

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

Aaron Chalfin¹, David Mitre-Becerril¹ and Morgan C. Williams, Jr.*²

¹Department of Criminology, University of Pennsylvania

²Department of Economics, Barnard College, Columbia University

November 13, 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. 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. Contrary to the remainder of the literature, the authors conclude that proactive policing can cause an increase in crime. 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., Barnard College, Columbia University, 1018 Milstein Center, New York, NY, 10027, E-Mail: mcwillia@barnard.edu.

1 Introduction

For decades, activists, policymakers, and social scientists have debated the role of police presence, particularly in lower income neighborhoods where crime tends to be most prevalent. While there is a consensus that the number of police officers (Evans and Owens, 2007; Chalfin and McCrary, 2018; Mello, 2019; Weisburst, 2019; Kaplan and Chalfin, 2019) combined with their presence and visibility (Sherman and Weisburd, 1995; Klick and Tabarrok, 2005; Braga and Bond, 2008; Draca et al., 2011; Machin and Marie, 2011; Braga et al., 2014; Groff et al., 2015; MacDonald et al., 2016; Mastrobuoni, 2019; Lovett and Xue, 2021; Weisburd, 2021) reduces serious crime, the strategies that the police have sometimes deployed, including directed patrol and the intensive use of field interrogations are also thought to create high costs for disadvantaged communities (Weitzer et al., 2008; Howell, 2009; Edwards et al., 2019; Ang, 2020; Shjarback and Nix, 2020; Bacher-Hicks and de la Campa, 2021). Research shows that while larger police forces save lives in the United States, they also make more arrests for low-level “quality-of-life” offenses, crimes which often do not have a specific victim and for which arrests accrue disproportionately to Black Americans (Chalfin et al., 2020).

Since the 1982 publication of *Broken Windows* by Wilson and Kelling and the subsequent expansion of proactive, order maintenance policing tactics in many U.S. cities, there has been considerable debate about the public safety value of making large numbers of arrests for low-level quality of life crimes. While the majority of early research, much of it using data from NYC, suggests that today’s misdemeanor arrests prevent tomorrow’s felony crimes (Kelling and Sousa, 2001; Corman and Mocan, 2005; Messner et al., 2007), a persuasive re-analysis by Harcourt and Ludwig (2006) and a litany of more recent scholarship (Braga and Bond, 2008; MacDonald et al., 2016; Caetano and Maheshri, 2018; Cho et al., 2021) calls such a conclusion into question.¹ At

¹These findings are likewise buttressed by recent research that documents the harmful effects of pre-trial detention

the same time, there is conflicting evidence on the effect of police “pullbacks” where police officers choose to engage in less proactive policing, with some research finding that police pullbacks do not substantially compromise public safety (Chandrasekher, 2016; Shjarback et al., 2017) and other research finding that when the police pull back, increases in crime (Mas, 2006; Shi, 2008) and gun violence (Cheng and Long, 2018; Devi and Fryer Jr, 2020) can follow.²

In a 2017 paper published in the highly influential general interest journal, *Nature Human Behaviour*, Sullivan and O’Keeffe (hereafter referred to as “S-O’K”) assess the effect of a large and discrete reduction in proactive policing in NYC during a seven week period. Their research design is motivated by a coordinated “slowdown” of police work undertaken by NYC police officers in response to a series of events that shook their faith in the city’s political leadership. S-O’K 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.”³ This is an extraordinary claim insofar as it does not merely suggest that order maintenance policing has no discernible impact on public safety but that it directly compromises public safety. So far as we are aware, this is the only empirical paper published in the last twenty years that reports such a finding. Given the novelty and importance of the finding, it is perhaps unsurprising that it is published in a prominent journal.

In this paper, we re-evaluate this claim and point out several fatal weaknesses in the authors’ analysis that calls their principal finding into question. We begin by noting that the authors’

(Leslie and Pope, 2017; Dobbie et al., 2018) and prosecution for low-level misdemeanor crimes, especially for first-time offenders (Agan et al., 2021).

²Examining a well-documented *nine-month* NYPD slowdown motivated by a collective bargaining dispute, Chandrasekher (2016) finds evidence of a sharp decline in ticket-writing by NYPD officers and mixed evidence suggestive of an increase in felonious assault and larceny.

³The authors further state: “Our results imply not only that these tactics fail at their stated objective of reducing major legal violations, but also that the initial deployment of proactive policing can inspire additional crimes that later provide justification for further increasing police stops, summonses and so on. The vicious feedback between proactive policing and major crime can exacerbate political and economic inequality across communities.”

“differences-in-differences” analysis uses only two years of data, comparing whether crime was unusual during the slowdown period relative to the same period in the prior year, accounting for annual crime trends. Since the entire city was subject to the police slowdown, the authors’ research design does not follow the canonical difference-in-differences framework which requires both a time dimension and a unit that is unexposed (or at least less exposed) to an intervention. The S-O’K analysis is instead a variant of a simple pre-post study design.

We extend the authors’ analysis, collecting the same public microdata they used in their 2017 paper for the 14-year period between 2006 and 2019. We then note a number of troubling issues that point to critical flaws in their analysis. While we confirm S-O’K’s finding that the decline in major crimes during the 2014-15 slowdown period is statistically significant at conventional levels, a closer inspection of the data suggests that the result may be spurious. While the estimate for the 2014-15 slowdown period is the most negative estimate among the twelve year-over-year pairs in the data, 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. In other words, the result does not survive a simple placebo test.

Most troubling though is that the authors have ignored the considerable variation in the intensity of the work slowdown across NYC communities. Leveraging police precinct-level as well as U.S. Census tract-level variation, we show that while some communities experienced an enormous decline in police enforcement during the slowdown period, other NYC communities experienced only a modest slowdown.⁴ Critically and contrary to the paper’s central claim — that the slowdown

⁴While we cannot be certain about the underlying reasons for this variation, we note that communities that were more exposed to the slowdown do not differ from less exposed communities with respect to baseline crime or

caused major crimes to decline — the communities which experienced a much more pronounced reduction in police enforcement *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.

Finally, we point to several additional testable implications of the authors’ analysis that are inconsistent with the conclusion reached in the paper. In particular, 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. We show that the slowdown does not appear to have re-allocated police attention to clearing major crimes. Moreover, the largest reductions in crime observed during the slowdown occurred for crimes committed inside people’s homes, areas which police officers are unable to actively patrol. These findings are inconsistent with two of the most likely mechanisms for a slowdown in proactive policing to have beneficial effects on public safety.

Having re-analyzed the data, we conclude that there is little evidence that the NYPD’s 2014-2015 work slowdown caused a decline in major crimes. Instead, the evidence is more consistent with the finding that major crimes did not change appreciably as a function of the slowing down of police work. While the degree to which this is a useful natural experiment to understand the longer-term effects of a reduction in so-called proactive policing remains uncertain, it is worth noting that during a seven-week period in which the scale of policing was dramatically reduced, the sky did not fall. This is, in our view, a considerably more evidence-based interpretation of the authors’ own data.

The remainder of this paper is organized as follows. In Section 2, we provide a brief description

pre-intervention crime trends.

of the natural experiment exploited in S-O’K’s original paper. In Section 3, we describe the public microdata used to study the effects of the NYPD work slowdown and summarize our empirical strategy which, to maintain consistency, is a replication of that of S-O’K. In Section 4, we present the results of our new investigation and note three curiosities which call into question S-O’K’s central finding — that a short-term reduction in the use of proactive policing curtailed major crimes. Section 5 concludes.

2 Institutional Background

The natural experiment studied in S-O’K’s original contribution and our re-analysis is motivated by a coordinated “slowdown” of police work undertaken by NYC police officers in response to a series of events that shook their faith in the city’s political leadership.⁵ On December 4th, 2014 a New York grand jury declined to indict Daniel Pantaleo, the NYPD officer who used a choke hold to kill an unarmed black man named Eric Garner in July 2014. Following the grand jury’s decision, there was a series of large-scale protests against police brutality at the Brooklyn Bridge, Manhattan’s West Side Highway, and several other sites throughout the city. In response to the perception among some police officers that these protests were supported by NYC’s Mayor Bill de Blasio, NYPD officers began to systematically scale back their enforcement activity, issuing fewer summons and making fewer arrests for misdemeanor crimes. Later that month, two NYPD officers, Wenjian Liu and Rafael Ramos, were murdered by a gunman in a targeted attack, leading to concern among officers for their safety and further entrenching officers in their position. The night that Officers Ramos and Liu were killed, there are reports that an e-mail was circulated among NYPD officers urging them not to make arrests or issue summonses “unless absolutely necessary”

⁵Research does not support the idea that danger has increased for police officers in recent years (Maguire et al., 2017) though there is evidence that some police officers may pull back when they perceive that public opinion has changed (Wolfe and Nix, 2016). For an excellent discussion of police officers’ perceptions regarding whether there is a “war on police” see Nix et al. (2018).

