The Effects of Chicago's Teacher Strike on Crime*

Abigail R. Banan† Mary Kate Batistich‡ Jillian B. Carr§ Clint Harris¶
Kendall J. Kennedy‖

July 27, 2022

Abstract

On September 10, 2012, Chicago's teachers began a strike that forced over 400,000 students out of school for 8 consecutive school days. This study examines how this interruption in schooling affected crime in Chicago, both during and immediately after the strike. Using a variety of methods – interrupted time series analysis, synthetic controls, and place-based intensity of treatment analysis – we find that crime rates in Chicago neither significantly nor substantially changed after the strike, suggesting incapacitation has little role in the effect of education on crime.

JEL Codes: I24, J08, K42

Keywords: Education and Crime, Incapacitation, Synthetic Controls

*We thank Gelsey Guerra and Sugandh Khobragade for invaluable research assistance on this project. Researchers' own analyses calculated based in part on NIBRS data and Chicago crime data.

†Department of Economics, Purdue University. 403 W. State St, West Lafayette, IN 47907. Email: abanan@purdue.edu.

‡University of Notre Dame, 3060 Jenkins-Nanovic Hall, Notre Dame, IN 46556, USA; email:mbatisti@nd.edu

§Department of Economics, Purdue University. 403 W. State St, West Lafayette, IN 47907. Email: carr56@purdue.edu.

¶Wisconsin School of Business and Wisconsin Institute for Discovery, University of Wisconsin-Madison, 330 North Orchard Street, Madison, WI 53715, USA; email: clint.harris@wisc.edu

‖Department of Finance and Economics, Mississippi State University, PO Box 9580, Mississippi State, MS 39762; email: kkennedy@business.msstate.edu

1 Introduction

The relationship between education and crime is widely studied and multifaceted. Becker (1968) was one of the first to link the two topics, suggesting that increasing human capital, and thereby wages, reduces incentives to participate in illegal activities. However, the effects of education on crime go beyond long-term human capital acquisition. Schooling itself causes children to gather together at schools on weekdays, simultaneously incapacitating schoolchildren and changing their access to different types of crimes. A pure test of the incapacitation and access effects of education on crime is difficult however, because education increases human capital and most changes to the typical school schedule, such as summer break or modification to the Monday to Friday school week, are anticipated by the community far in advance.

In this study, we consider a disruption to the school schedule that was large, lengthy, and occurred with little warning to isolate the incapacitation effect of education on youth criminal activity and victimization – the Chicago teacher strike of 2012. The strike, in response to failed negotiation between the Chicago Teachers' Union and Chicago Public Schools, resulted in the closing of every public school in Chicago from September 10–19, 2012. Importantly, the timing of the strike was only announced 10 days in advance (due to an Illinois law requiring 10 days' notice before a teacher strike), and was not fully confirmed to occur until negotiations on September 8 and 9 broke down. This event effectively shuttered schools for over 400,000 students in Chicago Public Schools, providing an ideal natural experiment to test the incapacitation effect of education on crime without contamination from anticipation effects or large human capital changes.

We overwhelmingly find that the 2012 Chicago teacher strike had little to no effect on crime, both during the strike and in the weeks following the strike. Using reported crime data from the Chicago Data Portal 2010–14, we begin by estimating interrupted time series models for a descriptive exploration of changes in crime rates around the teacher strike. Testing all crime, domestic crime, personal crime, property crime, and violent crime, we find no statistically significant discontinuous change in the post-strike crime rate, and only a small, marginally significant decrease in the overall trend in the property crime rate following the teacher strike. Even when separately testing all 25 types of crimes for which there were at least 70 occurrences during our sample period, we only find extremely small significant discontinuous changes in the crime rate (< 0.35 weekly crimes per 100,000) for exceedingly rare types of crime (< 2 weekly crimes per 100,000).

We then proceed to estimation using various synthetic controls methods. We construct a synthetic Chicago using information on crimes in the National Incident Based Reporting System (NIBRS) dataset for other cities, as well as 2010-14 ACS data on demographics and National Centers for Environmental Information data on weather. Though our constructed synthetic Chicago does not perfectly match the "true" Chicago in the pre-strike period, we find at most a decrease in crimes of 37 per 100,000 three weeks after the beginning of the teacher strike (or 1.5 weeks after the conclusion of the strike), corresponding to a 14.4% reduction in overall crime compared to the pre-strike mean.[1] We consistently find null effects on crime in the first week following the beginning of the strike, and find null effects in the second week in four of our six model specifications. Overall, we take this as additional evidence that reducing incapacitation due to the teacher strike had little to no effect on crime in Chicago, with the only statistically significant estimates occurring after schools reopened.

Finally, we consider how different geographical areas may be exposed to a teacher strike differently, depending on the child population present in those areas. At the Census block group level, we interact an indicator for the beginning of the strike with the share of the population in the Census block group that is school-aged (ages 5–19). In our main specifications controlling for indicators for year, week, and day of week, we again find no effect of the strike on crime rates at the Census block group level, and no differential effect depending on the youth population of those block groups. Taken together, we find that the 2012 Chicago teacher strike generally had no effect on crime, and that any effects that did occur were small and negative.

Our study conflicts to some extent with the previous literature on incapacitation effects of education on crime. Several prior studies have used exogenous variation in school attendance to test its link to crime involvement among youth. Akee, Halliday, and Kwak (2014) use teacher furlough days in Hawaii to test the effect of cancelled classes on juvenile crime. They find a decrease in juvenile assault and drug-related arrests, but no change to other crimes such as burglaries. Jacob and Lefgren (2003) use teacher in-service days, Luallen (2006) uses unexpected days off school caused by teacher strikes, and Fischer and Argyle (2018) use the adoption of the four-day school week across schools in Colorado. All three studies link being in

[1] This almost certainly overstates the effect of the Chicago teacher strike on crime. Figure 1 shows that our synthetic Chicago matches the actual data in Chicago fairly well until about 1-6 months before the teacher strike, depending on the method. After that point, synthetic Chicago has substantially higher projected crime rates than actual Chicago, and this pattern does not appear to change after the teacher strike. As we add more and better data and refine this synthetic control method, we fully expect this estimated effect to tend toward zero.

school to a reduction in property crime. Jacob and Lefgren (2003) and Luallen (2006) find an increase in violent crime, while Fischer and Argyle (2018) find no change to violent crime. The teacher strike we examine differs substantially from the contexts studied by Jacob and Lefgren (2003); Akee, Halliday, and Kwak (2014); Fischer and Argyle (2018), as their school closures were frequent and were easily anticipated well in advance. Compared to Luallen (2006), our methods and outcomes differ substantially, even though the core question remains the same. Luallen (2006) focuses only on juvenile crime, and studies multiple fairly short and small-scale strikes occurring in Washington state between 1980–2001 using a negative binomial model with time and location fixed effects. Our analysis implements modern synthetic controls methods to study a relatively long strike in the third largest school district in the US – our context is substantially different from that of Luallen (2006), and should not be seen as refuting or even directly conflicting with his findings.

Other papers use variation in minimum school dropout age laws to test for effects of staying in school on crime. Anderson (2014) using state-level variation in dropout age laws finds that staying in school reduces juvenile arrest rates for violent and property crimes. Anderson, Hansen, and Walker (2013) present evidence that increasing the school dropout age displaces crime from the streets to schools.

Distinct from incapacitation effects, Modestino (2019) tests for the effects of a summer youth employment program on criminal justice system involvement 17 months after participation, and points to improved social skills as a mechanism. Davis and Heller (2020) in contrast find reduced violent crime arrests but possibly increased property crime arrests for youth randomized into supported summer jobs in Chicago. They describe a more complicated picture with heterogeneous effects by different types of youth.

2 Data

We restrict our attention to crimes that are severe enough to warrant police involvement. Because we intend to compare crime before and after the strike (which lasted ten days), we also require crime data aggregated to, at most, the weekly level. We also intend to compare locations within Chicago over time to identify types of neighborhoods that are especially affected by the teacher strike. To these ends, we make use of incident-level data for the city of Chicago.

The specific dataset we primarily use is the incident-level administrative crime data from the Chicago Data Portal, from January 1, 2010 to January 1, 2014, which encompasses the teacher strike in September 2012. These data contain time stamps for crimes at the minute level, making them sufficient for time series analyses for the city of Chicago as a whole. They also include latitude and longitude for the reported crime. Finally, crimes are coded by type both in terms of the offense committed and in the location where they were committed. We include all reported crimes, regardless of whether an arrest was made.

Data on crime at the minute level is quite noisy. In our preferred specifications, we aggregate crime to the daily or weekly level by summing over crime by type. Not only does this substantially smooth temporal trends in the data, which is helpful for validating the identifying assumptions of our empirical strategies, it also permits comparison to other cities in the National Incident Based Reporting System (NIBRS) dataset, which aggregates crimes at the daily level.

