percentiles of the control group before taking the difference. To estimate the QTEs, we take the difference between the observed and counterfactual distributions at each ventile. We construct bootstrapped confidence intervals using the percentile method, drawing 1,000 bootstrap samples of schools from our sample with clusters at the district level. We estimate FH bounds separately for the 2013 and 2014 treatment groups, but exclude the 2015 and 2016 treatment groups due to small sample size. The bounds are estimated for disciplinary outcomes such as *suspen-*

sions, *expulsions*, and *law enforcement actions* as *gun offenses* and *other offenses* first appear in the 2015-16 CRDC and *chronic absenteeism* first appears in the 2013-14 CRDC.

5.2 Results

Table 6 contains FH bounds estimates at the 5th, 25th, 50th, 75th, and 95th percentiles in addition to the Pearson outcome correlation, impact standard deviation, and the fraction of schools in the treated group that had a positive, negative, or zero treatment effect. Estimates are presented for both the rank preservation and rank inversion cases. First, we find that for expulsions, FH bounds up to the 75th percentile are estimated at zero because most schools (86.3%) do not report any counts of expulsions. Law enforcement actions follow a similar pattern as 71.1% of schools do not report any counts of law enforcement actions. We also find larger treatment effects associated with higher percentiles, illustrating the variation in treatment effect across the outcome distribution.

Next, we examine treatment effect heterogeneity captured by differences in the direction of treatment effect. For suspensions, a majority of treated schools had a positive treatment effect, meaning suspensions increased due to SRO hiring, at the least 50.5% of schools under the rank inversion case for the 2014 treatment group and at most 90.5% of schools under the rank preservation case for the 2013 treatment group. A smaller proportion of schools had a negative treatment effect (between 0-42.4% of schools in the 2013 treatment group and 24.2-49.5% in the 2014 treatment group), but none or up to 1% of schools had a zero treatment effect. For expulsions, as suggested by the treatment effects on the ventiles, most treated schools had a zero treatment effect (77.8-88.9% of schools in the 2013 treatment group and 86.9-91.9% in the 2014 treatment group) with slightly more positive than negative effects in the 2013 treatment group but the opposite for the 2014 treatment group. There is less of an overwhelming trend for law enforcement actions: for the 2013 treatment group, around a third of schools have an increase in law enforcement actions following treatment and between 1-32.3% of schools experience a decrease while between 35.4-67.7% of schools have a zero treatment effect. However, for the 2014 treatment group, schools are more likely to experience a decrease in law enforcement

actions (between 41.4-51.5% of schools) with 11.1-48.5% of schools with a positive treatment effect. Anywhere between 0-47.5% of schools experienced a zero treatment effect. Lastly, for grade retention, a majority of schools had an increase (52.5-72.7% for the 2013 treatment group and 54.5-83.8% for the 2014 treatment group) while it is not clear whether the second largest group are schools that had a negative or zero effect. This reflects the upward trend in grade retentions observed in Figure 1 and may not be from SRO hiring as the conditional parallel trends assumption could not be satisfied.

Turning to the impact standard deviation, suspensions have the largest bounds, with (8.962, 23,719) in the 2013 treatment group and (1.923, 16.470) in the 2014 treatment group. This is consistent with the fact that suspensions have the smallest mass point at zero with 85.5% of schools in the analysis sample having a non-zero count of suspensions and having the widest range of values. Expulsions have the most modest bounds, with $(0.144, 0.408)$ in the 2013 treatment group and $(0.021, 0.130)$ in the 2014 treatment group as only 13.7% of schools in the analysis sample have a non-zero count of expulsions. Law enforcement action and grade retention have more comparable bounds with law enforcement action having bounds of $(0.607, 1.738)$ and $(0.206, 2.026)$ and grade retention having bounds of $(0.709, 3.944)$ and $(0.892, 5.680)$ for the 2013 and 2014 treatment groups, respectively. Comparing across treatment groups, the bounds on the impact standard deviation for the 2013 treatment group has a wider range than those for the 2014 treatment group.

We note that the lower bound on the impact standard deviation is bounded away from zero for all outcomes in both 2013 and 2014 treatment groups. However, when we compare the lower bound estimate against the Monte Carlo cutoff values constructed using the method from Appendix E of Heckman et al. (1997), we find that we can reject the null of a common treatment effect for suspensions and grade retention in both treatment groups, and expulsions and law enforcement actions in only the 2013 treatment group, all with a p-value of < 0.001. The inability to reject the null for expulsions and law enforcement actions in the 2014 treatment group may be due to the large proportion of zeroes within these outcomes, which reduces the range of potential values. These results provide evidence for treatment effect heterogeneity within the study setting by showing that there is variance in the effect of SRO hiring on student outcomes

across schools in our sample.

6 Sources of heterogeneity

6.1 Conceptual framework

How do SROs impact school discipline and why would we expect their impact to be heterogeneous across schools? At first glance, we would expect SROs to have no effect as they do not formally engage in school disciplinary decisions: the National Association of School Resource Officers (NASRO), a professional organization for SROs, recommends as a best practice that SROs be prohibited from being directly involved in student discipline NASRO (2015) and, when surveyed, SROs, schools administrators and staff, parents, and students agree with NASRO's stance (Curran et al., 2019; McKenna et al., 2016). While our results corroborate this viewpoint, past literature has instead shown that, despite their lack of a formal role, having SROs in schools increase rates of suspension and expulsion (Sorensen et al., 2021, 2023; Weisburst, 2019). This suggests that there must exist channels through which SROs indirectly impact school discipline.

One potential channel is the increased detection of rule-breaking activity. According to the 2019 Survey of Law Enforcement Personnel in Schools, nearly all (97.8%) SROs patrolled school grounds and 82.9% conducted video surveillance or monitoring of student behavior (Davis, 2023). The increased likelihood that student misbehavior is detected means school administrators also have a higher chance of being informed of incidents that are serious enough to warrant exclusionary discipline (Curran et al., 2019). In conjunction with increased detection, the literature on SROs also theorizes that SROs escalate the response to student misbehavior either by increasing the severity of the response (e.g. out-of-school suspension instead of in-school suspension) or moving from within-school disciplinary action to formal law enforcement action such as police referrals or arrests (Hirschfield, 2008; Thurau and Wald, 2010). Previous empirical studies have provided evidence suggesting that SRO presence increases the number of referrals to police for minor offenses (Theriot, 2009; Sorensen et al., 2021). Increased detection working in tandem with escalation is likely what drives the increase in exclusionary discipline noted in past

work: Weisburst (2019) finds that receiving grant funding for an SRO results in more students who experience disciplinary action for school code of conduct violations, with out-of-school suspension being the most common sanction.

Another potential channel which drives disciplinary action in the opposite direction is a deterrence effect. Previous literature shows that increased police presence, especially in "hot spots" or small local areas where crime is likely to occur, reduces the crime rate (Chalfin and McCrary, 2017; Owens, 2013). As the main responsibilities of the SRO require them to interact often with students in a highly visible way (e.g. patrolling the hallways, teaching a class), SRO presence on school grounds may deter students from misbehaving, thus reducing the need for discipline.

These channels can be used to think through how the variation across SRO programs and their local contexts translate into heterogeneity in SRO impact. For instance, consider the variation across individual school administrators, who all have their own approaches to school discipline. Heise and Nance (2021) observe that principals play a role in mediating students' exposure to law enforcement by choosing whether to report or omit certain disciplinary incidents. Then, schools whose principal is more flexible about engaging SROs may have a higher likelihood that incidents escalate to formal law enforcement action compared to schools whose principal is more restrictive. Considering there are a wide variety of factors about SRO programs that can differ across schools, we use the ecological systems model outlined in Curran et al. (2019) to systematically guide our untangling of these differences. The model lays out how the actions of the SRO are nested within broader contexts (e.g., school, district, state) and the influence that comes from interactions with different actors within each context (e.g., principal, school board, state legislation). We pair this with a review of the qualitative literature to determine factors that may be salient to SRO impact, which are outlined in the next section.

6.2 Moderators

Table 8 lists the moderators of interest, or school characteristics that proxy for the variation in program and school context that may explain the treatment effect heterogeneity observed in our

analysis. The first set of moderators are related to program context, or how the law enforcement agency approached the implementation of the SRO program and what issues within the school they set out to address. This may have an effect on school discipline by influencing how the SRO interprets his role within the school. For example, an SRO assigned to a school that has an issue with bullying may emphasize his role as a mentor or counselor compared to an SRO assigned to a school experiencing multiple drug or alcohol offenses in which the SRO may focus more on his law enforcement role. The moderators are identified via a text analysis that found five common themes in the grant narrative where law enforcement agencies were asked to explain what problems were facing the school. Section 6.3 details the text analysis method and explains the themes that were identified.

We also consider the school's racial composition. Previous literature shows that SRO programs have a disproportionately negative impact on students of color (Anyon et al., 2014; Sorensen et al., 2023; Weisburst, 2019). This disparity may arise because the racial composition of a school can impact an SRO's perceptions of whether potential threats are internal or external to the school, which may subsequently affect his approach to student discipline (Fisher et al., 2022). We use the CRDC to identify schools that are comprised of a majority of black or Hispanic students (more than 75% of students are black or Hispanic in the 2011-12 school year) and schools that are comprised of a majority of white students (more than 75% of students are white in the 2011-12 school year). Lastly, we consider the interactions between SROs and school staff, specifically a school counselor. Having a counselor on school grounds may mediate the effect that SROs have on school discipline as counselors may be able to provide the context behind students' misbehavior, which could prevent incidents from escalating into formal disciplinary action. We use the 2011-12 CRDC to identify the presence of counselors in schools.

6.3 Text-based analysis on grant narratives

To understand how SRO programs may be implemented differently across different schools, we train text models on CHP grant narratives to identify common themes. These narratives vary across applicants, but often contain descriptions of the problems that motivate the need for

SRO hiring, what activities the SROs would perform, and how they would be allocated across schools. As an outcome of the text-based analysis, we can categorize schools into groups based on their SRO program attributes for treatment effect heterogeneity analysis.

We used the Correlated Topic Model (CTM) to identify themes in the grant application texts (Blei and Lafferty, 2007). The CTM treats input texts as collections of words, and assigns different topics to texts based on the frequency at which words appear using a variational Expectation-Maximization (EM) algorithm. See Appendix Section C for a more detailed description of the CTM and how it is estimated. When using the CTM, researchers need to select the optimal number of topics K for their model. Using cross-validation, we search over a range of possible number of latent topics ($K \in \{2, ..., 20\}$). The main trade-off in selecting K is between topic semantic coherence, whether a topic represents a set of words that appear together frequently, and topic exclusivity, whether the words under a topic are distinct from those in other topics. Based on the metrics generated from testing a range of topic numbers K (Appendix Figure A3), we select $K = 5$ topics to proceed with the CTM estimation.

The resulting topics from the CTM estimation are presented in Table 9. The model assigns a score indicating the prevalence of each topic in a given text, and we consider a grant application as discussing a given topic k if the prevalence score is above the 75th percentile. Based on the most common words under each topic, it is clear that the grant narrative texts capture what applicants perceive as the key motivations for the SROs, not necessarily the exact details of their SRO program implementation. However, previous qualitative studies on SROs theorize and describe associations between perceptions of SRO roles and how SRO programs are implemented. To give an illustrating example, Fisher et al. (2022) find that SROs vary in terms of how they perceive sources of threat (students vs. intruders). Using the methods outlined above, we can construct sets of keywords that indicate whether SROs would be hired primarily to discipline students or protect students from outsiders. For instance, schools in the former category may be more likely to mention words such as "interpersonal" and "bullying," while those in the latter may be more likely to mention "shooter" and "outsider." It is important to acknowledge that though the grant narratives enrich our understanding of how the applicants *intend* to implement their SRO programs, they do not perfectly predict what happens in reality. However, we

believe the narrative data still provide useful information on SRO program heterogeneity.