(Lind, 2015).⁶ While the source of the e-mail appears to remain unresolved, on December 22nd, two days after the killing of the officers, the number of summonses issued and misdemeanor arrests fell even further.⁷

In **Figure 1** we plot the daily number of criminal summonses issued city-wide before, during and after the slowdown period. In Panel A, we focus on 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 18, 2015 represent the slowdown period as defined by S-O’K. The dotted line at December 22, 2014 denotes an alternative slowdown period defined using 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, peaking in the summer and falling during the winter months, reaching a global minimum during the 2014-15 slowdown period. Second, referring to Panel B, while summonses do decline in early December, the majority of the decline in summonses actually occurs starting on December 22nd, corresponding with the circulation of the union memo. The slowdown wanes beginning in mid-January, with enforcement activity returning to pre-slowdown levels on January 18th. As **Figure 1** calls into question whether December 1st is the best marker of the beginning of the slowdown, we also provide an auxiliary analysis using the December 22, 2014 - January 18, 2015 alternative slowdown period.

⁶As documented by [Ba and Rivera \(2019\)](#), such memos can indeed have large effects on police behavior.

⁷Some sources attribute the e-mail to the NYPD Patrolmen’s Benevolent Association (PBA). However, the PBA denied the allegation.

3 Data and Methods

3.1 Data

This research uses publicly-available administrative records obtained from New York City’s Open Data portal.⁸ 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’K, 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. 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.

⁸See <https://opendata.cityofnewyork.us/>

3.2 Econometric Strategy

We begin by replicating the main results from S-O’K, 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 1, 2014 - January 18, 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. Critically, there is no control group that is unexposed (or less exposed) to the treatment.

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

$$E[Y_{it}|S_t, T_t, X_t] = r_{it} = \exp(\alpha + \gamma S_t + \lambda T_t + \delta(S_t \times T_t) + X_t' \beta) \quad (1)$$

In (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_t 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_t 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. Standard errors are clustered at the precinct level.⁹

4 Re-Analysis

In our re-analysis, we note three curiosities which call into question S-O’K’s central finding — that a short-term reduction in the use of proactive policing curtailed major crimes. We detail each of these items below.

4.1 Inconsistent Evidence and Failed Placebo Tests

Using the approach described in equation (1), S-O’K find that major crimes declined by approximately 6% during the 2014-15 slowdown relative to the same period in the previous year. In **Figure 2** we replicate their analysis for all $t, t-1$ pairs during the 2007-2019 period.¹⁰ The figure reports incident rate ratios, standard errors clustered by police precinct and associated 95% confidence intervals from the negative binomial regression equation outlined in (1). Estimates for the 2014-15 period are presented using red circles. For the other periods, estimates which are significant at the $\alpha = 0.05$ level are presented using green triangles; estimates that are not statistically significant are presented using gray rhombuses.

In Panel A, we consider the issuance of criminal summonses, the primary marker of the work slowdown, though we reproduce the same figure using arrests for non-major crimes in **Appendix Figure A.1**. In Panel B, we consider S-O’K’s primary outcome: major crimes known to law enforcement.¹¹ Consistent with S-O’K, the 2014-15 slowdown appears to have reduced criminal summonses by approximately 50% when compared to the previous year. Happily, we are able to replicate S-

⁹Following S-O’K, we also exclude the Central Park Precinct from the regression models, but this decision is inconsequential to the magnitude of the point estimates.

¹⁰We exclude 2020 due to the large disruptions to social and economic life caused by the COVID-19 pandemic.

¹¹Auxiliary results for non-major crimes are available in **Appendix Figure A.1, Panel B**.

O’K’s finding that major crimes declined by approximately 7% during the 2014-15 slowdown period relative to the same period in the prior year, an estimate that is statistically significant at conventional levels.

A more careful review of Figure 2 however reveals several troubling issues. We begin by noting that substantively important year-over-year changes in summonses issued by the NYPD 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 year-over-year changes in summonses issued in 10 out of the 11 years for which we have data. Some of these changes are qualitatively large. For example, relative to the same period in the prior year, there was a 29% decrease in summonses issued during the December 1st, 2017-January 18th, 2018 period and a 56% *increase* in summonses issued during the December 1st, 2018-January 18th, 2019 period. We likewise observe a 91% increase in summonses issued during the 2015-16 period which is itself an artifact of the 2014-15 slowdown which depressed police activity during the prior year.

The fact that large year-over-year changes in summonses are common motivates a more systematic inquiry into the relationship between police enforcement and major crimes during the December 1st-January 18th period. Referring to Figure 2, we observe that when summonses fell by 29% during the 2017-18 period, there was, in fact, no change in major crimes. Similarly, when summonses issued increased by more than 90% in 2015-16, we also do not see any resulting changes in major crimes in NYC. In both cases, standard errors are small, easily allowing us to rule out an 7% decline in major crimes. Contrary to the claim made by S-O’K, the figure 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’K’s result. While the decline in major crimes during the 2014-15 slowdown period is statistically significant at conventional levels ($p <$

0.01), a closer inspection of the figure reveals a troubling issue. While the estimate for the 2014-15 slowdown period is the most negative estimate among the twelve year-over-year 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 may just as easily be an artifact of ordinary year-to-year variation in major crime rates in a single city during a seven-week period. These results call into the question the identification strategy employed by S-O'K while also raising additional concerns regarding the authors' ability to credibly disentangle alternative determinants of summons activity from the slowdown itself.¹²

4.2 An Ecological Fallacy

The authors' analysis is notable in that it draws a city-level inference, ignoring substantial sub-city variation in the intensity of the work slowdown. In panel A of **Figure 3**, we document variation in the intensity of the slowdown using a police precinct-level heat map of the five boroughs of NYC.¹³ 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. We see that while some communities experienced a sizable 80% year-over-year reduction in summonses issued, other communities experienced only a modest reduction of 10-20%. In **Appendix Figure A.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

¹²For example, [Chandrasekher \(2016\)](#) notes that there is some evidence that both administrative changes to the NYPD (e.g., the formal merging of the New York City Transit and Housing Authority Police Departments into the NYPD) and the introduction of broken windows policing under the Giuliani administration led to important shifts in enforcement activity preceding the 1997 slowdown. The absence of a valid control group within the S-O'K empirical strategy makes identification in this setting less clear.

¹³The same heat maps are presented for non-major crime arrests and non-major crimes in **Appendix Figure A.2**.

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. Using an event study, we likewise show in **Appendix Figure A.6** that communities with higher and lower than median exposure to the slowdown experienced similar crime trends prior to the slowdown.

Given the considerable spatial variation in the intensity of the slowdown, we note that, in analyzing the slowdown at the city-level, 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. If the authors' central claim — that the slowdown caused major crimes to decline — is correct, then we should observe that the police precincts which experienced the most intense slowdowns should also have experienced the largest declines in crime. We present a more systematic analysis of the relationship between the intensity of the slowdown and the change in major crimes in **Figure 3**, Panel B which, using precinct-level data, plots the year-over-year change in major crimes during the 2013-2014 to 2014-2015 SO study period (y -axis) against the year-over-year change in summonses issued (x -axis). A best-fit line is drawn through the 76 data points.¹⁴

If crime had fallen most dramatically in the police 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.¹⁵ Indeed, contrary to the paper's central claim — that the slowdown caused major crimes to decline — the communities which experienced a more pronounced reduction in police proactivity did not

¹⁴In keeping with prior literature using NYPD data, we exclude the NYPD's un-numbered Central Park precinct as it reports very few crimes.

¹⁵The figure is plotted using percentage changes rather than in raw numbers of crime per thousand residents. We do so because there is considerable variation in the applicability of population as a denominator in NYC. In particular, NYC has a number of neighborhoods with few residents but which receive an enormous number of daily visitors/tourists. One in particular, Times Square, becomes an outlier in an analysis that focuses on December and January, the peak season for tourism in NYC. To see this, we present **Appendix Figure A.4** which shows that excluding the precinct that includes Times Square with around 20,000 residents but a daily pedestrian count of more than 450,000 people, makes a non-existent correlation between criminal summons and major crime rates appear to be large and significant.

experience the largest reductions in major crime. Instead, the almost zero correlation between these two variables means that annual precinct-level changes in criminal summonses explains very little the variation in the change in major crimes.^{16,17}

4.3 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, if it does, why it would be causal. While numerous papers have noted that issuing large numbers of summonses and low-level arrests may be relatively unproductive (Harcourt and Ludwig, 2006; MacDonald et al., 2016), how is it that proactive policing would lead to an *increase* in major crimes? 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 longstanding “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 creates frustration and induces citizens to want to commit crimes. While this is theoretically possible, there is considerable year-over-year variation in the issuing of summonses in NYC with no clear relationship between

¹⁶We present the same analysis for non-major crimes as well as major crimes, disaggregated into indoor and outdoor crimes in **Appendix Figure A.5**.

¹⁷**Appendix Figure A.6** shows that precincts with larger versus smaller than median changes in criminal summonses during the 2014-15 slowdown follow parallel crime trends prior to the slowdown, providing some reassurance that the results presented in the scatterplots are not an artifact of gravitation to the mean.