The NIBRS dataset is our primary resource for comparing Chicago's crime trends, before and after the teacher strike, to those of other cities. Our primary inclusion criteria for cities in the NIBRS data is that they have populations above 100,000, in the interest of omitting small cities that are likely less comparable to Chicago. The NIBRS dataset includes a selected sample of cities throughout the US, with detailed incident-level data that is comparable to our data for Chicago, with the exception that the data is aggregated at the day, rather than minute, level.[2] We also merge in data from the 2010-2014 5-year American Community Survey on unemployment, racial demographics, gender demographics, poverty, and per capita income for the metropolitan statistical area that encompasses each city, including Chicago.

In addition to crime data for all cities, we include daily information on temperature and precipitation from airports in each city obtained from the National Centers for Environmental Information's climate data archive. Weather explains a large amount of daily and seasonal variation in crime, making these controls valuable for trend comparisons between Chicago and other cities. We drop one city from the NIBRS that has poor weather data coverage for 2010-2012, and another with an abnormal (discontinuous at zero) daily reported crime distribution, leaving us with 25 comparison cities, which are listed in Table A.2.

In addition to detailed time stamps, the Chicago data contain longitude and latitude (to nine

[2]Uniform Crime Reports data, a natural alternative to the NIBRS data, contains a larger sample of geographic locations, but aggregates crime to the monthly level. As the teacher strike only lasted 10 days, comparisons of daily/weekly Chicago crime to monthly crime data for other cities require judgments about the dispersion of crime within each month that we prefer to avoid making.

decimal places) for reported crimes. This allows us to match crime incidents to census blocks within-Chicago neighborhood measures using Census Shape files based on 2010 boundaries. We connect the geographic coordinates of the crime data to census blocks through the geospatial processing program ArcMap (part of Esri's ArcGIS software). We aggregate our data to Census block groups which typically have between 600 and 3,000 people. We link the crime data to demographic and socioeconomic information at the census block group level from the 2010-2014 5 year ACS data obtained from the National Historical Geographic Information System (NHGIS; Manson, Schroeder, Van Riper, Kugler, and Ruggles, 2021). Specifically, we include information on population age shares, median household income, college graduate population share, and racial population shares. These variables are used both to investigate heterogeneous effects, and to establish control groups (census block groups with few or no school-age children) for the intensity of treatment specification described in Section 3.3. We define school-age share as the share of population aged 5-19. We also define four racial and ethnic categories: Hispanic, non-Hispanic white, non-Hispanic black, and other.

An additional advantage of our Chicago crime data is the level of detail provided in terms of crime type. The data contain over 350 Illinois Uniform Crime Reporting (IUCR) codes that identify types of offenses. The NIBRS data similarly identify offenses at the Uniform Crime Reporting code level. We consider all crime, domestic crime, battery, criminal damage, narcotics, offenses involving children, sex offenses, theft, and weapons violations.

There is a trade off between using adminstrative reported crime data rather than survey data on crimes. Administrative data likely includes more severe crimes because individuals sought police involvement near the time at which the crime took place, whereas survey evidence may include less-severe crimes. Furthermore, administrative crime data is less subject to measurement error in type or intensity that may occur with ex-post descriptions from surveys. On the other hand, it is possible that teacher strikes may affect reported crime by reducing the observation of crime by those likely to report it (such as school personnel for crimes committed at or near school). For this reason, we acknowledge that some of our estimated effects of the teacher strike on crime could be due to effects on reporting propensity, rather than actual crime occurrences.

3 Empirical Strategies

We employ multiple identification strategies to determine the effect of the teacher strike on crime. The overarching theme of these strategies is to compared observed crime in Chicago after the beginning of the teacher strike to estimates of crime in the counterfactual Chicago in which no teacher strike occurred. The event study relies on extrapolation of Chicago's pre-strike crime trend into the strike period and assigns deviations from extrapolated crime to estimated treatment effects. The synthetic controls and related strategies attribute changes in crime that occur in post-treatment synthetic Chicago (a weighted average of control cities) to aggregate shocks, and assigns deviations in actual Chicago from post-treatment synthetic Chicago to estimated treatment effects. The intensity of treatment strategy which tests for effects across Census block groups using school-aged children share of population as a measure of exposure to the strike. This approach attributes changes in crime that occur in plausibly untreated areas within Chicago (block groups with no school-aged children, such as college campuses) to aggregate shocks, and assigns deviations in treated locations in Chicago from untreated locations in Chicago to estimated treatment effects.

3.1 Interrupted Time Series

We begin our analysis by estimating interrupted time series models, following Simonton (1977) and Huitema and Mckean (2000) by estimating a linear regression of the following general form:

$$Y_t = \beta_0 + \beta_1\ t + \beta_2 D_t + \beta_3\ t \times D_t + \varepsilon_t \tag{1}$$

In Equation (1), Y_t is per-capita crime, t is a linear time trend, and D_t an indicator equal to one if the observation occurs after the beginning of Chicago's teacher strike. The main parameters of interest are β_2, the shift in the mean level of per-capita crime, and β_3, the change in the time trend after the teacher strike. This model is effectively identical to a typical regression discontinuity model with time as the running variable. The effects estimated by the model, then, are highly dependent on any additional included control variables and the length of time included in the data pre- and post-strike, analogous to bandwidth selection in a regression discontinuity framework. However, we hesitate to interpret this model causally; the identifying assumption is that pre-strike trends would continue in the absence of the teacher strike, and that no other shocks happened contemporaneously. Given that crime data is extremely noisy and easily affected by many external factors (e.g. weather, policing, etc.), we prefer to interpret these interrupted time series regressions as descriptive of the trends in crime before and after the 2012 teacher strike, saving causal interpretation for our later analysis including comparison counterfactual groups.

3.2 Synthetic Controls

To construct counterfactual estimates of Chicago that are robust to contemporaneous aggregate shocks, we first employ synthetic controls, described by Abadie and Gardeazabal (2003) and Abadie, Diamond, and Hainmueller (2010), and two closely related methods. These methods all involve constructing a synthetic Chicago from other cities using data on crime and predictors of crime such as socio-economic indicators. To deal with noise in daily crime levels, we convert daily data to weekly data (taking sums for crime and precipitation and averages for temperature), starting each week on Monday so that the beginning of the teachers' strike is the first day of a week. Abadie (2021) emphasizes the importance of minimizing the pre-treatment root mean squared prediction error (RMSPE) for the credibility of results, so we employ a residualized synthetic controls estimator (with reference to the suggestion in Doudchenko and Imbens (2016)) as well as a version of the constrained regression method described by Doudchenko and Imbens (2016), both of which leverage our large number of pre-treatment periods to reduce pre-treatment RMSP. The details of these methods are discussed below, with results across methods summarized in Figure 1 and Table 5.

3.2.1 Standard Synthetic Controls

We begin by employing the standard synthetic control method, following Abadie, Diamond, and Hainmueller (2010). This method calculates weights on donor cities that minimize distances from researcher-chosen predictor variables, while also calculating importance weights on predictors for minimizing pre-period residual mean-squared prediction error. If weights that help synthetic Chicago match a particular predictor variable in actual Chicago do not help synthetic Chicago match crime in actual Chicago, that predictor will effectively be discarded. However, if there are weights on donor cities that would help synthetic Chicago match crime in actual Chicago that do not contribute to matching any of the predictors, this method will have no mechanism for choosing these weights. Along these lines, the weights are also restricted to be between zero and one, while summing to one. These combinations of restrictions, along with there being many more time periods than donor cities, prevent synthetic Chicago from matching actual Chicago by construction, as a regression-based method would do if the number of donor cities exceeded the number of time periods. The comparison cities used are shown in Table A.2, while the predictor variables are shown in Table A.1.

Treatment effects are estimated as the difference between Chicago and synthetic Chicago in the treatment time periods. To test hypotheses that treatment effects are significantly different from zero, we employ the placebo inference technique from Abadie, Diamond, and Hainmueller (2010). This procedure involves calculating pre-treatment and post-treatment RMSPE from the differences between actual and synthetic crime for all cities in the donor pool, and comparing their post-RMSPE to pre-RMSPE ratios to that of Chicago. If there is no effect of treatment on crime (as we assume for the placebo units) and no contemporaneous shock, we expect this ratio to be near one. Because the true null distribution is unknown in the presence of contemporaneous shocks, we use the empirical distribution of placebo units as a stand-in and define the p-value for Chicago's treatment effect as its RMSPE ratio rank divided by the number of units.