Looking at Table 9, we find that the keywords associated with each model-estimated topic represent distinct approaches to SRO programs. Topic 1 is focused on problems with student behavior, such as fighting, truancy, and bullying, while topics 2 and 3 feature words associated with responding to outside threats. Meanwhile, the keywords in topic 4 relate to how SROs could form positive relationships with students, build trust, and work as role models. Finally, topic 5 address calls made to local law enforcement and potential budget constraints that prevent officers from being assigned to patrol within or around schools.

7 How moderators explain treatment effect heterogeneity

We have shown evidence of treatment effect heterogeneity using the FH bounding distributions; next, we consider how the moderators of interest described in Section 6.2 moderate the impact of SRO hiring on student outcomes. To do so, we divide our analysis sample into subgroups based on a binary measure of each moderator and estimate the ATT separately for the moderator and comparison subgroups. We test the null of equal subgroup ATTs and construct 95% confidence intervals using the Wild bootstrap. To estimate the subgroup ATTs for *chronic absenteeism*, which appears only in the 2013-14 the 2015-16 CRDCs, we use the 2014 and 2015 treatment groups. We do not estimate subgroup ATTs for *gun offenses* and *other offenses* as these outcomes appear only in the 2015-16 and 2017-18 CRDCs, and the 2016 treatment group has a small sample size.

Table 11 contains the post-treatment ATT estimates by moderator. For *suspensions*, schools whose agency's grant narrative contained themes that focused on building community and trust within schools (Topic 4) have a decrease of 0.855 per 100 students who experienced suspensions while schools whose grant narrative did not include this theme have an increase of 1.297 per 100 students who experienced suspensions. Schools that had a grant narrative that focused on the SRO's traditional law enforcement responsibilities (Topic 5) similarly have a decrease in suspensions (-0.317 per 100 students) while schools whose narratives did not include this theme have

an increase (0.935 per 100 students). However, schools that focused on community have a much larger decrease; one explanation may be that building positive relationships with students either deters students from misbehavior through mentorship or reduces the use of exclusionary discipline.

Next, we examine how moderators impacted *expulsions*. Schools whose agency's grant narrative contained themes that viewed threats to school safety as being internal (Topic 1) have a decrease of 0.157 per 100 students who experienced expulsions while schools whose grant narrative did not include this theme have an increase of 0.074 per 100 students who experienced expulsions. In comparison, schools that viewed threats to school safety as being external (Topic 3) have an increase of 0.159 per 100 students who experienced expulsions while schools that did not hold such a view had a smaller increase of 0.027 per 100 students. One possible explanation for these results is that schools whose SRO is focused on reducing within-school student misbehavior may focus more on duties such as patrolling hallways that makes him more visible to students, possibly deterring misbehavior. In terms of schools that view threats as external, Topic 3 is associated with words about crime occurring outside the school, and these schools may be more proactive in expelling students who are involved in criminal activity or having students who are involved in external criminal activity in school may make it more likely that their offenses are severe enough to warrant expulsion. Lastly, schools who had counselors had an additional 0.122 per 100 students who were expelled compared to schools that did not have a counselor.

For *law enforcement actions*, schools whose agency's grant narrative contained themes that focused on community (Topic 4) have a decrease of 0.373 per 100 students who experienced law enforcement actions while schools whose grant narrative did not include this theme have a smaller decrease of 0.106 per 100 students. Schools that had a grant narrative that focused on law enforcement (Topic 5) also have a decrease of 0.406 per 100 students who experienced law enforcement actions compared to a decrease of 0.135 per 100 students in schools whose narratives did not include this theme. While the differences in ATTs is similar between Topics 4 and 5, schools that discussed the theme from Topic 5 had a larger decrease: it could be the case that SROs in schools focused on traditional law enforcement responsibilities are more visible to

students, which may have a deterrence effect on students. In addition, schools with a majority of minority students had 0.218 per 100 less students who experienced law enforcement action than schools with more mixed student demographics.

Looking at *grade retention*, schools that had a grant narrative that focused on law enforcement (Topic 5) have an additional 0.363 per 100 students that retained a grade compared to schools whose narratives did not include this theme. Schools with a majority of minority students also have an additional 2.542 per 100 students that retained a grade than schools with more mixed student demographics. This is an 88.0% increase from the pre-treatment mean. These results may stem from differences in students' sense of safety at school after SROs are stationed, which may impact their learning. Theriot and Orme (2015) find that black students felt less safe at school when there was an SRO, which may explain the increase in grade retention in majority minority schools. In addition, Curran et al. (2021) find that interactions with SROs increased students' sense of danger; it may be the case that in schools that focus on traditional law enforcement responsibilities, the SRO is more visible to students and also interacts with students more frequently.

Lastly, we consider *chronic absenteeism*. Schools that had a grant narrative that viewed threats to school safety being internal (Topic 1) have a large decrease of 12.596 per 100 students who were chronically absent compared to a decrease of 0.332 per 100 in schools whose grant narrative did not contain this theme. We also see that schools that focused on community (Topic 4) have a large decrease of 10.402 per 100 students who were chronically absent compared to a decrease of 0.747 per 100 in schools whose grant narrative did not contain this theme. These results may be because schools that have the SRO focus on reducing internal threats such as fighting or bullying or have the SRO build a sense of community and trust within school may succeed in creating a better school environment for students, reducing the number of students who are chronically absent. In line with the idea of creating a better environment for students, we also see that schools with counselors have a decrease of 4.874 per 100 students who were chronically absent while schools that did not have a decrease of 2.488 per 100.

8 Conclusion

In this paper, we explore treatment effect heterogeneity in SRO hiring on school disciplinary and academic outcomes and examine how the variation in treatment effects can be explained by SRO program and school contexts. Following the framework in Callaway and Sant'Anna (2021), we use a staggered difference-in-differences design that leverages a federal grant for SRO hiring and compares schools that received the grant against schools that applied but did not receive the grant. We estimate mean effects that corroborate findings in past works about violent offenses, male students, and students with disabilities (Weisburst, 2019; Sorensen et al., 2021, 2023) although we differ in finding decreases in law enforcement actions.

Using classical statistical bounds called FH bounds, we identify the bounds on the variance of the distribution of treatment effects and test the null hypothesis of a common treatment effect across schools. This is done by estimating the counterfactual distribution of the treated group using the change over time in the control group, following the Changes-in-Changes method (Athey and Imbens, 2006). We find evidence of meaningful treatment effect heterogeneity at a p-value < 0.001 for suspensions and grade retention in schools treated in 2013 and 2014, and for expulsions and law enforcement actions in schools treated in 2013. Then, having established there is treatment effect heterogeneity to dig into, we use novel grant narrative data and observable school characteristics to proxy for the variation in program and school contexts. We estimate the treatment effect separately for schools that were associated with a given moderator and for schools that were not to examine how these moderators mediate the effect of SRO hiring. We find program and school context do explain aspects of observed treatment effect heterogeneity.

Our paper contributes to the existing literature on the impact of SROs on student discipline by systematically examining the treatment effect heterogeneity in SRO hiring and how student outcomes are impacted as a result. Our findings can help policymakers better understand what SRO program and school contexts may have a magnifying or dampening effect on the impact of SROs, which may help inform future efforts on minimizing the role that SROs may play in facilitating student exposure to the criminal justice system.

Tables

Table 1: Summary statistics of main analysis sample by treatment status

	Never applied	Applied but not awarded	Awarded	Difference	P-value
Outcomes per 100 students (winsorized)					
Suspensions	11.59	11.77	15.77	4.00	0.00
	(14.98)	(15.83)	(20.80)		
Expulsions	0.21	0.17	0.10	-0.07	0.00
	(0.86)	(0.71)	(0.53)		
Law enforcement actions	0.55	0.48	0.93	0.45	0.00
	(1.66)	(1.38)	(1.89)		
Chronic absenteeism	12.37	12.80	14.69	1.89	0.00
	(11.58)	(10.32)	(10.31)		
Grade retention	2.21	2.82	3.29	0.47	0.00
	(3.43)	(6.36)	(5.87)		
Gun offenses	0.02	0.02	0.01	-0.01	0.00
	(0.20)	(0.13)	(0.08)		
Other offenses	1.97	1.72	2.79	1.08	0.00
	(5.26)	(3.99)	(6.11)		
Student characteristics					
Prop. Male	0.52	0.52	0.52	0.00	0.25
	(0.05)	(0.05)	(0.06)		
Prop. White	0.63	0.55	0.46	-0.09	0.00
	(0.31)	(0.32)	(0.28)		
Prop. Black	0.14	0.18	0.22	0.05	0.00
	(0.23)	(0.24)	(0.24)		
Prop. Hispanic	0.15	0.19	0.23	0.04	0.00
	(0.22)	(0.25)	(0.22)		
Prop. Other race	0.08	0.08	0.08	0.00	0.05
	(0.10)	(0.08)	(0.09)		
Prop. Disabled	0.15	0.14	0.16	0.01	0.00
	(0.11)	(0.10)	(0.13)		
Prop. LEP	0.07	0.09	0.10	0.02	0.00
	(0.12)	(0.14)	(0.13)		
Prop. Free or reduced price lunch	0.49	0.50	0.59	0.09	0.00
	(0.26)	(0.27)	(0.25)		
Student enrollment	528.56	626.59	696.42	69.83	0.00
	(403.18)	(425.07)	(485.92)		
Teacher-to-pupil ratio	15.44	16.20	16.06	-0.14	0.34
	(6.86)	(5.54)	(7.97)		
N	52,249	8,118	2,468		

	Never applied	Applied but not awarded	Awarded	Difference	P-value
School characteristics					
Elementary school	0.62	0.66	0.67	0.00	0.66
	(0.49)	(0.47)	(0.47)		
Middle school	0.31	0.27	0.25	-0.02	0.06
	(0.46)	(0.44)	(0.43)		
High school	0.23	0.17	0.18	0.00	0.82
	(0.42)	(0.38)	(0.38)		
Urban location	0.20	0.21	0.34	0.13	0.00
	(0.40)	(0.41)	(0.48)		
Suburban location	0.27	0.39	0.38	-0.02	0.13
	(0.44)	(0.49)	(0.48)		
Town location	0.15	0.12	0.10	-0.02	0.01
	(0.35)	(0.33)	(0.30)		
Rural location	0.38	0.28	0.18	-0.10	0.00
	(0.49)	(0.45)	(0.38)		
Northeast region	0.18	0.13	0.05	-0.08	0.00
	(0.38)	(0.33)	(0.22)		
South region	0.36	0.59	0.68	0.09	0.00
	(0.48)	(0.49)	(0.47)		
Midwest region	0.36	0.19	0.18	-0.01	0.45
	(0.48)	(0.39)	(0.39)		
West region	0.10	0.10	0.09	-0.01	0.13
	(0.31)	(0.30)	(0.28)		
School staff					
Has counselor	0.91	0.91	0.87	-0.04	0.02
	(0.29)	(0.29)	(0.34)		
N	52,249	8,118	2,468		

Notes: Columns 1, 2, and 3 contain the unweighted mean and standard deviation for schools that never applied for the CHP grant between 2013 and 2016, schools that applied but did not receive the CHP award and the schools that applied for and received the CHP award, respectively. Column 4 is the difference between the means of schools that applied for the grant but did not receive it and schools that received the grant and Column 5 presents the p-value from a t-test of the null of equal population means. Chronic absenteeism and whether the school has an SRO were measured from the 2013-14 CRDC; gun offenses and other offenses were measured from the 2015-16 CRDC as these questions were not asked in earlier waves; all other variables are measured from the 2011-12 CRDC. All outcomes are winsorized at the 99th percentile and are normalized per 100 students.