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 serious crime via routine patrols and maintaining visibility (Sherman and Weisburd, 1995; Braga et al., 2014; MacDonald et al., 2016; Weisburd, 2016). 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 test this mechanism by separately considering indoor versus outdoor crimes in **Figure 4** which, like Figure 2, plots incidence rate ratios and 95% confidence intervals using equation (1). Contrary to the hypothesis that outdoor crimes would 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 police were able to successfully deter more outdoor offending during the slowdown.

4.4 Robustness

Finally, we consider the several additional issues that shed light on the general robustness of S-O’K’s analysis of this natural experiment. We begin by re-considering the treatment period studied by S-O’K. Next, recognizing the importance of the considerable degree of sub-city variation in the intensity of the police slowdown, we consider a much more granular unit of analysis than the police precinct: the U.S. Census tract.

4.4.1 Robustness to Alternative Treatment Period

In Figure 1, we established that the majority of the slowdown’s impact, in fact, occurred after December 22nd, 2014. Accordingly, in **Figure 5**, 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 authors’ 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 A.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.

4.4.2 Robustness to an Alternative Unit of Analysis

In NYC, police precincts are responsible for maintaining public safety within a designated area, making them a reasonable unit of analysis for this research. Still, there could be concerns about whether the results of our analyses would be different if we were to use a more granular unit of geography. This concern, which is sometimes referred to as the “modifiable areal unit problem,” is common in studying spatial crime data ([Bernasco and Elffers, 2010](#); [Ratcliffe, 2010](#)).

To address this concern, we repeat our analysis at the U.S. Census tract level, leveraging 2,163 geographic units (each of which is home to approximately 4,000 people) rather than 76 units using police precincts. Following our main analysis, **Appendix Figure A.8** presents estimates for each

year t , year $t-1$ pair at the Census tract level. Both the magnitude of the estimate and the number of statistical significant estimates are similar to the main results at the precinct-level. Similarly, **Appendix Figure A.9** exhibits the variation of the police slowdown and crimes using Census tract percent changes. As in our main results, the relationship between criminal summons and major crimes is very close to zero ($R^2 = 0.0005$).

5 Discussion

Since the 2017 publication of S-O’K’s article, there has been a proliferation of public concern about police use of force and the collateral consequences of police enforcement in low-income, predominantly Black communities. At the same time, since the beginning of the COVID-19 pandemic and the resulting protests in the aftermath of the murder of George Floyd by a Minneapolis police officer, there has been a large increase in gun violence in U.S. cities which has, in some cases, coincided with a dramatic reduction in police proactivity. Lawmakers and members of the public are actively engaged in an ongoing debate about the public safety returns to police resources and whether bringing those resources to bear to make low-level arrests is worthwhile given that both the opportunity cost of a police officer’s time and the collateral costs of this style of policing may be high.¹⁸ As such, the importance of studying the public safety value of order maintenance policing remains as high as ever.

The public safety value of the marginal low-level arrest is a key policy estimand that has powerful implications for how police resources should be deployed and which styles of policing should be incentivized by public officials. However, because the volume of low-level arrests typically rises and falls for reasons that are endogenous to public safety, good research designs remain difficult to come

¹⁸Recent debates have referred to the possibility that other policy inputs could potentially be substitutes for policing including investments in mental healthcare (Deza et al., 2020; Jácome, 2020), improvements to the built environment (Chalfin et al., 2021; Branas et al., 2016, 2018; Macdonald et al., 2021) and improvements to social programs (Heller, 2014; Heller et al., 2017; Sharkey et al., 2017; Kessler et al., 2021).

by. Because police slowdowns sometimes happen for reasons that are unrelated to public safety and because they can lead to extraordinarily large changes in police behavior, they have been used to study the public safety value of order maintenance policing.

In NYC, the police slowdown was brought about, in large part, by the killing of two police officers by a man from Baltimore, MD who traveled to NYC with the intention of attacking police. The slowdown dramatically changed the intensity of policing in some areas of the city over a seven-week period and therefore it is instructive to ask what happened to major crimes during this period. S-O’K’s analysis suggests that order maintenance policing compromises public safety, a fact which, if true, would provide compelling evidence in favor of the argument to curtail this style of policing — and possibly police resources in general — in U.S. cities. Unfortunately this claim, which is new to the literature and which therefore has attracted considerable attention among scholars and in the popular media, is built upon dubious evidence.¹⁹

Our re-analysis shows that while NYPD’s use of criminal summonses and “quality of life” arrests was dramatically reduced in some communities during this time period, other communities sustained relatively little exposure to the slowdown. Communities with greater exposure and lesser exposure to the slowdown had similar levels of baseline crime and, prior to the slowdown, experienced extraordinarily similar crime trends. During the slowdown period, there is no evidence that highly exposed communities and less exposed communities experienced different levels of major crimes. We likewise note that S-O’K’s estimates are not robust to basic placebo checks which study the same time period in other years in which a slowdown was not present. The best available evidence thus suggests that there does not seem to have been a detectable change in serious crimes in either direction as a function of the 2014-15 slowdown. Such a finding comports with new and

¹⁹After writing their paper, the authors wrote an op-ed which appeared in the *Washington Post* in which they highlight their central claim — that reducing proactive policing led to a decrease in major crimes in NYC. See: <https://www.washingtonpost.com/news/monkey-cage/wp/2016/07/25/does-more-policing-lead-to-less-crime-or-just-more-racial-resentment/>.

related work on the public safety value of marginal low-level arrests by [Cho et al. \(2021\)](#) who leverage the natural experiment created by the killing of a police officer which creates a short-term reduction in police proactivity, [Bacher-Hicks and de la Campa \(2021\)](#) which studies variation in proactive policing arising from precinct-level management changes and [Chandrasekher \(2016\)](#) who studies the effects of an earlier police slowdown which occurred in NYC in the 1990s. Each of these studies reports little evidence of changes in major crimes in either direction.

Future research on police slowdowns has the potential to contribute to the growing body of work on the value of proactive policing. However, it is crucial to employ a comparison group — either unexposed or lesser exposed neighborhoods within the same city or a comparison panel of other cities which arguably had no exposure to the policing shock. Indeed the existence of a credibly selected comparison group has been noted as a core feature of research credibility by the Campbell Collaboration.²⁰ In addition to a comparison group, a credible differences-in-differences analysis also requires an analysis of parallel trends to provide some assurance that the natural experiment in question is not confounded by pre-treatment changes in the crime environment. Finally, placebo tests should be employed to provide assurances that the reported results are unlikely to be spurious. Absent these basic features (which are necessary but, by no means, sufficient to ensure credibility), findings from research on police slowdowns should be viewed with skepticism.

²⁰See: https://campbellcollaboration.org/media/k2/attachments/0366_IDCG_Vollmer_Protocol.pdf.