3.2.2 Residualized Synthetic Controls

As described above, statistical inference for estimated treatment effects using synthetic controls relies on the ratio of post-treatment and pre-treatment RMSPE. We can conceptually decompose these into a generic noise component, which affects both the pre-treatment period and the post-treatment period, and a treatment effect which affects only the post-treatment period. It follows that by reducing the noise component, we will increase the post to pre RMSPE ratio for a given treatment effect, increasing our confidence in treatment effect estimates. Furthermore, a low pre-treatment RMSPE is important for reducing the bias of synthetic controls treatment effect estimates (Abadie, 2021). To this end, we employ a residualized synthetic controls estimator, as suggested by Doudchenko and Imbens (2016), that extracts variation in crime driven by weather and constructs a synthetic Chicago to match the residual crime (such as that driven by socioeconomic factors).

Intuitively, this estimator identifies sources of variation in crime that may contribute to optimal donor weights that conflict with optimal donor from other sources of variation in crime, and conditions out these nuisance sources of variation. Then, it applies standard synthetic controls on the estimated residuals. For instance, if optimal weights under the standard implementation would attempt to match both unobserved weather patterns and socioeconomic factors (insofar as both drive crime), the weights will do a poor job of constructing a useful synthetic Chicago if the optimal weights for matching crime as driven by weather patterns are very different from those for matching crime as driven by socioeconomic factors. By running synthetic controls only on "socioeconomic-factors" crime, synthetic "socioeconomic Chicago" may do a better job of fitting actual "socioeconomic Chicago" than synthetic Chicago does of fitting actual Chicago, without losing any of the variation in outcomes that are actually due to treatment effects of the teacher strike.

To be more precise, we revisit the motivating example model from Abadie, Diamond, and Hainmueller (2010). We define observed crime for unit i in time t as $Y_{it} = Y_{it}^N + \alpha_{it} D_{it}$, with Y_{it}^N giving crime in the absence of the binary policy, D_{it}, and Y_{it}^I giving crime when the policy is in place. The policy affects only Chicago ($i = 1$) in time periods after T_0, so we have

$$D_{it} = \begin{cases} 1 \text{ if } i = 1 \text{ and } t > T_0, \\ 0 \text{ otherwise.} \end{cases}$$

It follows that the treatment effect of interest for $t > T_0$ is given by

$$\alpha_{it} = Y_{1t}^I - Y_{1t}^N = Y_{1t} - Y_{1t}^N,$$

with Y_{it}^I giving crime in the presence of the intervention, which satisfies $Y_{it}^I = Y_{it}$ for $t > T_0$. Our departure from the motivating model in Abadie, Diamond, and Hainmueller (2010) is to augment Y_{it}^N with time-varying observed exogenous determinants, X_{it}, yielding

$$Y_{it}^N = \delta_t + \beta_i X_{it} + \theta_t Z_i + \lambda_t u_i + \epsilon_{it}, \tag{2}$$

where we allow for the additional determinants to have unit-specific effects on crime, β_i. As in ADH, δ_t is a time-varying unobserved common factor with common effects across units, Z_i is a vector of observed predictors of crime with time-varying effects θ_t that are common across units, and u_i is a vector of unobserved predictors of crime with time-varying effects λ_t that are common across units.

We can decompose crime into two unobserved components, defining $Y_{it}^N = Y_{it}^{N1} + Y_{it}^{N2}$, where

$$Y_{it}^{N1} = \delta_t + \theta_t Z_i + \lambda_t u_i + \epsilon_{it}$$
$$Y_{it}^{N2} = \beta_i X_{it}, \tag{3}$$

with $Y_{it}^1 = Y_{it}^{N1} + \alpha_{it} D_{it}$, $Y_{it}^{I1} = Y_{it}^{N1} + \alpha_{it}$, and $Y_{it}^2 = Y_{it}^{N2} = Y_{it}^{I2}$ (recalling that the component of crime driven by X_{it} is exogenous to the treatment). We point out here that the model for Y_{it}^{N1} is identical to the representation of the outcome Y_{it}^N in the motivating model presented in Abadie, Diamond, and Hainmueller (2010). It follows that their results for Y_{it}^N will apply here to Y_{it}^{N1}. It follows that the treatment effects for $t > T_0$ are given by,

$$\alpha_{1t} = Y_{it}^I - Y_{it}^N = Y_{1t}^{I1} - Y_{1t}^{N1} = Y_{1t}^1 - Y_{1t}^{N1}. \tag{4}$$

Following the derivations in Abadie, Diamond, and Hainmueller (2010) yields the residualized synthetic controls estimator,

$$\hat{\alpha}_{1t}^{RSC} = Y_{1t}^1 - \sum_{j=2}^{J+1} w_j^{*,RSC} Y_{jt}^1,$$

with donor weights $w_j^{*,RSC}$ chosen to approximate the unobserved counterfactual crime in Chicago, Y_{1t}^{N1}, using observed crime in donor cities.

Unfortunately, we do not observe Y_{jt}^1 for any j or t. We do, however, observe Y_{jt} and X_{jt} for all j and t. An estimate of Y_{jt}^1 is given by

$$\hat{Y}_{jt}^1 = Y_{jt} - \hat{\beta}_j X_{jt}, \tag{5}$$

where $\hat{\beta}_j$ is obtained from an ordinary least squares regression of Y_{jt} on X_{jt}, and $\hat{\beta}_j$ is a consistent estimate of β_j if X_{jt} is uncorrelated with the composite residual, $\theta_t Z_i + \lambda_t u_j + \epsilon_{jt}$.[3] In our application, X_{it} includes precipitation and temperature. If temporal variation in weather is uncorrelated with socio-economic conditions or other drivers of crime, our exclusion restriction will hold. Given estimates of the residual outcomes, $\hat{Y}_{jt}^1$, we define the feasible residualized

[3] Note that the values in Z_i and u_i are time-invariant, such that their average effects will be absorbed by the constant.

synthetic controls estimator as

$$\hat{\alpha}_{1t}^{FRSC} = \hat{Y}_{1t}^1 - \sum_{j=2}^{J+1} w_j^{*,FRSC} \hat{Y}_{jt}^1.$$

Given these estimates, hypothesis testing is performed as described in Section 3.2.1.

We conclude the discussion of this method by emphasizing the importance of the exogeneity of the predictors in X_{it}. In general, if values of the predictor variables that contribute to the synthetic treatment unit are themselves determined by the treatment itself, the treatment effect will work its way in to the synthetic treated unit, biasing estimates. With predictor variables, Z_i, that are exclusive to time periods prior to treatment (such as background demographic information), it is implausible that they would be affected by treatment if there are no anticipation effects, so this assumption is likely to hold in common synthetic controls applications. The predictors in X_{it}, however, vary over time with the outcome of interest, and we use their values to construct residuals in the post-treatment period. In our application, we include temperature and rainfall in X_{it}, which we assume do not respond to treatment. If they did, the procedure described here would introduce bias into our treatment effect estimates.[4]

3.2.3 Constrained Regression

In addition to the synthetic controls methods described above, we employ a version of the constrained regression method of Doudchenko and Imbens (2016). This method is similar to synthetic controls while allowing (depending on constraints) for a level shift between treated and untreated units, negative weights on donor cities, and weights that do not sum to one. Doudchenko and Imbens (2016) note specifically that these relaxations of restrictions are possible in settings with a large number of pre-treatment periods, which we have in our application.[5]

Our large number of time periods allow us to relax most of the constraints described Doudchenko and Imbens (2016). We maintain only the no-intercept constraint, such that we require Chicago to be produced by a convex combination of donor cities. We implement this method with an ordinary least squares regression of Chicago crime on treatment indicators for weeks

[4] Consider, for instance, the alternative of using daily policing levels as a predictor of crime. If the teacher strike caused, for example, a decrease in effective policing, increasing crime (e.g. via a union solidarity channel), and also increased crime independently of policing, the procedure described here would understate the total effect of the teacher strike on crime by removing the policing channel.

[5] In our application, this regression includes 31 covariates and 139 observations (time periods) prior to treatment. If we had fewer than 31 pre-treatment time periods, our regression model would be underidentified without the inclusion of additional constraints.

one, two, and three, same-week crime in all donor cities, temperature, rainfall, one-week lags of weather variables, and a five-week lag of Chicago crime, with no constant. The lag of Chicago crime is chosen carefully to avoid any post-treatment periods of interest including a lag that takes place after treatment, as this would effectively control for early effects of treatment in later post-treatment periods.

3.3 Intensity of Treatment

Another approach to detecting effects of the teacher strike is to compare neighborhoods across the city who are differentially exposed to the strike based on their concentration of school-age children.