Table 2: Treatment group by CHP grant application year

	No CHP award		CHP award	
	N schools	Prop. schools	N schools	Prop. schools
2013	4,586	0.56	1,034	0.42
2014	1,035	0.13	1,198	0.49
2015	1,221	0.15	179	0.07
2016	1,276	0.16	57	0.02
Total	8,118		2,468	

Notes: Schools in the "No CHP Award" group are those in the "never treated" group.

Table 3: Pre-treatment means of student outcomes

	Suspensions	Expulsions	Law enforcement actions	Chronic absenteeism	Grade retention	Gun offenses	Other offenses
Overall	12.787	0.148	0.578	13.278	2.890	0.019	1.974
	(17.531)	(0.654)	(1.511)	(10.329)	(5.924)	(0.136)	(4.716)
Male	9.317	0.113	0.424	7.505	1.860	-	-
	(13.128)	(0.530)	(1.148)	(7.724)	(4.690)	-	-
Female	4.112	0.039	0.172	6.935	1.305	-	-
	(7.585)	(0.261)	(0.628)	(7.090)	(4.120)	-	-
Black	5.054	0.053	0.184	2.773	1.678	-	-
	(11.211)	(0.342)	(0.703)	(5.833)	(4.591)	-	-
Hispanic	2.777	0.024	0.136	3.029	1.540	-	-
	(7.273)	(0.216)	(0.616)	(4.761)	(5.099)	-	-
Other	0.823	0.011	0.040	1.133	0.442	-	-
	(2.367)	(0.105)	(0.250)	(2.138)	(1.590)	-	-
White	4.842	0.067	0.237	6.486	2.134	-	-
	(8.435)	(0.406)	(0.845)	(6.975)	(4.981)	-	-
Has disability status	3.425	0.037	0.177	3.074	0.722	-	-
	(6.389)	(0.308)	(0.708)	(4.938)	(2.227)	-	-
No disability status	10.033	0.117	0.419	10.272	2.275	-	-
	(15.159)	(0.555)	(1.167)	(8.384)	(5.363)	-	-
Has LEP status	0.703	0.006	0.037	1.082	0.070	-	-
	(2.297)	(0.063)	(0.264)	(2.117)	(0.299)	-	-
No LEP status	12.402	0.142	0.548	12.229	2.829	-	-
	(17.631)	(0.642)	(1.466)	(9.971)	(5.912)	-	-

Notes: This table presents the mean outcome in the pre-treatment period for all schools in our sample. Standard deviations are reported in the parentheses. The means are measured from the 2013-14 school year for chronic absenteeism, 2015-16 school year for gun and other offenses, and 2011-12 for the remaining outcomes. All outcomes are defined as counts per 100 students and are winsorized at the 99th percentile.

Table 4: Post-treatment ATT estimates for student disciplinary outcomes

	Suspensions	Expulsions	Law enforcement actions
Overall	0.429	0.058	−0.186
	[-1.240, 2.098]	[-0.015, 0.130]	[-0.365, -0.006]
Male	0.326	0.047	−0.135
	[-0.995, 1.647]	[-0.004, 0.098]	[-0.259, -0.012]
Female	0.208	0.017	−0.054
	[-0.366, 0.782]	[-0.009, 0.042]	[-0.125, 0.017]
Black	0.454	0.025	−0.062
	[-0.524, 1.432]	[-0.011, 0.061]	[-0.134, 0.010]
Hispanic	0.212	0.008	−0.084
	[-0.311, 0.735]	[-0.010, 0.026]	[-0.182, 0.014]
Other	−0.052	0.014	−0.025
	[-0.309, 0.205]	[-0.001, 0.029]	[-0.072, 0.021]
White	0.106	0.020	−0.019
	[-0.536, 0.747]	[-0.008, 0.048]	[-0.074, 0.036]
Has disability status	0.134	0.016	−0.069
	[-0.510, 0.778]	[-0.005, 0.037]	[-0.136, -0.003]
No disability status	0.514	0.046	−0.121
	[-0.906, 1.933]	[-0.013, 0.106]	[-0.256, 0.013]
Has LEP status	0.032	0.003	−0.027
	[-0.122, 0.187]	[-0.002, 0.007]	[-0.061, 0.008]
No LEP status	0.462	0.055	−0.161
	[-1.196, 2.120]	[-0.015, 0.124]	[-0.319, -0.002]

Notes: This table presents the aggregate post-treatment ATT estimates for each student disciplinary outcome for all treated schools. All outcomes are defined as counts per 100 students and are winsorized at the 99th percentile. The 95% confidence intervals are reported in brackets. The standard errors are estimated using Wild bootstrap and clustered at the school district level.

Table 5: Post-treatment ATT estimates for other student outcomes

	Absenteeism	Grade Retention	Gun Offenses	Other Offenses
Overall	-2.495	-0.212	-0.005	-1.506
	[-7.056, 2.065]	[-0.927, 0.503]	[-0.032, 0.022]	[-2.224, -0.789]
Male	-1.369	0.572	-	-
	[-3.935, 1.197]	[-1.233, 2.378]	-	-
Female	-1.334	0.594	-	-
	[-3.835, 1.167]	[-1.202, 2.390]	-	-
Black	-1.264	0.263	-	-
	[-3.969, 1.442]	[-0.434, 0.959]	-	-
Hispanic	-0.608	0.702	-	-
	[-2.117, 0.900]	[-1.316, 2.721]	-	-
Other	-0.085	0.108	-	-
	[-0.351, 0.181]	[-0.402, 0.618]	-	-
White	-0.529	0.552	-	-
	[-1.273, 0.215]	[-0.795, 1.898]	-	-
Has disability status	-0.468	0.194	-	-
	[-1.341, 0.405]	[-0.330, 0.717]	-	-
No disability status	-2.000	0.615	-	-
	[-5.936, 1.936]	[-1.324, 2.554]	-	-
Has LEP status	0.083	-0.048	-	-
	[-0.248, 0.414]	[-0.194, 0.099]	-	-
No LEP status	-2.584	0.676	-	-
	[-7.095, 1.927]	[-1.268, 2.619]	-	-

Notes: This table presents the aggregate post-treatment ATT estimates for other student outcomes for all treated schools. All outcomes are defined as counts per 100 students and are winsorized at the 99th percentile. The 95% confidence intervals are reported in brackets. The standard errors are estimated using Wild bootstrap and clustered at the school district level. Note that for chronic absenteeism, data is only available for the 2013-14 and 2015-16 school years and so only schools treated in 2014 and 2015 are compared against the never treated group. For gun and other offenses, data is only available for the 2015-16 and 2017-18 school years and so, only schools treated in 2016 are compared against the never treated group.

Table 6: Fréchet-Höffding (FH) bounds

	2013 treatment group		2014 treatment group	
	Rank preservation	Rank inversion	Rank preservation	Rank inversion
Suspensions (per 100 students)	(1)	(2)	(3)	(4)
5th	0.000	54.068	0.000	30.820
	[0.000,0.000]	[41.940,64.332]	[0.000,0.000]	[24.707,38.376]
25th	0.968	18.238	-0.085	10.520
	[0.600,1.480]	[12.854,23.112]	[-0.231,0.027]	[8.737,13.014]
50th	2.610	2.610	0.070	0.070
	[1.763,3.442]	[1.763,3.442]	[-0.043,0.521]	[-0.043,0.521]
75th	8.672	-8.599	0.523	-10.082
	[6.311,11.208]	[-10.974,-5.800]	[0.341,1.660]	[-11.605,-8.083]
95th	26.537	-27.531	4.278	-26.541
	[19.374,32.740]	[-31.719,-22.964]	[2.122,7.520]	[-30.780,-21.853]
Outcome Correlation	0.997	-0.532	0.995	-0.593
	[0.991,0.999]	[-0.609,-0.460]	[0.992,0.999]	[-0.657,-0.545]
Impact Standard Deviation	8.962	23.719	1.923	16.470
	[7.212,10.833]	[19.510,27.598]	[1.036,2.864]	[13.771,19.054]
Fraction Positive	0.909	0.576	0.677	0.505
	[0.869,0.939]	[0.556,0.596]	[0.556,0.747]	[0.495,0.515]
Fraction Negative	0.000	0.424	0.242	0.495
	[0.000,0.000]	[0.404,0.444]	[0.172,0.364]	[0.485,0.505]
Fraction Zero	0.091	0.000	0.081	0.000
	[0.061,0.131]	[0.000,0.000]	[0.061,0.101]	[0.000,0.000]
Expulsions (per 100 students)	(1)	(2)	(3)	(4)
5th	0.000	0.642	0.000	0.085
	[0.000,0.000]	[0.250,1.260]	[0.000,0.000]	[0.000,0.428]
25th	0.000	0.000	0.000	0.000
	[0.000,0.000]	[0.000,0.000]	[0.000,0.000]	[0.000,0.000]
50th	0.000	0.000	0.000	0.000
	[0.000,0.000]	[0.000,0.000]	[0.000,0.000]	[0.000,0.000]
75th	0.000	0.000	0.000	0.000
	[0.000,0.000]	[0.000,0.000]	[0.000,0.000]	[0.000,0.000]
95th	0.157	-0.485	-0.052	-0.137
	[0.018,0.466]	[-0.795,-0.227]	[-0.137,0.064]	[-0.416,-0.094]
Outcome Correlation	0.992	-0.072	0.972	-0.046
	[0.977,0.996]	[-0.107,-0.042]	[0.861,0.993]	[-0.077,-0.026]
Impact Standard Deviation	0.144	0.408	0.021	0.130
	[0.076,0.257]	[0.256,0.595]	[0.019,0.044]	[0.061,0.322]
Fraction Positive	0.091	0.111	0.030	0.051
	[0.051,0.141]	[0.066,0.162]	[0.000,0.071]	[0.020,0.091]
Fraction Negative	0.020	0.111	0.051	0.081
	[0.000,0.040]	[0.066,0.162]	[0.030,0.091]	[0.051,0.121]
Fraction Zero	0.889	0.778	0.919	0.869
	[0.838,0.939]	[0.677,0.869]	[0.879,0.949]	[0.788,0.929]