References

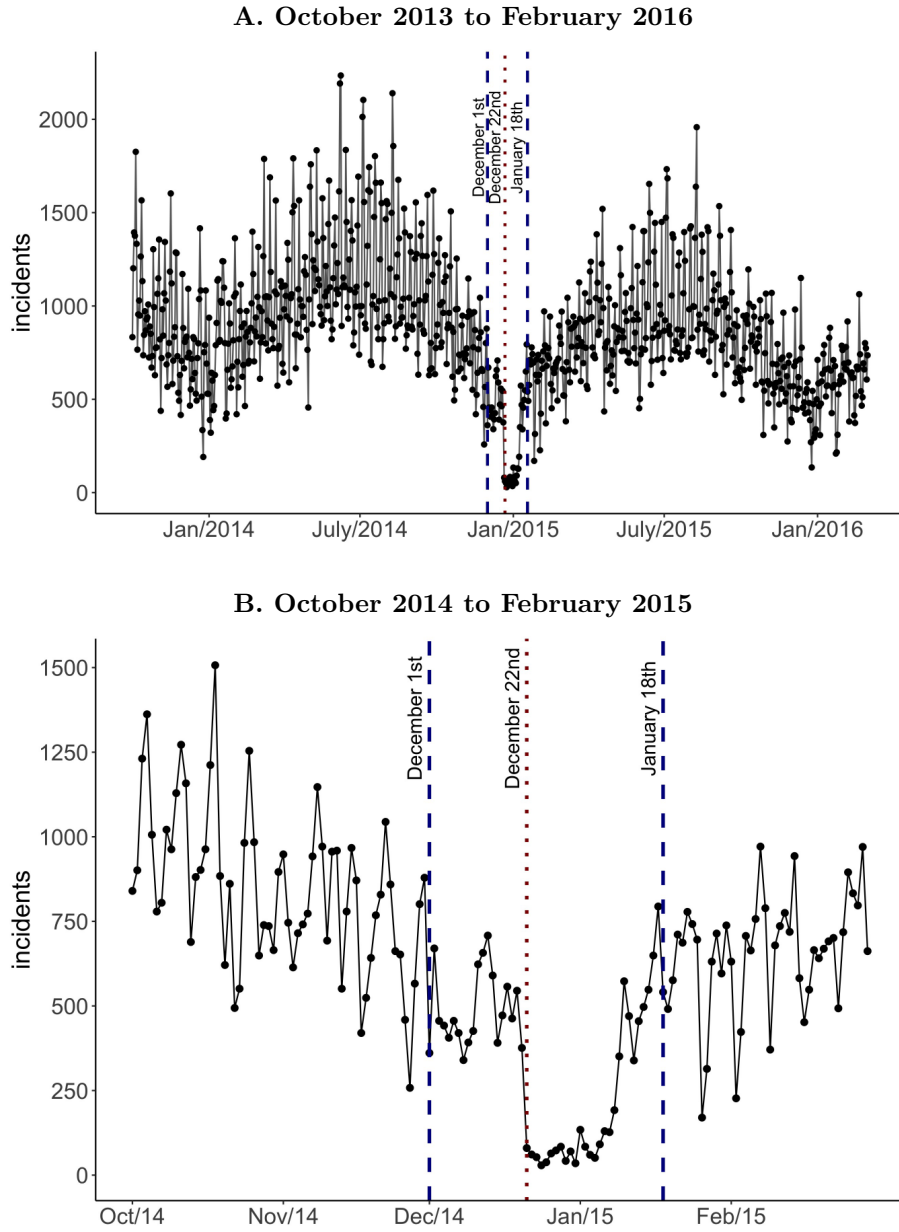
- Agan, A. Y., J. L. Doleac, and A. Harvey (2021). Misdemeanor prosecution. Technical report, National Bureau of Economic Research.
- Ang, D. (2020). The effects of police violence on inner-city students.
- Ba, B. A. and R. Rivera (2019). The effect of police oversight on crime and allegations of misconduct: Evidence from Chicago. *Working Paper* (19-42).
- Bacher-Hicks, A. and E. de la Campa (2021). The impact of New York City’s stop and frisk program on crime: The case of police commanders. Technical report, Working Paper.
- Bernasco, W. and H. Elffers (2010). Statistical analysis of spatial crime data. In *Handbook of Quantitative Criminology*, pp. 699–724. Springer.
- Braga, A. A. and B. J. Bond (2008). Policing crime and disorder hot spots: A randomized controlled trial. *Criminology* 46(3), 577–607.
- Braga, A. A., D. M. Hureau, and A. V. Papachristos (2014). Deterring gang-involved gun violence: Measuring the impact of Boston’s operation ceasefire on street gang behavior. *Journal of Quantitative Criminology* 30(1), 113–139.
- Braga, A. A., A. V. Papachristos, and D. M. Hureau (2014). The effects of hot spots policing on crime: An updated systematic review and meta-analysis. *Justice Quarterly* 31(4), 633–663.
- Branas, C. C., M. C. Kondo, S. M. Murphy, E. C. South, D. Polsky, and J. M. MacDonald (2016). Urban blight remediation as a cost-beneficial solution to firearm violence. *American Journal of Public Health* 106(12), 2158–2164.
- Branas, C. C., E. South, M. C. Kondo, B. C. Hohl, P. Bourgois, D. J. Wiebe, and J. M. MacDonald (2018). Citywide cluster randomized trial to restore blighted vacant land and its effects on violence, crime, and fear. *Proceedings of the National Academy of Sciences* 115(12), 2946–2951.
- Caetano, G. and V. Maheshri (2018). Identifying dynamic spillovers of crime with a causal approach to model selection. *Quantitative Economics* 9(1), 343–394.
- Chalfin, A., B. Hansen, J. Lerner, and L. Parker (2021). Reducing crime through environmental design: Evidence from a randomized experiment of street lighting in New York City. *Journal of Quantitative Criminology*, 1–31.
- Chalfin, A., B. Hansen, E. K. Weisburst, and M. C. Williams (2020). Police force size and civilian race. Technical report, National Bureau of Economic Research.
- Chalfin, A. and J. McCrary (2018). Are US cities underpoliced? Theory and evidence. *The Review of Economics and Statistics* 100(1), 167–186.
- Chandrasekher, A. C. (2016). The effect of police slowdowns on crime. *American Law and Economics Review* 18(2), 385–437.
- Cheng, C. and W. Long (2018). The effect of highly publicized police-related deaths on policing and crime: Evidence from large US cities. Technical report, Working Paper.

- Cho, S., F. Gonçalves, and E. Weisburst (2021). Do police make too many arrests? *University of California, Los Angeles Working Paper*.
- Corman, H. and N. Mocan (2005). Carrots, sticks, and broken windows. *The Journal of Law and Economics* 48(1), 235–266.
- Devi, T. and R. G. Fryer Jr (2020). Policing the police: The impact of “pattern-or-practice” investigations on crime. Technical report, National Bureau of Economic Research.
- Deza, M., J. C. Maclean, and K. T. Solomon (2020). Local access to mental healthcare and crime. Technical report, National Bureau of Economic Research.
- Dobbie, W., J. Goldin, and C. S. Yang (2018). The effects of pretrial detention on conviction, future crime, and employment: Evidence from randomly assigned judges. *The American Economic Review* 108(2), 201–40.
- Draca, M., S. Machin, and R. Witt (2011). Panic on the streets of london: Police, crime, and the july 2005 terror attacks. *The American Economic Review* 101(5), 2157–81.
- Edwards, F., H. Lee, and M. Esposito (2019). Risk of being killed by police use of force in the united states by age, race–ethnicity, and sex. *Proceedings of the National Academy of Sciences* 116(34), 16793–16798.
- Evans, W. N. and E. G. Owens (2007). Cops and crime. *Journal of Public Economics* 91(1-2), 181–201.
- Groff, E. R., J. H. Ratcliffe, C. P. Haberman, E. T. Sorg, N. M. Joyce, and R. B. Taylor (2015). Does what police do at hot spots matter? The Philadelphia policing tactics experiment. *Criminology* 53(1), 23–53.
- 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.
- Heller, S. B. (2014). Summer jobs reduce violence among disadvantaged youth. *Science* 346(6214), 1219–1223.
- Heller, S. B., A. K. Shah, J. Guryan, J. Ludwig, S. Mullainathan, and H. A. Pollack (2017). Thinking, fast and slow? some field experiments to reduce crime and dropout in chicago. *The Quarterly Journal of Economics* 132(1), 1–54.
- Howell, K. B. (2009). Broken lives from broken windows: The hidden costs of aggressive order-maintenance policing. *NYU Review of Law & Social Change* 33, 271.
- Jácome, E. (2020). Mental health and criminal involvement: Evidence from losing medicaid eligibility.
- Kaplan, J. and A. Chalfin (2019). More cops, fewer prisoners? *Criminology & Public Policy* 18(1), 171–200.
- Kelling, G. L. and W. H. Sousa (2001). *Do police matter?: An analysis of the impact of new york city’s police reforms*. CCI Center for Civic Innovation at the Manhattan Institute.

- Kessler, J. B., S. Tahamont, A. M. Gelber, and A. Isen (2021). The effects of youth employment on crime: Evidence from new york city lotteries. Technical report, National Bureau of Economic Research.
- Klick, J. and A. Tabarrok (2005). Using terror alert levels to estimate the effect of police on crime. *The Journal of Law and Economics* 48(1), 267–279.
- Leslie, E. and N. G. Pope (2017). The unintended impact of pretrial detention on case outcomes: Evidence from new york city arraignments. *The Journal of Law and Economics* 60(3), 529–557.
- Lind, D. (2015). The nypd “slowdown” that’s cut arrests in new york by half, explained. *Vox*, January 6, 2015.
- Lovett, N. and Y. Xue (2021). Rare homicides, criminal behavior, and the returns to police labor. *Working 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., V. Nguyen, S. T. Jensen, and C. C. Branas (2021). Reducing crime by remediating vacant lots: the moderating effect of nearby land uses. *Journal of Experimental Criminology*, 1–26.
- 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.
- Machin, S. and O. Marie (2011). Crime and police resources: The street crime initiative. *Journal of the European Economic Association* 9(4), 678–701.
- Maguire, E. R., J. Nix, and B. A. Campbell (2017). A war on cops? the effects of ferguson on the number of us police officers murdered in the line of duty. *Justice Quarterly* 34(5), 739–758.
- Mas, A. (2006). Pay, reference points, and police performance. *The Quarterly Journal of Economics* 121(3), 783–821.
- Mastrobuoni, G. (2019). Police disruption and performance: Evidence from recurrent redeployments within a city. *Journal of Public Economics* 176, 18–31.
- Mello, S. (2019). More cops, less crime. *Journal of Public Economics* 172, 174–200.
- Messner, S. F., S. Galea, K. J. Tardiff, M. Tracy, A. Bucciarelli, T. M. Piper, V. Frye, and D. Vlahov (2007). Policing, drugs, and the homicide decline in new york city in the 1990s. *Criminology* 45(2), 385–414.
- Nix, J., S. E. Wolfe, and B. A. Campbell (2018). Command-level police officers’ perceptions of the “war on cops” and de-policing. *Justice Quarterly* 35(1), 33–54.
- Ratcliffe, J. (2010). Crime mapping: spatial and temporal challenges. In *Handbook of Quantitative Criminology*, pp. 5–24. Springer.
- Rosenfeld, R. (2016). Documenting and explaining the 2015 homicide rise: Research directions. *Washington, DC*.


