Does Proactive Policing Really Increase Major Crime? Accounting for An Ecological Fallacy

Aaron Chalfin[1], David Mitre-Becerril[1] and Morgan C. Williams, Jr.*[2]

[1]Department of Criminology, University of Pennsylvania
[2]Robert F. Wagner Graduate School of Public Service, New York University

May 27, 2021

Abstract

In December 2014 and January 2015, police officers in New York City engaged in an organized slowdown of police work to protest the murder of two police officers who were targeted by a gunman while sitting in their patrol car and in response to a perceived lack of political support from NYC Mayor Bill de Blasio. An influential 2017 article in *Nature Human Behaviour* studies the effect of the NYPD's work slowdown on major crimes and concludes that the slowdown led to a significant *improvement* in public safety. We re-evaluate this claim and point out several fatal weaknesses in the authors' analysis that call this finding into question. In particular, we note that there was considerable variation in the intensity of the slowdown across NYC communities and that the communities which experienced a more pronounced reduction in police proactivity did not experience the largest reductions in major crime. The authors' analysis constitutes a quintessential example of an ecological fallacy in statistical reasoning, a logical miscalculation in which inferences from aggregated data are mistakenly applied to a more granular phenomenon. We raise several additional and equally compelling concerns regarding the tests presented in the paper and conclude that there is little evidence that the slowdown led to short-term changes in major crimes in either direction.

Keywords: Proactive Policing, Crimes, Ecological Fallacy

*Correspondence: Morgan C. Williams, Jr., Robert F. Wagner Graduate School of Public Service, New York University, 295 Lafayette St., New York, NY, 10012-9604, E-Mail: mcw394@nyu.edu.

In a 2017 paper published in *Nature Human Behaviour*, Sullivan and O'Keeffe (hereafter S-O) identify a potentially promising natural experiment to study the effect of a large and discrete reduction in proactive policing in NYC. The natural experiment is motivated by a coordinated "slowdown" of police work undertaken by NYC police officers in December 2014 and January 2015 in response to a series of events that shook their faith in the city's political leadership. S-O study major crimes known to law enforcement during the slowdown period and, accounting for annual crime trends, document a 6% decline relative to the same period in the previous year. Their conclusion is that "curtailing proactive policing can reduce major crime." We re-evaluate this claim and point out several fatal weaknesses in the authors' analysis that calls their finding into question.

Inconsistent Evidence and Failed Placebo Tests

In Figure 1 we use the statistical approach described in equation (1) of S-O's paper to replicate their analysis for all t, t-1 year pairs during the 2007-2019 period.[1] The analysis identifies the effect of the slowdown by differencing the year-over-year change in the number of major crimes experienced during the December 1st-January 18th slowdown period from the year-over-year change in major crimes experienced during the remainder of the year. For each pair of years, the figure reports incident rate ratios, standard errors clustered by police precinct and 95% confidence intervals from a series of negative binomial regressions. Panel A considers the issuing of criminal summonses, the primary marker of the work slowdown. Panel B considers S-O's primary outcome, major "index" crimes known to law enforcement. We are able to substantively replicate S-O's analysis, finding that criminal summonses declined by 50% and major crimes declined by 7% during the 2014-15 slowdown period as compared to the previous year.

A more careful review of Figure 1 reveals several troubling issues. We begin by noting that

[1] We replicate S-O's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as S-O note, therefore do not influence the resulting estimates.

sizable annual changes in summonses are actually very common and so the large reduction in summonses during the 2014-15 slowdown period is not unique. Indeed we observe statistically significant annual changes in summonses issued in 10 out of the 11 additional periods for which we have data. Some of these changes are qualitatively important and include changes as large as 56% and 91%.[2] Referring to the figure, we observe that when summonses fell by 29% during the 2017-18 period and when summonses increased by 91% in the 2015-16 period, there was, in fact, no change in major crimes. Contrary to the claim made by S-O, Figure 1 suggests that the relationship between police proactivity and major crimes is highly variable and remains far from clear.

Next, we focus on the statistical significance of S-O's result. While the estimate for the 2014-15 slowdown period is the second largest among the twelve year t, t-1 pairs, 6 out of the 11 remaining estimates are statistically significant at conventional levels. A number of the estimates have confidence intervals which overlap substantially with that of the 2014-15 slowdown period. Since there were no work slowdowns during the other periods, the authors' significant finding for the 2014-15 period appears to be an artifact of ordinary year-to-year variation in major crime rates in a single city during a seven-week period. Given that 6 out of 11 tests yield a significant coefficient, we suggest that the implicit p-value on the authors' finding is 0.55 rather than < 0.05.

Ecological Fallacy

The authors' analysis draws a city-level inference, ignoring sub-city variation in the intensity of the work slowdown. If there is considerable spatial variation in the intensity of the slowdown, the authors' analysis is vulnerable to an ecological fallacy in statistical reasoning, a logical miscalculation in which inferences from aggregated data are mistakenly applied to a more granular phenomenon. We analyze the relationship between the intensity of the slowdown and the change in major crimes

[2]The 91% increase in summonses issued during the 2015-16 period is an artifact of the 2014-15 slowdown.

at the police precinct level in Figure 2. In Panel A, we present a precinct-level heat map of NYC and document enormous variation in the intensity of the slowdown across NYC communities. For each of NYC's 76 police precincts (excluding the Central Park precinct), the intensity of the color in the heat map corresponds with the year-over-year change in summonses issued during the December 1st-January 18th period. While some communities experienced an 80% reduction in summonses issued, other communities experienced a more modest reduction of 20-25%.

In Panel B, we exploit the precinct-level variation and plot each precinct's percentage change in major crimes between the December 1st, 2013 - January 18th, 2014 and the December 1st, 2014 - January 18th, 2015 periods (y-axis) against the same percentage change in summonses issued (x-axis). A best-fit line is drawn through the data. If crime had fallen most dramatically in the precincts that experienced the largest slowdowns, we would expect this regression line to have a positive slope — that is, the larger the slowdown, the larger the drop in crime. Referring to the figure, we see little evidence that this is the case. Indeed, the correlation between these two variables is very close to zero indicating that precinct-level changes in criminal summonses during the slowdown period explains very little of the variation in major crimes across communities, a finding which is wholly inconsistent with the authors' claim.

Discussion

S-O claim that order maintenance policing compromises public safety, a fact which, if true, provides extraordinarily compelling evidence against the deployment of this style of policing in U.S. cities. Unfortunately this claim is built upon dubious evidence. Having re-analyzed the data, we conclude that there is little evidence that the NYPD's work slowdown caused major crimes to decline. While causal claims require strong assumptions, the evidence is instead largely consistent with the finding that major crimes did not change appreciably as a function of the 2014-2015 slowdown nor did they change appreciably in response to other changes in summonses issued during other periods.

Similar to an earlier analysis of a 1997 work slowdown by NYPD officers (Chandrasekher, 2016), the natural experiment provided by the NYPD's 2014-15 work slowdown provides a useful data point on the short-term effects of a decline in police proactivity, at least in one relatively safe U.S. city. However, there is conflicting evidence from Chicago which experienced a similar work slowdown in response to the October 2015 release of body-worn camera video of the death of Laquan McDonald and a subsequent legal settlement with the ACLU that increased the administrative burden of a street stop. In that case, the work slowdown was followed by a large increase in gun violence (Kapustin et al., 2017; Cassell and Fowles, 2018). Ultimately, it is difficult to infer causality from a single event. We thus urge caution in drawing global conclusions from data collected over a short time period in a single city. We likewise note that whether a more prolonged work slowdown might have different effects remains unknown.

References

Braga, A. A., B. S. Turchan, A. V. Papachristos, and D. M. Hureau (2019). Hot spots policing and crime reduction: an update of an ongoing systematic review and meta-analysis. *Journal of Experimental Criminology 15*(3), 289–311.

Cassell, P. G. and R. Fowles (2018). What caused the 2016 chicago homicide spike: An empirical examination of the aclu effect and the role of stop and frisks in preventing gun violence. *University of Illinois Law Review*, 1581.

Chandrasekher, A. C. (2016). The effect of police slowdowns on crime. *American Law and Economics Review 18*(2), 385–437.

Cho, S., F. Gonçalves, and E. Weisburst. Do police make too many arrests? *University of California, Los Angeles Working Paper*.

Harcourt, B. E. and J. Ludwig (2006). Broken windows: New evidence from new york city and a five-city social experiment. *The University of Chicago Law Review*, 271–320.

Kapustin, M., J. Ludwig, M. Punkay, K. Smith, L. Speigel, and D. Welgus (2017). Gun violence in chicago, 2016. *University of Chicago Crime Lab White Paper*.