	2013 treatment group		2014 treatment group	
	Rank preservation	Rank inversion	Rank preservation	Rank inversion
LE actions (per 100 students)	(1)	(2)	(3)	(4)
5th	0.000	3.577	0.000	3.782
	[0.000,0.000]	[2.410,4.812]	[0.000,0.000]	[3.142,4.545]
25th	0.000	0.506	0.000	1.076
	[0.000,0.000]	[0.000,0.950]	[0.000,0.000]	[0.761,1.313]
50th	0.000	0.000	-0.054	-0.054
	[0.000,0.000]	[0.000,0.000]	[-0.120,0.000]	[-0.120,0.000]
75th	0.142	-0.365	-0.191	-1.267
	[-0.010,0.349]	[-0.626,0.000]	[-0.298,-0.007]	[-1.444,-0.902]
95th	1.417	-2.160	0.406	-3.376
	[0.765,2.091]	[-2.913,-1.553]	[0.294,0.656]	[-3.972,-2.673]
Outcome Correlation	0.994	-0.197	0.991	-0.438
	[0.985,0.999]	[-0.277,-0.142]	[0.954,0.995]	[-0.519,-0.256]
Impact Standard Deviation	0.607	1.738	0.206	2.026
	[0.349,0.907]	[1.141,2.360]	[0.162,0.575]	[1.752,2.473]
Fraction Positive	0.313	0.323	0.111	0.485
	[0.212,0.414]	[0.227,0.424]	[0.081,0.242]	[0.384,0.495]
Fraction Negative	0.010	0.323	0.414	0.515
	[0.000,0.030]	[0.232,0.424]	[0.263,0.505]	[0.384,0.525]
Fraction Zero	0.677	0.354	0.475	0.000
	[0.576,0.768]	[0.152,0.540]	[0.384,0.616]	[0.000,0.232]
Grade retention (per 100 students)	(1)	(2)	(3)	(4)
5th	0.000	8.000	0.000	10.511
	[0.000,0.000]	[6.486,9.016]	[0.000,0.000]	[9.148,11.272]
25th	-0.006	2.677	0.231	4.681
	[-0.041,0.119]	[1.702,3.327]	[-0.004,0.429]	[4.008,5.285]
50th	0.183	0.183	0.802	0.802
	[0.075,0.259]	[0.075,0.259]	[0.384,0.906]	[0.384,0.906]
75th	0.541	-2.142	1.245	-3.205
	[0.319,0.717]	[-2.607,-1.374]	[1.029,1.408]	[-3.890,-2.553]
95th	2.208	-5.792	2.594	-7.917
	[1.636,2.760]	[-6.853,-4.808]	[2.113,3.115]	[-8.635,-6.933]
Outcome Correlation	0.998	-0.531	0.999	-0.773
	[0.989,0.999]	[-0.657,-0.390]	[0.994,1.000]	[-0.850,-0.589]
Impact Standard Deviation	0.709	3.944	0.892	5.680
	[0.602,0.988]	[3.243,4.711]	[0.748,1.036]	[4.934,6.057]
Fraction Positive	0.727	0.525	0.838	0.545
	[0.616,0.808]	[0.515,0.535]	[0.616,0.919]	[0.525,0.566]
Fraction Negative	0.030	0.475	0.010	0.455
	[0.020,0.051]	[0.465,0.485]	[0.000,0.040]	[0.434,0.475]
Fraction Zero	0.242	0.000	0.152	0.000
	[0.162,0.343]	[0.000,0.000]	[0.071,0.359]	[0.000,0.000]

Notes: Columns 1 and 3 contain results estimated under rank preservation, or when percentiles of the treated distribution are matched to corresponding percentiles of the comparison distribution (i.e. 5th percentile of the treated distribution matched to 5th percentile of the comparison distribution). Columns 2 and 4 contains results estimated under rank inversion, or when percentiles of the treated distribution are matched to the inverse percentile of the comparison distribution (i.e. 5th percentile of the treated distribution matched to 95th percentile of the comparison distribution). Bootstrapped 95% confidence intervals (using 1,000 replications) are in brackets.

Table 7: Simulated cutoff values for the test of the null of a common treatment effect

	Cutoff value for			
P-value	Suspensions	Expulsions	LE action	Grade retention
Panel A: 2013 treatment group				
0.001	1.933	0.108	0.252	0.405
0.010	1.535	0.089	0.217	0.343
0.050	1.206	0.068	0.182	0.275
0.100	1.070	0.060	0.159	0.249
0.200	0.918	0.050	0.136	0.211
0.300	0.839	0.044	0.121	0.190
0.400	0.768	0.040	0.110	0.173
0.500	0.704	0.036	0.098	0.158
Panel B: 2014 treatment group				
0.001	1.474	0.105	0.623	0.626
0.010	1.139	0.088	0.476	0.474
0.050	0.949	0.065	0.218	0.394
0.100	0.831	0.058	0.164	0.370
0.200	0.716	0.048	0.134	0.331
0.300	0.640	0.036	0.117	0.301
0.400	0.585	0.030	0.108	0.241
0.500	0.544	0.025	0.095	0.217

Notes: This table contains cutoff values that correspond to p-values. The values are estimated using the method outlined in Appendix E of Heckman et al. (1997), which involves drawing observations from the control group to construct a faux treatment group and then estimating the impact standard deviation. As draws are made from the control group, the expected impact standard deviation is 0. The distribution of the impact standard deviation from 1,000 draws is then used to construct the cutoff values. Panel A contains values constructed using the 2013 treatment group and Panel B contains values for the 2014 treatment group.

Table 8: Moderators of interest

Category	**Moderator**	**How it will be measured**
Program context	Topic 1 (internal threat)	Text analysis using CHP grant narrative data (see Section 6.3)
	Topic 2 (external threat)	Same as above
	Topic 3 (external threat)	Same as above
	Topic 4 (focused on community)	Same as above
	Topic 5 (focused on law enforcement)	Same as above
Student demographics	Majority black/Hispanic	75%+ of students are black or Hispanic
	Majority white	75%+ of students are white
School staff	Any counselors	School has a counselor

Notes: This table lists the moderators of interest and how they are measured.

Table 9: Grant narrative topics from CTM

Topic	Theme	Top associated words	Median prevalence
1	Internal threat	fight, assault, truancy, abuse, violent, report, bullying	0.13
2	External threat	plan, administrator, meeting, inform, emergency, security, staff	0.17
3	External threat	crime, prevent, address, awareness, conflict, expand, problem	0.15
4	Focused on community	positive, relationship, trust, commit, interact, model, communities	0.18
5	Focused on law enforcement	call, patrol, deputies, assign, budget, location, full	0.16

Notes: This table shows the topics generated by the Correlated Topic Model (CTM) trained on CHP applicant narrative responses to the prompt "description of problem/focus area." The "top associated words" column presents some of the most common vocabulary classified into each topic, and the column "median prevalence" shows the median prevalence score of each topic across all the applicant narratives. For instance, the prevalence score of topic 1 is defined as the proportion of words in a given text that fall under topic 1.

Table 10: Proportion of schools in each moderator subgroup by CHP award status

	No CHP award	CHP award
Program context		
Topic 1 (internal threat)	0.18	0.26
Topic 2 (external threat)	0.26	0.20
Topic 3 (external threat)	0.25	0.22
Topic 4 (focused on community)	0.30	0.31
Topic 5 (focused on LE)	0.23	0.25
Student demographics		
Majority minority students	0.31	0.26
Majority white students	0.47	0.25
School staff		
Any counselors	0.17	0.16
N	8,111	2,466

Notes: This table shows the proportion of schools in the analysis sample that belong to the moderator subgroup, which is determined by using binary measures of the moderators of interest. Table 9 contains further information about each narrative topic and Section 6.3 further details how these narrative topics were determined using the CHP grant narrative data. Student demographics and presence of counselors is measured using the CRDC.

Table 11: Post-treatment ATT estimates: By moderators

	Suspensions			Expulsions		
	Moderator ATT	Comparison ATT	Diff in ATT	Moderator ATT	Comparison ATT	Diff in ATT
	(1)	(2)	(3)	(1)	(2)	(3)
Program context						
Topic 1 (internal threat)	1.489	0.871	0.618	-0.157	0.074	-0.230
	[-0.367, 3.346]	[0.324, 1.418]	[-1.078, 2.315]	[-0.272, -0.041]	[0.048, 0.099]	[-0.343, -0.117]
Topic 2 (external threat)	1.691	0.656	1.035	0.031	0.067	-0.036
	[0.296, 3.086]	[0.088, 1.225]	[-0.301, 2.372]	[-0.049, 0.111]	[0.041, 0.093]	[-0.113, 0.041]
Topic 3 (external threat)	-0.474	0.568	-1.042	0.159	0.027	0.132
	[-1.601, 0.653]	[0.075, 1.061]	[-2.025, -0.059]	[0.074, 0.245]	[0.007, 0.047]	[0.062, 0.203]
Topic 4 (focused on community)	-0.855	1.297	-2.152	0.066	0.055	0.011
	[-1.635, -0.075]	[0.602, 1.992]	[-3.530, -0.774]	[0.023, 0.110]	[0.020, 0.090]	[-0.065, 0.087]
Topic 5 (focused on LE)	-0.317	0.935	-1.253	0.105	0.055	0.050
	[-1.400, 0.766]	[0.390, 1.481]	[-2.414, -0.091]	[0.053, 0.157]	[0.023, 0.086]	[-0.024, 0.124]
Student demographics						
Majority minority students	1.347	0.966	0.381	0.046	0.082	-0.035
	[0.008, 2.687]	[0.159, 1.774]	[-1.161, 1.923]	[-0.010, 0.103]	[0.032, 0.132]	[-0.157, 0.086]
Majority white students	0.134	0.966	-0.832	0.038	0.082	-0.044
	[-0.436, 0.704]	[0.159, 1.774]	[-1.897, 0.233]	[0.000, 0.076]	[0.032, 0.132]	[-0.108, 0.021]
School staff						
Any counselors	0.471	0.756	-0.285	0.167	0.044	0.122
	[-1.232, 2.175]	[0.285, 1.227]	[-1.178, 0.609]	[0.071, 0.262]	[0.020, 0.069]	[0.057, 0.187]

	Law enforcement actions			Grade retention		
	Moderator ATT (1)	Comparison ATT (2)	Diff in ATT (3)	Moderator ATT (1)	Comparison ATT (2)	Diff in ATT (3)
Program context						
Topic 1 (internal threat)	-0.046	-0.135	0.089	0.169	0.300	-0.131
	[-0.232, 0.140]	[-0.202, -0.068]	[-0.058, 0.235]	[-0.281, 0.619]	[0.109, 0.490]	[-0.550, 0.288]
Topic 2 (external threat)	0.001	-0.221	0.222	-0.185	0.842	-1.027
	[-0.130, 0.132]	[-0.292, -0.150]	[0.076, 0.367]	[-0.438, 0.068]	[0.557, 1.126]	[-1.521, -0.532]
Topic 3 (external threat)	-0.009	-0.220	0.211	0.366	0.650	-0.283
	[-0.156, 0.137]	[-0.290, -0.151]	[0.082, 0.340]	[0.105, 0.627]	[0.372, 0.927]	[-0.703, 0.136]
Topic 4 (focused on community)	-0.373	-0.106	-0.267	0.210	0.300	-0.090
	[-0.477, -0.270]	[-0.186, -0.026]	[-0.414, -0.120]	[-0.009, 0.429]	[0.074, 0.525]	[-0.522, 0.343]
Topic 5 (focused on LE)	-0.406	-0.135	-0.270	0.511	0.148	0.363
	[-0.576, -0.235]	[-0.199, -0.072]	[-0.412, -0.129]	[0.115, 0.907]	[0.018, 0.279]	[0.034, 0.692]
Student demographics						
Majority minority students	-0.374	-0.156	-0.218	2.904	0.362	2.542
	[-0.531, -0.216]	[-0.244, -0.068]	[-0.406, -0.030]	[1.493, 4.315]	[0.143, 0.581]	[1.684, 3.400]
Majority white students	0.027	-0.156	0.183	-0.096	0.362	-0.457
	[-0.062, 0.117]	[-0.244, -0.068]	[0.053, 0.314]	[-0.289, 0.098]	[0.143, 0.581]	[-0.781, -0.134]
School staff						
Any counselors	-0.112	-0.185	0.073	0.272	0.597	-0.325
	[-0.311, 0.087]	[-0.245, -0.125]	[-0.053, 0.199]	[-0.201, 0.745]	[0.375, 0.820]	[-0.703, 0.053]