We estimate the following econometric specification by ordinary least-squares:

$$Y_{jt} = \beta_0 + \beta_1 * \text{Strike}_t * \text{School-Age Share}_j + \beta_2 * \text{School-Age Share}_j + \beta_3 * \text{Strike}_t + \mathbf{X}_{jt} * \gamma + \epsilon_{jt} \quad (6)$$

Y_{ijt} is the outcome crimes per 100,000 population in Census block group j on day t. The coefficient of interest is β_1 on the interaction between an indicator for presence of strike on day t and the school-age share of the population in Census block group j.[6] Our sample includes each September day for the years 2010 through 2014. We include year, week of year, and day of week fixed effects. We also include controls for median household income, share of population aged 25 and over with a college degree, and share of population that is black non-Hispanic, white non-Hispanic, and Hispanic.

This approach does not detect aggregate effects of the teacher strike on crime and instead detects variation across Census block groups that can be explained by the concentration of school-aged children in that block group. This strategy assumes that crime involving school-age children will take place in the child's neighborhood.

4 Results

4.1 Interrupted Time Series

Table 1 shows interrupted time series estimates of the effect of the 2012 teacher walkout on weekly crimes per 100,000 population. The term $\mathbf{1}(t > \textit{Strike Start})$ is the estimated discon-

[6] For school-age share and the other demographic variables in this analysis, we use the 2010-2014 5-year ACS.

tinuity in the weekly crime rate, and the term $t \times \mathbf{1}(t > Strike\ Start)$ is the estimated change in the trend in weekly crime rates. Columns are labelled with the category of crime included in each outcome. In all 5 columns, we see no significant discontinuous change in crime rates associated with the 2012 teacher strike, and we only see a marginally significant change in the trend in crime rates for property crimes. Our estimated effects in all cases are small, with estimated discontinuities of less than 10% of the pre-strike mean and estimated slope changes of less than 0.2% of the pre-strike mean.[7] Overall, we see no indication that crime in any category was substantially affected by the 2012 teacher strike in the short- or long-term.

Table 2 further separates the categories presented in Table 1 into 25 types of categories as coded by Chicago PD.[8] In most cases, we see small, statistically insignificant effects of the 2012 teacher strike on both the discontinuity and the trend in crime rates. The five statistically significant discontinuities we observe are all for rare crimes – criminal sexual assault, homicide, interference with a public officer, kidnapping, and prostitution – all of which decrease except for kidnapping. The post-strike trend is statistically significant for seven types of crimes: decreases in criminal burglary, motor vehicle theft, and weapons violations, and increases in liquor law violations, offenses involving children, other offenses, and prostitution. The decrease in the trends for burglaries and motor vehicle thefts is particularly interesting here. Previous work using similar methods (Luallen, 2006; Akee, Halliday, and Kwak, 2014) found null effects or increases in property crimes due to school closings. Our estimated decreases in the post-strike trend in these two types of property crimes combined with null effects in the post-strike discontinuity suggest that the effect of a strike in aggregate may actually be zero, and any increases in property crimes during a teacher strike may only be due to a change in when these property crimes occur.

Table 3 replicates Table 1 including linear controls for weather. As in Table 1, none of the discontinuities are significant or substantial. However, now two of the post-strike trends are statistically significant: all crimes and property crimes both decline after the teacher strike. However, both of the changes in the trend are fairly small: 0.05% and 0.14% of the pre-

[7]Given that our sample extends for two years after the teacher strike and that we use weekly data, our largest estimated effect of the trend change in Column (4) would project a decrease of about 23 crimes per 100,000 at the end of our sample, which is 17.9% of the pre-strike mean.

[8]We exclude exceedingly rare types of crimes from this analysis; concealed carry license violations, human trafficking, non-criminal offenses, obscenity, other narcotic violations, and public indecency all had less than 70 total occurrences in our sample time period.

strike mean, respectively. Including controls for weather affects the point estimates, but still generally results in null findings. Finally, Table 4 considers using different time windows for the interrupted time series estimation, analogous to the choice of bandwidth in a regression discontinuity. The previous results in Tables 1-3 included all data from the beginning of 2010 through the end of 2014. Here, we consider ± 3 weeks around the strike, ± 1 month, all of 2012, and ± 1 year, to ensure that the choice of starting and ending points in our analysis are not driving our null results. In Table 4, changing the time window does not substantially change the interpretation of our findings. The discontinuities in columns (3)-(5) are larger than our main result in column (1), but still remain less than 5% of the pre-strike mean. The post-strike trends in columns (2) and (4) are extremely large, but, given the short post-strike time window, correspond to an increase of 18.9 crimes per 100,000 and a decrease of 54.6 crimes per 100,000, respectively. These effects are both also offset by large pre-strike trends, further dampening the estimated effects. Overall, though we hesitate to interpret these interrupted time series estimates as causal, our descriptive analysis suggests that the 2012 Chicago teacher strike had little to no effect on crime.

4.2 Synthetic Controls and Constrained Regression

As discussed above, the interrupted time series results for total crime show no evidence of substantial effects of the teacher strike on crime. However, it is possible that this is due to an aggregate shock that affects Chicago as well as other cities. The synthetic controls and constrained regression methods described in Section 3.2 address this concern by using post-treatment variation in donor cities to project the likely trajectory of crime in Chicago in the absence of the strike.

Estimates of the effects of the teacher strike on total crime per capita for Synthetic Controls and related methods are given in Table 5. Recalling that the strike lasted 10 days, we present estimates for each method for the first, second, and third week after commencement of the strike. If the strike displaced crime, we should see opposite sign estimates in the first and third weeks. We also include specifications for each method that use the entire pre-period and that use only the first half of the pre-period. If a specification can fit the second half of the pre-period without explicitly targeting it, we will gain confidence in its validity for out-of-sample extrapolation for untreated Chicago in the post-treatment time periods.

Figure 1 shows total crime over time for Chicago as well as Synthetic Chicago constructed from donor states for each specification, with Figure 2 showing the gaps between actual and synthetic Chicago for total crime for each specification. The predictor variables in actual and synthetic Chicago are shown for each specification in Table A.1, with the weights on donor cities for each specification shown in Table A.2. P-values for effects are derived from permutation inference for the synthetic controls estimators, and from robust standard errors for the constrained regressions. The empirical distributions of post to pre RMSPE for each treatment period for the synthetic controls specifications are shown in Figure A.1. In general, our estimates show that actual crime in Chicago drops relative to estimated counterfactuals, but these estimates are only statistically significant at conventional levels for the third week after the strike began (and second week for the constrained regressions).

Importantly, we note that the pre-period fit is imperfect, with especially poor fit just prior to the treatment, for all methods. The residualized synthetic controls specification does have slightly better fit than the standard method, but nonetheless fails to match actual Chicago just prior to the treatment. Interestingly, the residualization places less weight on Rockford, IL (0.194) than standard synthetic controls (0.329), which makes sense given the goal of the residualization. Rockford is 90 miles west of Chicago, and likely does a good job of predicting Chicago's weather and crime-as-driven-by-weather. The residualization effectively extracts this source of variation and places weights on cities that are better matches for Chicago in ways unrelated to weather. Finally, the constrained regression using the entire pre-period exhibits the best fit with an RMSPE of 0.086, which is unsurprising given that it enforces the fewest restrictions on the donor weights. However, the constrained regression actually does the worst job of matching the pre-period when only explicitly fitting the first half of the pre-period, with an RMSPE of 0.138.

Because the pre-period fit is poor just prior to the teacher strike, and the identification of causal effects rests on accurate extrapolation, we hesitate to interpret the results from these methods as causal at this time. One hopeful avenue is searching for additional data from other sources to obtain a richer population of donor cities. It is possible that cities not included in our current dataset will do a better job of matching Chicago's crime trends. Another is estimation of synthetic difference-in-differences, which infers treatment effects from the difference between crime in actual and synthetic Chicago just before and just after treatment.

4.3 Intensity of Treatment

We show results for the intensity of treatment estimation in table 6. In column (1) we include no controls and find that on strike days there are 0.03 fewer crimes per 100,000 people across Census block groups. In column (2) we include the school-age share of population and its interaction with strike days. Here we see a positive coefficient on the interaction, but it is imprecise. In column (3) we add demographic characteristics of the neighborhood, including median household income, share of population 25 and older who have a college degree, share black, share Hispanic, and share white. We see that adding these controls changes the sign of the coefficient on strike day interacted with school-age share. The magnitude is also reduced.

In column (4) we include indicators for year, week of year, and day of week, with little change to the result. Finally, in column (5) we add a time trend, or a continuous variable representing the date. Once including demographic characteristics, we find that the effect of having a 10 percent increase in school-age share (roughly one standard deviation) decreases the crime rate on strike days by 0.9 crimes per 100,000 population, or 22%, but the result is not significant at conventional levels.