- Sharkey, P., G. Torrats-Espinosa, and D. Takyar (2017). Community and the crime decline: The causal effect of local nonprofits on violent crime. *The American Sociological Review* 82(6), 1214–1240.
- 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.
- Shi, L. (2008). Does oversight reduce policing? evidence from the Cincinnati police department after the April 2001 riot. *Journal of Public Economics*.
- Shjarback, J. A. and J. Nix (2020). Considering violence against police by citizen race/ethnicity to contextualize representation in officer-involved shootings. *Journal of Criminal Justice* 66, 101653.
- Shjarback, J. A., D. C. Pyrooz, S. E. Wolfe, and S. H. Decker (2017). De-policing and crime in the wake of Ferguson: Racialized changes in the quantity and quality of policing among Missouri police departments. *Journal of Criminal Justice* 50, 42–52.
- 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, D. (2016). Does hot spots policing inevitably lead to unfair and abusive police practices, or can we maximize both fairness and effectiveness in the new proactive policing. *University of Chicago Legal Forum*, 661.
- Weisburd, S. (2021). Police presence, rapid response rates, and crime prevention. *The Review of Economics and Statistics* 103(2), 280–293.
- Weisburst, E. K. (2019). Police use of force as an extension of arrests: Examining disparities across civilian and officer race. In *The American Economic Review*, Volume 109, pp. 152–56.
- Weitzer, R., S. A. Tuch, and W. G. Skogan (2008). Police–community relations in a majority-black city. *Journal of Research in Crime and Delinquency* 45(4), 398–428.
- Wilson, J. Q. and G. L. Kelling (1982). Broken windows. *Atlantic Monthly* 249(3), 29–38.
- Wolfe, S. E. and J. Nix (2016). The alleged “Ferguson effect” and police willingness to engage in community partnership. *Law and Human Behavior* 40(1), 1.

Figure 1: Daily city-wide criminal summonses



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 summonses.


Figure 2: The 2014-2015 NYPD slowdown and its effect on major crimes



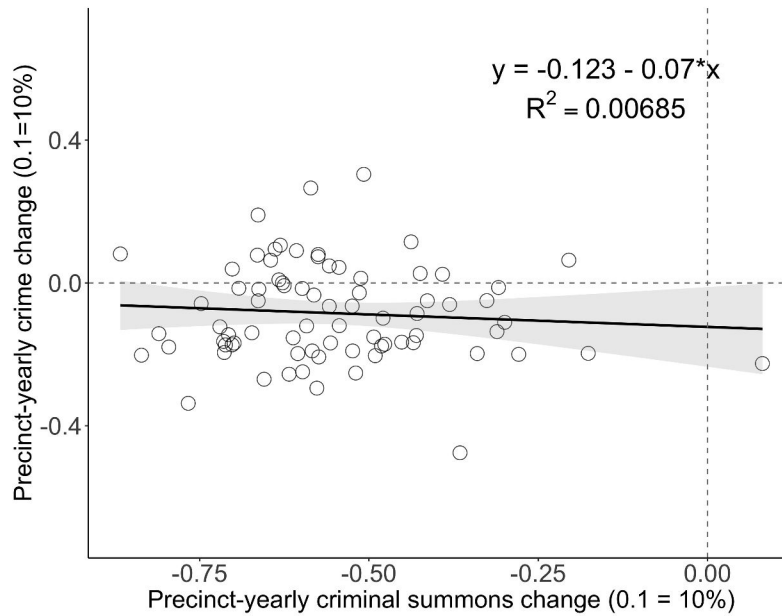
Note: Each row presents estimates from SO's differences-in-differences estimator. The analysis identifies the effect of the slowdown by differencing the year-over-year change in the number of major crimes experienced during the December 1st-January 18th slowdown period from the year-over-year change in major crimes experienced during the remainder of the year. For each year t , year $t-1$ pair, we replicate SO's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as SO 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 SO'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 3: 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



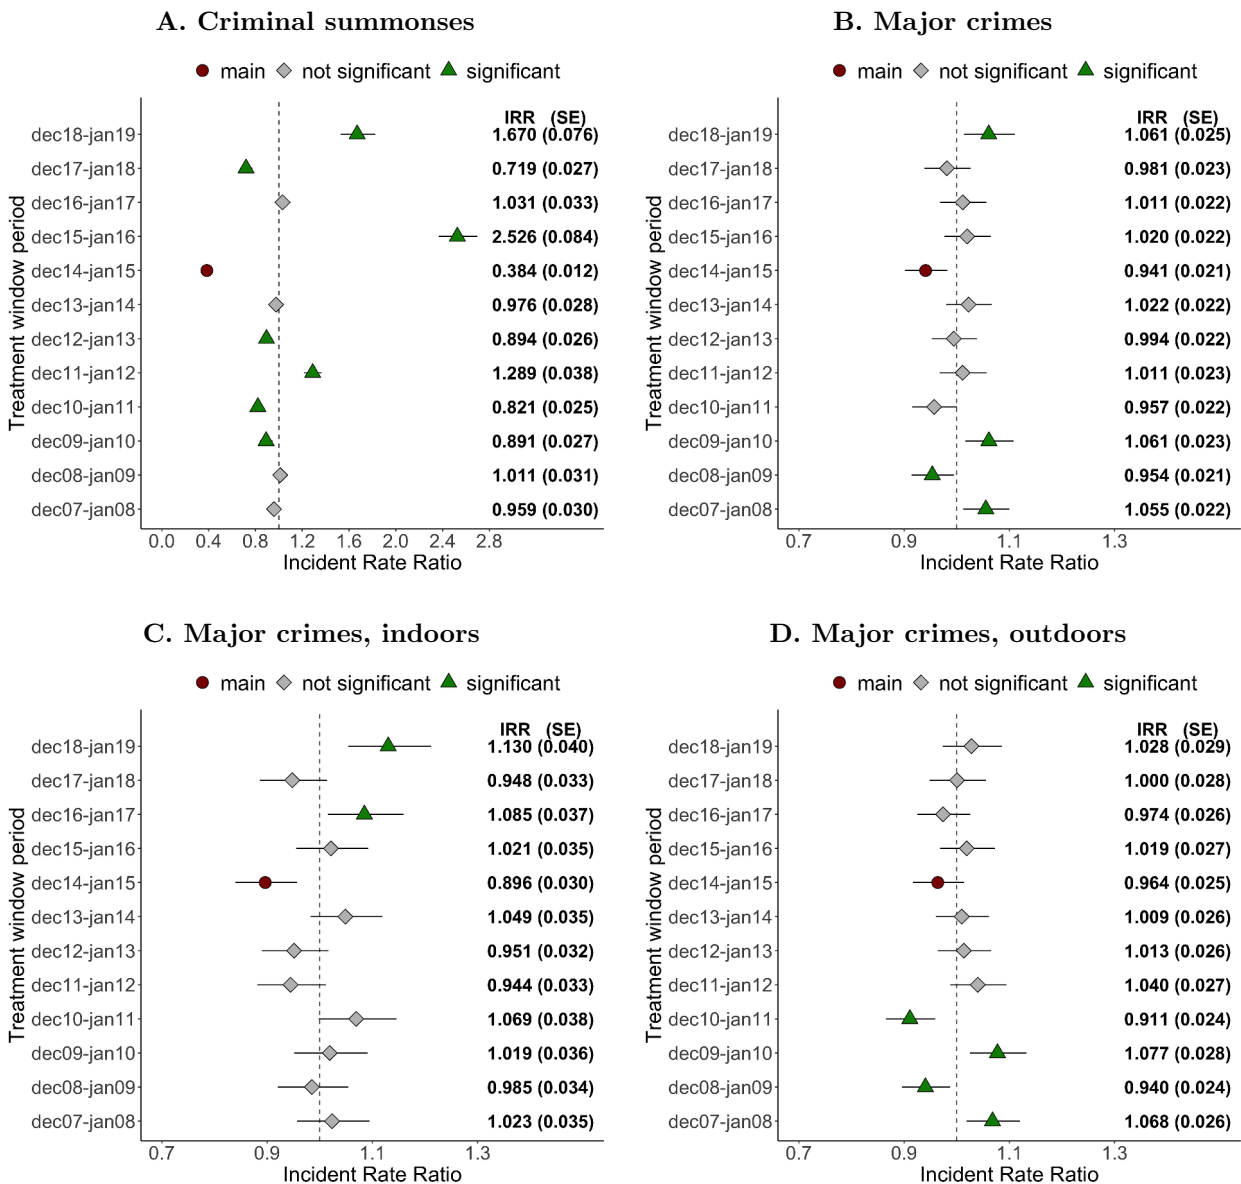
Note: Panel A presents a precinct-level heat map of NYC and documents enormous variation in the intensity of the slowdown across NYC communities. For each of NYC's 76 police precincts (excluding the Central Park precinct), the intensity of the color in the heat map corresponds with the year-over-year change in summonses issued during the December 1st-January 18th period. Panel B exploits the precinct-level variation and plots each precinct's percentage-level 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.