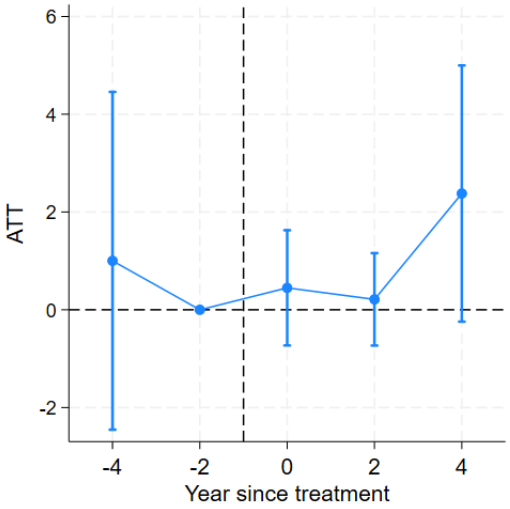
	Chronic absenteeism		
	Moderator ATT	Comparison ATT	Diff in ATT
	(1)	(2)	(3)
Program context			
Topic 1 (internal threat)	-12.596	-0.332	-12.264
	[-15.546, -9.646]	[-1.044, 0.380]	[-15.652, -8.876]
Topic 2 (external threat)	-0.575	-4.750	4.175
	[-1.946, 0.796]	[-5.860, -3.639]	[1.633, 6.716]
Topic 3 (external threat)	0.016	-3.393	3.408
	[-1.868, 1.900]	[-4.238, -2.547]	[1.228, 5.588]
Topic 4 (focused on community)	-10.402	-0.747	-9.655
	[-12.796, -8.008]	[-1.513, 0.018]	[-12.303, -7.006]
Topic 5 (focused on LE)	-2.038	-3.653	1.616
	[-3.909, -0.166]	[-4.636, -2.671]	[-0.539, 3.770]
Student demographics			
Majority minority students	-4.778	-1.824	-2.953
	[-8.493, -1.063]	[-2.821, -0.828]	[-6.528, 0.621]
Majority white students	-1.104	-1.824	0.721
	[-2.372, 0.164]	[-2.821, -0.828]	[-0.671, 2.112]
School staff			
Any counselors	-4.874	-2.488	-2.386
	[-7.251, -2.498]	[-3.348, -1.629]	[-3.837, -0.935]

Notes: Column 1 contains the ATT of schools who satisfied the binary measure of the moderator (moderator ATT) and Column 2 contains the ATT of schools who did not (comparison ATT). ATTs were estimated by dividing the analysis sample into two subgroups based on a binary measure of the moderator and estimating the ATT separately for the moderator and comparison subgroups. Column 3 is the difference between the moderator ATT and the comparison ATT. 95% confidence intervals are constructed using Wild bootstrap. To estimate the subgroup ATTs for chronic absenteeism, which appears only in the 2013-14 the 2015-16 CRDCs, we use the 2014 and 2015 treatment groups. See Section 6.2 for additional details about the moderators of interest.

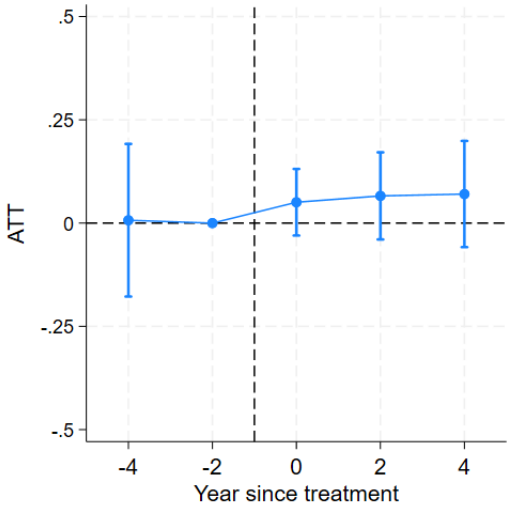
Figures

Figure 1: Dynamic ATT estimates, conditional on pre-treatment covariates

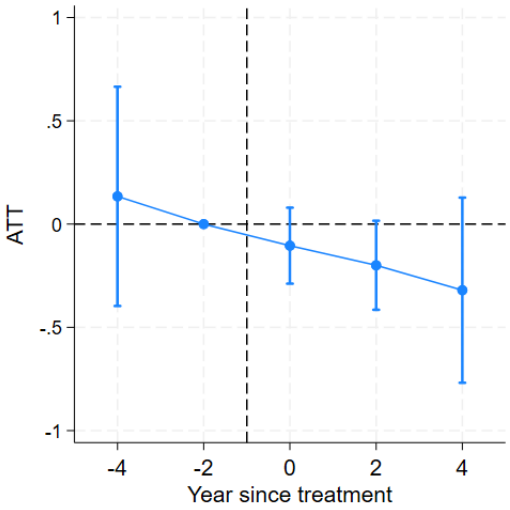
(a) Suspensions per 100 students



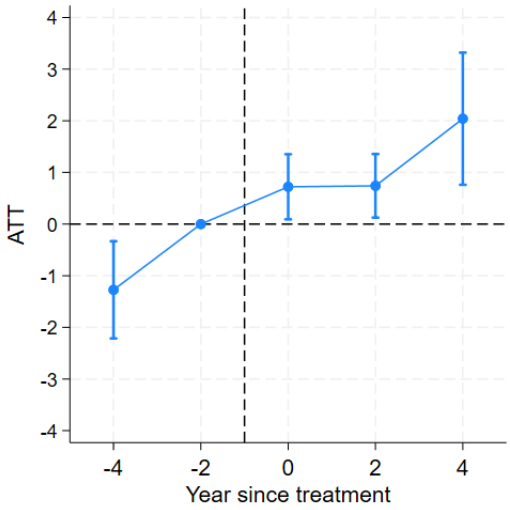
(b) Expulsions per 100 students



(c) Law enforcement action per 100 students



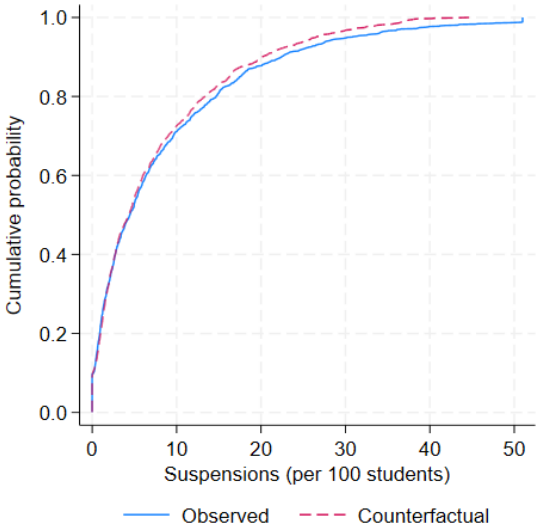
(d) Grade retention per 100 students



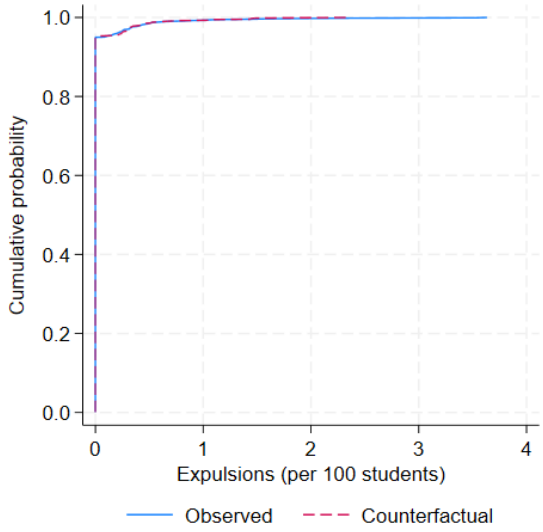
Notes: This figure shows the Average Treatment Effect on the Treated (ATT) estimates of receiving CHP grant funding for SRO hiring on student outcome counts per 100 students. Estimation is implemented separately for each combination of treatment group g and post-treatment year t, then aggregated by the number of years of exposure to the treatment. The dynamic ATT estimates progress biennially, meaning one- and two-year treatment effects are pooled together. The estimates are conditional on pre-treatment school characteristics such as the proportion of enrolled students who are male, black, Hispanic, disabled, have Limited English Proficiency (LEP), receiving free or reduced price lunch; school level (elementary, middle, or high); school locale (urban, suburban, town, or rural); and school region (Northeast, South, Midwest, West). Standard errors are clustered at the school district level and estimated by Wild bootstrap.

Figure 2: Observed and counterfactual cumulative distribution of the 2014 treatment group

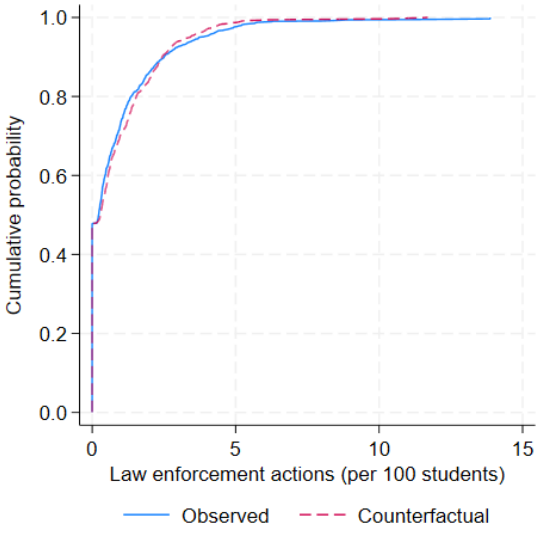
(a) Suspensions per 100 students



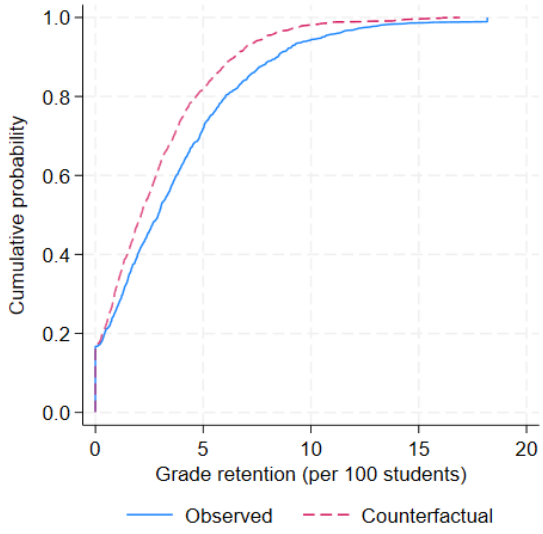
(b) Expulsions per 100 students



(c) Law enforcement action per 100 students



(d) Grade retention per 100 students



Notes: This figure shows the observed post-treatment cumulative distribution of the 2014 treatment group and the counterfactual cumulative distribution estimated using the Changes-in-Changes method (Athey and Imbens, 2006). The outcomes are the average of outcomes across post-treatment time periods, which come from the 2015-16 and 2017-18 CRDC for the 2014 treatment group. The estimates are conditional on pre-treatment school characteristics such as the proportion of enrolled students who are male, black, Hispanic, disabled, have Limited English Proficiency (LEP), receiving free or reduced price lunch; school level (elementary, middle, or high); school locale (urban, suburban, town, or rural); school region (Northeast, South, Midwest, West); and the population size that the law enforcement agency serves (more or less than 150,000).

References

Aizer, Anna and Joseph J. Doyle Jr., "Juvenile Incarceration, Human Capital, and Future Crime: Evidence from Randomly Assigned Judges," *The Quarterly Journal of Economics*, 2015, *130* (2), 759–803.

Anyon, Yolanda, Jeffrey M Jenson, Inna Altschul, Jordan Farrar, Jeanette McQueen, Eldridge Greer, Barbara Downing, and John Simmons, "The persistent effect of race and the promise of alternatives to suspension in school discipline outcomes," *Children and Youth Services Review*, 2014, *44*, 379–386.