MacDonald, J., J. Fagan, and A. Geller (2016). The effects of local police surges on crime and arrests in new york city. *PLoS one 11*(6).

MacDonald, J. M., J. Klick, and B. Grunwald (2016). The effect of private police on crime: evidence from a geographic regression discontinuity design. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 179*(3), 831–846.

Rosenfeld, R. (2016). Documenting and explaining the 2015 homicide rise: Research directions. *Washington, DC*.

Sherman, L. W. and D. Weisburd (1995). General deterrent effects of police patrol in crime "hot spots": A randomized, controlled trial. *Justice Quarterly 12*(4), 625–648.

Sullivan, C. M. and Z. P. O'Keeffe (2017). Evidence that curtailing proactive policing can reduce major crime. *Nature Human Behaviour 1*(10), 730–737.

Weisburd, S. (2016). Police presence, rapid response rates, and crime prevention. *The Review of Economics and Statistics*, 1–45.

Figure 1: The 2014-2015 NYPD slowdown and its effect on major crimes

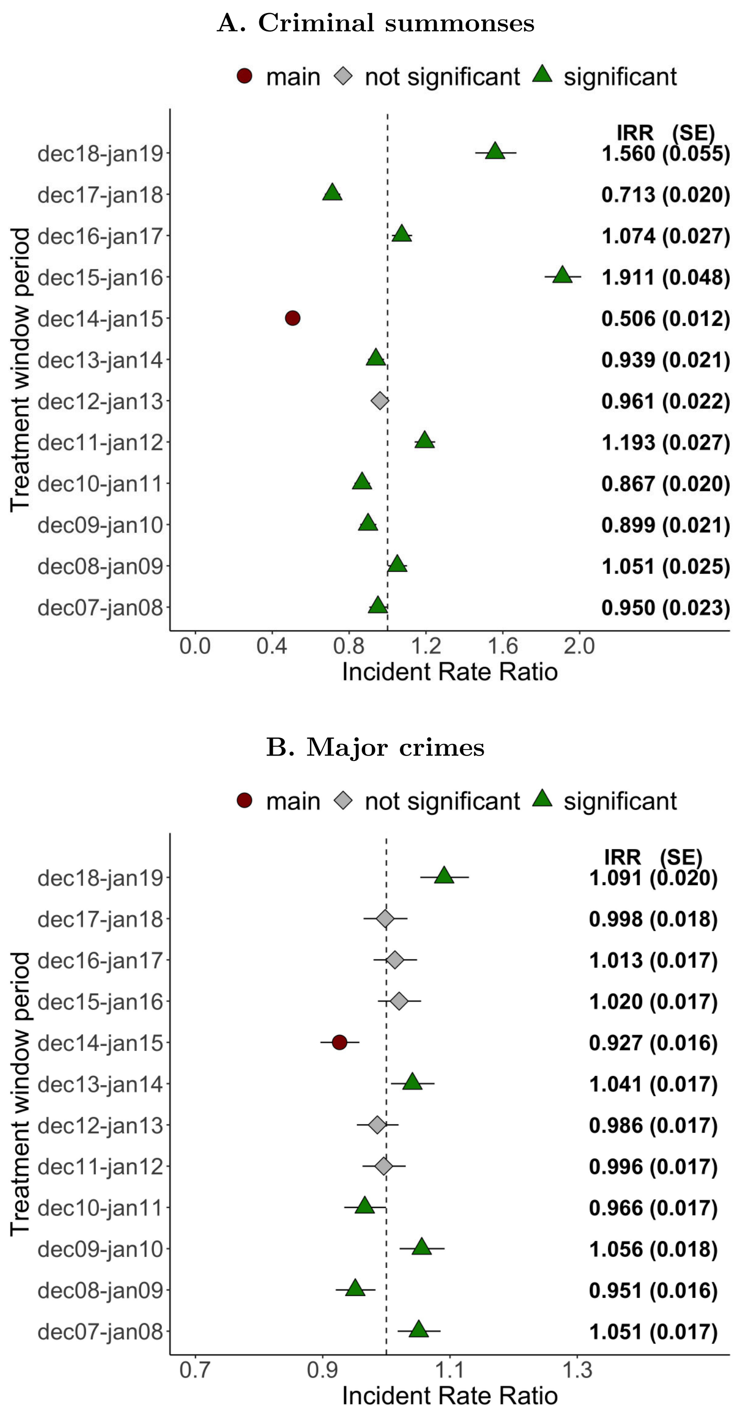

Note: Each row presents estimates from S-O's differences-in-differences estimator. The analysis identifies the effect of the slowdown by differencing the year-over-year change in the number of major crimes experienced during the December 1st-January 18th slowdown period from the year-over-year change in major crimes experienced during the remainder of the year. For each year t, year t-1 pair, we replicate S-O's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as S-O note, therefore do not influence the resulting estimates. For a given period, the estimating equation is as follows: $y_{it} = \alpha + \beta_1 S_{it} + \beta_2 T_{it} + \beta_3 S_{it} T_{it} + \gamma X_{it} + e_{it}$, where y_{it} is the outcome variable at the precinct i, day-month-year t, $S_{it} = 1$ for period t between January/19/YEAR+1 and January/18/YEAR+2, zero otherwise, $T_{it} = 1$ is the treatment window for period t between December/01/YEAR and January/18/YEAR+1 as well as December/01/YEAR+1 and January/18/YEAR+2, zero otherwise, and X_{it} represents weather data (snow, precipitation, and temperature) at the city-day-month-year level. Each model only includes data between January/19/YEAR to January/18/YEAR+2. For each pair of years, the figure reports incident rate ratios, standard errors clustered by police precinct and 95% confidence intervals from a series of negative binomial regressions. Panel A considers the issuing of criminal summonses, the primary marker of the work slowdown. Panel B considers S-O's primary outcome, major crimes known to law enforcement, which includes the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft.