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



Note: Each row presents estimates from SO'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 SO's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as SO 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 5: Effects of NYPD slowdown on public safety using alternative September 22nd, 2014 - January 18th, 2015 treatment period

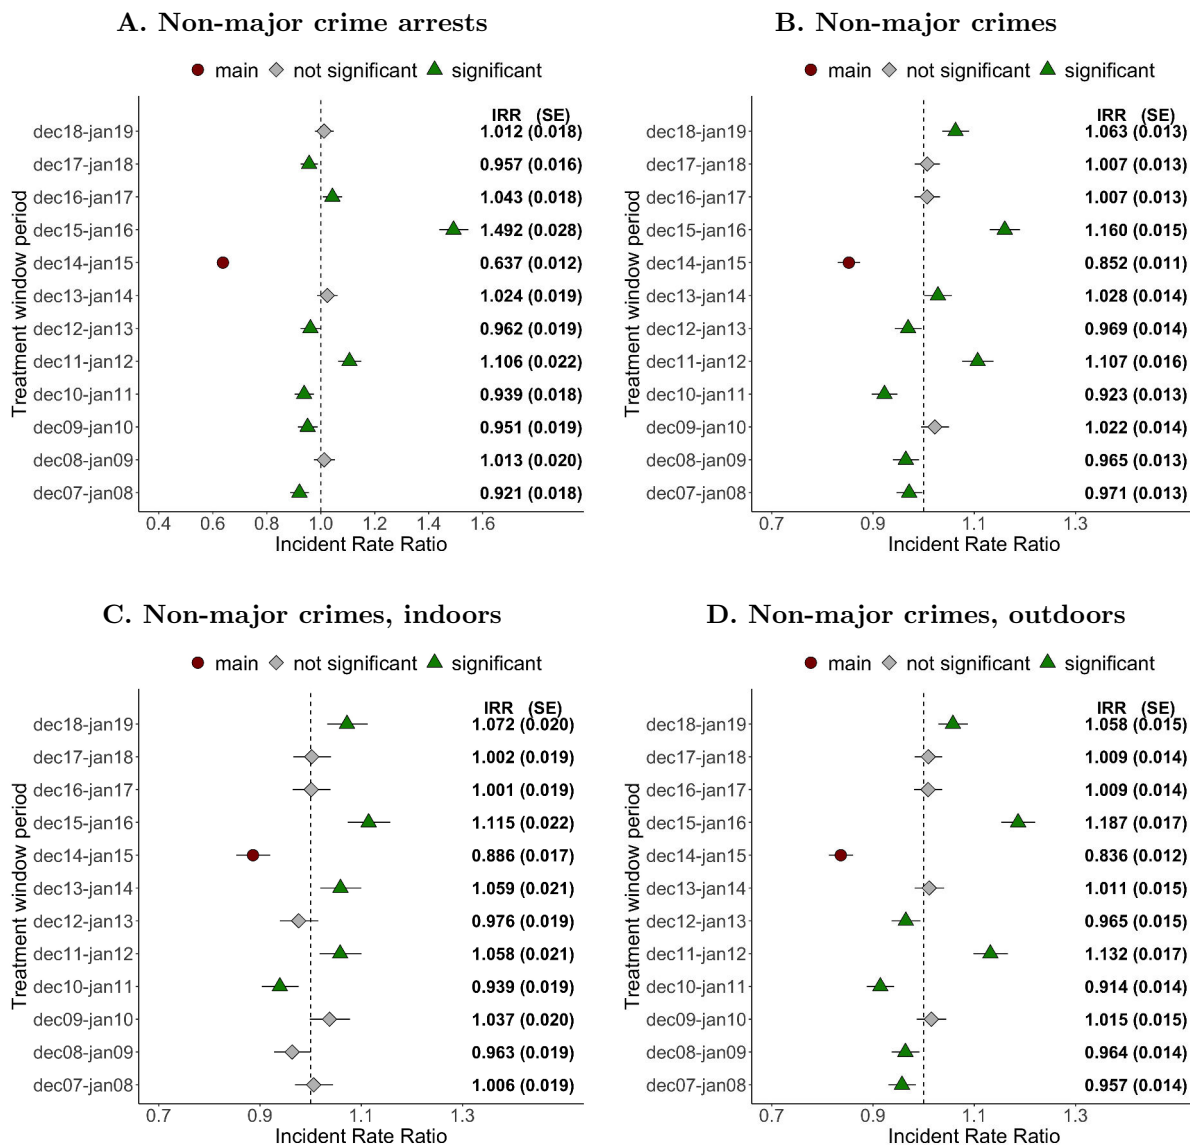


Note: Each row presents estimates from SO's differences-in-differences estimator. For each year t , year $t-1$ pair, we replicate SO's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as SO note, therefore do not 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.