Athey, Susan and Guido W. Imbens, "Identification and Inference in Nonlinear Difference-in-Differences Models," *Econometrica*, 2006, *74* (2), 431–497.

Bacher-Hicks, Andrew, Stephen B. Billings, and David J. Deming, "The School-to-Prison Pipeline: Long-Run Impacts of School Suspensions on Adult Crime," *American Economic Journal: Economic Policy*, 2024, *16* (4), 165–193.

Blei, David M. and John D. Lafferty, "A correlated topic model of Science," *The Annals of Applied Statistics*, 2007, *1* (1), 17 – 35.

Blei, David M, Andrew Y Ng, and Michael I Jordan, "Latent dirichlet allocation," *Journal of Machine Learning Research*, 2003, *3* (Jan), 993–1022.

Brown, Ben, "Controlling Crime and Delinquency in the Schools: An Exploratory Study of Student Perceptions of School Security Measures," *Journal of School Violence*, 2005, *4* (4).

Buhl-Wiggers, Julie, Jason T. Kerwin, Juan Muñoz-Morales, Jeffrey Smith, and Rebecca Thornton, "Some children left behind: Variation in the effects of an educational intervention," *Journal of Econometrics*, 2024, *243* (1), 105256.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller, "Bootstrap-based improvements for inference with clustered errors," *The Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Chalfin, Aaron and Justin McCrary, "Criminal Deterrence: A Review of the Literature," *Journal of Economic Literature*, 2017, *55* (1), 5–48.

Counts, Jennifer, Kristina N. Randall, Joseph B. Ryan, and Antonis Katsiyannis, "School Resource Officers in Public Schools: A National Review," *Education and Treatment of Children*, 2018, *41* (4), 405–430.

Cuellar, Matthew J, Susan E Elswick, and Matthew T Theriot, "n Investigation of School Social Worker Perceptions toward School Security Personnel," *School Social Work Journal*, 2017, *41* (2), 41–60.

Curran, F. Chris, Benjamin W. Fisher, Samantha Viano, and Aaron Kupchik, "Why and When Do School Resource Officers Engage in School Discipline? The Role of Context in Shaping Disciplinary Involvement," *American Journal of Education*, 2019, *126* (1), 33–63.

_ , Samantha Viano, Aaron Kupchik, and Benjamin W. Fisher, "Do Interactions With School Resource Officers Predict Students' Likelihood of Being Disciplined and Feelings of Safety? Mixed-Methods Evidence From Two School Districts," *Educational Evaluation and Policy Analysis*, 2021, *43* (2), 200–232.

Davis, Elizabeth J., "School Resource Officers, 2019–2020," Report, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics 2023. NCJ 307334.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, September 2020, *110* (9), 2964–96.

Djebbari, Habiba and Jeffrey Smith, "Heterogeneous impacts in PROGRESA," *Journal of Econometrics*, 2008, *145* (1-2), 64–80.

Fisher, Benjamin W, Ethan M Higgins, Aaron Kupchik, Samantha Viano, F Chris Curran, Suzanne Overstreet, Bryant Plumlee, and Brandon Coffey, "Protecting the Flock or Policing the Sheep? Differences in School Resource Officers' Perceptions of Threats by School Racial Composition," *Social Problems*, 5 2022, *69* (2), 316–334.

Fréchet, M., "Les tableaux de corrélation dont les marges sont données," *Ann. L'Univer. Lyon. Sec. A: Sci., Math. Et Astron.*, 1951, *14*, 53–77.

Ghavami, Negin, Bryan E Thornton, and Sandra Graham, "School police officers' roles: The influence of social, developmental and historical contexts," *Journal of Criminal Justice*, 2021, *72*, 101724.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Gottfredson, Denise C., Scott Crosse, Zhiqun Tang, Harmon Bauer Erin L., Carol A. Michele A. Hagen, and Angela D. Greene, "Effects of school resource officers on school crime and responses to school crime," *Criminology Public Policy*, 2020, *19*, 905–940.

Gottlieb, Aaron and Robert Wilson, "The effect of direct and vicarious police contact on the educational achivement of urban teens," *Children and Youth Services Review*, 2019, *103*, 190–199.

Heckman, James, Jeffrey Smith, and Nancy Clements, "Making The Most Out Of Programme Evaluations and Social Experiments: Accounting For Heterogeneity in Programme Impacts," *The Review of Economic Studies*, 1997, *64* (4), 487–535.

Heise, Michael and Jason P Nance, "To Report or Not to Report: Data on School Law Enforcement, Student Discipline, Race, and the "School-to-Prison Pipeline"," *UC Davis Law Review*, 2021, *55*, 209.

Hirschfield, Paul J, "Preparing for prison? The criminalization of school discipline in the USA," *Theoretical Criminology*, 2008, *12* (1), 79–101.

Hjalmarsson, Randi, "Criminal justice involvement and high school completion," *Journal of Urban Economics*, 2008, *63* (2), 613–630.

Höffding, W., "Masstabinvariante korrelationsmasse für diskontinuierliche verteilungen," *Arkiv Mat. Wirtsch. Sozialforschung*, 1970, *7*, 49–70.

James, Nathan and Gail McCallion, "School Resource Officers: Law Enforcement Officers in Schools," *Congressional Research Service*, 2013.

Kirk, David S. and Robert J. Sampson, "Juvenile Arrest and Collateral Educational Damage in the Transition to Adulthood," *Sociology of Education*, 2013, *86* (1), 36–62.

Lynch, Caitlin G., Randy R. Gainey, and Allison T. Chappell, "The effects of social and educational disadvantage on the roles and functions of school resource officers," *Policing: An International Journal*, 2016, *39* (3), 521–535.

McKenna, Joseph M, Kathy Martinez-Prather, and Scott W Bowman, "The roles of school-based law enforcement officers and how these roles are established: A qualitative study," *Criminal Justice Policy Review*, 2016, *27* (4), 420–443.

Melly, Blaise and Giulia Santangelo, "The changes-in-changes model with covariates," *Unpublished manuscript*, 2015.

Monahan, Kathryn C., Susan VanDerhei, Jordan Bechtold, and Elizabeth Cauffman, "From the school yard to the squad car: school discipline, truancy, and arrest," *Journal of Youth and Adolescence*, 2014, *43* (7), 1110–1122.

Mowen, Thomas and John Brent, "School Discipline as a Turning Point: The Cumulative Effect of Suspension on Arrest," *Journal of Research in Crime and Delinquency*, 2016, *53* (5), 628–653.

NASRO, "NASRO Position Statement on Police Involvement in Student Discipline," 2015.

Owens, Emily, "COPS and Cuffs," in Philip Cook, Stephen Machin, Olivier Marie, and Giovanni Mastrobouni, eds., *Lessons from the Economics of Crime: What Works in Reducing Offending?*, Cambridge: MIT Press, 2013, pp. 17–44.

Owens, Emily G., "Testing the School-to-Prison Pipeline," *Journal of Policy Analysis and Management*, 2017, *36* (1), 11–37.

Roberts, Margaret E, Brandon M Stewart, and Dustin Tingley, "Stm: An R package for structural topic models," *Journal of Statistical Software*, 2019, *91*, 1–40.

Ryan, Joseph B., Antonis Katsiyannis, Jennifer M. Counts, and Jill C. Shelnut, "The Growing Concerns Regarding School Resource Officers," *Intervention in School and Clinic*, 2018, *53* (3), 188–192.

Sorensen, Lucy C, Montserrat Avila-Acosta, John B Engberg, and Shawn D Bushway, "The thin blue line in schools: New evidence on school-based policing across the US," *Journal of Policy Analysis and Management*, 2023, *42* (4), 941–970.

_ , Yinzhi Shen, and Shawn D Bushway, "Making schools safer and/or escalating disciplinary response: A study of police officers in North Carolina schools," *Educational Evaluation and Policy Analysis*, 2021, *43* (3), 495–519.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Swartz, Kristin, Dustin L. Osborne, Cherie Dawson-Edwards, and George E. Higgins, "Policing Schools: Examining the Impact of Place Management Activities on School Violence," *American Journal of Criminal Justice*, 2016, *41* (3), 465–483.

Theriot, Matthew T, "School resource officers and the criminalization of student behavior," *Journal of criminal justice*, 2009, *37* (3), 280–287.

Theriot, Matthew T. and John G. Orme, "School Resource Officers and Students' Feelings of Safety at School," *Youth Violence and Juvenile Justice*, 2015, *14* (2), 130–146.

Thurau, Lisa H and Johanna Wald, "Controlling partners: When law enforcement meets discipline in public schools," *New York Law School Review*, 2010, *54*, 977.

Wang, Ke, Yongqiu Chen, Jizhi Zhang, and Barbara A Oudekerk, "Indicators of School Crime and Safety: 2019," *National Center for Education Statistics, U.S. Department of Education, and Bureau of Justice Statistics, Office of Justice Programs, U.S. Department of Justice*, 2020.

Weisburst, Emily K., "Patrolling Public Schools: The Impact of Funding for School Police on Student Discipline and Long-term Education Outcomes," *Journal of Policy Analysis and Management*, 2019, *38* (2), 338–365.

Wood, Brandon J. and Eric Hampton, "The Influence of School Resource Officer Presence on Teacher Perceptions of School Safety and Security," *School Psychology Review*, 2021, *50* (2-3).

Wu, Chien-Fu Jeff, "Jackknife, bootstrap and other resampling methods in regression analysis," *The Annals of Statistics*, 1986, *14* (4), 1261–1295.

Zhang, Gary, "The Effects of a School Policing Program on Crime, Discipline, and Disorder: A Quasi-Experimental Evaluation," *American Journal of Criminal Justice*, 2019, *44* (1), 45–62.

Appendix A Tables and Figures

Table A1: Summary statistics of raw outcomes by treatment status

	Never applied	Applied but not awarded	Awarded	Difference	P-value
Outcomes per 100 students (raw)					
Suspensions	11.92	12.14	16.69	4.54	0.00
	(20.02)	(16.83)	(23.06)		
Expulsions	0.22	0.21	0.09	-0.12	0.00
	(1.30)	(1.42)	(0.46)		
Formal law enforcement actions	0.55	0.48	0.92	-0.44	0.00
	(1.65)	(1.37)	(1.87)		
Chronic absenteeism	12.54	13.18	15.36	2.18	0.00
	(12.53)	(11.99)	(11.82)		
Grade retention	2.33	3.03	3.31	0.27	0.14
	(4.75)	(8.89)	(6.79)		
Gun offenses	0.02	0.02	0.02	-0.00	0.65
	(0.31)	(0.12)	(0.12)		
Other offenses	2.18	1.76	3.52	1.76	0.00
	(15.10)	(4.98)	(26.46)		
N	52,249	8,944	2,641		

Notes: Column 1, 2, and 3 contain the unweighted variable mean and standard deviation for schools that never applied for the CHP grant between 2013 and 2016, schools that applied but did not receive the CHP award and the schools that applied for and received the CHP award, respectively. Column 4 is the difference between the variable means of schools that applied for the grant but did not receive it and schools that received the grant and Column 5 presents the p-value from the t-test. The variable for chronic absenteeism was measured from the 2013-14 CRDC and gun offenses and other offenseswere measured from the 2015-16 CRDC as these questions were not asked in earlier waves; all other variables are measured from the 2011-12 CRDC. All outcomes are raw and are normalized per 100 students.