Figure 2: Precinct-level variation in the 2014-2015 NYPD slowdown

A. Geographic variation in the annual change in criminal summonses

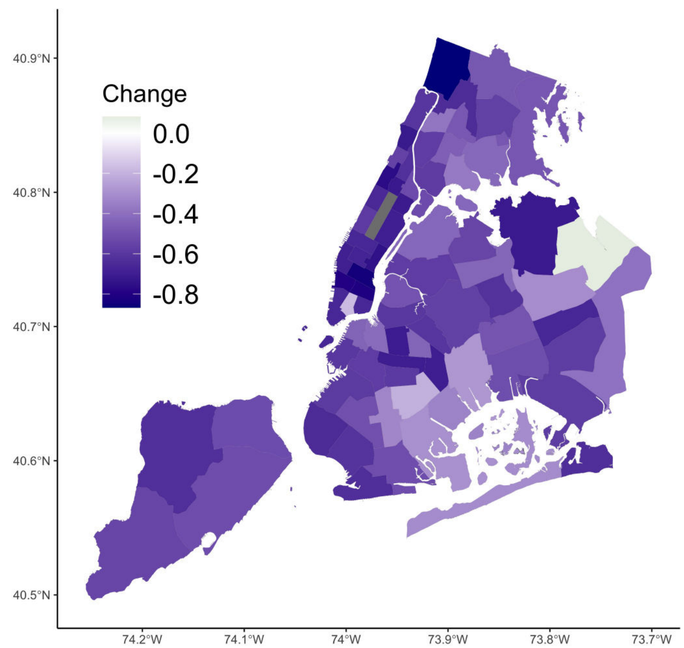

B. Annual change in major crimes against the annual change in criminal summonses

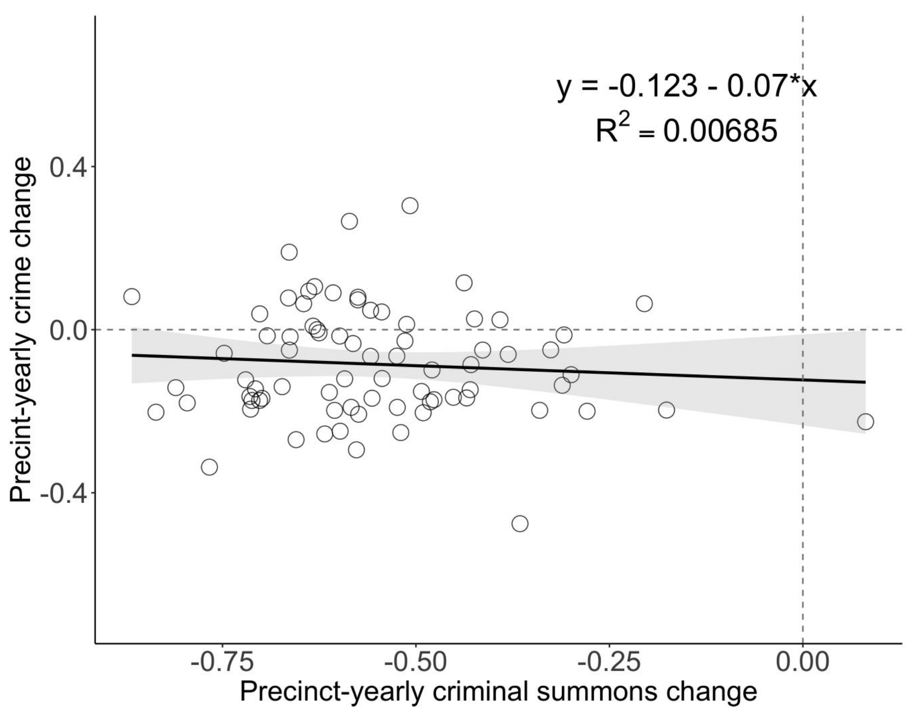

Note: Panel A presents a precinct-level heat map of NYC and documents enormous variation in the intensity of the slowdown across NYC communities. For each of NYC's 76 police precincts (excluding the Central Park precinct), the intensity of the color in the heat map corresponds with the year-over-year change in summonses issued during the December 1st-January 18th period. Panel B exploits the precinct-level variation and plots each precinct's percentage change in major crimes between the December 1st, 2013 - January 18th, 2014 and the December 1st, 2014 - January 18th, 2015 periods (y-axis) against the same percentage change in summonses issued (x-axis). A best-fit line is drawn through the data. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft.

ONLINE APPENDIX

In this supplementary appendix we present additional detail on our data and methods along with a series of auxiliary findings which are not central to our critique of S-O's paper but which present supporting evidence against their interpretation of the available data.

A Data and Methods

A.1 Data

We utilize publicly-available administrative records obtained from New York City's Open Data portal.[3] We use data on criminal court summonses and arrests to obtain measures of the extent of the work slowdown and crime complaint data to study the effect of the slowdown on crimes known to law enforcement. The crime, arrest and summons data contain time-stamped incident-level information that includes information on the type of incident and its geographic location. Following S-O, we classify crimes into "major" and "non-major" incidents. Major crimes include the seven Part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault (all of which are classified as violent crimes), and burglary, theft, and motor vehicle theft (defined as property crimes). Non-major crimes refer to all other offenses known to the New York City Police Department. The complaint dataset also has information about the specific location of occurrence and whether the crime occurred in a residential or commercial location or on the street. We use these variables to categorize the crimes as indoor (offenses committed inside a residential premise that could be either an apartment, house, or public housing) or outdoor offenses. To construct an analytic file, the data were aggregated at the precinct-day level to better capture changes in law enforcement activity during the slowdown window.

Finally, we collected weather information (including data on snowfall, precipitation, and temperature) from the National Oceanic and Atmospheric Administration Local Climatological dataset.

[3] See https://opendata.cityofnewyork.us/

We average the information from New York's LaGuardia and John F. Kennedy International Airports, and Central Park weather stations to estimate city-wide daily level data.

A.2 Econometric Strategy

We replicate the main results from S-O, supplementing their analysis with several auxiliary analyses which point to critical flaws in their analysis. The authors run an unusual differences-in-differences model in which the change in crime during the December 1st, 2014 - January 18th, 2015 slowdown period relative to the remainder of the calendar year is compared to the same change in crimes during in the previous year. Simply put, the authors identify whether crime was unusual during the slowdown period, relative to the same period in the prior year, accounting for annual changes in crime during the 2013-2015 period.

Following S-O, we fit the following regression equation, estimated using a negative binomial model, a count regression model which is often used to study count data. For ease of comparison, we describe the equation using the authors' original notation:

$$E[Y_{it}|S_{it}, T_{it}, X_{it}] = r_{it} = exp(\alpha + \gamma S_{it} + \lambda T_{it} + \delta(S_{it} \times T_{it}) + X_{it}^{'}\beta \quad (1)$$

In equation (1), r_{it} represents an outcome such as the count of summonses issued or the count of major crimes in a given daily-precinct. S_{it} is an indicator for the "series" — that is, whether a given year was treated by the slowdown. This variable is equal to 1 if a given day occurred during January 19th, 2014 to January 18th, 2015 period and 0 if otherwise. T_{it} is an indicator for the treatment window and is equal to 1 if a given day occurs during the December 1st-January 18th period. The coefficient on the interaction between the series and treatment window indicators, δ, provides an estimate of the differences-in-differences treatment effect. Formally this estimate tells us whether or not the difference in a given outcome between the treatment period and the remainder of the year for the 2014-2015 period differs from the difference in that outcome between the treatment period and the remainder of the year for the 2013-2014 period. We clustered the standard errors at the precinct level. Following S-O, we also exclude the Central Park Precinct from the regression models, but this decision is inconsequential to the magnitude of the point estimates.

B Supplementary Analyses

B.1 Spatial Variation in Treatment Intensity and Crime

In Figure 2 we documented considerable spatial variation in the intensity of the 2014-15 work slowdown. Here, we document variation in the annual change in crimes among NYC police precincts during the slowdown period. In Appendix Figure C.1, we document variation in the change in major "index" crimes. The same heat maps are presented for non-index crime arrests as non-index crimes in Appendix Figure C.2. A casual glance at the heat maps suggests that there is an imperfect correspondence between the intensity of the slowdown and the change in crimes. For instance, the community which experienced the largest slowdown — the Riverdale section of the Bronx — actually experienced one of the largest year-to-year *increases* in crime during the slowdown period. Likewise, the 123rd precinct located on Staten Island South Shore neighborhood, which experienced among the more intensive slowdowns also faced an *increase* in major crimes.[4]

B.2 Failure to Consider Testable Implications

The authors document a reduction in major crimes in NYC during the slowdown but the paper offers little explanation for why such a relationship might exist or why it would be causal. While a number of papers have noted that issuing large numbers of summonses and low-level arrests may

[4] In Appendix Figure C.3, we explore the nature of the variation by considering whether the intensity of the slowdown is related to pre-intervention crimes. Overall, there is little evidence that the slowdown was concentrated in either the safest or least safe areas of the city, with pre-intervention crime rates explaining under 2% of the variation in the year-over-year decline in summonses.

be relatively unproductive (Harcourt and Ludwig, 2006; MacDonald et al., 2016; Cho, Gonçalves, and Weisburst, Cho et al.), how is it that proactive policing would lead to an *increase* crime? We can think of several potential mechanisms. First, to the extent that proactive policing is a source of tension between police and the community, scaling back enforcement of low-level crimes may encourage greater cooperation with police investigations. While this is a possibility, given the long-standing "reservoir of discontent" (Rosenfeld, 2016) that has existed between police departments and lower-income Black communities for many years, we are skeptical that a seven-week slowdown is sufficient to measurably repair police-community relations. We further note that in NYC such a mechanism does not seem to have been at play as arrests for major crimes declined by 20% during the slowdown as compared to an 7% decline in major crimes. As such, it does not appear as though police were, through enhanced community cooperation, able to clear more crimes during this period. Second, it is possible that acrimonious contact with a police officer itself induces citizens to want to commit crimes. While this is theoretically possible, we have noted via Figure 1 there is considerable year-over-year variation in the issuing of summonses in NYC with no clear relationship between the yearly change in summonses and the yearly change in major crimes, a feature of the data that calls such a mechanism into question.

Finally, it is possible that when police officers are not spending their time issuing tickets for incivilities and booking arrestees charged with low-level "quality-of-life" crimes, they will spend more time deterring crime via routine patrols (Sherman and Weisburd, 1995; MacDonald et al., 2016; Weisburd, 2016; Braga et al., 2019) or by arresting offenders who have committed serious crimes. This mechanism suggests that scaling back proactive policing will have a disproportionate effect on crimes that occur outdoors or in a commercial setting as these are areas in which police are able to actively patrol. We separately consider indoor versus outdoor crimes in Appendix Figure C.4 which, like Figure 1, plots incidence rate ratios and 95% confidence intervals using equation

(1). Contrary to the hypothesis that outdoor crimes will have been more sensitive to the slowdown than crimes committed in areas which police cannot surveil, we see that the crime decline during the 2014-15 slowdown is, in fact, disproportionately driven by declines in *indoor* crimes. We view this evidence as being inconsistent with the hypothesis that the police slowdown is responsible for the decline in major crimes during the 2014-15 period.