5 Conclusion

In this paper we use three distinct approaches to test the effect of being in school on short term crime, considering the case of the 2012 teacher strike in Chicago and the resulting unexpected 8 consecutive days off school for all Chicago Public School students. First we estimate an interrupted time series model. We then employ synthetic control methods. Finally we consider an intensity of treatment design based on heterogeneous exposure across Chicago neighborhoods from their concentration of school-age children. In all three analyses, we detect no change in the crime rate. These findings suggest incapacitation at most plays a minor role in the overall effects of education on crime.

References

ABADIE, A. (2021): "Using synthetic controls: Feasibility, data requirements, and methodological aspects," *Journal of Economic Literature*, 59(2), 391–425.

ABADIE, A., A. DIAMOND, AND J. HAINMUELLER (2010): "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program," *Journal of the American statistical Association*, 105(490), 493–505.

ABADIE, A., AND J. GARDEAZABAL (2003): "The economic costs of conflict: A case study of the Basque Country," *American economic review*, 93(1), 113–132.

AKEE, R. Q., T. J. HALLIDAY, AND S. KWAK (2014): "Investigating the effects of furloughing public school teachers on juvenile crime in Hawaii," *Economics of Education Review*, 42, 1–11.

ANDERSON, D. M. (2014): "In school and out of trouble? The minimum dropout age and juvenile crime," *Review of Economics and Statistics*, 96(2), 318–331.

ANDERSON, D. M., B. HANSEN, AND M. B. WALKER (2013): "The minimum dropout age and student victimization," *Economics of Education Review*, 35, 66–74.

BECKER, G. S. (1968): "Crime and punishment: An economic approach," in *The economic dimensions of crime*, pp. 13–68. Springer.

DAVIS, J. M., AND S. B. HELLER (2020): "Rethinking the benefits of youth employment programs: The heterogeneous effects of summer jobs," *Review of Economics and Statistics*, 102(4), 664–677.

DOUDCHENKO, N., AND G. W. IMBENS (2016): "Balancing, regression, difference-in-differences and synthetic control methods: A synthesis," Discussion paper, National Bureau of Economic Research.

FISCHER, S., AND D. ARGYLE (2018): "Juvenile crime and the four-day school week," *Economics of education Review*, 64, 31–39.

HUITEMA, B. E., AND J. W. MCKEAN (2000): "Design specification issues in time-series intervention models," *Educational and Psychological Measurement*, 60(1), 38–58.

JACOB, B. A., AND L. LEFGREN (2003): "Are idle hands the devil's workshop? Incapacitation, concentration, and juvenile crime," *American economic review*, 93(5), 1560–1577.

Luallen, J. (2006): "School's out… forever: A study of juvenile crime, at-risk youths and teacher strikes," *Journal of urban economics*, 59(1), 75–103.

Manson, S., J. Schroeder, D. Van Riper, T. Kugler, and S. Ruggles (2021): "IPUMS National Historical Geographic Information System: Version 16.0 [dataset]," Technical report, Minneapolis, MN: IPUMS. http://doi.org/10.18128/D050.V16.0.

Modestino, A. S. (2019): "How do summer youth employment programs improve criminal justice outcomes, and for whom?," *Journal of Policy Analysis and Management*, 38(3), 600–628.

Simonton, D. K. (1977): "Cross-sectional time-series experiments: Some suggested statistical analyses.," *Psychological Bulletin*, 84(3), 489.

Tables and Figures

Table 1: Interrupted Time Series: Effect of Strike on Weekly Crimes per 100,000 Population

	(1)	(2)	(3)	(4)	(5)
	All Crimes	Domestic Crimes	Personal Crimes	Property Crimes	Violent Crimes
t	-0.166	-0.00913	-0.0327	-0.0625	-0.0298
	(0.111)	(0.0178)	(0.0399)	(0.0736)	(0.0451)
$\mathbf{1}(t > \textit{Strike Start})$	-7.104	-2.421	-4.928	-1.088	-4.494
	(9.869)	(1.838)	(3.932)	(5.588)	(4.213)
$t \times \mathbf{1}(t > \textit{Strike Start})$	-0.224	-0.0343	-0.0240	-0.222*	-0.0562
	(0.196)	(0.0304)	(0.0688)	(0.134)	(0.0788)
Pre-Strike Mean of DV	253.91	36.05	65.8	129.1	71.8
N	255	255	255	255	255

Prais-Winston standard errors corrected for autocorrelation of residuals in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Interrupted time series estimates presented, allowing for a change in the post-treatment trend. Outcome variable in Column (1) is all crimes in Chicago, and Column (2) is Chicago's categorization of "Domestic" crimes. Column (3) is all personal crimes, including assault, battery, criminal sexual assault, homicide, human trafficking, interference with public officer, intimidation, kidnapping, offenses involving children, prostitution, sex offenses, and stalking. Column (4) is all property crimes, including arson, burglary, criminal damage, criminal trespass, motor vehicle theft, robbery, and theft. Column (5) is all violent crimes, including arson, assault, battery, criminal sexual assault, homicide, human trafficking, intimidation, kidnapping, robbery, and sex offenses.

Table 2: Interrupted Time Series: Effect of Strike on Weekly Crimes by Type

	Crimes A-C								
	All Crimes	All Domestic Crimes	Arson	Assault	Battery	Criminal Burglary	Criminal Sexual Assault	Criminal Damage	Trespass
t	-0.166	-0.00913	0.0000797	-0.00735	-0.0221	-0.0151	0.000744	-0.0228	-0.00828***
	(0.111)	(0.0178)	(0.000395)	(0.00933)	(0.0285)	(0.0135)	(0.000706)	(0.0165)	(0.00182)
$\mathbf{1}(t > \textit{Strike Start})$	-7.104	-2.421	-0.0663	-0.588	-3.659	-0.768	-0.261***	-1.085	0.161
	(9.869)	(1.838)	(0.0473)	(0.980)	(2.809)	(1.192)	(0.0847)	(1.581)	(0.215)
$t \times \mathbf{1}(t > \textit{Strike Start})$	-0.224	-0.0343	-0.000518	-0.0102	-0.0178	-0.0408*	0.000523	-0.0223	0.00243
	(0.196)	(0.0304)	(0.000642)	(0.0158)	(0.0491)	(0.0238)	(0.00115)	(0.0286)	(0.00296)
Pre-Strike Mean of DV	253.91	36.05	0.39	14.90	44.63	18.09	1.02	27.32	6.34
N	255	255	255	255	255	255	255	255	255

	Crimes D-N								
	Deceptive Practice	Gambling	Homicide	Interference with Public Officer	Intimidation	Kidnapping	Liquor Law Violation	Motor Vehicle Theft	Narcotics
t	0.00459	0.000524	0.000917**	0.00406***	-0.000299*	-0.000935***	-0.00263***	-0.0142**	-0.0506***
	(0.00307)	(0.00166)	(0.000457)	(0.000733)	(0.000161)	(0.000256)	(0.000467)	(0.00695)	(0.0101)
$\mathbf{1}(t > \textit{Strike Start})$	0.520	-0.247	-0.112**	-0.259***	0.0157	0.0817***	0.0326	-0.410	-0.615
	(0.366)	(0.183)	(0.0531)	(0.0873)	(0.0189)	(0.0305)	(0.0559)	(0.714)	(1.143)
$t \times \mathbf{1}(t > \textit{Strike Start})$	0.00529	-0.000789	-0.000774	-0.00108	0.000234	0.000297	0.00177**	-0.0373***	0.0123
	(0.00499)	(0.00283)	(0.000728)	(0.00119)	(0.000258)	(0.000420)	(0.000765)	(0.0118)	(0.0166)
Pre-Strike Mean of DV	9.07	0.70	0.35	0.75	0.16	0.23	0.49	13.12	28.70
N	255	255	255	255	255	255	255	255	255

	Crimes O-Z								
	Offenses Involving Children	Other Offenses	Prostitution	Public Peace Violations	Robbery	Sex Offenses	Stalking	Theft	Weapons Violations
t	-0.00298***	-0.0325***	-0.00342***	-0.00355**	0.00219	0.000556	0.0000405	0.000602	0.00189
	(0.00106)	(0.00361)	(0.00130)	(0.00160)	(0.00787)	(0.000738)	(0.000190)	(0.0341)	(0.00174)
$\mathbf{1}(t > \textit{Strike Start})$	0.152	0.557	-0.340**	0.147	-0.733	-0.130	-0.00270	-0.112	-0.301
	(0.127)	(0.425)	(0.156)	(0.187)	(0.736)	(0.0884)	(0.0228)	(2.896)	(0.202)
$t \times \mathbf{1}(t > \textit{Strike Start})$	0.00305*	0.0253***	0.00420**	0.00244	-0.0222	-0.000569	-0.000102	-0.0930	-0.00601**
	(0.00173)	(0.00588)	(0.00212)	(0.00261)	(0.0137)	(0.00120)	(0.000309)	(0.0609)	(0.00286)
Pre-Strike Mean of DV	1.72	14.47	1.75	2.33	9.84	0.81	0.17	54.12	2.74
N	255	255	255	255	255	255	255	255	255

Prais-Winston standard errors corrected for autocorrelation of residuals in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Interrupted time series estimates presented, allowing for a change in the post-treatment trend. Outcomes are presented by type of crime, using Chicago's categorization of all crime types. All types of crime with at least 70 occurrences in the sample period are presented in this table; this excludes concealed carry license violations, human trafficking, non-criminal offenses, obscenity, other narcotic violations, and public indecency.