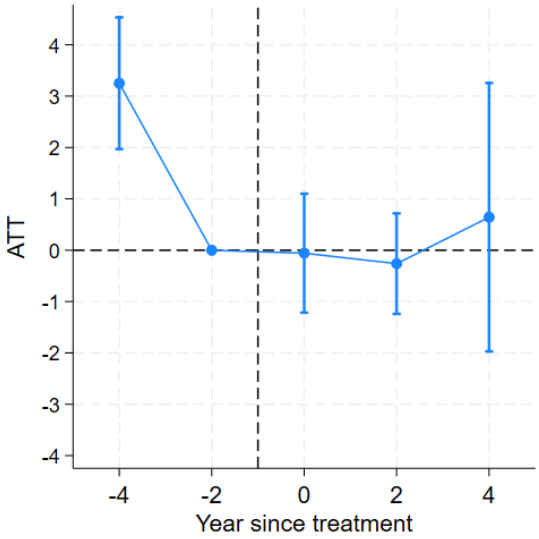
Table A2: Post-treatment ATT estimates by treatment group

	Suspensions	Expulsions	Law enforcement actions	Chronic absenteeism	Grade retention	Gun offenses	Other offenses
2013	0.904	0.071	-0.256	-	1.761	-	-
	[-2.152, 3.960]	[-0.029, 0.172]	[-0.560, 0.049]	-	[-2.399, 5.922]	-	-
2014	-0.369	0.029	-0.124	-3.531	-0.049	-	-
	[-1.609, 0.870]	[-0.058, 0.117]	[-0.361, 0.113]	[-9.069, 2.006]	[-0.949, 0.851]	-	-
2015	1.107	0.128	0.069	-1.466	0.128	-	-
	[-1.134, 3.348]	[-0.071, 0.327]	[-0.121, 0.259]	[-6.767, 3.835]	[-0.306, 0.561]	-	-
2016	0.467	0.248	0.090	-	-0.823	-0.005	-1.511
	[-1.136, 2.070]	[0.066, 0.429]	[-0.046, 0.226]	-	[-1.973, 0.327]	[-0.030, 0.020]	[-2.242, -0.781]

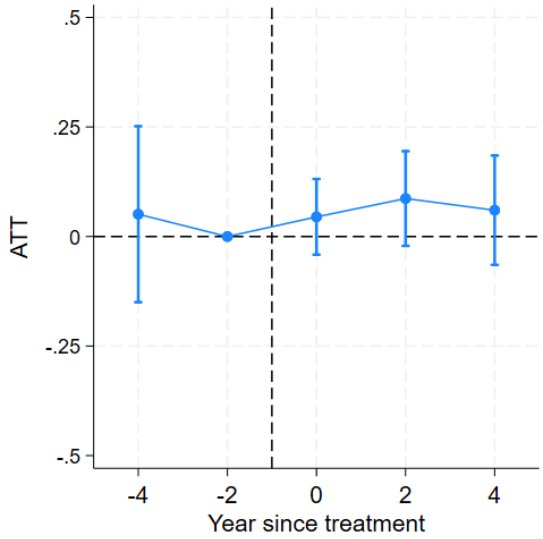
Notes: This table presents the aggregate post-treatment ATT estimates for each student disciplinary outcome for all treated schools, aggregated by treatment group. Note that for chronic absenteeism, data is only available for the 2013-14 and 2015-16 school years and so only schools treated in 2014 and 2015 are compared against the never treated group. For gun and other offenses, data is only available for the 2015-16 and 2017-18 school years and so, only schools treated in 2016 are compared against the never treated group. All outcomes are defined as counts per 100 students and are winsorized at the 99th percentile. The 95% confidence intervals are reported in brackets. The standard errors are estimated using Wild bootstrap and clustered at the school district level.

Figure A1: Dynamic ATT estimates, unconditional on pre-treatment covariates

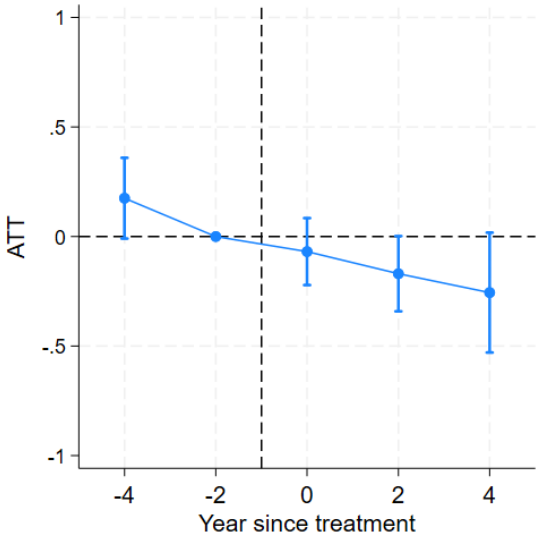
(a) Suspensions per 100 students



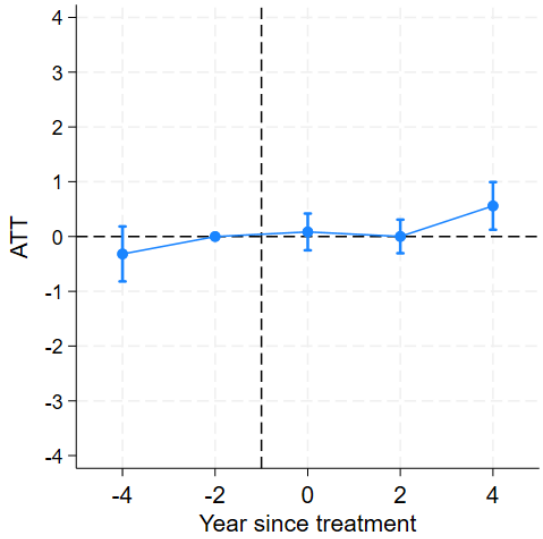
(b) Expulsions per 100 students



(c) Law enforcement action per 100 students



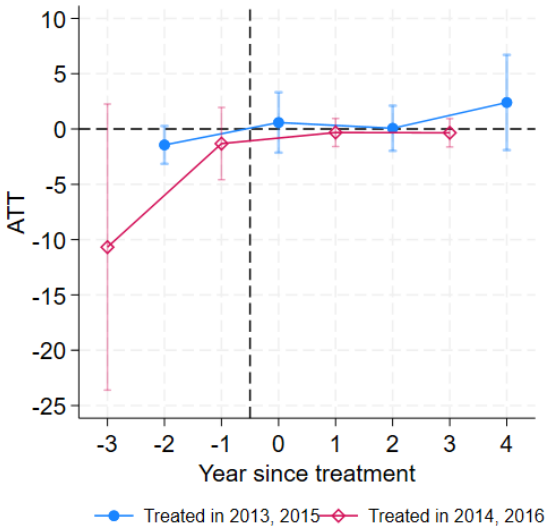
(d) Grade retention per 100 students



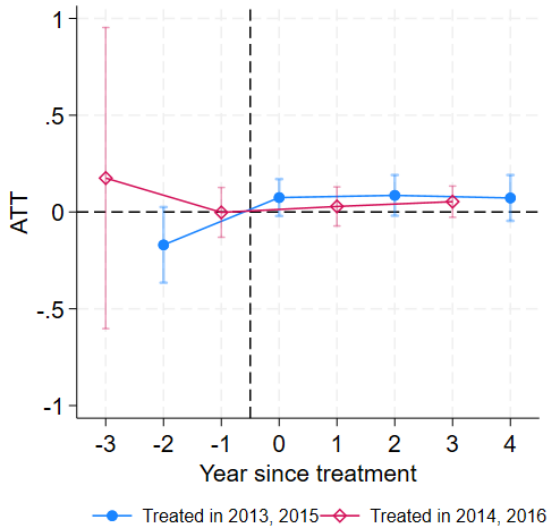
Notes: This figure shows the Average Treatment Effect on the Treated (ATT) estimates of receiving CHP grant funding for SRO hiring on student outcome counts per 100 students. Estimation is implemented separately for each 2×2 combination of treatment group g and year t, then aggregated by the number of years of exposure to the treatment. The dynamic ATT estimates progress biennially, meaning one- and two-year treatment effects are pooled together. The estimates are compared against the last year before treatment. The estimates are unconditional on pre-treatment school characteristics. Standard errors are clustered at the school district level and estimated by Wild bootstrap.

Figure A2: Dynamic ATT estimates, unpooled and conditional on pre-treatment covariates