B.3 Robustness to Alternative Treatment Period

In Appendix Figure C.5 we plot the daily number of criminal summonses issued city-wide before during and after the slowdown period. In Panel A, we focus the period between October 2013 and February 2016. In Panel B, we zoom in on the October 2014 - February 2015 period. In each panel, the dashed lines at December 1st, 2014 and January 18th, 2015 represent the slowdown period as defined by S-O. The dotted line at December 22nd, 2014 denotes an alternative slowdown period defined as the date of the alleged internal NYPD memo that encouraged officers to step back from proactive police work. Several features are worth noting. First, referring to Panel A, though there is predictable seasonal variation in summonses issues with summonses peaking in the summer and falling during the winter months, summonses issued reach a global minimum during the 2014-15 slowdown period. Second, referring to Panel B, while summonses continue to decline in early December, the majority of the decline in summonses actually occurs starting on December 22nd, there is a sharp decline in summonses issued corresponding with the circulation of the memo. The slowdown wanes beginning in mid-January, with enforcement activity returning to pre-slowdown levels on January 18th.

Appendix Figure C.5 calls into question whether December 1st is the best marker of the beginning of the slowdown. Accordingly, in Appendix Figure C.6, we re-analyze the data denoting the December 22nd-2014 - January 18th, 2015 period as the slowdown period. In Panel A, we verify that the year-over-year decline in summonses issued is even larger (IRR = 0.38) for this period. In the remaining panels, we consider the impact of the slowdown on overall major crimes (Panel B), indoor major crimes (Panel C) and outdoor major crimes (Panel D). Using the differences-in-differences analysis in (1), major crimes declined by 6% during the alternative slowdown period. However, five other year-over-year changes in major crimes were greater than 5% (four of which were significant at conventional levels), even though there were no other identified slowdowns during the same time period. Focusing on outdoor major crimes, the estimate for the slowdown period (IRR = 0.964) is no longer statistically significant and is only the 6th largest impact among 12 time periods studied. We likewise show, in Appendix Figure C.7 that, for the alternative slowdown period, there is no relationship between the intensity of the slowdown and the change in major crimes at the precinct level.

B.4 Results for Non-Index Crimes

Finally, we present auxiliary estimates for non-index crimes and non-index crime arrests. In Appendix Figure C.8 we present the analog to Figure 1 in the main text for non-index crime arrests (Panel A), non-index crimes (Panel B) and non-index crimes disaggregated to outdoor versus indoor locations (Panels C and D). In Appendix Figure C.9, we present the analog to Figure 2, in which we plot the change in each quantity at the police precinct level against the intensity of the slowdown, as measured by the annual change in the number of summonses issued and non-major crime arrests. Overall, precinct-level changes in criminal summonses during the alternative slowdown period explains very little of the variation in major crimes across communities, which goes against S-Os' central claim.

C Tables and Figures

Figure C.1: Annual change in major crimes by police precinct

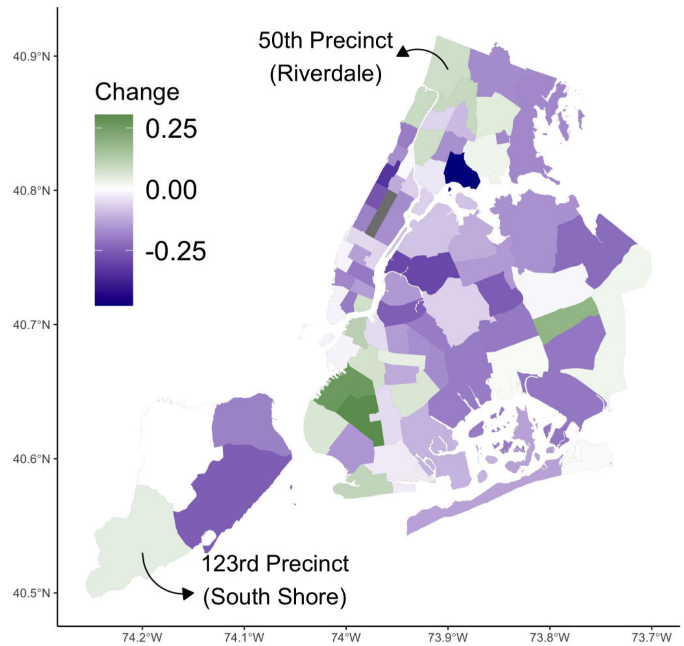

Note: The Figure presents the annual precinct-level change in major crimes between December 1st, 2014 to January 18th, 2015 and December 1st, 2013 to January 18th, 2014. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft.