Table 3: Interrupted Time Series: Effect of Strike on Weekly Crimes per 100,000 Population with Weather Controls

	(1) All Crimes	(2) All Domestic Crimes	(3) Personal Crimes	(4) Property Crimes	(5) Violent Crimes
t	-0.276***	-0.0279***	-0.0806***	-0.102***	-0.0808***
	(0.0251)	(0.00769)	(0.0121)	(0.0343)	(0.0115)
$\mathbf{1}(t > \textit{Strike Start})$	3.554	-0.218	0.701	1.338	1.077
	(2.976)	(0.901)	(1.428)	(3.604)	(1.355)
$t \times \mathbf{1}(t > \textit{Strike Start})$	-0.117***	-0.0201	0.0124	-0.175***	-0.0162
	(0.0403)	(0.0124)	(0.0195)	(0.0578)	(0.0185)
Precip. (in)	-3.169***	-0.000896	-0.677**	-1.579***	-0.720**
	(0.699)	(0.174)	(0.298)	(0.389)	(0.290)
Avg. High Temp.	0.728***	0.122***	0.357***	0.266***	0.373***
	(0.130)	(0.0333)	(0.0565)	(0.0790)	(0.0548)
Avg. Low Temp.	0.289*	0.0393	0.0256	0.201**	0.0436
	(0.148)	(0.0376)	(0.0641)	(0.0868)	(0.0622)
Pre-Strike Mean of DV	253.91	36.05	65.8	129.1	71.8
N	255	255	255	255	255

Prais-Winston standard errors corrected for autocorrelation of residuals in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Interrupted time series estimates presented, allowing for a change in the post-treatment trend. All models include linear controls for weather: precipitation in inches, the weekly average high temperature, and the weekly average low temperature. Outcome variable in Column (1) is all crimes in Chicago, and Column (2) is Chicago's categorization of "Domestic" crimes. Column (3) is all personal crimes, including assault, battery, criminal sexual assault, homicide, human trafficking, interference with public officer, intimidation, kidnapping, offenses involving children, prostitution, sex offenses, and stalking. Column (4) is all property crimes, including arson, burglary, criminal damage, criminal trespass, motor vehicle theft, robbery, and theft. Column (5) is all violent crimes, including arson, assault, battery, criminal sexual assault, homicide, human trafficking, intimidation, kidnapping, robbery, and sex offenses.

Table 4: Interrupted Time Series: Effect of Strike on Weekly Crimes per 100,000 Population with Different Timing

	(1) 2010-2014	(2) ± 3 Weeks	(3) ± 1 Month	(4) All of 2012	(5) ± 1 Year
t	-0.166	-7.939	-1.527	1.307***	-0.0456
	(0.111)	(3.682)	(1.180)	(0.388)	(0.332)
$\mathbf{1}(t > \textit{Strike Start})$	-7.104	2.663	-9.041*	-12.09	-10.11
	(9.869)	(5.834)	(3.762)	(9.407)	(9.693)
$t \times \mathbf{1}(t > \textit{Strike Start})$	-0.224	6.321	0.828	-4.210***	-0.0618
	(0.196)	(3.583)	(1.237)	(1.272)	(0.529)
Pre-Strike Mean of DV	253.91	247.07	248.49	243.46	243.65
N	255	7	10	50	103

Prais-Winston standard errors corrected for autocorrelation of residuals in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Interrupted time series estimates presented, allowing for a change in the post-treatment trend. Models estimate different time windows around the strike; all outcome variables are total crimes in Chicago. Column (1) includes all crimes from Jan 1, 2010 through Dec. 31, 2014. Column (2) only includes crimes from 3 weeks before the 2012 strike through 3 weeks after the 2012 strike: August 20, 2012 through October 10, 2012. Column (3) only includes crimes from 1 month before through 1 month after the 2012 strike: August 10, 2012 through October 19, 2012. Column (4) includes all of 2012. Column (5) includes crimes from 1 year before through 1 year after the 2012 strike: September 10, 2011 through September 19, 2013.