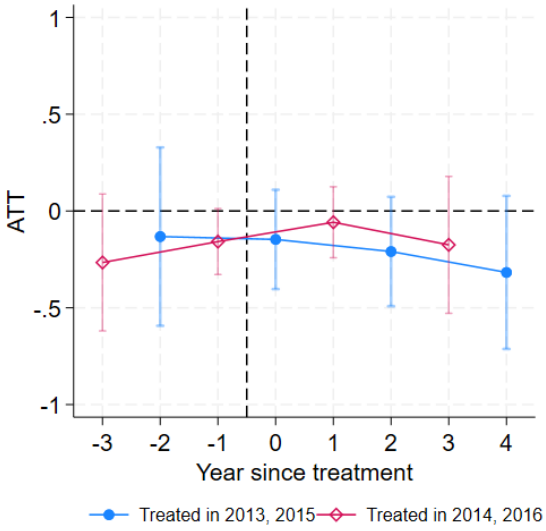
(a) Suspensions per 100 students



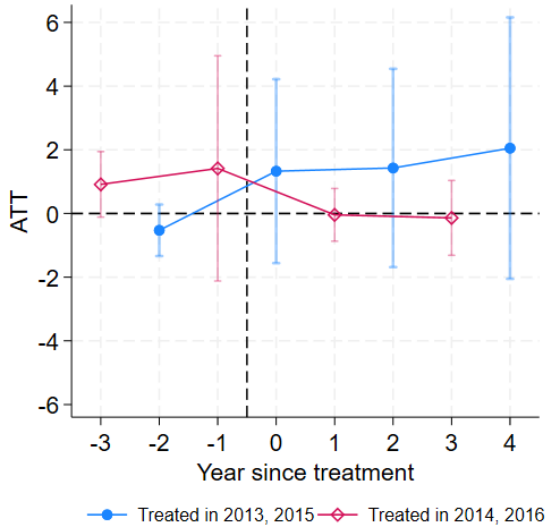
(b) Expulsions per 100 students



(c) Law enforcement action per 100 students

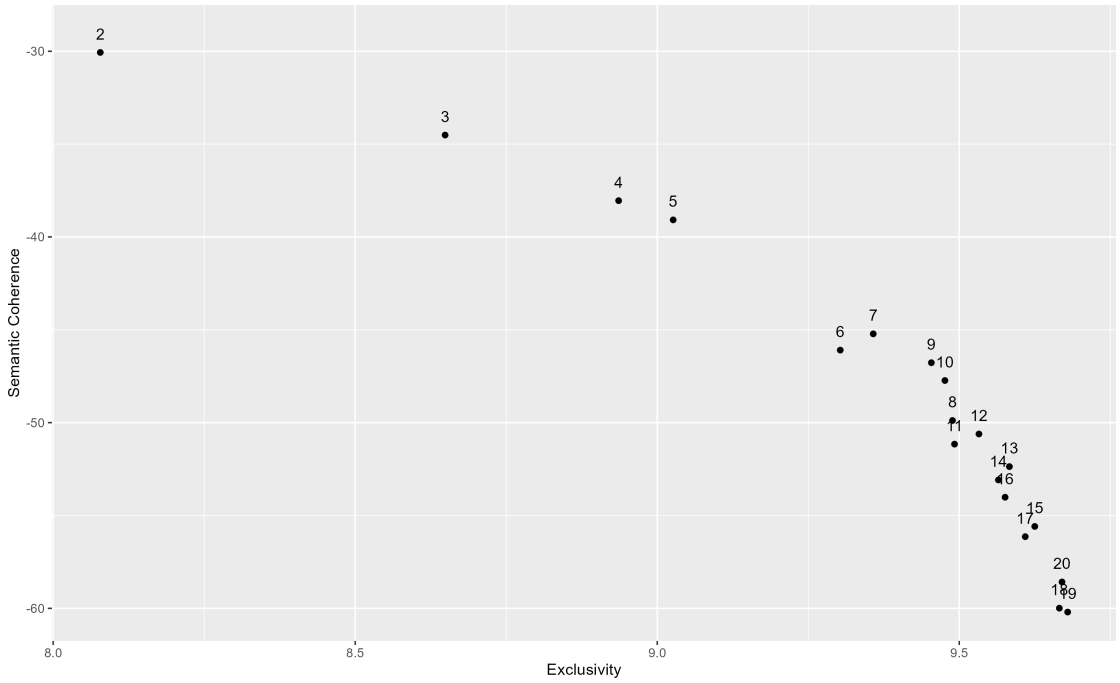


(d) Grade retention per 100 students



Notes: This figure shows the Average Treatment Effect on the Treated (ATT) estimates of receiving CHP grant funding for SRO hiring on student outcome counts per 100 students. Estimation is implemented separately for each 2×2 combination of treatment group g and year t, then aggregated by the number of years of exposure to the treatment. The dynamic ATT estimates progress yearly, and treatment groups from even treatment years are presented separately from treatment groups from odd treatment years. The estimates are compared against the nearest pre-treatment year. The estimates are conditional on pre-treatment school characteristics such as the proportion of enrolled students who are male, black, Hispanic, disabled, have Limited English Proficiency (LEP), receiving free or reduced price lunch; school level (elementary, middle, or high); school locale (urban, suburban, town, or rural); school region (Northeast, South, Midwest, West); and the population size that the law enforcement agency serves (more or less than 150,000). Standard errors are clustered at the school district level and estimated by Wild bootstrap.

Figure A3: Diagnostic Metrics from CTM with Various Topics (K)



Notes: This figure shows the trade off between topic semantic coherence and exclusivity as the number of topics K ranges from 2 to 20. Topic semantic coherence is a measure of whether topics contain words that frequently appear together, and exclusivity is measure of how distinct the words under a topic are from words under other topics.

Appendix B Data Appendix

B.1 Sample

To construct our analysis sample, we start with law enforcement agencies that applied for the CHP grant for school-based policing between 2013 and 2016. We exclude agencies located in U.S. territories, tribal agencies, agencies that only serve a post-secondary institution, and state agencies. We also exclude agencies whose applications remained pending at the time we received the data, withdrew their application, or declined the CHP grant award. As there are law enforcement agencies that applied for the CHP grant more than once within our study period (441 agencies, or 27.6% of our unique sample), this results in 1,596 grant applications, or 1,155 unique law enforcement agencies, that are included in our sample. To match these law enforcement agencies to schools, we used the grant narrative data (the process is described in Appendix Section B.3).

After linking law enforcement agencies to schools, we further narrow the sample at the school level. First, we exclude alternative schools, which are schools designed to fulfill the needs of students who are experiencing academic difficulties or discipline problems. Based on the population of students that alternative schools serve, they may have more punitive discipline policies or higher rates of student misbehavior than traditional public schools. We exclude 1,334 alternative schools. Second, we exclude 103 schools, which are juvenile justice programs, pre-K schools, special education schools, and transitional schools. Lastly, we exclude 283 schools in Hawaii from our sample due to suspected data entry errors. Specifically, the number of referrals reported by schools in Hawaii in the 2013-14 CRDC was extremely large compared to the number of referrals reported in the other CRDC survey years. One extreme example is a high school in Hawaii that reported no referrals to the police in the 2011-12 CRDC, 509 in the 2013-14 CRDC, 22 in the 2015-16 CRDC, and 13 in the 2017-18 CRDC, but many schools reported referrals in the hundreds in the 2013-14 CRDC compared to tens in all other CRDC survey years.

B.2 Transforming outcomes

We topcode the outcomes in our study at the 99th percentile. We do this because each outcome distribution contains a large number of zeros, ranging from 11.9% of the distribution for suspensions to 90.3% of the distribution for school-based arrests, and an extended right tail, where the minimum non-zero value ranges from 2 for expulsions and 12 for suspensions and the maximum non-zero value ranges from 483 for referrals and 2,773 for suspensions. This allows us to smooth the more extreme non-zero values without removing too much of the variation in non-zero values, which would occur if we topcoded at a lower percentile of the distribution. We also topcode the outcomes that are disaggregated by student subgroups using the value at the 99th percentile of the distribution of non-zero values of the aggregate outcome.

B.3 Matching law enforcement agencies to schools

To ensure that the treatment effect is estimated at the school level, we use the CHP grant narrative data to identify the specific schools that would receive the SRO placement should the law enforcement agency be awarded the grant. The grant narrative data are text responses to the following questions on the application: first, describe the problem the agency will address with the grant funds and the agency's approach to the problem; second, explain the agency's inability to address the need without federal assistance. While these prompts do not directly request information about the schools that will be the beneficiaries of the grant, in the process of responding to the questions, the majority of grant applicants either directly name the school that will receive the SRO placement or indirectly allude to the school (e.g. stating the SRO will be placed in elementary schools within the district).

We begin by identifying the school district that the law enforcement agency has partnered with. For the majority of applications, the agency directly names the local school district on the list of partnering institutions. In cases in which no school district is listed as a partnering institution, many of the grant narratives include the name of the district. Otherwise, we use the city name of the law enforcement agency to identify the name of the local municipality (city,

town, etc.), which is then used to identify the school districts located within or adjacent to that municipality. This results in a fuzzier match between agencies and school districts than using the grant narrative, but this is the case for a limited amount of grant applications.

To identify the individual schools, we read through the grant narratives for each application and use a coding scheme to keep track of the following information: name of the school(s) that will receive an SRO placement should the agency receive the grant; name of the school district that currently has SROs provided by the agency; name of the school(s) that currently have SROs provided by the agency; and the number of SROs currently working within the school district. We also track whether the agency provided SROs to the schools or school district in the past although not at the time of application; whether the grant narrative data directly states the name of the school that will receive an SRO placement; and whether the grant narrative data does not contain any school information. We also track information that we later use for sample exclusion such as whether the agency is a tribal agency, a state agency, an agency located in U.S. territories, or an agency serving a postsecondary institution.

In cases in which the grant narrative does not directly state the name of the school, they still usually contain clues that can be used to identify the specific school. For example, some narratives note that the SRO(s) will be placed in only the elementary schools in the district. Then, we refer to the Common Core of Data from the 2013-14 school year to find the names of the elementary schools in the given district. Another example is if the grant narrative states that the grant will be used to place SRO(s) in elementary schools that currently do not have one. If the narrative does not explicitly state the name of the schools that currently have an SRO, we attempted to identify these schools by searching through local news sites or by looking through Facebook pages of the given school district or law enforcement agency. If we were unable to definitively identify the specific schools or if the grant narrative did not contain any information about the school,

Appendix C Grant Narrative Text Analysis

In this paper, we utilize text modeling algorithms to discover latent topics in the CHP grant application narrative texts. In particular, we train a Correlated Topic Model (CTM) on applicants' responses to the prompt, "please describe the selected problem/focus area that you wish to address with COPS Office funding in 4,000 characters or less."

Correlated Topic Models (CTM) are defined by several key features. First, they are mixed membership models, which mean they assume a given text can embody multiple latent themes (Blei and Lafferty, 2007). Second, building on the latent Dirichlet allocation (LDA) model (Blei et al., 2003), CTMs consider each set of text as a collection of words ("bag of words"), and makes inferences about the correlation between words and topics based on the frequency at which words appear in the same text. Notably, this means CTMs do not account for the sequencing of words. Finally, unlike their predecessor (LDA), CTMs capture more complexity in text data by allowing the latent topics to be correlated.

Below, we define key components in the CTM and provide a brief summary of the model estimation procedure, which utilizes variational expectation-maximization (VEM). In this paper, we use the R package "stm" for the model estimation (Roberts et al., 2019). Following the convention in the prior literature, we call an application a "document" indexed by d, and there are K topics in the model (Blei and Lafferty, 2007; Roberts et al., 2019). The model parameters are defined as follows: once the narrative data is prepared, we then estimate the CTM with the following steps (Blei and Lafferty, 2007; Roberts et al., 2019).

1. Draw $\overrightarrow{\theta_d}$, a vector of topic prevalence in a document (i.e. application) from a log-normal distribution. Topic prevalence is the proportion of words in a given document that are classified under a particular topic. In practice, we draw $\eta_d \sim \mathcal{N}(\mu, \Sigma)$, and then let

$$\theta = \frac{\exp(\eta)}{\sum_i \exp(\eta_i)}$$

where μ is a K-length vector and Σ is a K by K covariance matrix (Blei and Lafferty, 2007).

2. Define $\beta_{k,w}$, a topic-specific distribution of word w, such that

$$\beta_{k,w} = \frac{\exp(m_w + \kappa_{k,w})}{\sum_w \exp(m_w + \kappa_{k,w})}$$

where m is the overall word distribution and κ_k captures how the distribution of word w deviates for topic k. So, $\beta_{k,w}$ tells us how likely we are to observe each word in a document that is classified under topic k.

3. For each word (indexed by $n \in 1, ..., N_d$) in document d:

   - Draw the topic of word n in document d, $z_{d,n} \sim$ Multinomial$(\overrightarrow{\theta_d})$.
   - Conditional on the topic $z_{d,n} = k$, we draw a word that is classified under the topic k, $w_{d,n} \sim$ Multinomial$(\beta_{k,w})$.

We use the variational expectation-maximization algorithm to iteratively estimate the model parameters until convergence. Since inference on the actual posterior distribution is intractable, Blei and Lafferty (2007) propose using the variational Gaussian distribution. The E-step involves updating variational posterior probabilities of model parameters λ_d and ν_d by minimizing the Kullback-Leibler divergence between the variational and true posterior distributions. The M-step involves maximizing the approximate evidence lower bound (ELBO) with respect to μ and Σ. To establish the initial points of the model parameters for estimation, we use the spectral decomposition method described in **?**.

A key parameter that is set by the researcher is the number of topics, K. To set the number of topics, we use a data-driven approach by testing models with a range of different topic numbers (K between 2 - 20). For each number of topics K, we randomly split the model in half, train the model on half of the data, and then compute diagnostic measures on the held-out half of the data. The results are summarized in Figure A3. We see that exclusivity (the likelihood that words are exclusively associated with a topic) goes up as we increase the number of topics, but semantic coherence (the likelihood that words classified in the same topic show up in the same application) goes down. We choose $K = 5$, as that choice balances the trade-offs between having many differentiated topics and topics that are coherent.

It is important to note that we take several steps to prepare the narrative text data for analysis, following previous literature in natural language processing and text modeling. The main goal of the data preparation is dimension reduction: we separate the sentences into words, remove punctuation and white-space, stem (only keeping root forms of words), remove common vocabulary that are likely not informative (e.g. "the", "a", "and"), and remove rare vocabulary. The final step is done in a data-driven manner; we plot the number of words and applications that would be removed depending on our threshold for inclusion (i.e. across how many applications must a given vocabulary word appear to be included). The key trade-off we make when setting the threshold is between reducing noise and keeping sufficient variation in words between applications. We decide on a threshold of 90 applications (i.e. a word must have appeared in at least 90 applications to be kept in the data set), which removes roughly 40% of the unique words and leaves us with 105,137 unique words across all applications.