Figure C.2: Annual change in non-major crimes and arrests by police precinct

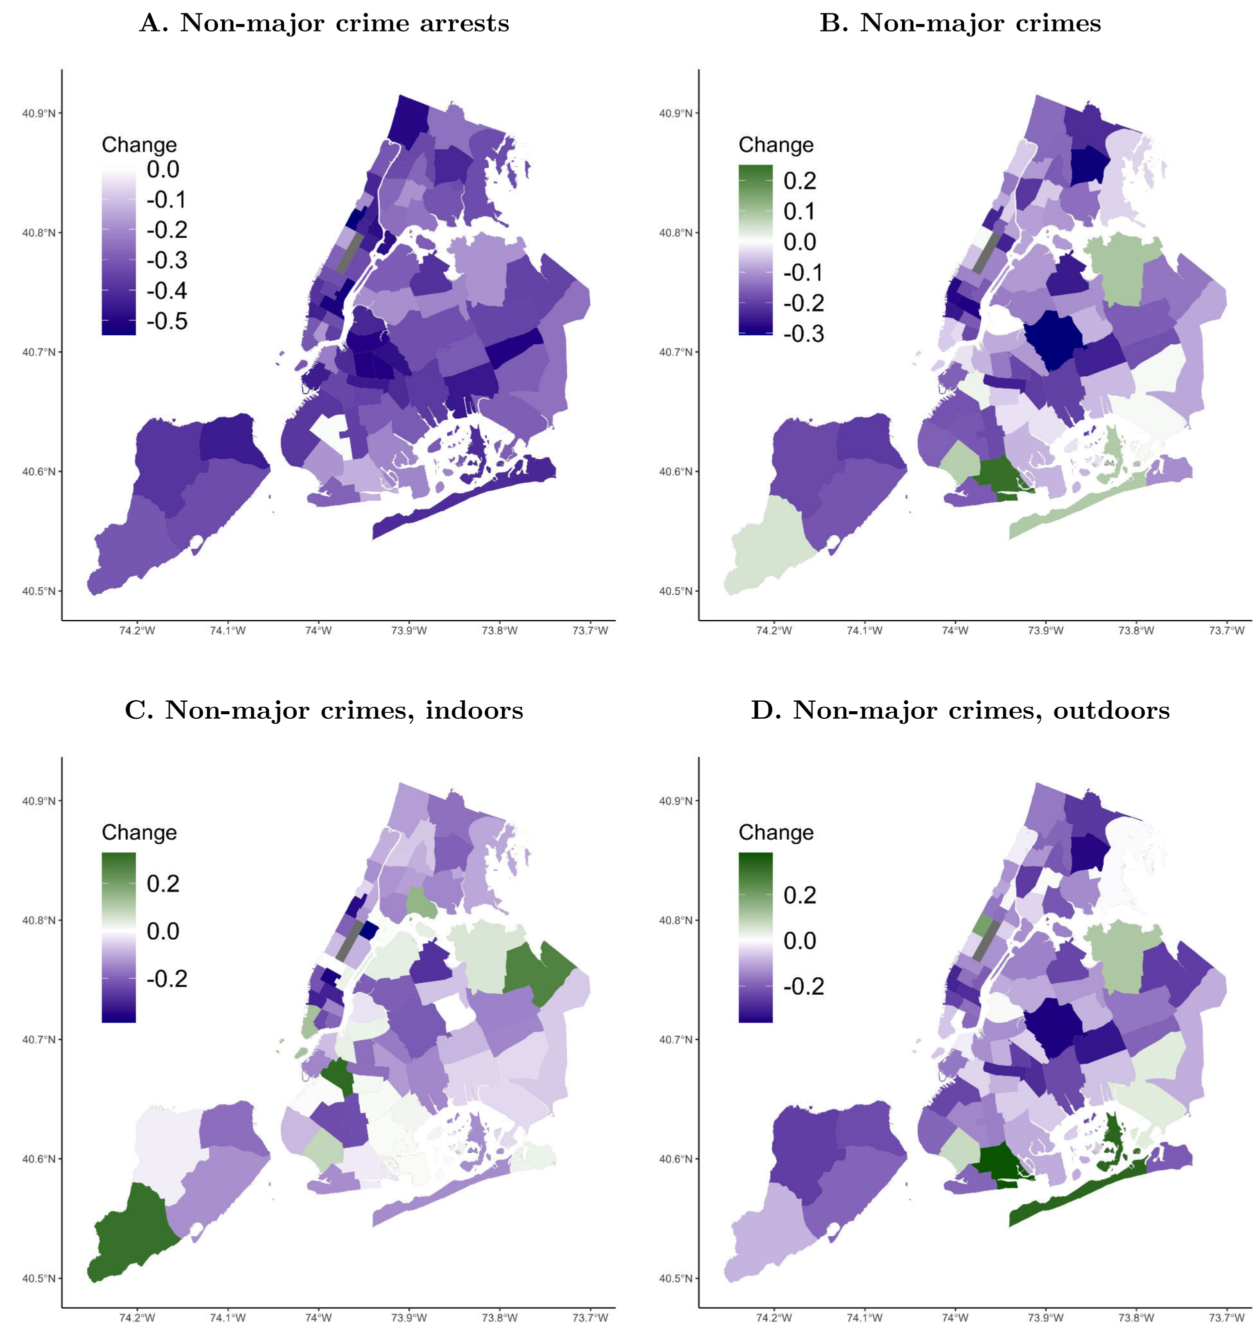

Note: The Figure presents the annual precinct-level change in non-major crime arrests (Panel A), non-major crimes (Panel B), indoor non-major crimes (Panel C) and outdoor non-major crimes (Panel D) between December 1st, 2014 to January 18th, 2015 and December 1st, 2013 to January 18th, 2014. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft. Non-major crimes refer to all the other crimes reported to the New York Police Department. Indoor crimes comprise offenses committed inside a residential premise (apartment, house, or public housing); outdoor crimes include all the other criminal offenses.

Figure C.3: Police slowdown and pre-intervention yearly crime rate per 1,000 people

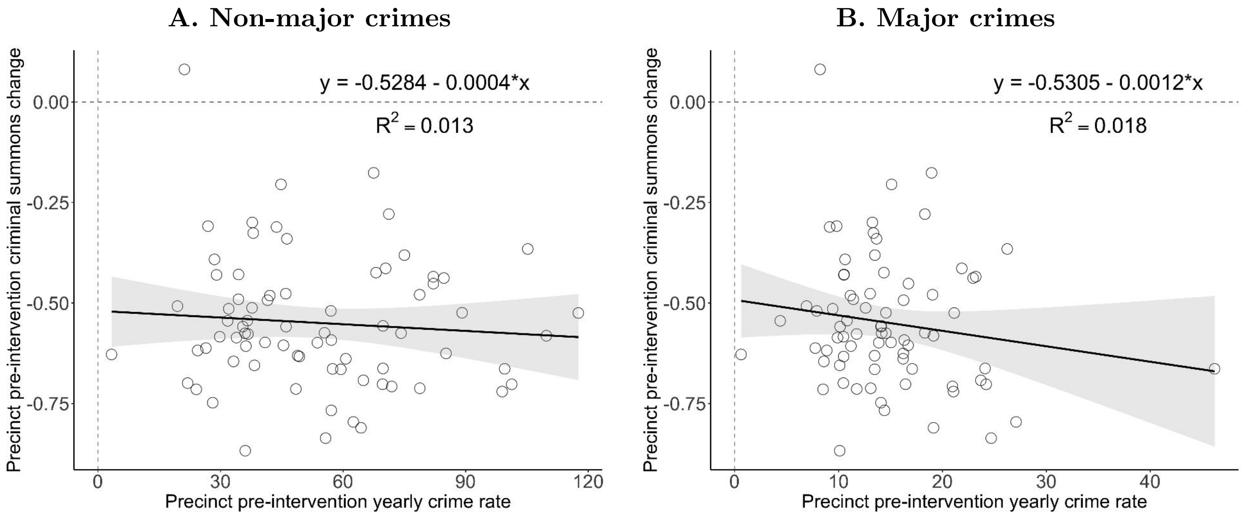

Note: Scatterplots plot the annual precinct-level change in criminal summonses (0.1 = 10%) between December 1st, 2014 to January 18th, 2015 and December 1st, 2013 to January 18th, 2014 against the mean pre-intervention (2006-2013) annual crimes per 1,000 persons in each police precinct. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft. Non-major crimes refer to all other crimes reported to the New York Police Department. Both panels exclude the data point of precinct 14 as it is an outlier (a 162.4 and 428.9 crime rate for major and non-major crimes, respectively, and a 67.2 percent criminal summons change), but it is consider in the regression line. We also exclude the Central Park Precinct following Sullivan and O'Keeffe (2017).

Figure C.4: Effects of NYPD slowdown on outdoor and indoor major crimes

A. Major crimes, indoors

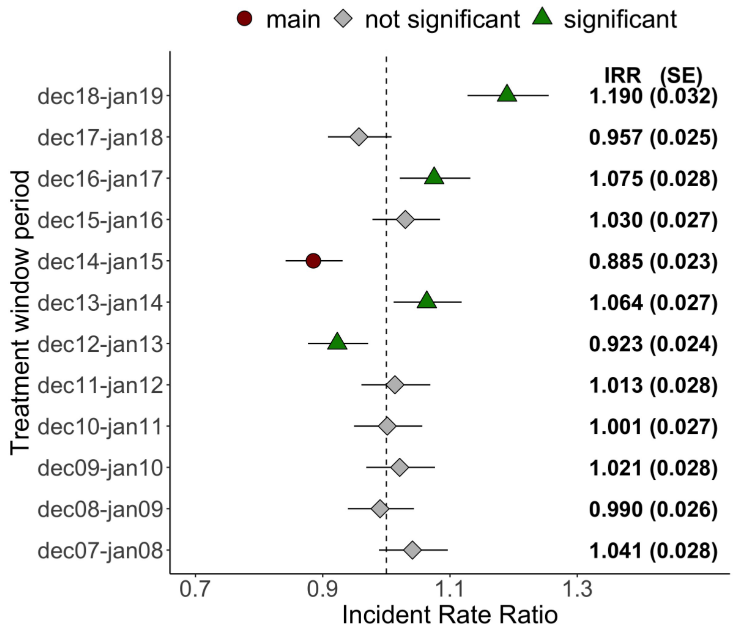

B. Major crimes, outdoors

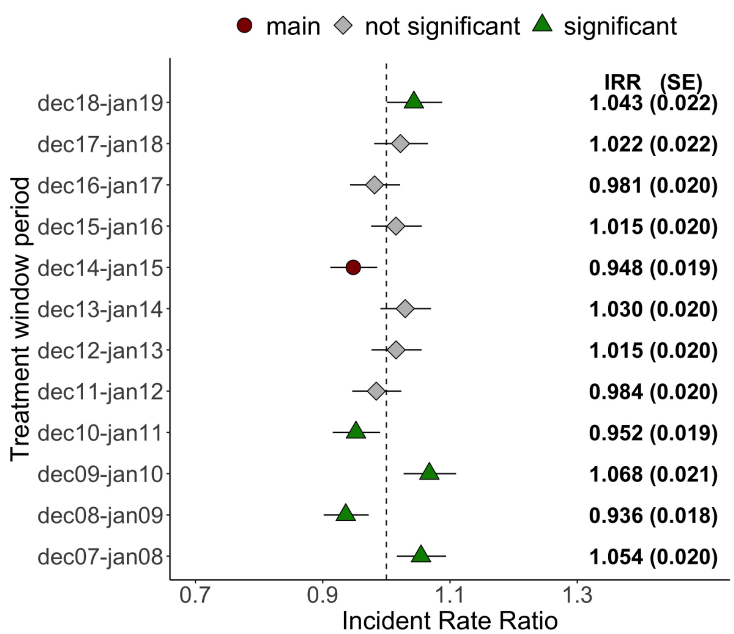

Note: Each row presents estimates from S-O's differences-in-differences estimator. The analysis identifies the effect of the slowdown by differencing the year-over-year change in the number of major crimes experienced during the December 1st-January 18th slowdown period from the year-over-year change in major crimes experienced during the remainder of the year. For each year t, year t-1 pair, we replicate S-O's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as S-O note, therefore do not influence the resulting estimates. For a given period, the estimating equation is as follows: $y_{it} = \alpha + \beta_1 S_{it} + \beta_2 T_{it} + \beta_3 S_{it} T_{it} + \gamma X_{it} + e_{it}$, where y_{it} is the outcome variable at the precinct i, day-month-year t, $S_{it} = 1$ for period t between January/19/YEAR+1 and January/18/YEAR+2, zero otherwise, $T_{it} = 1$ is the treatment window for period t between December/01/YEAR and January/18/YEAR+1 as well as December/01/YEAR+1 and January/18/YEAR+2, zero otherwise, and X_{it} represents weather data (snow, precipitation, and temperature) at the city-day-month-year level. Each model only includes data between January/19/YEAR to January/18/YEAR+2. For each pair of years, the figure reports incident rate ratios, standard errors clustered by police precinct and 95% confidence intervals from a series of negative binomial regressions. Panel A considers major crimes committed indoors; Panel B considers major crimes, committed outdoors. Indoor crimes comprise offenses committed inside a residential premise (apartment, house, or public housing); outdoor crimes include all the other criminal offenses.