ONLINE APPENDIX

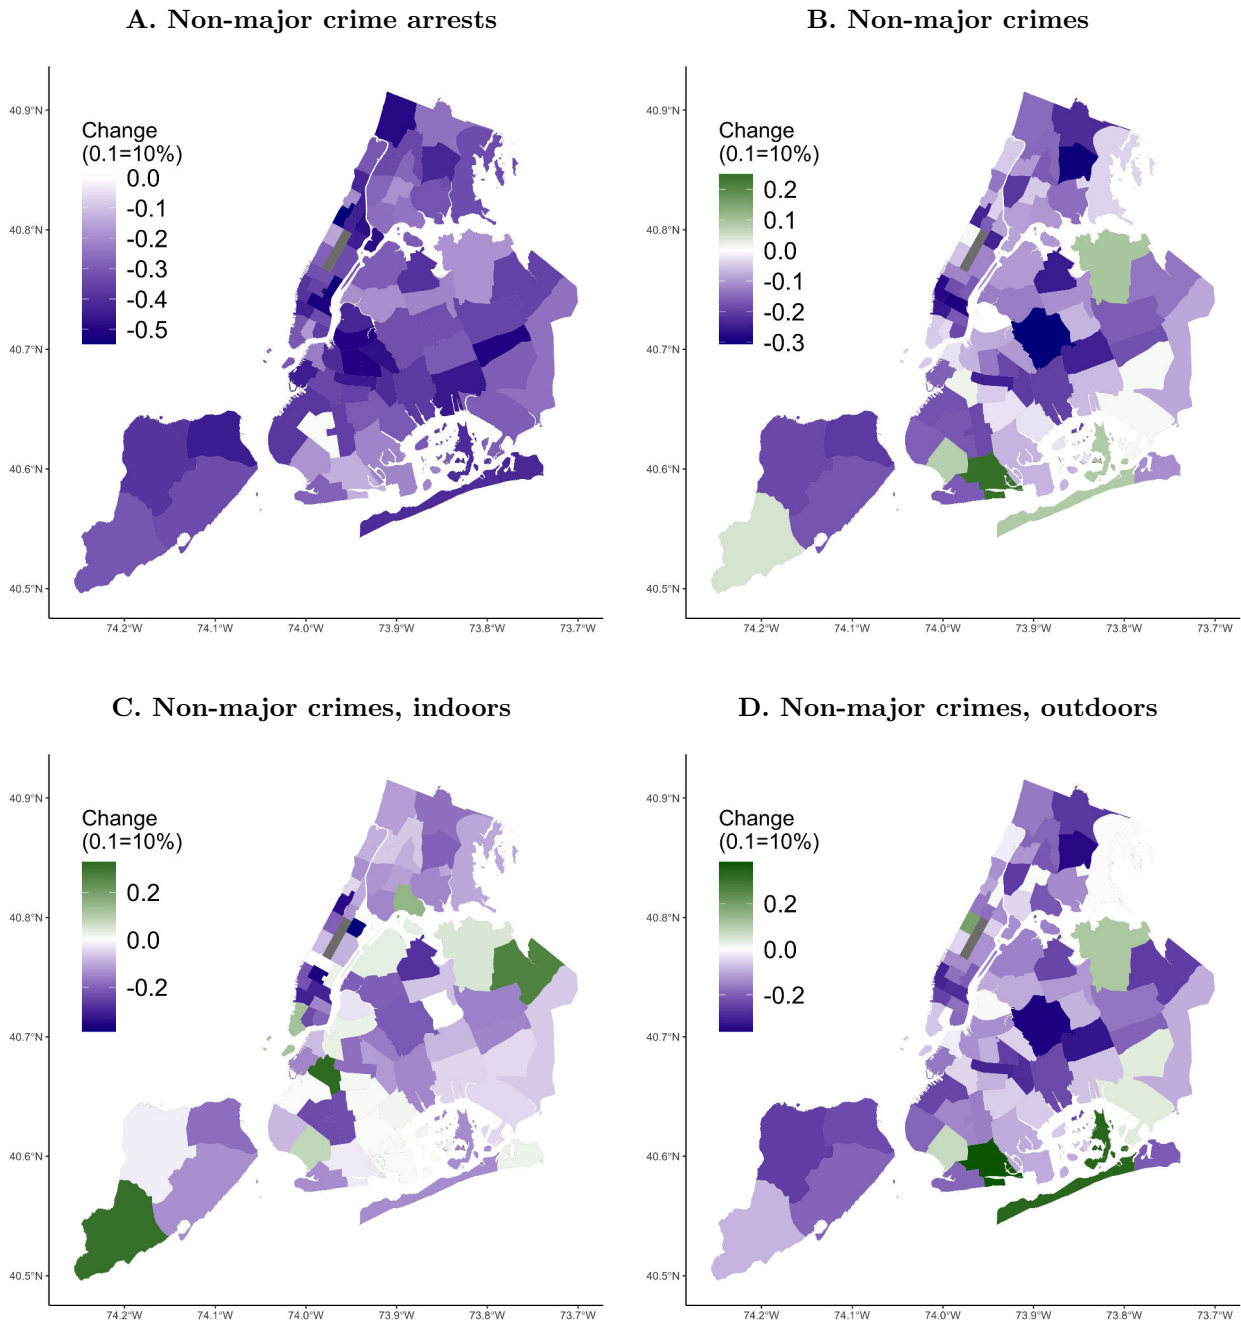
A Appendix: Tables and Figures

Figure A.1: Effects of slowdown on public safety



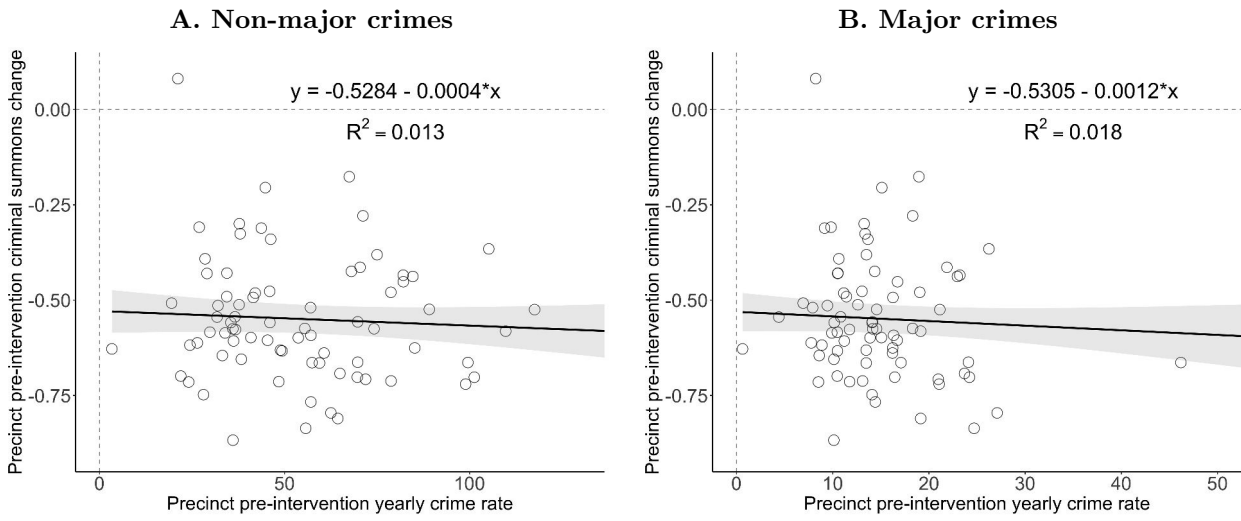
Note: Each row presents estimates from SO's differences-in-differences estimator. For each year t , year $t-1$ pair, we replicate SO's models, collapsing crime to the precinct-day level, excluding demographic control variables, which do not vary at a daily level and which, as SO note, therefore do not 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 A.2: Annual change in non-major crimes and arrests by police precinct




Note: The maps 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 A.3: Police slowdown and pre-intervention yearly crime rate per 1,000 people



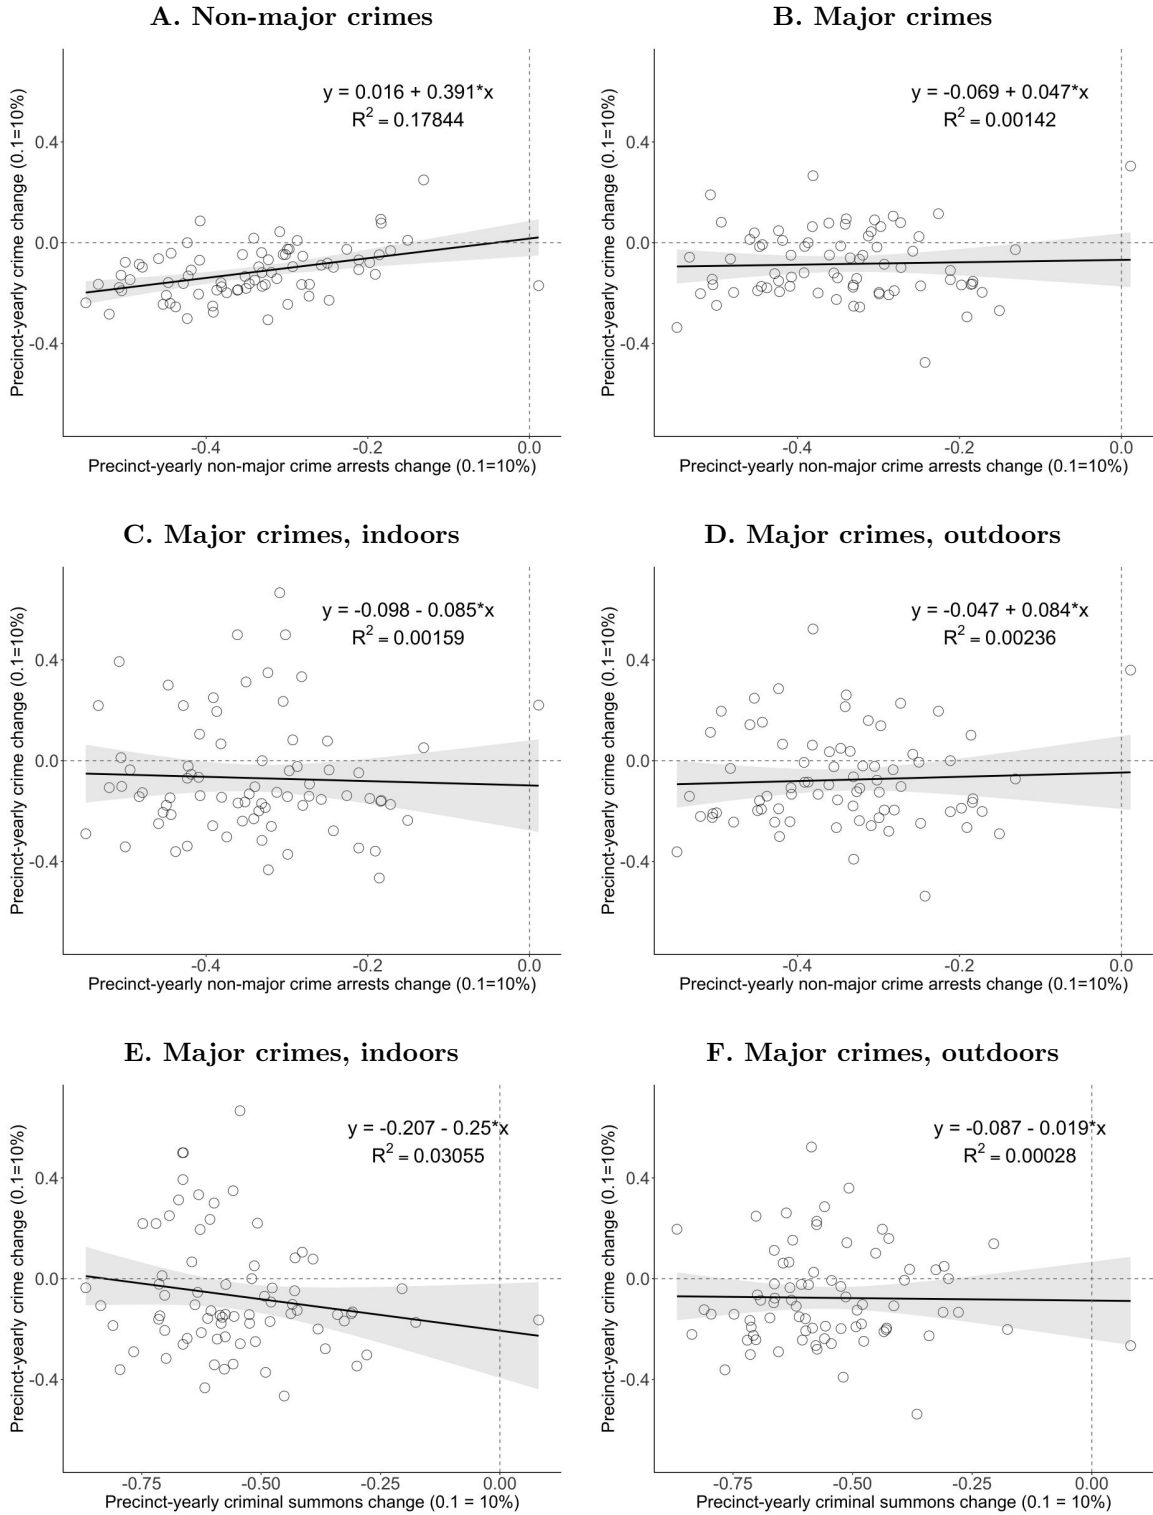
Note: The scatterplots exhibit the annual precinct-level change in criminal summonses (0.1 = 10%) between December 1st, 2014 to January 18th, 2015 and December 1st, 2013 to January 18th, 2014 against the mean pre-intervention (2006-2013) annual crimes per 1,000 persons in each police precinct. Major 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 Midtown South Precinct 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 summonses change), but it is considered in the regression line.