Figure 1: Total Crime in Actual and Synthetic Chicago, by Method

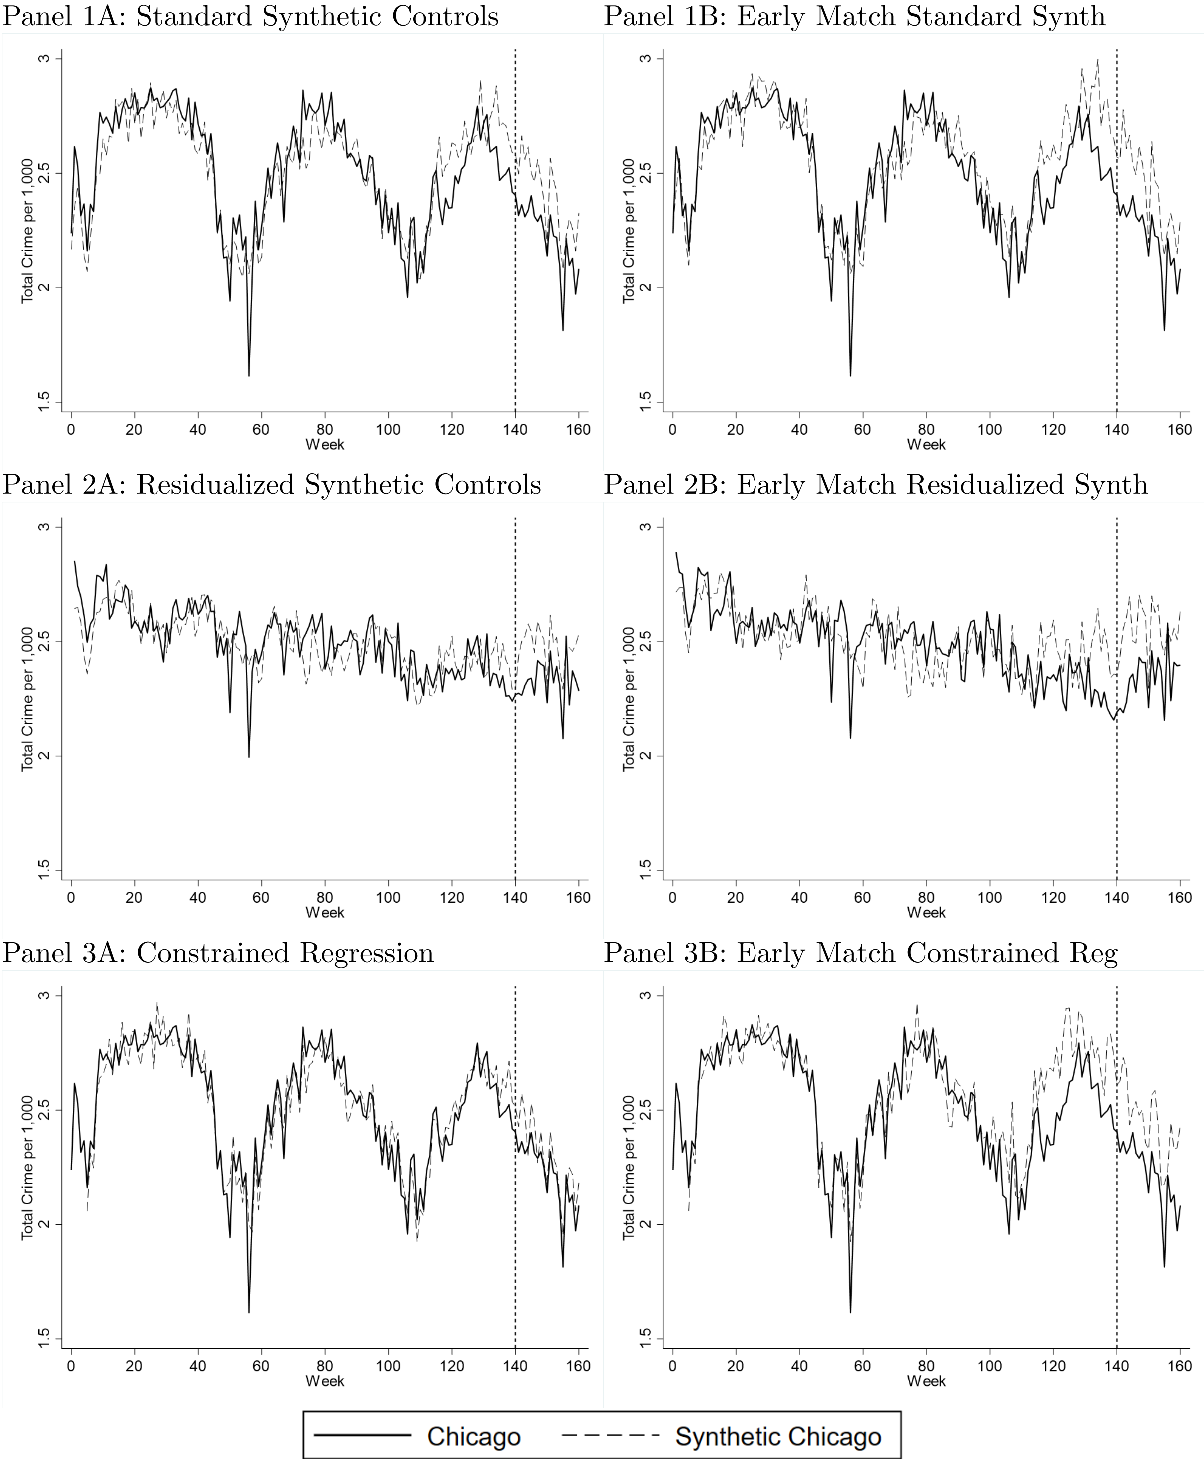

Notes: The top panel shows results for standard synthetic controls. The second panel shows results for residualized synthetic controls, which effectively conditions out crime as driven by weather. The third panel shows results for constrained regression, which removes the synthetic controls restrictions that weights be non-negative and sum to one. The 2nd column shows results for each method when minimizing RMSPE for only the first 70 weeks.

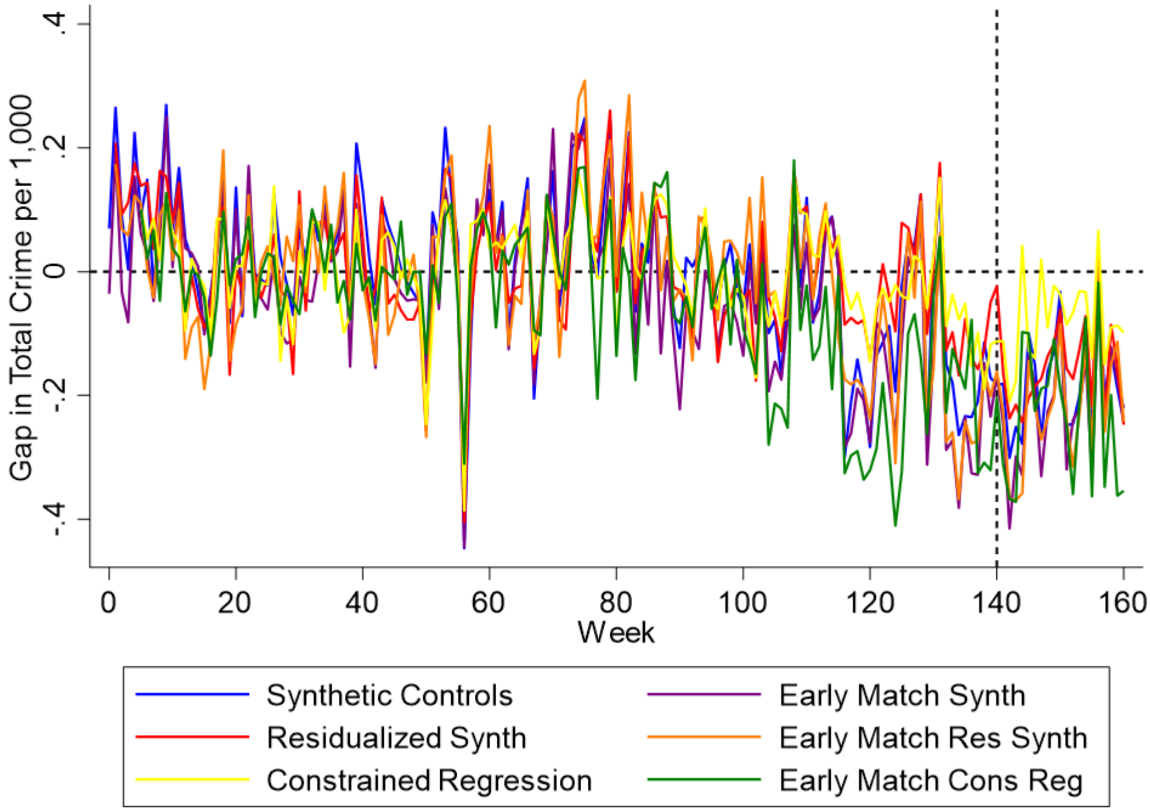

Figure 2: Actual-Minus-Predicted Crime Gaps by Method

Table 5: Effects of Teacher Strike on Weekly Crime per 1,000 Population, Synthetic Controls and Related Methods

	Synthetic Controls	Synth (Early)	Residualized Synth	Res Synth (Early)	Constrained Regression	Cons Reg (Early)
Effect (Week 1)	-0.182	-0.166	-0.024	-0.161	-0.112	-0.205
	(0.154)	(0.308)	(0.808)	(0.308)	(0.154)	(0.132)
Effect (Week 2)	-0.182	-0.246	-0.164	-0.227	-0.112	-0.312
	(0.231)	(0.115)	(0.231)	(0.231)	(0.066)	(0.024)
Effect (Week 3)	-0.185	-0.222	-0.159	-0.221	-0.210	-0.367
	(0.038)	(0.038)	(0.038)	(0.038)	(0.000)	(0.000)
Pre-Treatment RMSPE	0.127	0.137	0.106	0.131	0.086	0.138
Weight Estimation Periods	140	70	139	69	135	65

Notes: P-values in parentheses, obtained from permutation inference for synthetic controls and from robust standard errors for constrained regressions.

Table 6: Effect of Strike on Crimes per 100,000 Population

	(1)	(2)	(3)	(4)	(5)
Strike Day=1	-0.0328***	-0.101*	0.00523	0.00988	0.0103
	(0.0124)	(0.0546)	(0.0120)	(0.0127)	(0.0127)
School-Age Share		-0.134	-0.278***	-0.275***	-0.275***
		(0.102)	(0.0213)	(0.0214)	(0.0214)
Strike Day=1 × School-Age Share		0.375	-0.0876	-0.0896	-0.0896
		(0.241)	(0.0586)	(0.0586)	(0.0586)
Median Household Income, Thousands of Dollars			-0.00123***	-0.00121***	-0.00121***
			(0.0000572)	(0.0000576)	(0.0000576)
College Graduate Share			0.0183**	0.0133	0.0133
			(0.00838)	(0.00843)	(0.00843)
Black Non-Hispanic Share			0.524***	0.522***	0.522***
			(0.0113)	(0.0113)	(0.0113)
Hispanic Share			0.0861***	0.0826***	0.0826***
			(0.0113)	(0.0114)	(0.0114)
White Non-Hispanic Share			-0.00849	-0.00728	-0.00728
			(0.0121)	(0.0122)	(0.0122)
Day					0.000660
					(0.000677)
Constant	0.410***	0.435***	0.279***	0.367***	-11.84
	(0.00501)	(0.0233)	(0.0113)	(0.0132)	(12.53)
Year Indicators	No	No	No	Yes	Yes
Week of Year Indicators	No	No	No	Yes	Yes
Day of Week Indicators	No	No	No	Yes	Yes
Non-Strike Mean of DV	.41	.41	.41	.41	.41
N	337,595	337,595	336,994	335,250	335,250

Notes: Heteroskedasticity robust standard errors are in parentheses. $p \leq 0.10, **p \leq 0.05, ***p \leq 0.01$