Figure C.5: Daily city-wide criminal summonses

A. October 2013 to February 2016

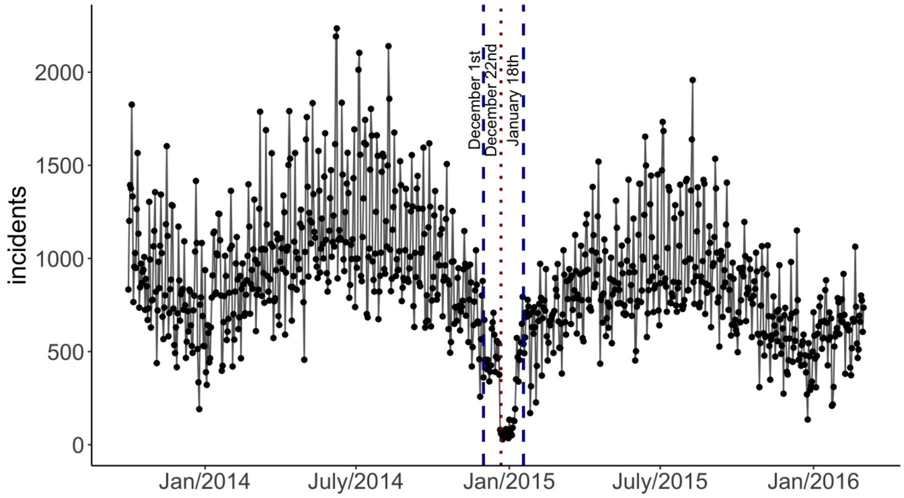

B. October 2014 to February 2015

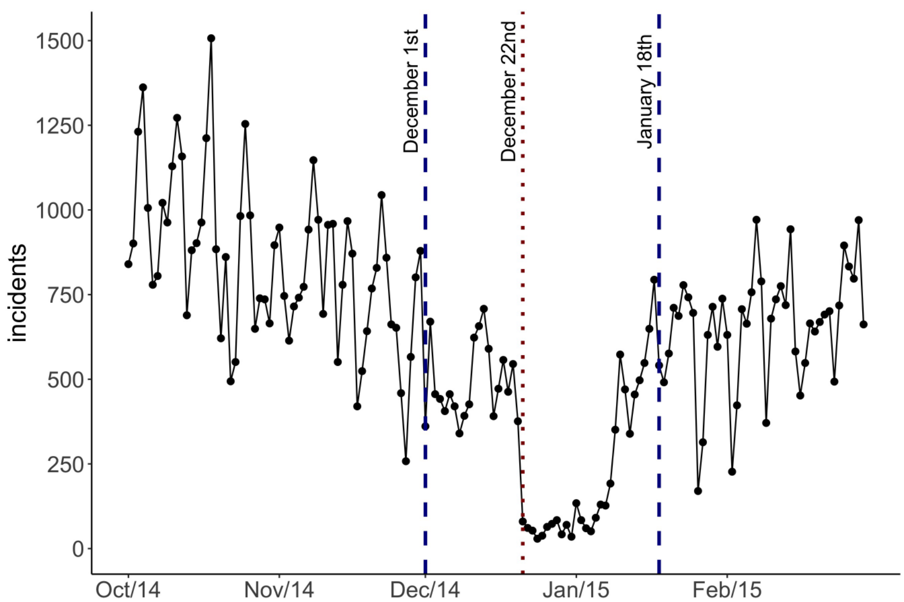

Note: Figures plot the daily number of criminal summonses issued by NYPD officers for a given time period. Panel A focuses on the October 2014-February 2016 period; Panel B focuses on the October 2014-February 2015 period. In each panel, the blue dashed lines mark the New York Police Department slowdown period used by Sullivan and O'Keeffe (2017), ranging from December 1st, 2014 to January 18th, 2015. The red dotted line highlights December 22nd, 2014, the date that several media outlets report as the starting period of the police slowdown, which coincides with a large decrease in criminal summons.

Figure C.6: Effects of NYPD slowdown on public safety using alternative September 22nd, 2014 - January 18th, 2015 treatment period