Figure A.4: Precinct-level scatterplot of the change in major crimes rates against the change in criminal summonses rates per thousand persons




Note: The figure exploits the precinct-level variation and plots each precinct's change in major crime rates per thousand persons between the December 1st, 2013 - January 18th, 2014 and the December 1st, 2014 - January 18th, 2015 periods (y -axis) against the change in criminal summonses rates (x -axis). Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft. A best-fit line is drawn through the data including (blue dashed line) and excluding (red solid line) the outlier of Midtown South Precinct. The regression equation, R^2 , correlation coefficient (r), and its p value (p) are shown for both best-fit lines.

Figure A.5: 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.

Figure A.6: Event study estimates on major crimes between high and low exposure precincts



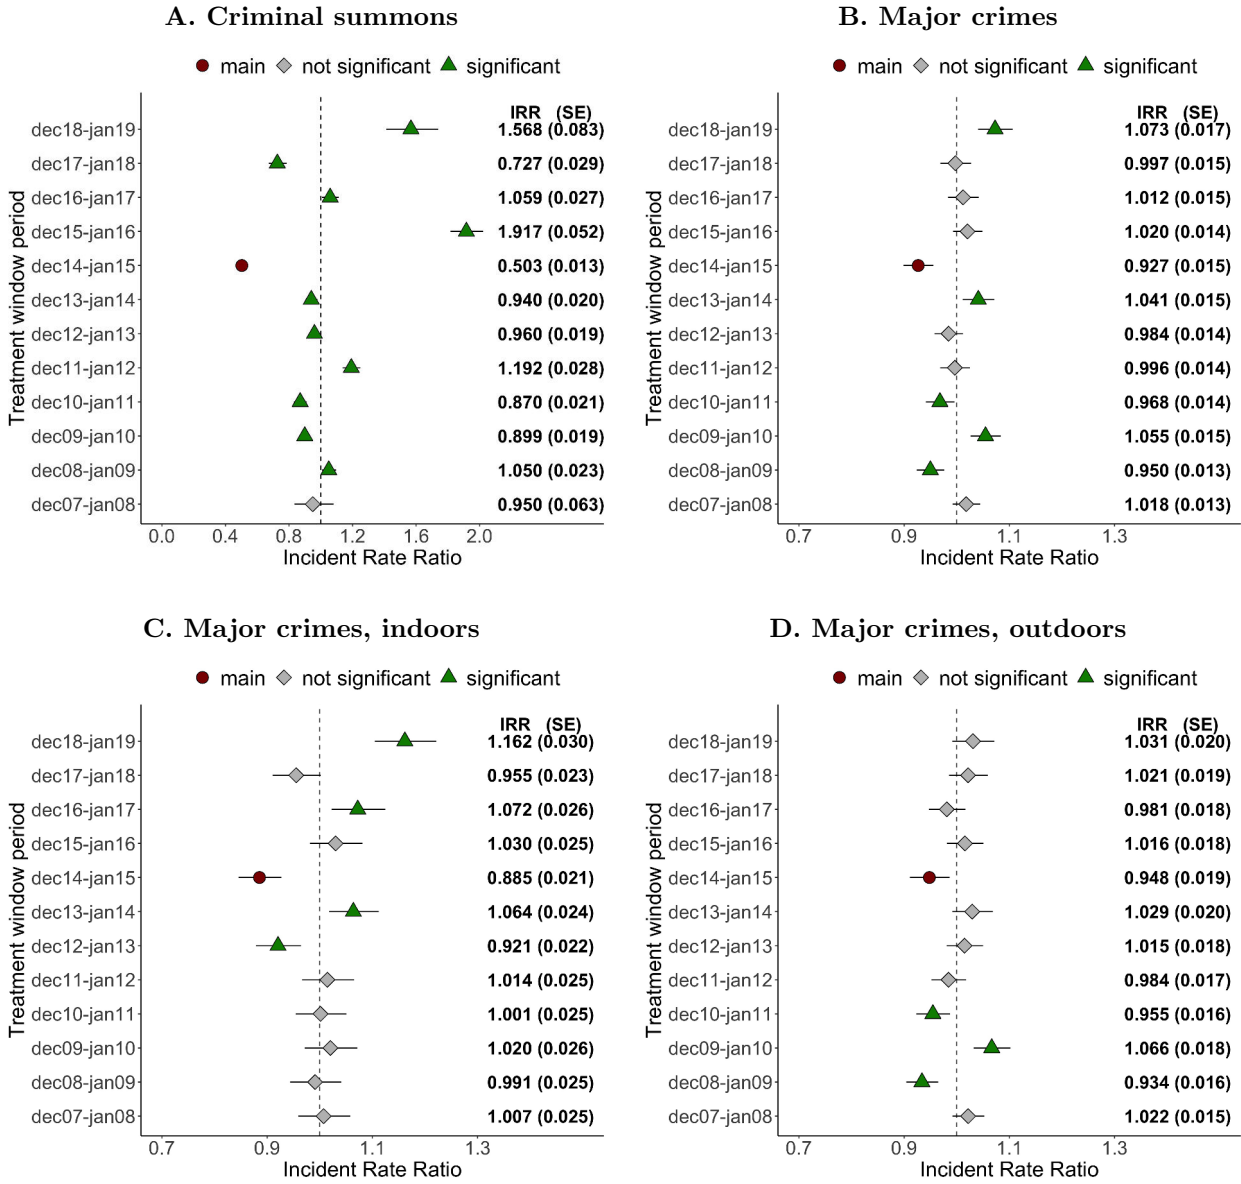
Note: Event study design estimates following: $y_{it} = \alpha + \beta_1 S_t + \beta_2 H_i + \beta_3 S_t H_i + \sum_{\tau=-q}^m \delta_\tau S_t 1[t - t_0 = \tau] + \sum_{\tau=-q}^m \gamma_\tau S_t H_i 1[t - t_0 = \tau] + X'_t \beta_6 + e_{it}$, where S_t is an indicator for the “series” — that is, whether a given year was treated by the slowdown—equal to one if a given day occurred during January 19th, 2014 to January 18th, 2015 period and zero if otherwise. H_i is an indicator variable of whether the precinct’s criminal summons percent change between December 1st, 2014 to January 18th, 2015 and December 1st, 2013 to January 18th, 2014 is above the median (57.6%), zero if otherwise. $1[t - t_0 = \tau]$ is an indicator variable, where τ is the one-week bin relative to the beginning of the treatment window (t_0 is week 49 that includes December 1st), specifically $\tau = \{-45, -44, \dots, -2, 0, 1, 2, \dots, 6\}$. The reference period is one week before the treatment window started. The figure plots γ_τ , estimating any differential trends between high and low exposure precincts. The Negative binomial regression clusters the standard errors at the precinct level and includes data from 01/19/2010 to 01/18/2015 (restricting the sample to 01/19/2013 to 01/18/2015 to following the year-pair analysis provides the same results).

Figure A.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 A.8: Effects of slowdown on public safety at the census tract level



Note: Each row presents estimates from SO's differences-in-differences estimator. For each year t , year $t-1$ pair, we collapsed crime to the census tract-day level, excluding demographic control variables, which do not vary at a daily level and which, as SO 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 census tract 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 at the census tract level and 95% confidence intervals from a series of negative binomial regressions. Major crimes include the seven-part I Uniform Crime Reporting categories: murder, rape, robbery, aggravated assault, burglary, theft, and motor vehicle theft. Indoor crimes comprise offenses committed inside a residential premise (apartment, house, or public housing); outdoor crimes include all the other criminal offenses.

Figure A.9: Census tract level scatterplot of the change in crimes against criminal summons



Note: The figure exploits the census tract-level variation and plots each tract'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). 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. A best-fit line is drawn through the data. Both panels exclude outliers in the visualization but they are included in the regression line. The sample sizes are different between panels because they exclude non-finite values (i.e., division over zero values).