Appendix A: Additional Figures and Tables

Table A.1: Actual and Synthetic Chicago Predictor Balance by Specification

Variable Names	Actual	Synthetic Controls	Synth (Early)	Residualized Synth	Res Synth (Early)
%Pop Age 15-19 White Male	2.211	2.395	2.362	2.374	2.382
%Pop Age 15-19 Black Male	0.739	0.702	0.715	0.700	0.711
%Pop Age 20-44 White Male	0.637	0.547	0.535	0.540	0.525
%Pop Age 20-44 Black Male	11.376	12.704	12.802	12.770	12.946
%Pop Age 15-19 Other Male	2.607	2.712	2.750	2.697	2.740
%Pop Age 20-44 Other Male	3.335	2.551	2.509	2.530	2.485
Poverty Rate	14.102	15.587	15.115	15.291	15.053
Average Income Per Capita	3.1e+04	2.6e+04	2.6e+04	2.6e+04	2.6e+04
Unemployment Rate	10.466	8.797	8.272	8.491	8.288
Total Crime per Capita (Week 2, September 2010)	2.729	2.653	2.738	.	.
Total Crime per Capita (Week 4, April 2011)	2.606	.	2.490	.	.
Total Crime per Capita (Week 2, September 2011)	2.588	2.470	.	.	.
Total Crime per Capita (Week 1, September 2012)	2.419	2.512	.	.	.
Residualized Total Crime per Capita (Week 2, September 2010)	2.548	.	.	2.514	2.553
Residualized Total Crime per Capita (Week 4, April 2011)	2.585	.	.	.	2.505
Residualized Total Crime per Capita (Week 2, September 2011)	2.550	.	.	2.383	.
Residualized Total Crime per Capita (Week 1, September 2012)	2.240	.	.	2.340	.

Notes: Crime variables are all in per capita terms. Actual crime for residualized synthetic controls are from Chicago's actual residualized crime. Predictor balance not shown for constrained regression, as that method does not choose weights to fit predictors.

Table A.2: Synthetic Chicago Donor Weights by Specification

Donor City	Synthetic Controls	Synth (Early)	Residualized Synth	Res Synth (Early)	Constrained Regression	Cons Reg (Early)
AKRON, OH	0	0	0	0	-0.065	-0.047
AMARILLO, TX	0	0	0	0	0.193	0.043
ANN ARBOR, MI	0	0	0	0	0.099	0.040
BOULDER, CO	0	0	0	0	-0.088	-0.105
CEDAR RAPIDS, IA	0	0	0	0	0.011	0.001
CHATTANOOGA, TN	0	0	0	0	-0.048	0.001
CINCINNATI, OH	0	0	0	0	0.025	-0.061
CLARKSVILLE, TN	0	0	0	0	-0.031	0.063
COLORADO SPRINGS, CO	0	0	0	0	0.038	0.253
COLUMBIA, SC	0	0	0	0	0.093	0.086
COLUMBUS, OH	0	0	.294	0	0.185	-0.066
FLINT, MI	0	.134	0	0	0.102	0.162
GREEN BAY, WI	0	0	0	0	0.102	0.206
KANSAS CITY, MO	.11	0	.027	0	0.277	0.309
KNOXVILLE, TN	0	0	0	0	0.053	-0.133
LONGVIEW, TX	0	0	0	0	0.063	0.070
LUBBOCK, TX	0	0	0	0	0.001	-0.010
MADISON, WI	0	0	0	0	0.045	-0.094
MEMPHIS, TN	.313	.206	.277	.263	-0.024	0.168
ROANOKE, VA	.069	.087	.049	.047	0.062	0.076
ROCKFORD, IL	.329	.29	.194	.363	0.132	0.156
SALEM, OR	0	0	0	0	-0.067	0.109
SALT LAKE CITY, UT	.179	.282	.158	.301	0.010	0.052
TULSA, OK	0	0	0	0	-0.054	-0.014
WORCESTER, MA	0	0	0	.025	0.183	0.282

Figure A.1: Ratio of post to pre RMSPE by method and post-treatment week

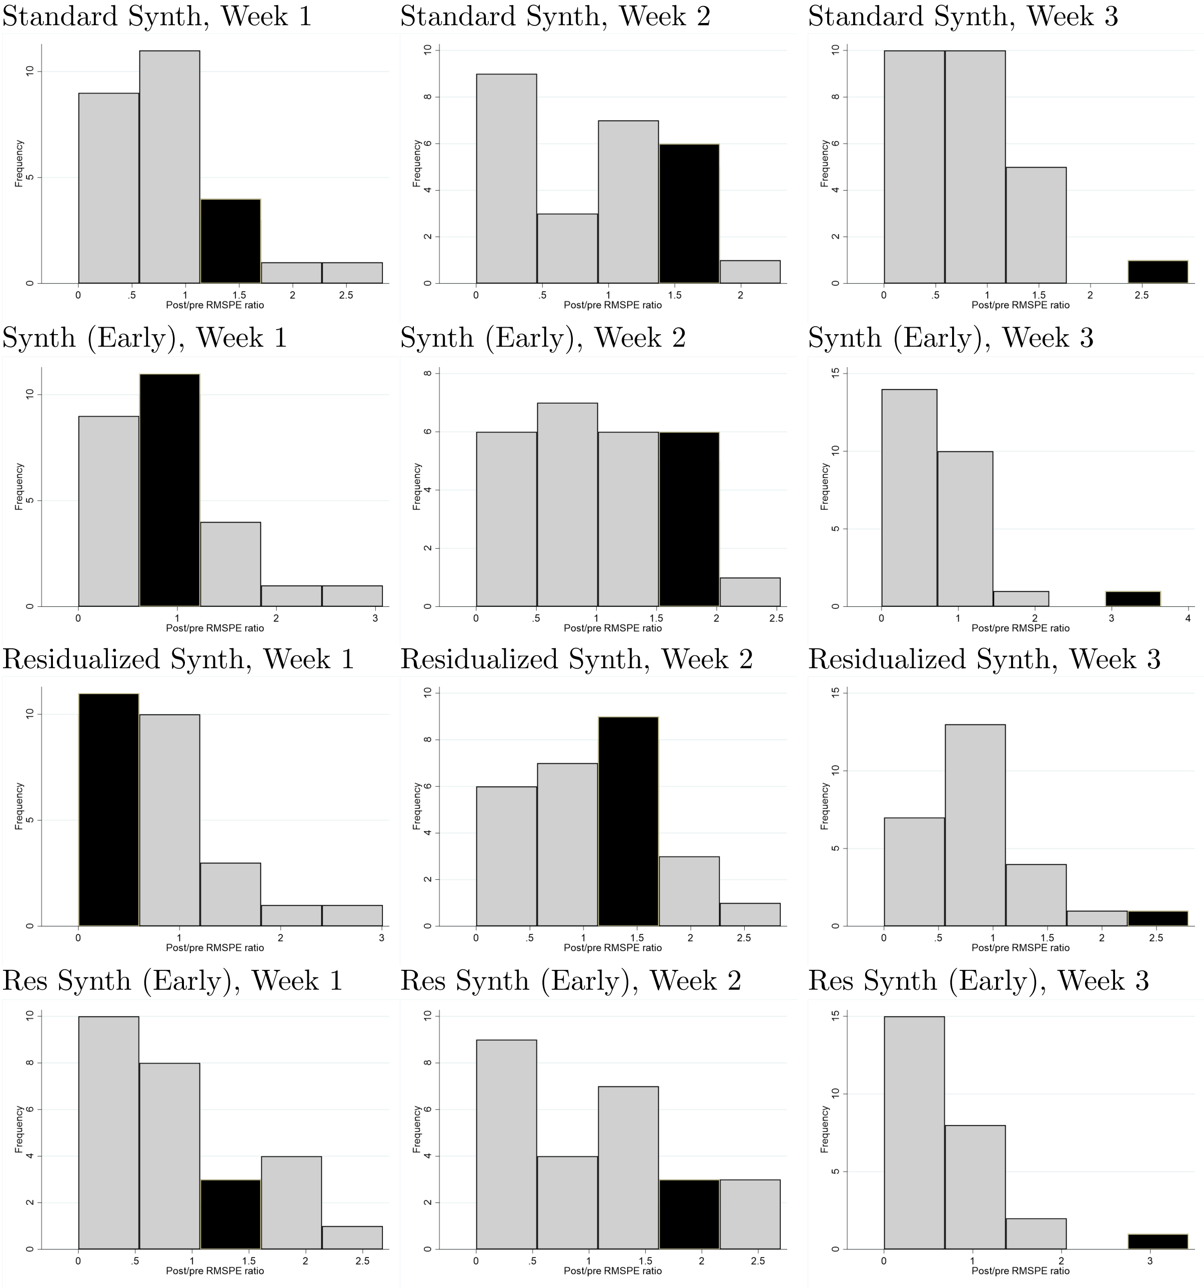

Notes: Post-to-pre RMSPE ratios shown for synthetic controls specifications, for 3 weeks after the beginning of the strike. Chicago's position in these distributions (dark bins) are sufficient to determine p-values for treatment effects using permutation inference.