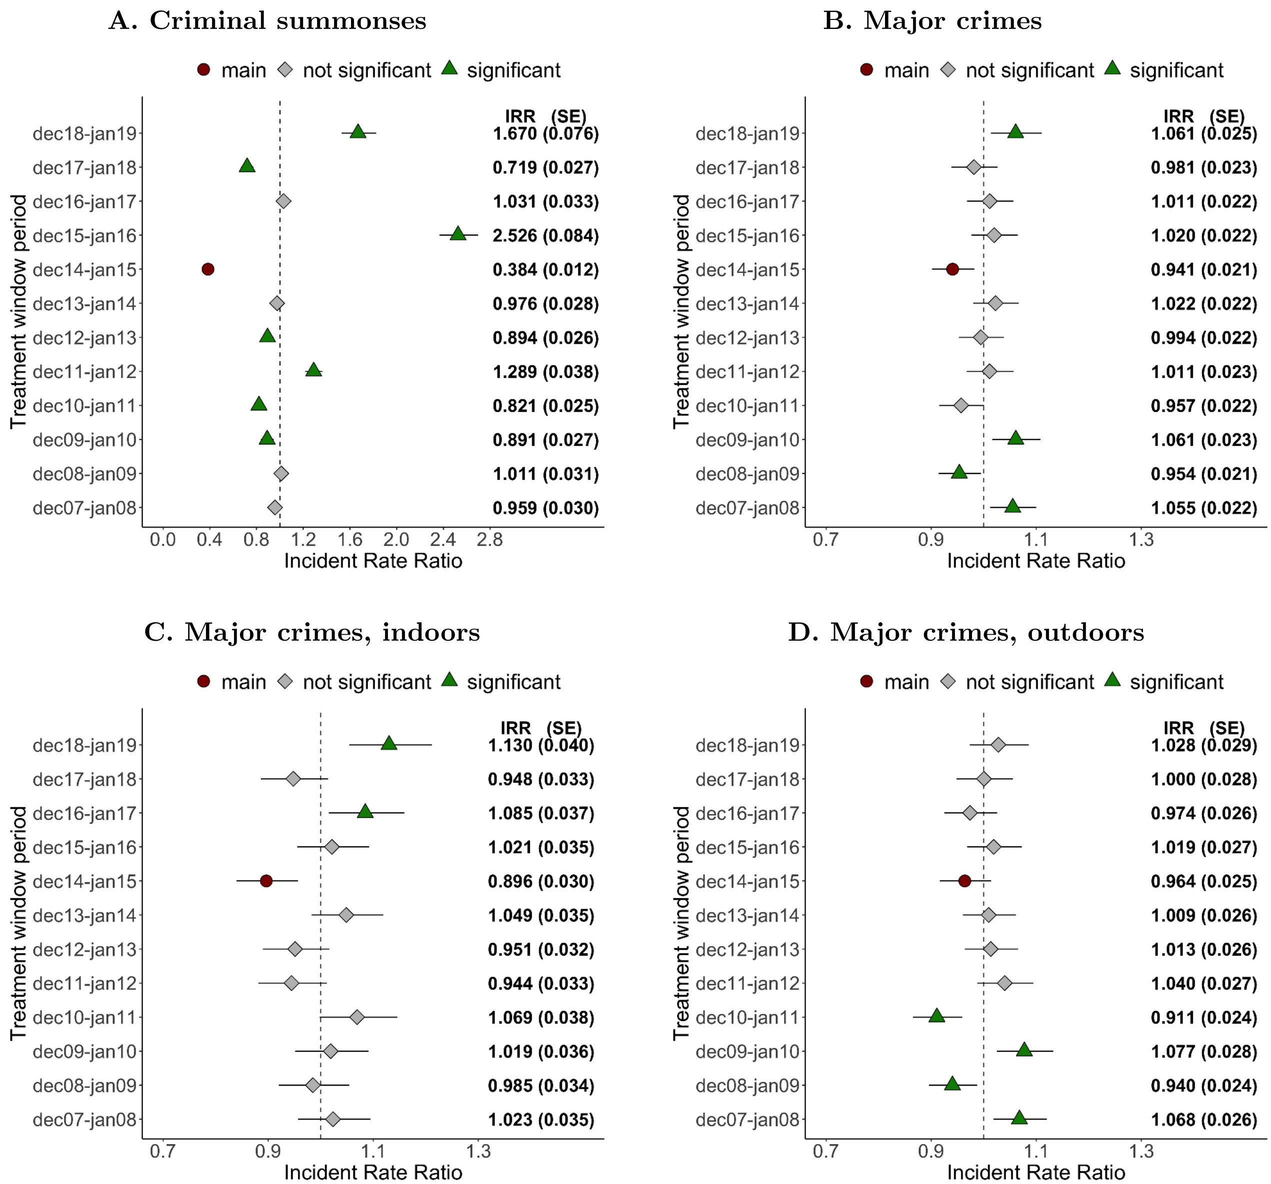

Note: Each row presents estimates from S-O's differences-in-differences estimator. For each year t, year t-1 pair, we replicate S-O's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as S-O note, therefore do not influence the resulting estimates. For a given period, the estimating equation is as follows: $y_{it} = \alpha + \beta_1 S_{it} + \beta_2 T_{it} + \beta_3 S_{it} T_{it} + \gamma X_{it} + e_{it}$, where y_{it} is the outcome variable at the precinct i, day-month-year t, $S_{it} = 1$ for period t between January/19/YEAR+1 and January/18/YEAR+2, zero otherwise, $T_{it} = 1$ is the treatment window for period t between December/01/YEAR and January/18/YEAR+1 as well as December/22/YEAR+1 and January/18/YEAR+2, zero otherwise, and X_{it} represents weather data (snow, precipitation, and temperature) at the city-day-month-year level. Each model only includes data between January/19/YEAR to January/18/YEAR+2. For each pair of years, the figure reports incident rate ratios, standard errors clustered by police precinct and 95% confidence intervals from a series of negative binomial regressions. Panel A considers criminal summonses, Panel B considers major crimes and Panels C and D consider major indoor and outdoor crimes, respectively. Indoor crimes comprise offenses committed inside a residential premise (apartment, house, or public housing); outdoor crimes include all the other criminal offenses.

Figure C.7: Precinct-level scatterplot of the change in major crimes against the change in criminal summonses using alternative September 22nd, 2014 - January 18th, 2015 treatment period

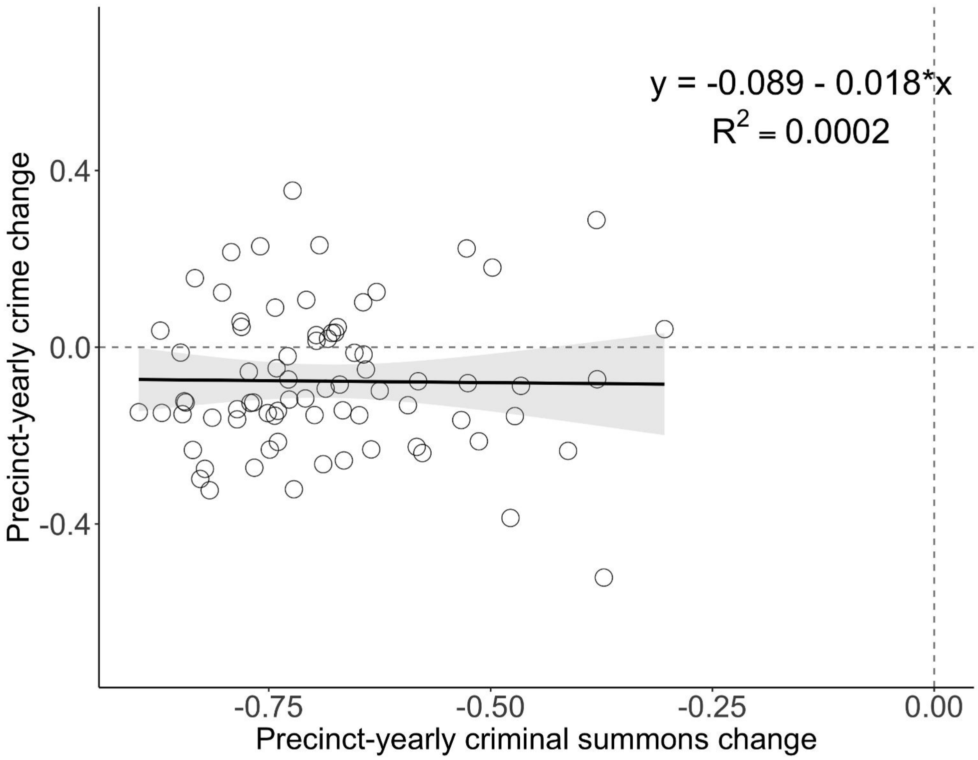

Note: Figure exploits the precinct-level variation and plots each precinct's percentage change in major crimes between the December 22nd, 2013 - January 18th, 2014 and the December 22nd, 2014 - January 18th, 2015 periods (y-axis) against the same percentage change in summonses issued (x-axis). A best-fit line is drawn through the data. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft.

Figure C.8: Effects of slowdown on public safety

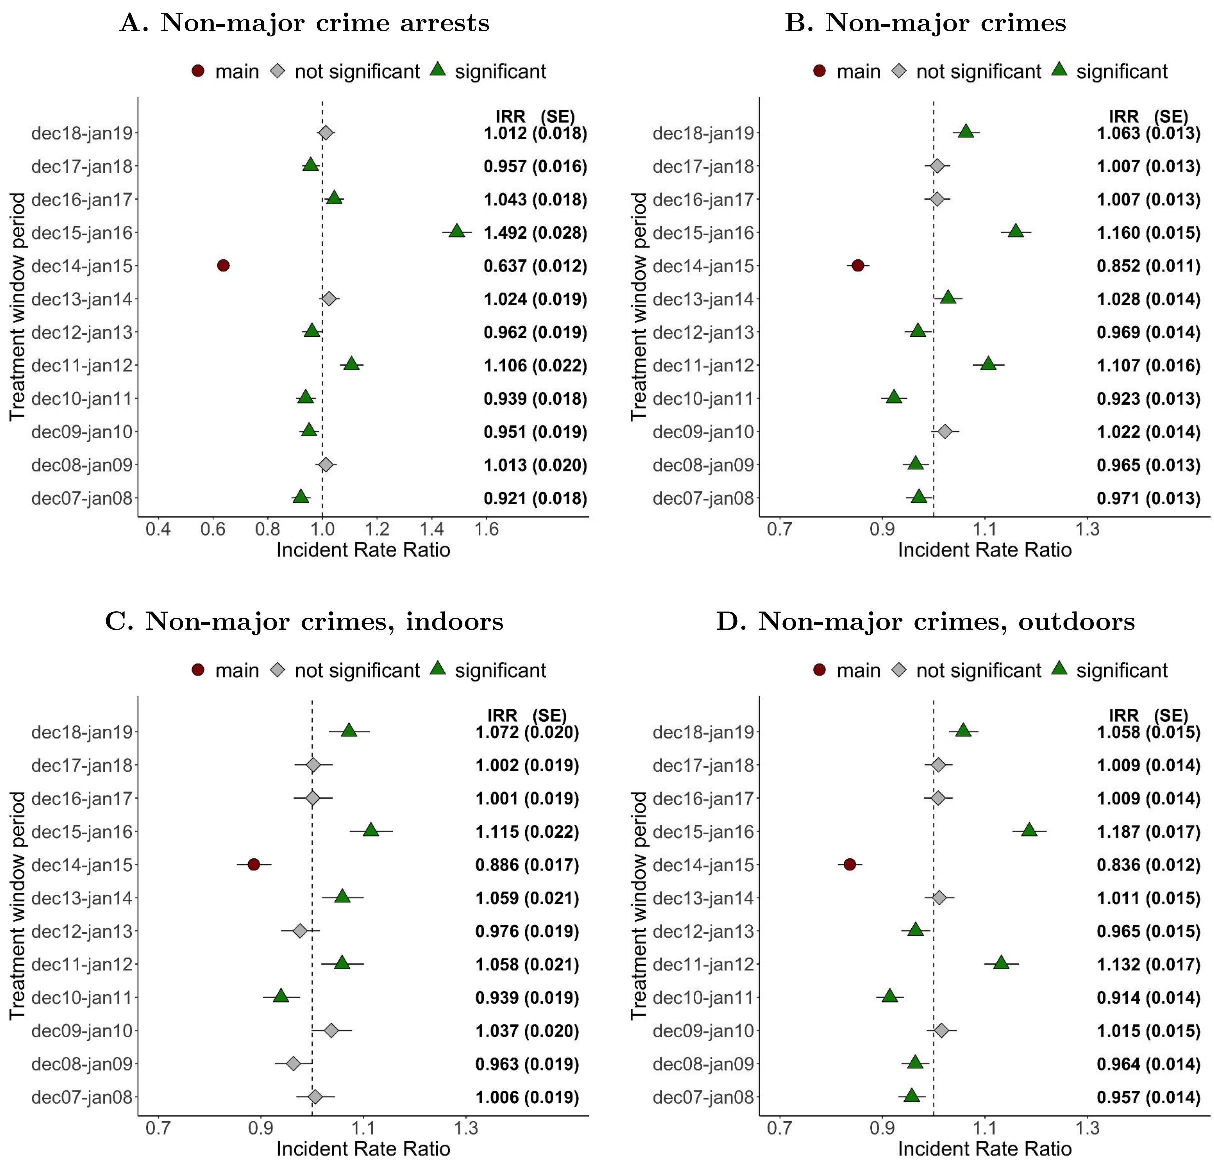

Note: Each row presents estimates from S-O's differences-in-differences estimator. For each year t, year t-1 pair, we replicate S-O's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as S-O note, therefore do not influence the resulting estimates. For a given period, the estimating equation is as follows: $y_{it} = \alpha + \beta_1 S_{it} + \beta_2 T_{it} + \beta_3 S_{it} T_{it} + \gamma X_{it} + e_{it}$, where y_{it} is the outcome variable at the precinct i, day-month-year t, $S_{it} = 1$ for period t between January/19/YEAR+1 and January/18/YEAR+2, zero otherwise, $T_{it} = 1$ is the treatment window for period t between December/01/YEAR and January/18/YEAR+1 as well as December/01/YEAR+1 and January/18/YEAR+2, zero otherwise, and X_{it} represents weather data (snow, precipitation, and temperature) at the city-day-month-year level. Each model only includes data between January/19/YEAR to January/18/YEAR+2. For each pair of years, the figure reports incident rate ratios, standard errors clustered by police precinct and 95% confidence intervals from a series of negative binomial regressions. Panel A considers non-major crime arrests, Panel B considers non-major crimes and Panels C and D consider indoor and outdoor non-major crimes, respectively. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft. Non-major crimes refer to all the other crimes reported to the New York Police Department. Indoor crimes comprise offenses committed inside a residential premise (apartment, house, or public housing); outdoor crimes include all the other criminal offenses.

Figure C.9: Precinct-yearly changes on non-major crime arrests and crimes

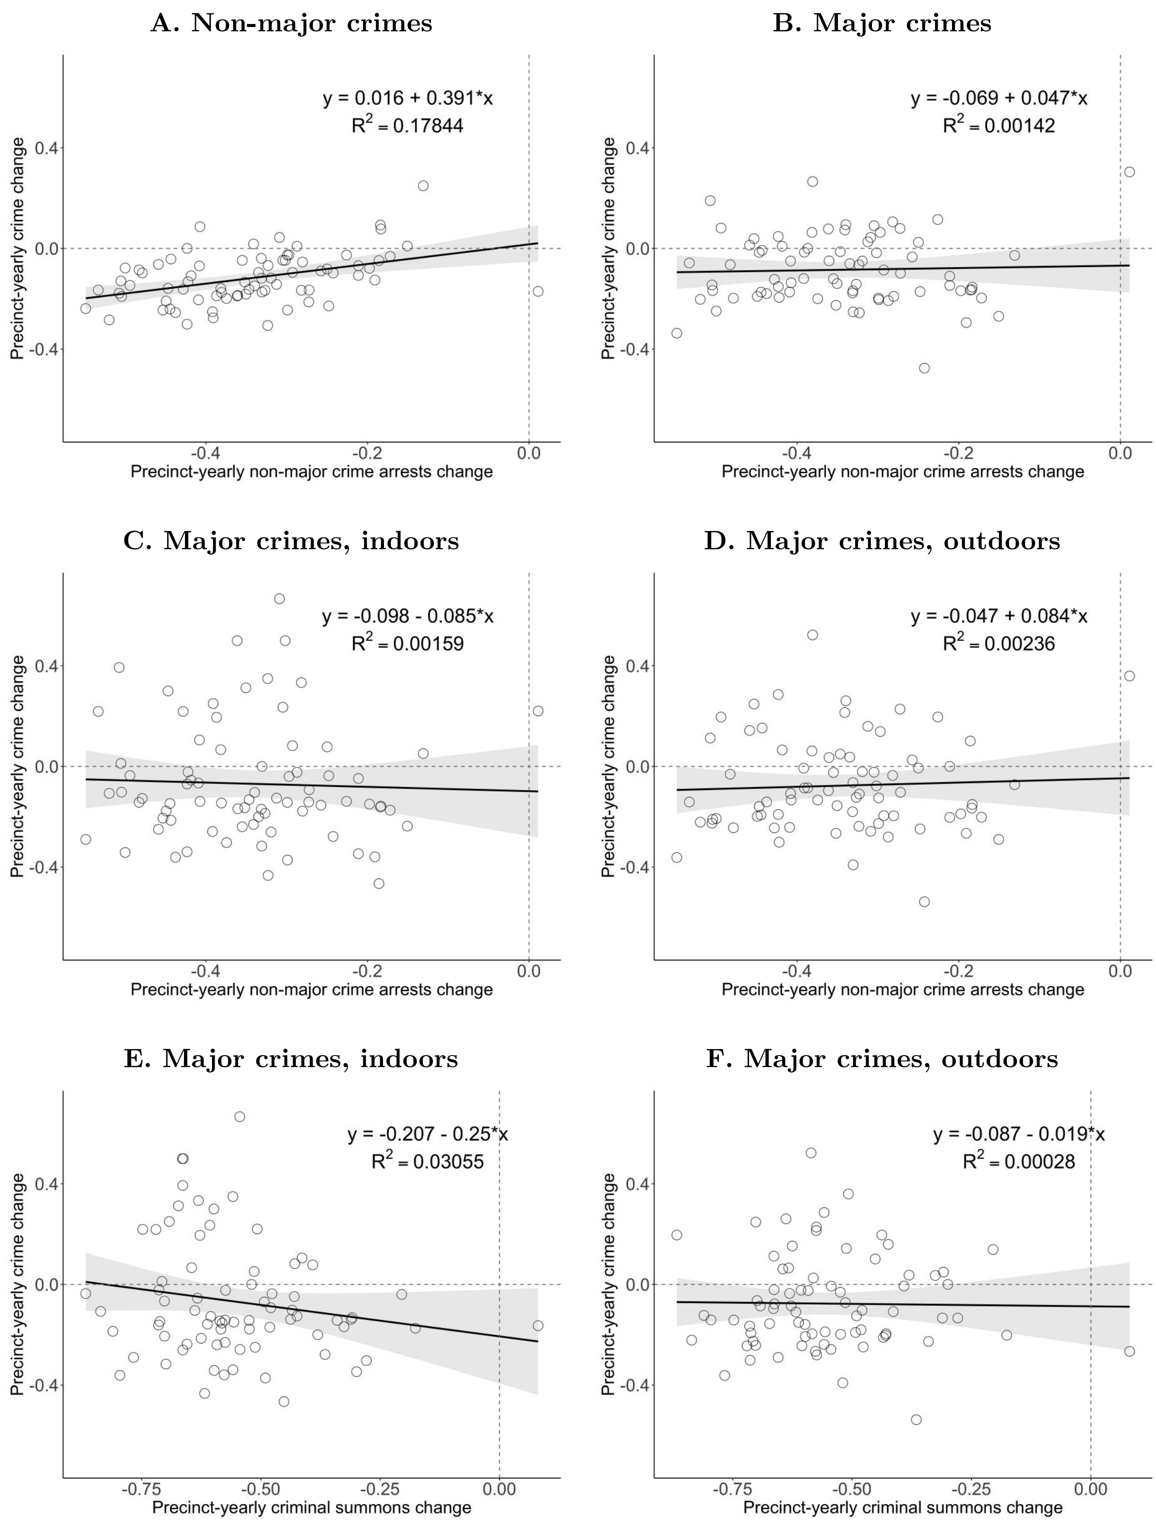

Note: Figure exploits the precinct-level variation and plots each precinct's percentage change in crimes between the December 1st, 2013 - January 18th, 2014 and the December 1st, 2014 - January 18th, 2015 periods (y-axis) against the same percentage change in summonses issued or arrests made (x-axis). A best-fit line is drawn through the data. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft. Panels A-D use the yearly change of non-major crime arrests as the horizontal variable, while Panels E-F use the change in criminal summonses. The vertical variable is the one presented in the panel subtitle.